Stochastic Choice and Noisy Beliefs in Games: an Experiment*

Evan Friedman[†] and Jeremy Ward

December 31, 2020

PRELIMINARY

Abstract

We conduct an experiment in which we elicit subjects' beliefs over opponents' behavior multiple times for a given game without feedback. We find that the large majority of individual subjects have stochastic belief reports, which we argue cannot be explained by learning or measurment error. Using both actions and beliefs data, we test directly the assumptions, or axioms, underlying equilibrium models with "noise in actions" (quantal response equilibrium) and "noise in beliefs" (noisy belief equilibrium). We find that, while both types of noise are important in explaining observed behaviors, there are systematic violations of the axioms. We discuss possible explanations and some implications for modelling stochastic choice in games.

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: ef20396@essex.ac.uk

1 Introduction

A large literature has documented a pattern of stochastic choice, or "noisy actions," in individual decision making. In games, in which payoffs depend on beliefs over opponents' behavior, another potentially important source of stochasticity is in the beliefs themselves. However, little is known about the empirical relevance of "noisy beliefs."

Recognizing that there may be both "noise in actions" and "noise in beliefs" in games, we run a laboratory experiment designed to answer three related questions. First, are beliefs noisy? Second, what is the relative importance of the two sources of noise for explaining data? Third, how valid are the assumptions underlying models of stochastic choice that have been applied to games?

To guide the experimental design, we formulate as a benchmark model a generalization of Nash equilibrium 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". Importantly, the model is parameter-free and will not be trivially rejected if there is even one failure of best response or one instance of incorrect beliefs.

The model is a hybrid model, borrowing from existing models of stochastic choice. It is defined by an *action-map,* which determines the mixed actions taken by players given their beliefs, and a *belief-map*, which determines the distribution of players' beliefs as a function of the opponents' mixed actions.[1] The action-map is restricted to satisfy the axioms of regular quantal response equilibrium (QRE) (McKelvey and Palfrey [1995] and 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 noisy belief equilibrium (NBE) (Friedman [2020b]), 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*).

Motivated by the idea of hybrid model, the goal of our experiment is to make observable the empirical action- and belief-maps. Noting that both mappings depend on both actions and beliefs, we naturally elicit beliefs and have subjects take actions. By doing this for a family of games with systematically varied payoffs, we observe multiple points on, i.e. "trace

[1] To be clear, for player i forming beliefs about opponent j, the belief-map maps from the simplex representing the space of j's mixed actions to a belief-distribution on the simplex. Hence, each draw from this belief-distribution is a belief over j's mixed action.

out", the empirical action- and belief-maps. With these empirical objects, we answer our three questions.

Central to our design are the 2×2 asymmetric matching pennies games whose payoffs are in Table 1. Indexed by different values of player 1's payoff parameter $X > 0$, these X-*games* have unique mixed strategy Nash equilibria. By varying X, the hybrid model predicts systematic variation in actions and beliefs for both players so that we may trace out the empirical action- and belief-maps.

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

Table 1: *Game X*. Player 1's payoff parameter $X > 0$ controls the asymmetry of payoffs.

Prior to previewing our results, we make two conceptual points that weigh heavily in our design and analysis. First, by "noise in beliefs," we mean within-subject variation in beliefs that is not due to predictable sources of variation. For this reason, unlike existing experiments, we have subjects state beliefs multiple times for each game without feedback, and so there is no scope for learning. The selection of games too is such that we would not expect the phenomenon of no-feedback learning.[2] Second, we hold the view that stated beliefs are simply measures of the unobservable true beliefs that subjects hold in their minds and guide their actions. As such, stated beliefs may be subject to random measurement error. While our main analysis takes stated beliefs as equal to the true beliefs, we argue that our results are robust to measurement error under the assumption that a stated belief is an unbiased signal of the true belief.

In answering the questions that motivated our study, we make three contributions. First, we establish that beliefs are, in fact, noisy. Second, we show, using a counterfactual exercise, that belief-noise is similarly important as action-noise for explaining data. Third, we study the validity of assumptions underlying models of stochastic choice by way of statistically testing the axioms of the hybrid model–and therefore of QRE and NBE as well. We find systematic violations for which we offer potential explanations.

[2] Weber [2003] documents, using relatively complicated dominance solvable games, that subjects' behavior drifts toward the Nash equilibrium, even without feedback. The X-games, on the other hand, are very simple and fully mixed.

The conclusion that beliefs are noisy is based on subject-level variation in stated beliefs for a given game. On average, each subject's beliefs have a range of 23 percentage points over five repetitions of a given game. These results cannot be driven by learning since we show that subjects typically have no trend in stated beliefs across appearances of a given game, and the reported ranges of beliefs are little affected by linearly detrending each subject's beliefs. We also argue that this noise in stated beliefs cannot be due entirely to random measurement error since within-subject variations in stated beliefs are strongly predictive of the actions subjects take. This implies that a high stated belief signals a high true belief, and thus noisy stated beliefs reflect noisy true beliefs.

To conclude that belief-noise is similarly important as action-noise for explaining the data, we conduct a counterfactual exercise. Specifically, we construct, for each subject, two counterfactual action frequencies that result from "turning off" just one source of noise; and we say that a source of noise is important if turning it off leads to large prediction errors relative to the observed frequency. To turn off action-noise, we calculate the action frequencies that we would have observed if subjects had best responded to stated beliefs. To turn off belief-noise, we first fit a random random utility model (RUM) for each subject. We then calculate the action frequency that would result from the best-fit RUMs if belief statements were replaced with the central tendency for each subject-game. We find that the two counterfactuals perform about equally well, suggesting that the two sources of noise are similarly important.

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 a subject's action (through variations in her stated beliefs) is associated with an increased probability that 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 stated beliefs to which subjects fail to best respond more often than not. These regions of violation, which are quite large and typically include the mode of stated beliefs, are consistent with concavity in the utilility function. Since matrix payoffs are in probability points in our experiment, this cannot be interpreted as risk aversion, so we discuss potential alternatives.

Consistent with *belief-responsiveness*, we find that player i's stated 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 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 documented (e.g. Huck and Weizsacker [2002]), we believe this asymmetric pattern of bias is novel. As a partial explanation for the pattern of bias, we provide suggestive evidence that player 1-subjects are causally induced, by merit of facing asymmetric payoffs, to approach the game with greater strategic sophistication (in the level k sense[3]) than player 2-subjects.

The paper is organized as follows. The remainder of this section discusses related literature. Section 2 presents the theory, Section 3 gives the experimental design, and Section 4 provides an overview of the data. Section 5 provides evidence that beliefs are noisy, and Section 6 presents the results from testing the axioms. Section 7 studies the relative importance of action- and belief-noise for explaining the data. Section 8 argues that our results, including the fact of noisy beliefs, are robust to measurment error, and Section 9 concludes.

Related literature. This paper contributes to the theory and empirical study of equilibrium models with stochastic elements. We directly test the assumptions, or axioms, underlying regular quantal response equilibrium (QRE) (Goeree et al. [2005] and McKelvey and Palfrey [1995]) and noisy belief equilibrium (NBE) (Friedman [2020b]), but these are closely related to other models found in the literature.[4] We are novel in studying a hybrid model that allows for both noisy actions and noisy beliefs. Because it is a non-parametric, axiomatic model that makes set-predictions, 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. [2018], Goeree and Louis [2018], and Friedman [2020a])). While the focus of our paper is testing

[3]See, for example, Nagel [1995] and Stahl and Wilson [1995]. For a review of the level k literature, see Crawford et al. [2013].

[4]Models with noisy actions include those of Chen et al. [1997] and Stahl [1990], and models with noisy beliefs include those of Friedman and Mezzetti [2005], Rubinstein and Osborne [2003], Goeree and Holt [2002], and Goncalves [2020]. NBE borrows the idea of the belief-map from the concept of random belief equilibrium (Friedman and Mezzetti [2005]), but imposes behavioral axioms to derive testable restrictions.

axioms using actions and beliefs data jointly, we do show that the hybrid model–despite being a generalization of regular QRE–is falsifiable using standard actions data.[5] Very few studies explictly test predictions using both actions and beliefs data jointly, though there are exceptions (e.g. Goeree and Louis [2018]).

By studying noisy beliefs as a driver of stochastic choice, we contribute to the experimental literature on the nature and determinants of stochastic choice (e.g. Tversky [1969], Agranov and Ortoleva [2017], and Agranov et al. [2020]). More directly related, Wolff and Bauer [2018] argue that strategic uncertainty is related to noise in beliefs, and Mauersberger [2018] models beliefs as draws from a bayesian posterior.

Our central tool is belief elicitation, so we engage with the growing literature on belief elicitation (see Schotter and Trevino [2014] and Schlag et al. [2015] for review articles). 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. This distinguishes us from experiments that elicit beliefs once for each game in a set of games spanning very different features (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]). Whereas one focus of the literature has been on documenting rates of best responses,[6] we study how rates of best reponse vary across every neighborhood of stated beliefs. We believe we are also unique in showing that, within-subject, actions vary with changes in stated beliefs.

2 Theory as the basis for experimental design

To guide the experimental design, we consider a benchmark model that replaces the deterministic assumptions of "best response" and "correct beliefs" with stochastic generalizations. Importantly, the model is parameter free and will not be trivially rejected if any noise is observed in the data.

The model is a hybrid model, defined by an action-map satisfying the axioms of regular QRE (Goeree et al. [2005]) and a belief-map satisfying the axioms of NBE (**?**). For clarity, we

[5]Since Haile et al. [2008] showed that *structural* QRE can rationalize the data from any one game without strong restrictions on the error distributions, people are often concerned that QRE models may be non-falsifiable. Goeree et al. [2005], however, show that *regular* QRE is falsifiable.

[6]An exception is Hyndman et al. [2013] who show using 3×3 games that beliefs in the corner of the simplex are best responded to more often.

review QRE and NBE, and provide results for the X-games, before introducing the hybrid. Our main takeaway: as long as there is noise in actions, beliefs, or both, the hybrid model predicts that the games played in the experiment will give rise to systematic variation in actions and beliefs.

Anticipating the experiment, we present the case of 2×2 games in which there are two players with two actions each, but QRE and NBE–and therefore their hybrid–generalize to all finite, normal form games. A game is defined by $\Gamma^{2\times 2} = \{N, A, u\}$ 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 σ_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.[7]

[7] Requiring that $v_{ik} = v_{il} \implies Q_{ik}(v_i) = \frac{1}{2}$ in (A4) is unnecessary since it is implied by $v_{ik} > v_{il} \implies$

QRE is the special case of the hybrid model defined by an action-map satisfying (A1)-(A4) and a belief-map that is the identity map imposing "correct beliefs".

Definition 1. Fix $\{\Gamma^{2\times 2}, Q\}$. A QRE is any mixed action profile σ such that $\sigma = Q(\bar{u}(\sigma))$.

2.2 Belief-map

As in NBE, 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 [2020b], 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 $\bar{\sigma}_j \in (0, 1)$.

(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.[8] (B4) imposes that belief distributions are correct *on median*. Both median- and mean-unbiasedness can be microfounded via a model of sampling (Friedman [2020b]). 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.

$Q_{ik}(v_i) > Q_{il}(v_i)$ and (A1). We added this condition to (A4) in order to have a clean division between technical and behavioral axioms.

[8] 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.

NBE is the special case of the hybrid model defined by a belief-map satisfying (B1)-(B4) and the perfect action-map of "best response".

Definition 2. Fix $\{\Gamma^{2\times 2}, \sigma^*\}$. An NBE is any 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.

2.3 X-games

We specialize theory for the family of X-games whose payoffs are in Table 1. The results derived here serve as our justification for using the X-games as the basis for the experiment.

The X-games have unique, mixed strategy 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.

Within any one X-game, QRE and NBE are unique for fixed primitives. Flexibility in choosing the primitives, however, gives rise to set-valued predictions. Friedman [2020b] shows that these sets of attainable QRE and NBE exactly coincide within any one X-game.[9] Assuming further that the primitives are held fixed as X-varies–an assumption we maintain throughout–gives rise to comparative static predictions.

Extending the results of Friedman [2020b], Proposition 1 characterizes QRE and NBE for any finite family of X-games. Again, the models' predictions coincide.

Proposition 1. *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 10.1. □

Figure 1 illustrates the proposition, showing the sets of attainable QRE and NBE as functions of σ_L^{NE}. 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

[9] QRE-NBE equivalence holds for any 2×2 game with a unique mixed strategy NE.

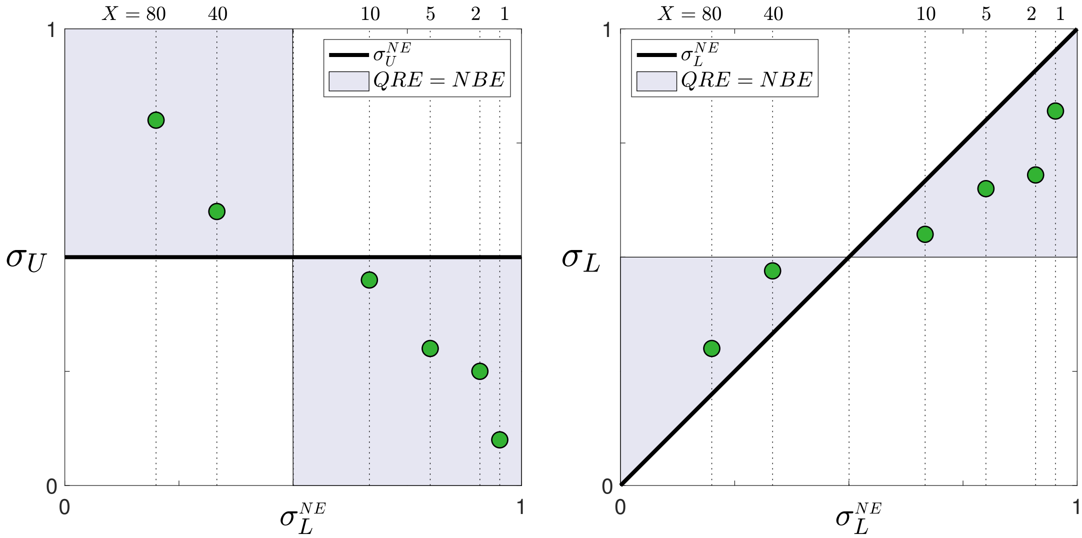

Figure 1: *QRE and NBE in the X-games as a function of* σ_L^{NE}. This figure plots a hypothetical dataset $\{\hat{\sigma}_U^X, \hat{\sigma}_L^X\}_X$ of mixed actions for the set of X-games with $X \in \{80, 40, 10, 5, 2, 1\}$. The left panel plots $\hat{\sigma}_U^X$ and the right panel plots $\hat{\sigma}_L^X$ (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.

$\hat{\sigma}_U^X$ and the right panel plots $\hat{\sigma}_L^X$. 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. Thus, varying X will cause systematic variation in mixed actions, and thus belief distributions also, for both players. This comparative static is essential in order to "trace out" the empirical action- and belief-maps.

2.4 Hybrid model

The hybrid model is defined in the obvious way as a generalization of both QRE and NBE. It is defined by an action-map satisfying (A1)-(A4) and a belief map satisfying (B1)-(B4). We provide the formal definition and characterization results in Appendix 10.2. We briefly mention a few key points below.

First, in terms of mixed action profiles, the set of hybrid equilibria is much larger than that of either QRE or NBE.[10] Importantly, however, the hybrid is rejected if and only if any of the axioms are rejected, which is to say that the model is much easier to falsify by

[10]For example, when $X = 80$, the Lebesgue measure of QRE and NBE outcomes is 15%, whereas the measure of hybrid equilibria is 51.25%, a more than 3-fold increase.

augmenting standard actions data with beliefs data, as we do in our experiment.

Second, by extending Proposition 1 to the hybrid model, the comparative static result in X carries over (i.e. parts *(iii)* and *(iv)*). Hence, the hybrid predicts that varying X will give rise to the necessary variation as long as there is noise in actions, beliefs, or both.

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 X-games. An important consideration guiding our design is being 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

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 experiment consisted of two treatments, which we label "[A,BA]" and "[A,A]". The main treatment is [A,BA], which we describe here. The treatment [A,A] is similar, but does not involve belief elicitation.

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 treatment name [A,BA] reflects the two stages: "A" for "action" and "BA" for "belief-action".

In each of the 20 rounds of the first stage, subjects were anonymously and randomly paired with subjects in the opposite role, and actions were taken simultaneously. Each of the 40 rounds of the second stage were similar to the first-stage rounds, except for two differences. First, prior to taking an action, each subject stated a belief over the matched subject's expected action choice.[11] Second, rather than being paired against a subject acting simul-

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

taneously, each subject faced a first-stage subject whose action had already been recorded. To be precise, each subject was presented with a payoff matrix that appeared in the first stage. Then, the subject's belief was elicited over the action taken by a randomly selected first-stage subject for that game and the subject's action was similarly paired against the action of a randomly selected subject for that game. Since second-stage subjects were not paired against other subjects acting simultaneously, they were not required to wait for all subjects to finish a round before moving on to the next. In both stages, however, subjects were required to wait for 10 seconds before submitting their answers. Screenshots of the experimental interface are given in Appendix 10.10.

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

Table 2: *Subjects in each treatment*

Before the start of the first stage, instructions (see Appendix 10.11) 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. Table 2 summarizes the number of subjects who participated in the experiment by treatment and player role.[12]

3.2 Procedures and incentives

We are interested in observing the stochasticity inherent in beliefs. Hence, we wished to eliminate predictable sources of variation in beliefs due to new information or learning. For the same reason, we also wished to minimize variation in stated beliefs due to measure-

[12]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.

ment error, resulting from either noisy introspection or physically entering beliefs. These considerations weighed heavily in the choice of procedures.

To avoid learning, 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. Furthermore, 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 also cannot be due to a higher-ordered belief that other subjects were learning.

To minimize measurement error, belief statements had to be entered as whole numbers into a box rather than via a slider. Of course, measurement error can never be fully eliminated, and we acknowledge that the act of elicitation itself may introduce some degree of error.

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 3.3.

In addition to a $10 show-up fee, subjects were paid according to one randomly selected round from the first stage (based on actions) and four randomly selected rounds from the second stage–two rounds based on actions and two rounds based on beliefs. 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.

To incentivize actions, if a round were selected for a subject's action payment, the subject was paid according to the outcome of the game. Each unit of payoff in the matrix corresponded to a probability point of earning $10. For example, a payoff of 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.[13] This is important for our purposes since several of our tests require that utilities are identified.

[13]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].

To incentive subjects to accurately report their beliefs, we used the random binary choice mechanism (Karni [2009]), also known as the "lottery method". We describe the mechanism and its implementation in greater detail in Appendix 10.5. The important feature is that reporting truthfully is incentive compatible, independent of risk attitude.[14] Our variant gives subjects a chance at $5 for each elicitation selected for payment.

To allay any hedging concerns, all five payments were based on different matrices, and this was emphasized to subjects. On average, the experiment took about 1 hour and 15 minutes, and the average subject payment was $19.5.

3.3 The games

As discussed in Section 10.3, the X-games take center-stage since they are predicted to give rise to systematic variation in actions and beliefs. Henceforth, we refer to the game with $X = 80$ as "$X80$" 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.

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.*

For the experiment, we chose the six values of X given in Table 3. These correspond to the vertical lines in Figure 1. 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 games $X40$ and $X10$ are symmetric-pairs in that $\sigma_L^{NE,X80} = 1 - \sigma_L^{NE,X5}$ and $\sigma_L^{NE,X40} = 1 - \sigma_L^{NE,X10}$.

[14] In fact, expected utility maximization is not required. Subjects need only have preferences that respect stochastic dominance.

This does not, however, imply the analogous symmetry relation for the hybrid model without additional conditions.[15]

Each of the X-games are played 2 times in the first stage and 5 times in the second stage. Since there are six X-games, they appear a total of 14 times in the first stage and 30 times in the second stage.

We also included a small number of additional 2×2 games, whose analysis will be the subject of a separate paper. These were included to break up the appearance of the X-games, similar to what has been done in experiments testing models of stochastic choice (e.g. Tversky [1969]). With these additional games, the first stage has 20 rounds in total and the second stage has 40 rounds in total. This implies that the second stage has twice the number of rounds as the first as desired (see Section 3.2) and that the X-games take up a similar fraction of total games in each stage.

The games appear in a random order, subject to some restrictions, such as ensuring that the same game does not appear more than once within 3 consecutive rounds. Appendix 10.4 shows all of the games and explains the precise randomization.

4 Overview of the data

We begin with an overview of the data. Since our procedures are novel, we also benchmark our findings against those from existing experiments.

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 special notation to indicate the data source. In particular, "$[\underline{\text{A}}, \circ\ ]$" refers to first-stage actions pooled across [A,BA] and [A,A], "$[\text{A},\text{B}\underline{\text{A}}]$" refers to second-stage actions from [A,BA], and we occasionally refer to other data sources in a similar fashion.

We are primarily interested in the aggregate action frequencies from $[\underline{\text{A}}, \circ\ ]$ 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 $[\underline{\text{A}}, \circ\ ]$.

[15] 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 $\sigma_L^{hybrid,X} = 1 - \sigma_L^{hybrid,X'}$ and $\sigma_U^{hybrid,X} = 1 - \sigma_U^{hybrid,X'}$ if $\sigma_L^{NE,X} = 1 - \sigma_L^{NE,X'}$.

Figure 2 plots the action frequencies from [A, ∘], which are also given in Appendix Table 12.

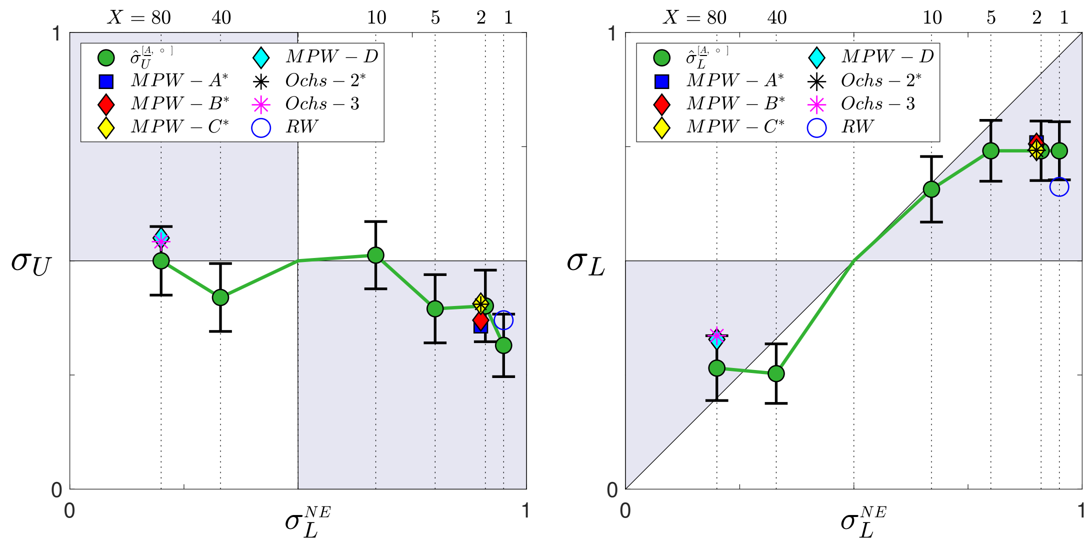

Figure 2: *Actions data.* This figure plots the first stage empirical frequencies from [A, ∘] with 90% confidence bands (clustered by subject), superimposed with the empirical frequencies from other studies.

We see that the data largely falls within the QRE-NBE region, and F-tests strongly reject that the data generating process is Nash equilibrium. The only surprise concerns $X40$ for 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 outcomes of the hybrid model, as we show in Appendix Figure 24.

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 that which is collected under more standard experimental conditions? Figure 2 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 that played games with "sparse" payoffs[16] 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. This being said,

[16] 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.

we cannot find precedents in the literature for games closely matching our more symmetric games–those with σ_L^{NE} relatively close to $\frac{1}{2}$.

We will also make use of actions data from [A,B$\underline{\text{A}}$] because, in testing axioms on the action-map, we must associate to each belief statement a corresponding action. We present the empirical action frequencies from [A,B$\underline{\text{A}}$] in Appendix Table 12, and make the following remark.

Remark 1. There is a small but significant difference in first- and second-stage action frequencies from [$\underline{\text{A}}$, ∘] and [A,B$\underline{\text{A}}$]. In Appendix 10.6, we show that this difference is caused by the process of belief elicitation itself.[17] This difference has no effect on our main conclusions in that they would still hold if the empirical action frequencies were replaced with any convex combination of those from [$\underline{\text{A}}$, ∘] and [A,B$\underline{\text{A}}$].

4.2 Beliefs

Figure 3 plots individual belief statements along with the median and quartiles of beliefs for each game. Appendix Table 12 reports both median and mean beliefs.

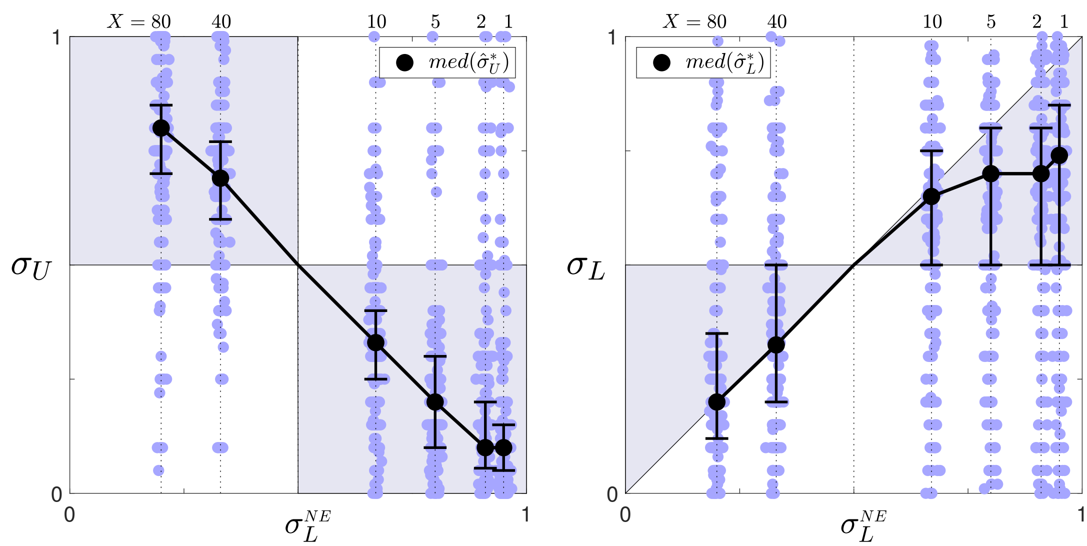

Figure 3: *Beliefs data.* This figure plots individual belief statements along with the median and quartiles of beliefs. The left panel gives beliefs over σ_U, and the right panel gives beliefs over σ_L.

[17] Studies report mixed results on the existence and magnitude of such belief elicitation effects, and there is no consensus on the conditions that evoke them. Schotter and Trevino [2014] provide a discussion.

Deferring a deeper discussion of the beliefs data to sections testing the axioms, we make two immediate observations. First, beliefs clearly respond to changes in X, with median and quartiles of beliefs varying monotonicially in X. Second, beliefs over σ_U tend to have much less dispersion than beliefs over σ_L, which have average interquartile ranges of 15 and 29 belief points, respectively. Figure 4 plots the entire histogram of beliefs for each player and game, and will be an important reference in subsequent sections.

4.3 Actions given beliefs

We begin by presenting rates of best response, which the experimental literature has suggested as a method for validating elicited beliefs (see Schotter and Trevino [2014] for a discussion of this view). Figure 5 plots histograms of subjects' rates of best response, calculated from all six 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 13 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.

The premise of quantal response is that beliefs determine actions only insofar as they pin down expected payoff vectors. This suggests to a convenient way of visualizing the whole of the data–through the empirical "quantal response surface."

The left panel of Figure 6 plots the convex hull of all expected payoff vectors that we may observe in the data. Vector (v_{i1}, v_{i2}) refers to the expected payoffs to (U, D) for player 1 or the expected payoffs to (L, R) for player 2. 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 the empirical quantal response surface that gives the expected action probability as a function of payoff vectors based on a local linear (lowess) regression.

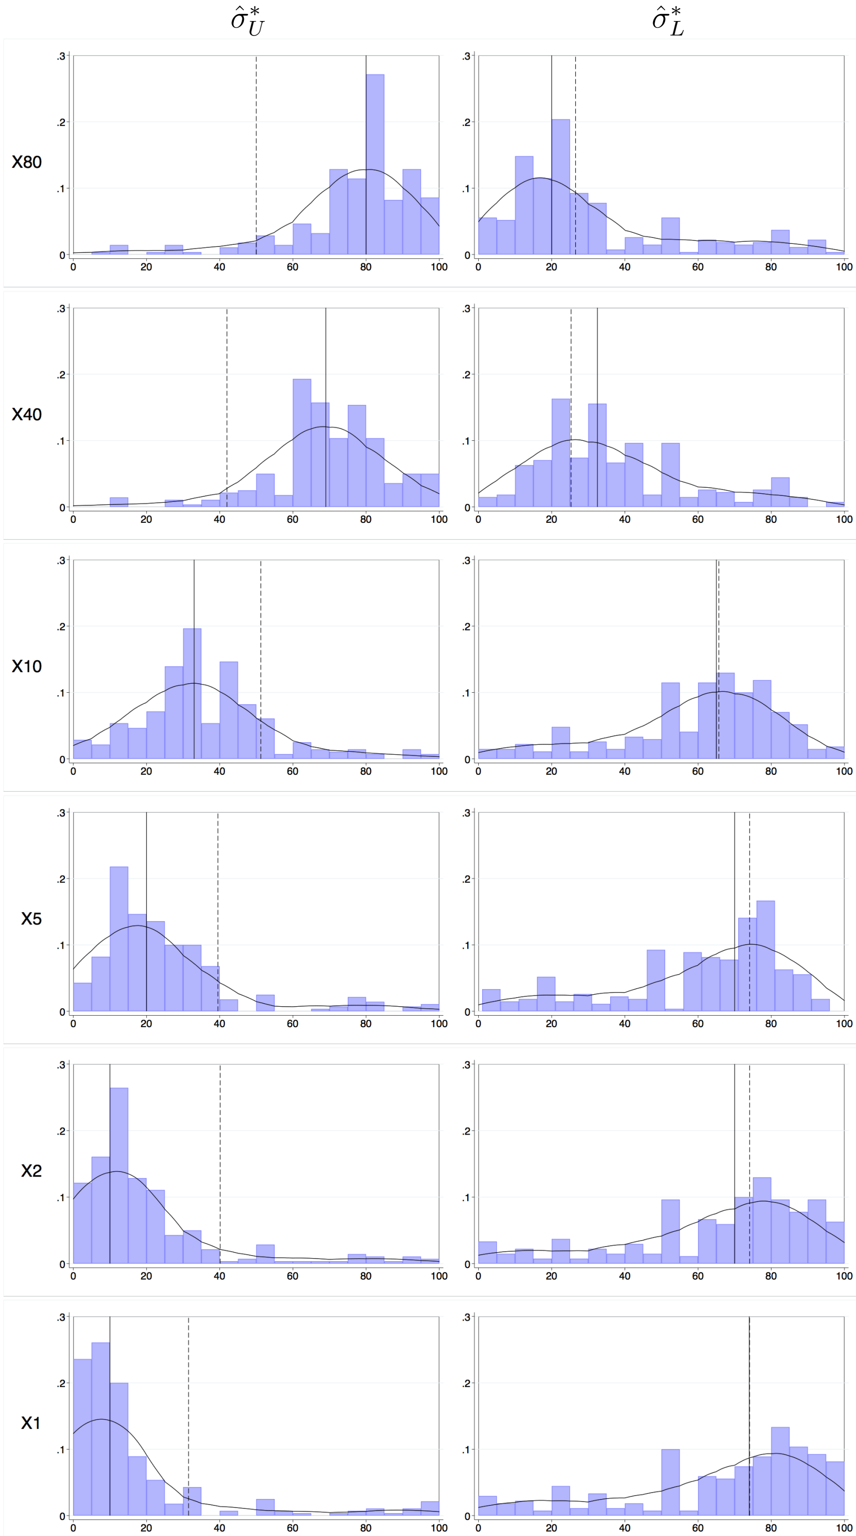

Figure 4: *Belief distributions.* The left panel is for player 2's beliefs about U, and the right panel is for player 1's beliefs about L. The solid lines mark the median of i's beliefs and the dashed line marks the empirical frequency of j's actions.

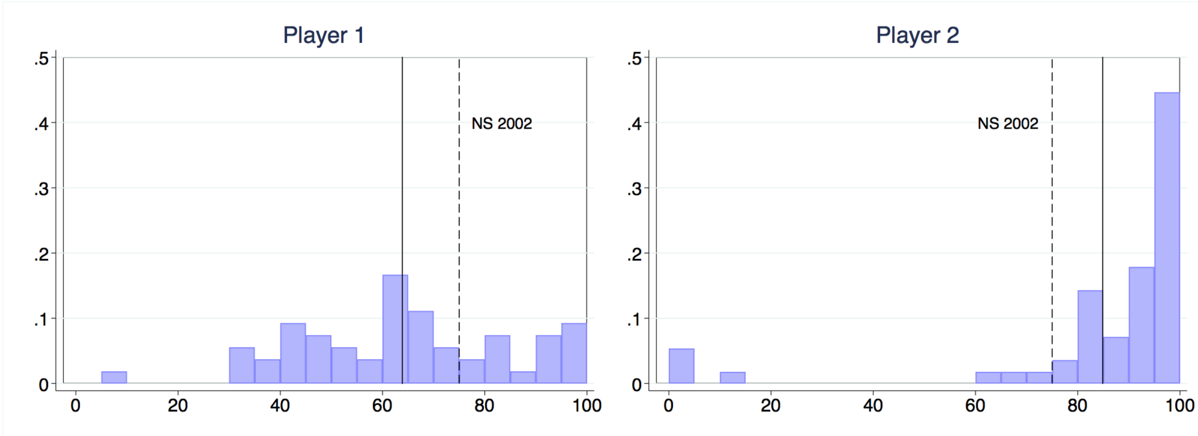

Figure 5: *Subjects' rates of best response.* This figure gives histograms of subjects' rates of best response across all X-games. The solid lines are averages, and the dashed lines in the bottom panel mark the average rate of best response from Nyarko and Schotter [2002].

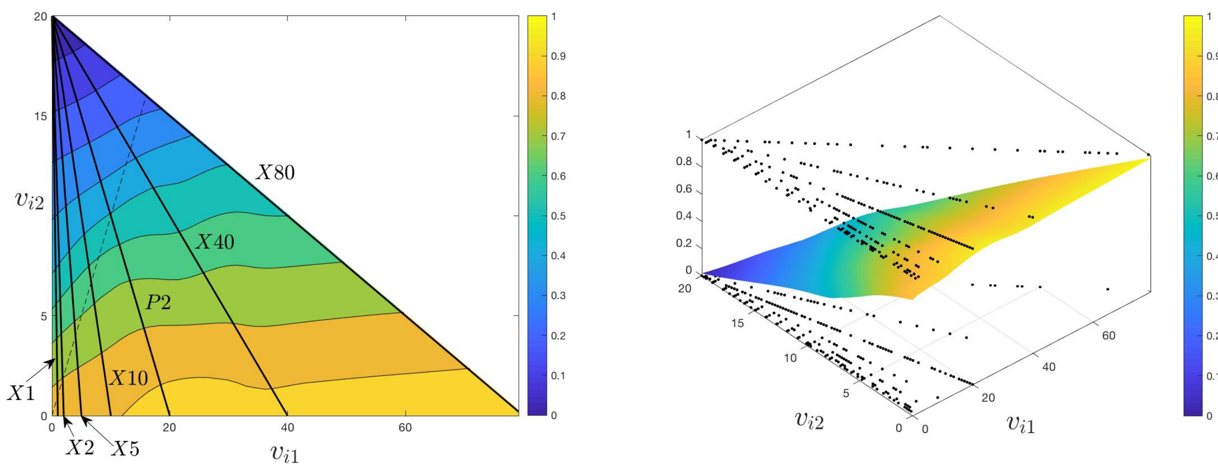

Figure 6: *Empirical quantal response surface.* The left panel gives the convex hull of all expected payoff vectors that may be observed in the data in any of the X-games. (v_{i1}, v_{i2}) refers to either the payoffs to (U, D) for player 1 or the payoffs to (L, R) for player 2. "X1", "X2", ... , and "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.

The left panel gives the associated level sets.

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

5 Are beliefs noisy?

To our knowledge, this is the first study to have multiple belief elicitations per subject-game without feedback. This allows us to answer the basic question: are beliefs noisy?

For each subject and X-game, we calculate the spread of her beliefs–the highest belief minus the lowest belief–across the five belief statements. We average this across the six X-games for each subject to get an average spread measure. Figure 7 plots histograms of subjects' spreads by player role. There is considerable heterogeneity in spreads, and there is a right tail of subjects with very high spreads. The average spreads are 25 and 21 belief-points for player 1- and player 2-subjects, respectively. Is this evidence for noisy beliefs? Or does it simply reflect learning or measurement error?

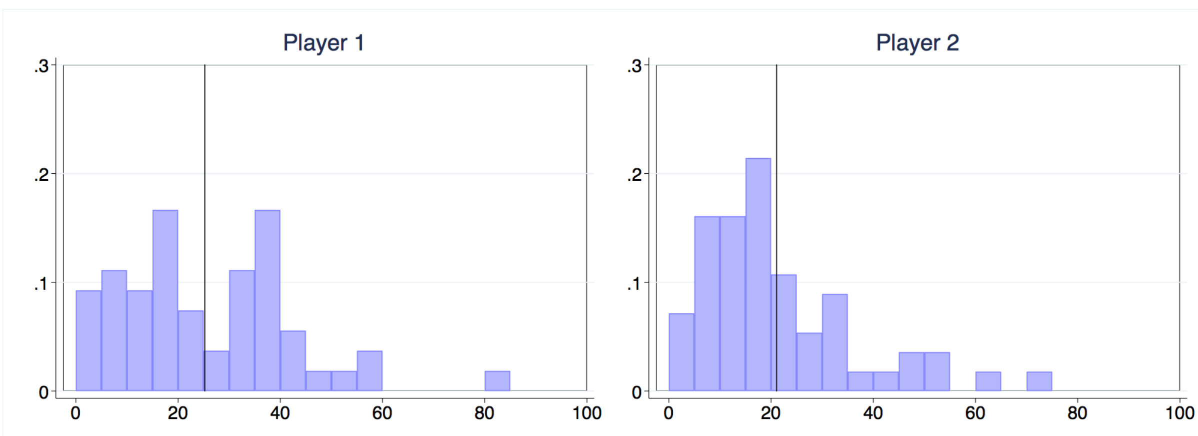

Figure 7: *Subjects' spreads of beliefs.* This figure gives histograms of subjects' spreads of beliefs, averaged across all X-games.

Our procedures, in particular the lack of feedback, were designed to minimize conventional learning due to new information (see Section 3.2). However, there may still be trends in beliefs across the five appearances of each game, which would indicate some form of no-feedback learning (Weber [2003]). Appendix Figure 25 shows that average beliefs are very stable throughout the experiment, suggesting there is no overall trend in beliefs. To account for subject-specific trends, we linearly detrend beliefs for each subject and game. Using these detrended beliefs, we recalculate for each subject an average spread. Appendix Figure 26 replicates Figure 7 using these detrended spreads. We find similar results, with only slightly smaller average spreads (21 and 18 for players 1 and 2, respectively). We conclude that learning does not drive variation in beliefs.

In Section 8, we argue formally that variation in stated beliefs cannot be due entirely to random measurement error, i.e., it must reflect variation in true underlying beliefs. The basic idea is that, within-subject, variations in stated beliefs are strongly predictive of the actions subjects take. Hence, a high stated belief signals a high true belief, and thus variation in stated beliefs cannot be due entirely to measurement error.

Result 1. *Beliefs are noisy.* The majority of subjects have stochastic belief reports, with an average spread of 21-25 percentage points. This is not the result of learning or measurement error.

6 Testing the axioms

We attempt to reject the four behavioral axioms: *responsiveness*, *monotonicity*, *belief-responsiveness*, and *unbiasedness*. To the extent that we cannot reject an axiom, we think of it as a plausible description of the data.

6.1 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 the second stage of [A,BA].

Since player 1's payoff parameter X is different in each game, 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.[18] Instead, we first consider tests game-by-game. Then, we consider player 2 only, whose payoff parameters are fixed across games, allowing us to pool data across the entire set of games.

[18]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 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*.

Since expected payoffs are one-to-one with beliefs within a game, *responsiveness* is easily translated in terms of beliefs. For player 1 and game X, *responsivness* requires that Q_U, the probability of taking U, is everywhere strictly increasing in belief σ_L'. Similarly, for player 2, *responsivness* requires that Q_L, the probability of taking L, is everywhere strictly decreasing in belief σ_U'.

We visualize the relevant data in Figure 8, which plots estimates of $\hat{Q}$ for games $X80$ and $X5$ for both players. Appendix Figure 27 gives the plots for all six games. These are simply the predicted action frequencies from regressing actions on beliefs using a flexible specification (see caption of Figure 8 for details). Recall that, for each game and player role, there are five observations per subject. The vertical dashed line gives the *indifferent belief* $\sigma_j' = \sigma_j^{NE}$ and the horizontal dashed line is set to one-half. The plots also include belief histograms and the average action within each of ten equally spaced bins (black dots).[19]

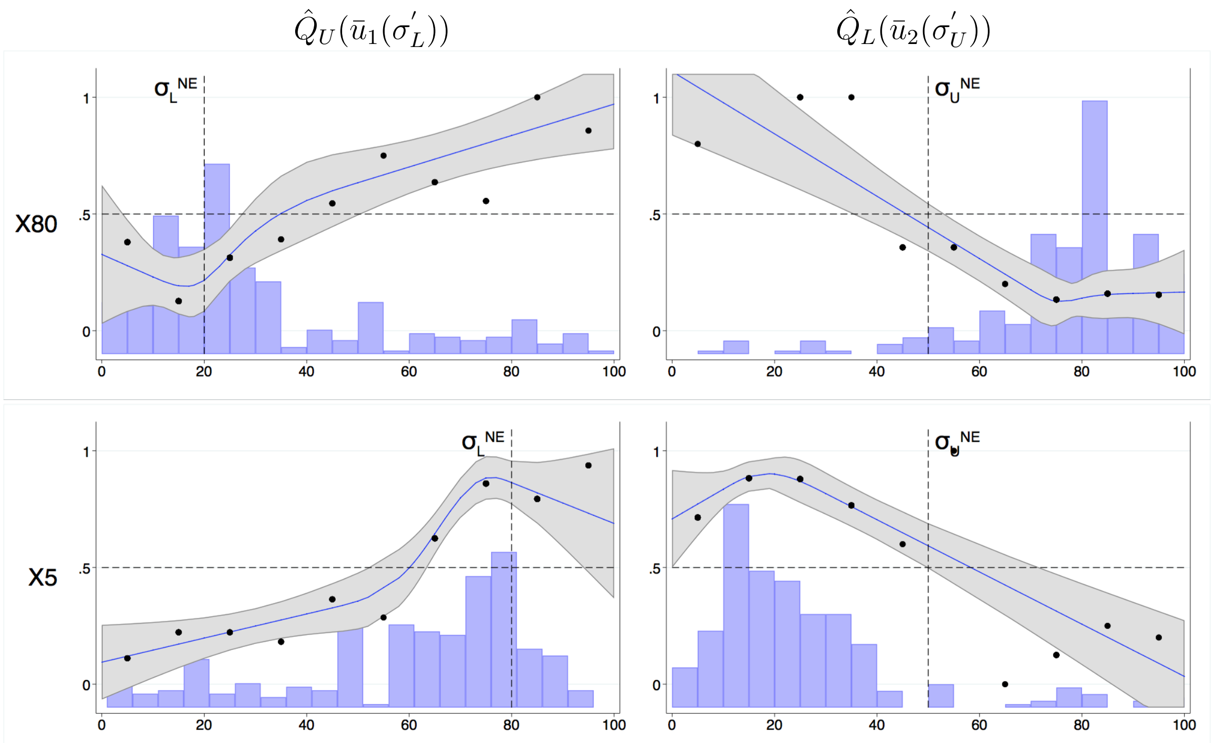

Figure 8: *Action frequencies predicted by beliefs.* For each player and games $X80$ and $X5$, we plot the predicted values (with 90% error bands) from restricted cubic spline regressions of actions on beliefs (4 knots at belief quintiles, standard errors clustered by subject). Belief histograms appear 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.

[19]Note that, for some bins, there is very little data.

Responsiveness is equivalent to an increasing slope for player 1 and a decreasing slope for player 2. Inspecting Figure 8, it appears there may be violations. However, since different subjects form different beliefs, the $\hat{Q}$-curves plotted in Figure 8 are patched together from different subjects representing different parts of the domain. Hence, any perceived violations 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. 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. Since *responsiveness* is a local property, we run separate regressions for different regions of stated beliefs. 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, actions 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:[20]

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

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 4. 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,[21] 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 9 reproduces Figure 8,

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

[21]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 4: *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.

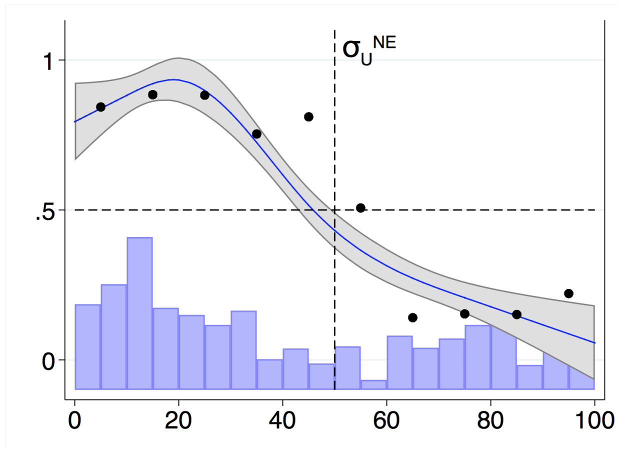

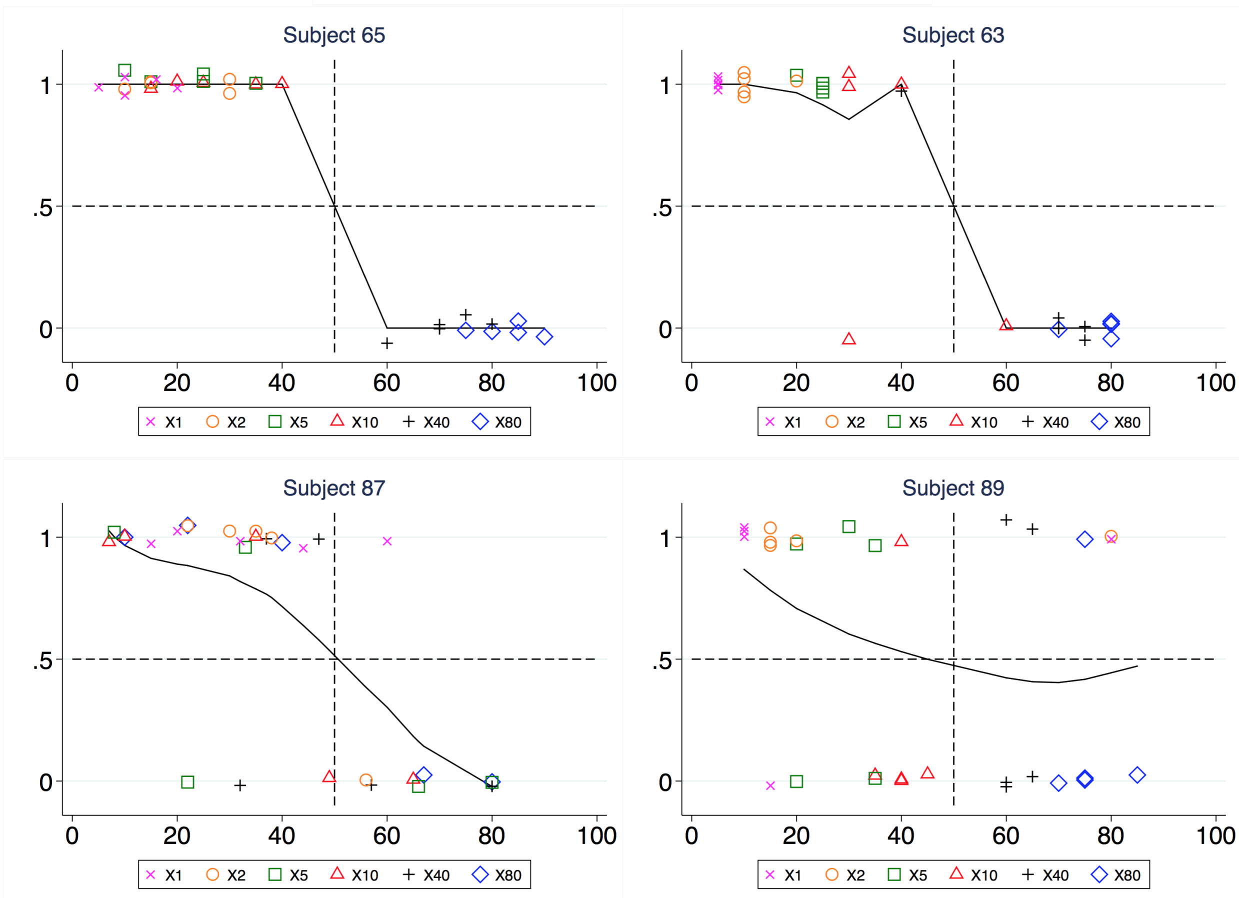

Figure 9: *Player 2–actions and beliefs, pooled across games.* All plots involve player 2-subjects whose data is pooled across all games. Action L is coded as 1, and action R is coded as 0. The top panel uses all player 2-subjects and gives the predicted action frequencies (with 90% error bands) from 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. Solid black curves are estimates from local linear regressions, 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.

and Appendix Table 14 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 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 14, 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.

To explore subject heterogeneity that is hidden in the regressions, Figure 9 plots the data, pooled across all six games, for four representative player 2-subjects. Subject 65 appears to have step function-like *responsiveness* and always best responds; subject 63 looks similar, but has a single "mistake"; subject 87 also appears responsive, but with action-probabilities that are more continuous in beliefs; subject 89 is mostly responsive, but very noisy.

Result 2. *Responsiveness cannot be rejected.* An increase in a subject's beliefs are associated with an increase in the probability of taking the action whose payoff is increasing in beliefs.

6.2 Monotonicity

Monotonicity is a weakening of best response which states that, *given belief*s, 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 the second stage of [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, *monotonicity* requires that Q_U, the probability of taking U, is greater than $\frac{1}{2}$ if and only if belief σ'_L is greater than $\sigma_L^{NE,X}$: $Q_U \gtreqless \frac{1}{2} \iff \sigma'_L \gtreqless \sigma_L^{NE,X}$. Similarly, for player 2, *monotonicity* requires $Q_L \gtreqless \frac{1}{2} \iff \sigma'_U \lesseqgtr \sigma_U^{NE,X}$.

In order to visualize potential *monotonicity* violations, we appeal once again to Figure 8, which plots estimates of $\hat{Q}$ for games $X80$ and $X5$ for both players (see figure caption for details; see Appendix Figure 27 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 8. 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 5 for five 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 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 8, it is clear that a large mass of beliefs (including the mode) fall in the regions with *monotonicity* violations.

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 5: *Testing monotonicity.* For each player and game, we test for *monotonicity* in the manner described in Section 6.2 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.

From Figure 8, 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. In Appendix 10.8, we show this pattern is consistent with concavity in the utility function. Since payoffs are in probability points, this cannot be interpreted as risk aversion, so we discuss possible alternatives.

Result 3. *Monotonicity is rejected.* There exists regions of beliefs for which player 1-subjects fail to best respond more often than not. The pattern is consistent with concavity in the utility function (over probability points).

Remark 2. A weak implication of *monotonicity* is that best responses will be taken with probability greater than one-half. As shown in Appendix Table 13, 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.

6.3 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,BA] and the actions data from [$\underline{\text{A}}$, ∘].

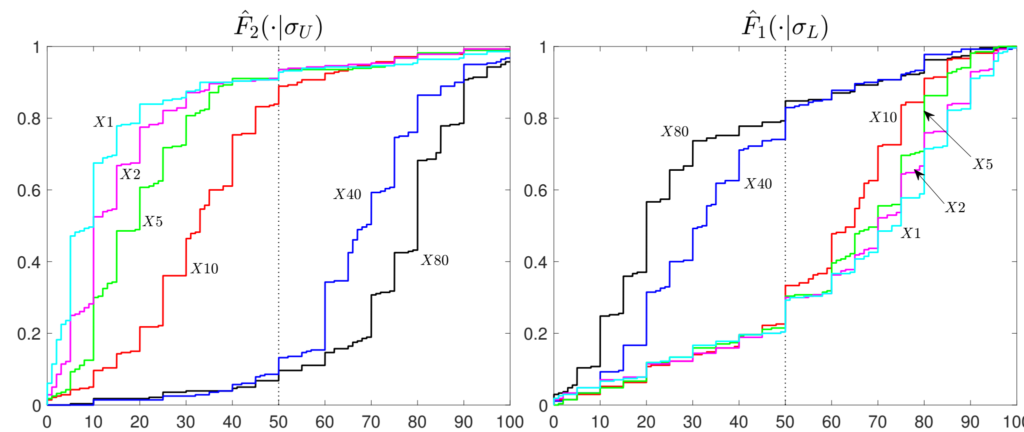

Figure 10: *CDFs of belief distributions.* We plot the empirical CDFs of belief distributions. The left panel is for player 2's beliefs about U, and the right panel is for player 1's beliefs about L.

Prior to testing, we visualize the data in Figure 4, which plots histograms of stated beliefs, superimposed with median beliefs (solid vertical lines) and the corresponding empirical frequencies of actions (dashed vertical lines). It appears that the distributions of beliefs shift monotonically in X in the direction predicted by the hybrid model: 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 10 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*.

A violation of *belief-responsiveness* occurs whenever, across two games x and y, $\sigma_j^x > \sigma_j^y$ and $F_i(\cdot|\sigma_j^x) \not\succ_{FOSD} F_i(\cdot|\sigma_j^y)$, meaning actions and beliefs do not go in the same direction (in terms of stochastic dominance) as the corresponding action frequencies. Hence, we perform one-sided tests of the null hypotheses $H_0 : \sigma_j^x > \sigma_j^y$ and $H_0 : F_i(\cdot|\sigma_j^y) \succ_{FOSD} F_i(\cdot|\sigma_j^x)$ for all games $x \neq y$ and players $i \neq j$. We say that *belief-responsiveness* is rejected whenever we reject *both* $\sigma_j^x > \sigma_j^y$ and $F_i(\cdot|\sigma_j^y) \succ_{FOSD} F_i(\cdot|\sigma_j^x)$.

Actions

		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
H_0 :	$X40$	0.42	–	0.00***	0.00***	0.00***	0.00***	0.11	–	**0.07***	0.65	0.61	0.96
σ_j^x	$X10$	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
σ_j^y	$X2$	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

		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
H_0 :	$X40$	0.80	–	0.00***	0.00***	0.00***	0.00***	0.00***	–	1.00	1.00	1.00	1.00
$F_i(\cdot\|\sigma_j^x)$	$X10$	0.96	0.93	–	0.00***	0.00***	0.00***	0.00***	**0.00*****	–	0.84	0.98	0.73
$\succ_{FOSD}$	$X5$	0.97	0.92	0.76	–	0.00***	0.00***	0.00***	0.00***	0.00***	–	0.95	0.77
$F_i(\cdot\|\sigma_j^y)$	$X2$	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 6: *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]. The entries in bold correspond to the only rejection, i.e. rejections of $\sigma_1^{X40} > \sigma_1^{X10}$ and $F_1(\cdot|\sigma_1^{X10}) \succ_{FOSD} F_1(\cdot|\sigma_1^{X40})$.

Table 6 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 all games $x \neq y$ and players $i \neq j$ (see table caption for details). 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.

Result 4. *Belief-Responsiveness cannot be rejected.* Beliefs distributions are ordered across games by stochastic dominance as predicted by the hybrid model. Empirical action frequencies are mostly ordered in the same way.

6.4 Unbiasedness

Unbiasedness states that beliefs are correct on median, i.e. $med(\sigma_j^{*,X}) = \sigma_j^X$ for all X. Once again, we use the beliefs data from [A,BA] and the actions data from $[\underline{A}, \circ\,]$.

Unbiasedness requires that beliefs are unbiased on median, so we plot in Figure 11 the aggregate action frequencies and median beliefs as well as the individual belief statements. Appendix Table 15 reports the bias in both median- and mean-beliefs with p-values of the hypothesis that beliefs are unbiased (see caption for details).

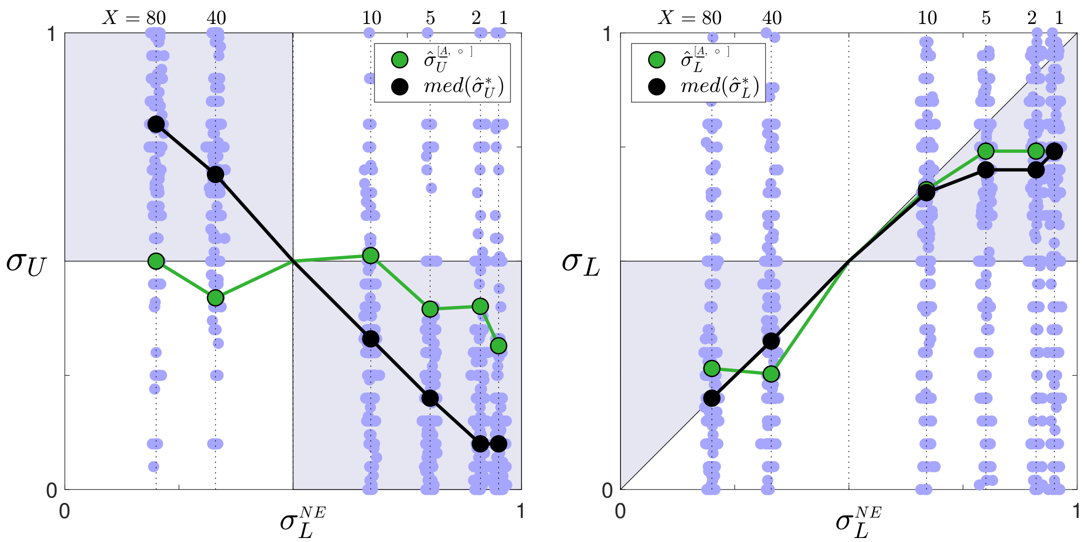

Figure 11: *Bias in beliefs.* The left panel gives player 1's action frequency from $[\underline{A}, \circ\,]$ and the median of player 2's beliefs about player 1. Blue circles are individual belief statements. The right panel gives player 2's action frequency from $[\underline{A}, \circ\,]$ and the median of player 1's beliefs about player 2. Red circles are individual belief statements.

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.

Result 5. *Unbiasedness is rejected.* Player 2-subjects have an "extreme" bias, exaggerating the direction of deviations from uniform play. Player 1-subjects have unbiased or slightly conservative beliefs.

Remark 3. To completely explain the observed pattern of bias would require an alternative theory of actions and beliefs jointly, which is beyond the scope of this paper. However, we provide in Appendix 10.9 evidence that player 1-subjects have longer response times and appear more sophisticated, in the level-k sense, in some of the non-X-games that were included in the experiment. Since all subjects saw the same games throughout the experiment and were randomized into their roles, we interpret this causally, i.e. the player role itself induces greater sophistication in player 1-subjects. We think of this as a potential "ingredient" in developing new models that can explain the pattern of bias.

7 Action-noise or belief-noise?

It is common practice to fit parametric models with stochastic elements to experimental data. Models with action-noise and correct beliefs, such as logit QRE, are particularly common. It is also common to augment level k (and other non-equilibrium models) with action-noise prior to fitting. Fitting models with belief-noise is uncommon, though there are a few exceptions. A priori, it is not clear what is lost in terms of explanatory power by not allowing for one of action- or belief-noise; and if an important type of noise is ignored, the misspecification may lead to biased estimates. For example, a model that only allows for action-noise may overestimate its importance and therefore overpredict the degree of mistakes in decision problems (for which subjects do not need to form beliefs). Following this motivation, we ask: which of action- or belief-noise is more important for explaining the data?

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, and we say a source of noise is important if turning it off leads to large prediction errors. For this, we use the second stage-data from [A,BA], for which we can associate actions with beliefs. Importantly, since we do not want to conflate noise with heterogeneity, we construct counterfactuals subject-by-subject.

To turn off belief-noise, we replace every belief statement with the median belief statement for the corresponding subject-game. For each subject and game, we then predict behavior based on this median belief and a best-fit random utility model for that subject. For player 1-subjects, based on evidence from Section 6.2, we also allow for curvature in the utility function.[22] We leave the details of the random utility estimation to Appendix 10.7, but emphasize that the estimation is done for each subject based on her data from all 6 games. We denote the counterfactual action frequency by $p_s^{iX}(\hat{b}_{s,med}^{iX}; \hat{\theta}_s^i)$. This is the predicted probability for subject s in role i of taking U if $i = 1$ (L if $i = 2$) in game X, where $\hat{\theta}_s^i = (\hat{\rho}_s^i, \hat{\mu}_s^i)$ are estimated curvature and noise parameters ($\hat{\rho}_s^i$ set to 0 if $i = 2$, corresponding to linearity) and $\hat{b}_{s,med}^{iX}$ is the median belief.

To turn off action-noise, we assume subjects best respond to every belief realization. Thus, the counterfactual action frequency is given by $q_s^{iX}(\{\hat{b}_{sl}^{iX}\}_l; \hat{\rho}_s^i) = \frac{1}{5}\sum_l BR^{iX}(\hat{b}_{sl}^{iX}; \hat{\rho}_s^i)$, where $\{\hat{b}_{sl}^{iX}\}_l$ are the five belief statements (indexed by $l = 1, .., 5$) for subject s in role i and game X; and BR^{iX} is the best response correspondence that, for player i, equals 1 if $i = 1$ and U is a best response to the belief ($i = 2$ and L is the best response), which depends on the estimated curvature $\hat{\rho}_s^i$ if $i = 1$.

For each subject-game and counterfactual, we calculate the absolute difference between empirical frequency and counterfactual frequency. Averaging within-subject across all six games gives the subject's average counterfactual prediction error:

$$\varepsilon_{si}^{\text{belief-noise}} = \frac{1}{6}\sum_X |\hat{\sigma}_{si}^X - p_s^{iX}(\hat{b}_{s,med}^{iX}; \hat{\theta}_s^i)|,$$
$$\varepsilon_{si}^{\text{action-noise}} = \frac{1}{6}\sum_X |\hat{\sigma}_{si}^X - q_s^{iX}(\{\hat{b}_{sl}^{iX}\}_l; \hat{\rho}_s^i)|.$$

Hence, each subject is associated with a pair $\varepsilon_{si} = \{\varepsilon_{si}^{\text{belief-noise}}, \varepsilon_{si}^{\text{action-noise}}\}$ which gives the prediction errors that result from turning off action- and belief-noise, respectively. In other words, $\varepsilon_{si}^{\text{belief-noise}}$ is the error of the counterfactual with only belief-noise, and $\varepsilon_{si}^{\text{action-noise}}$ is the error of the counterfactual with only action-noise. We say that one source of noise is important if its counterfactual leads to small errors.

Figure 12 plots all subjects' counterfactual prediction errors. Within each player role

[22] For player 2-subjects, due to symmetry of payoffs, this cannot be separately identified from the noise term.

i, subjects are sorted so that $\varepsilon_{si}^{\text{belief-noise}}$ is increasing. In addition to plotting $\varepsilon_{si}^{\text{belief-noise}}$ and $\varepsilon_{si}^{\text{action-noise}}$, we also give the errors that would result from predicting uniform random behavior.

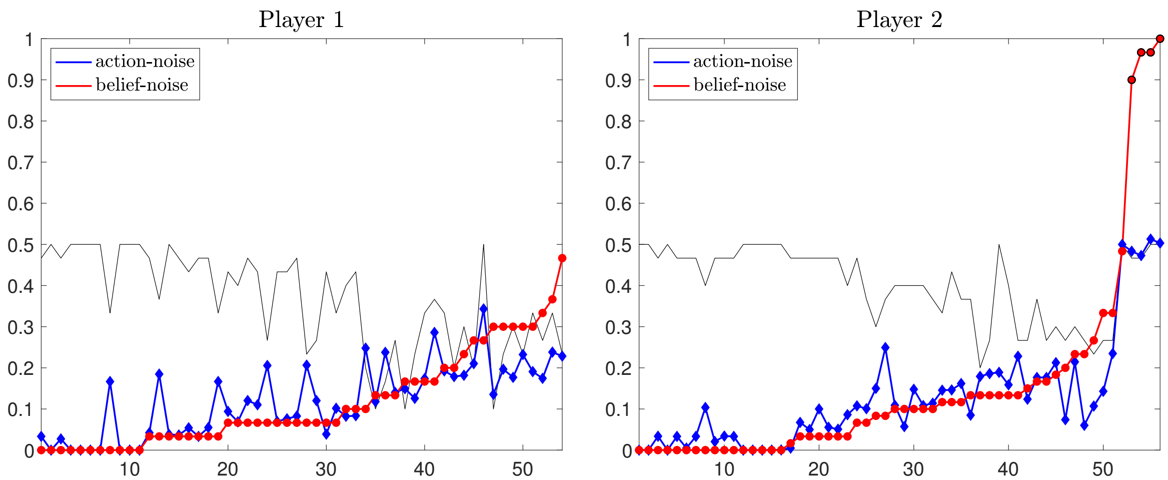

Figure 12: *Counterfactual prediction errors: all subjects.* For each player, we plot the prediction errors of all subjects from the counterfactuals with only belief-noise and action-noise, respectively. The thin line is the error resulting from predicting uniform random behavior.

Table 7 summarizes the average errors by player role, and also gives the numbers of subjects for whom each type of noise does a better job of explaining their behavior. Noting that the prediction errors range from 0 to 1, we see that both counterfactuals do rather well in absolute terms, and much better than the random benchmark.

	Player 1	Player 2	Player 2 (dropping "worst responders")
mean action-noise errors	0.12	0.13	0.10
mean belief-noise errors	0.12	0.16	0.09
uniform random benchmark	0.36	0.41	0.40
# (%) best expl. by action-noise	18 (33%)	14 (25%)	10 (19%)
# (%) best expl. by belief-noise	27 (50%)	34 (61%)	34 (65%)
# (%) equally well expl. by both	9 (17%)	8 (14%)	8 (15%)

Table 7: *Mean counterfactual prediction errors.*

From the table, both types of noise do equally well on average for player 1, and action-noise does slightly better for player 2. However, belief-noise does better for a larger fraction of subjects for both player roles, suggesting large belief-noise errors for a few subjects. In particular, as is clear from the right panel of the figure, four player 2-subjects have belief-noise errors greater than 0.9. These subjects are near perfect "worst responders" who systematically fail to best respond. The belief-noise counterfactual does so poorly for these subjects because the best response assumption trivially maximizes prediction error. For these subjects, the action-noise counterfactual predicts uniformly random behavior and so fairs much better. Since best responding for player 2-subjects should be easy given their symmetric payoffs, these subjects may simply be confused. For this reason, we also calculate average prediction errors after dropping these subjects and find that belief-noise performs slightly better than action-noise.

Result 7. *Both action-noise and belief-noise are similarly important.* Counterfactual action frequencies that result from "turning off" either action-noise or belief-noise lead to similar average prediction errors.

8 Measurement error

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 measurement erorr, resulting from errors in reporting or noisy introspection about one's beliefs. In this 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] or $\{0, ..., 100\}$ as convenient. 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 4, 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.

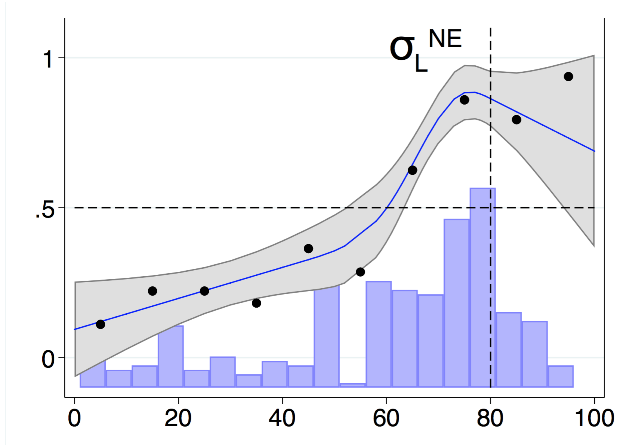

Figure 13: *Monotonicity violations and measurment error.*

Monotonicity. Consider player 1 in game $X5$, whose predicted probability of taking U as a function of stated beliefs is reproduced in Figure 13. The indifferent belief is 80, and the *monotonicity* violation occurs in the interval of stated beliefs $\{61, ..., 80\}$, i.e. to the left of the indifferent belief. Based on our framework, the observed *monotonicity* violation occurs when stated beliefs are in the interval $\{61, ..., 80\}$ and the associated true beliefs induce the action U more often than not. Hence, suppose that actions given true beliefs are governed by $Q_U(\bar{u}_1(b_0)) = \frac{1}{2}$ for $b_0 \leq 80$ and $Q_U(\bar{u}_1(b_0)) = 1$ for $b_0 > 80$, which is the *monotonic* quantal response function most likely to generate such violations. Under the assumption that $med[b_s^*(b_0)] = b_0$ for all b_0, we argue that the *monotonicity* violation we observe is very unlikely. To this end, suppose a given stated belief $b_s \in \{61, ..., 80\}$ is associated with $\hat{Q}_U(\bar{u}_1(b_s)) > \frac{1}{2}$, the predicted probability of taking U. It is straightforward to back out $\hat{\alpha}(b_s)$, the estimated probability that the associated true belief is strictly greater than 80:

$$
\begin{aligned}
(1 - \hat{\alpha}(b_s))Q_U(\bar{u}_1(b_0 \leq 80)) + \hat{\alpha}(b_s)Q_U(\bar{u}_1(b_0 > 80)) = \hat{Q}_U(\bar{u}_1(b_s)) &\iff \\
(1 - \hat{\alpha}(b_s))\frac{1}{2} + \hat{\alpha}(b_s) = \hat{Q}_U(\bar{u}_1(b_s)) &\iff \\
\hat{\alpha}(b_s) = 2\hat{Q}_U(\bar{u}_1(b_s)) - 1.&
\end{aligned}
$$

We make two observations. First, letting $\hat{N}(b_s)$ be the number of belief statements equal to b_s, the expected number of true belief realizations strictly greater than 80 associated with belief statement $b_s \in \{61, ..., 80\}$ is given by $\hat{\alpha}(b_s)\hat{N}(b_s)$. Second, since $med[b_s^*(b_0)] = b_0$ for every true belief b_0, whenever $b_0 > 80$, the probability of observing a stated belief strictly greater than 80 is at least $\frac{1}{2}$. Combining these two observations, it must be that

$$\frac{1}{2}\sum_{b_s\in\{61,...,80\}}\hat{\alpha}(b_s)\hat{N}(b_s)\leq\sum_{b_s\in\{81,...,100\}}\hat{N}(b_s),$$

where the left-hand side is a lower bound on the expected number of belief statements strictly greater than 80 based on data associated with belief statements in $\{61,...,80\}$, and the right-hand side is the observed number of belief statements strictly greater than 80. In the data, the left-hand-side equals 39.4 and the right-hand-side equals 37, meaning measurement error cannot explain away the observed *monotonicity* violation. To summarize the argument, it could in principle be that we observe what appears to be a violation of *monotonicity* due to measurement error if high true beliefs are often associated with low stated beliefs. However, in order to rationalize our data in this way, it would require that we observe more high belief statements than we do.

Unbiasedness. 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 4). 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 $\mathbb{P}(b_s^* > \hat{\sigma}_U) \leq \frac{3}{4}$,[23] 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*.

Result 6. *The results are robust to measurement error.* Assuming that actions depend on unobservable true belief realizations of which stated beliefs are noisy signals, we conclude that true beliefs are noisy and that *monotonicity* and *unbiasedness* are rejected with respect to true beliefs.

9 Conclusion

Equilibrium models with stochastic elements have had great success in explaining experimental data. While some such models have incorporated noisy beliefs, models with noisy actions have been much more prominent, and the empirical relevance of noisy beliefs has been little explored.

Our experiment shows that beliefs are, in fact, noisy. Moreover, belief-noise matters: it is similarly important as action-noise for explaining our data. These results suggest that

[23] 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 $\mathbb{P}(b_s^* > \hat{\sigma}_U) = \frac{3}{4}$.

ignoring belief-noise could lead to biased parameter estimates and poor out-of-sample predictions.

In addition to documenting belief-noise, we have uncovered a number of regularties in the failure of basic postulates about the relationships between beliefs and actions. Taken together, we hope our results stimulate the development of models with noisy beliefs.

References

Alberto Abadie. Bootstrap tests for distributional treatment effects in instrumental variable models. *Journal of the American Statistical Association*, 2002. 6

Marina Agranov and Pietro Ortoleva. Stochastic choice and preferences for randomization. *Journal of Political Economy*, 2017. 1

Marina Agranov, Paul Healy, and Kirby Nielsen. Stable randomization. *Working paper*, 2020. 1

Hsiao-Chi Chen, James W. Friedman, and Jacques-Francois Thisse. Boundedly rational nash equilibrium: A probabilistic choice approach. *Games and Economic Behavior*, 1997. 4

Miguel Costa-Gomes and Georg Weizsacker. Stated beliefs and play in normal-form games. *The Review of Economic Studies*, 2008. 1, 2, 6.4

Vincent Crawford, Miguel Costa-Gomes, and Nagore Iriberri. Structural models of nonequilibrium strategic thinking: Theory, evidence, and applications. *Journal of Economic Literature*, 2013. 3

Evan Friedman. Endogenous quantal response equilibrium. *Games and Economic Behavior*, 2020a. 1

Evan Friedman. Stochastic equilibria: Noise in actions or beliefs? *Working Paper*, 2020b. 1, 1, 2.2, 2.3, 10.1, 10.3

James Friedman and Claudio Mezzetti. Random belief equilibrium in normal form games. *Games and Economic Behavior*, 2005. 4

Jacob Goeree and Charles Holt. Ten little treasures of game theory and ten intuitive contradictions. *American Economic Review*, 2001. 16

Jacob Goeree and Charles Holt. Learning in economics experiments. *Encyclopedia of Cognitive Science*, 2002. 4

Jacob Goeree and Philippos Louis. M equilibrium. *Working paper*, 2018. 1

Jacob Goeree, Charles Holt, and Thomas Palfrey. Regular quantal response equilibrium. *Experimental Economics*, 2005. 1, 1, 2, 5, 2.1, 10.1, 10.3

Duarte Goncalves. Sequential sampling and equilibrium. *Working paper*, 2020. 4

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. 1, 5

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. 13

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, 6.4

Kyle Hyndman, Antoine Terracol, and Jonathan Vaksman. Beliefs and (in)stability in normal-form games. *Working paper*, 2013. 6

Edi Karni. A mechanism for eliciting probabilities. *Econometrica*, 2009. 3.2, 10.5

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, 1

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. 4.1

Emerson Melo, Kiril Pogorelskiy, and Matthew Shum. Testing the quantal response hypothesis. *International Economic Review*, 2018. 1

Rosemarie Nagel. Unraveling in guessing games: an experimental study. *American Economic Review*, 1995. 3, 10.9

Yaw Nyarko and Andrew Schotter. An experimental study of belief learning using elicited beliefs. *Econometrica*, 2002. 1, 4.3, 4.3, 5

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

Pedro Rey-Biel. Equilibrium play and best response to (stated) beliefs in normal form games. *Games and Economic Behavior*, 2009. 1, 2

Ariel Rubinstein and Martin Osborne. Sampling equilibrium, with an application to strategic voting. *Games and Economic Behavior*, 2003. 4

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

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

Andrew Schotter and Isabel Trevino. Belief elicitation in the laboratory. *Annual Review of Economics*, 2014. 1, 17, 4.3

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. 3, 10.9

Amos Tversky. Physical review. *Intransitivity of Preferences*, 1969. 1, 3.3

Roberto Weber. 'learning' with no feedback in a competitive guessing game. *Games and Economic Behavior*, 2003. 2, 3.3, 5

Irenaeus Wolff and Dominik Bauer. Elusive beliefs: Why uncertainty leads to stochastic choice and errors. *Working paper*, 2018. 1

10 Appendix

10.1 Proofs

Proof of Proposition 1. The *only if* direction can be found for very similar games for QRE in Goeree et al. [2005] and for NBE in Friedman [2020b] 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 14. 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 2. 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 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. The *only if* direction follows from Propositions 3 and the following comparative static in X. 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. We omit the *if* direction because it is very similar to that in the proof of Proposition 1 as it basically combines the results for QRE and NBE.

□

Proof of Proposition 5. *(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 $\tilde{\sigma}_L^{v,X} = \frac{20}{20+X} = \sigma_L^{NE,X}$.

□

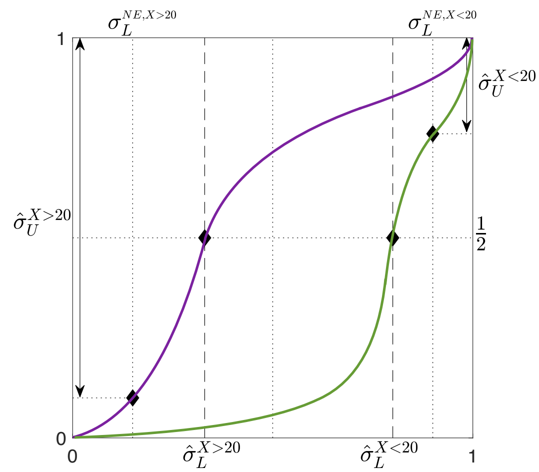

Figure 14: *Construction of belief-map in the proof of Proposition 1.* This illustrates the constructed CDFs of player 1's beliefs to rationalize as NBE the actions dataset $\{\hat{\sigma}_U^X, \hat{\sigma}_L^X\}_X$ for $X \in \{X', X''\}$ with $X' > 20$ and $X'' < 20$. The purple CDF is $F_1(\cdot|\hat{\sigma}_L^{X'})$ and the green CDF is $F_1(\cdot|\hat{\sigma}_L^{X''})$.

10.2 The hybrid model

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 $\Psi = (\Psi_1, \Psi_2) : [0,1]^2 \to [0,1]^2$. A quantal response-noisy belief equilibrium (QNBE) is defined as a mixed action profile that is a fixed point along with the supporting belief distributions.

Definition 3. Fix $\{\Gamma^{2\times 2}, Q, \sigma^*\}$. A QNBE is any pair $\{\sigma, \sigma^*(\sigma)\}$ where $\sigma = \Psi(\sigma; Q, \sigma^*)$.

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

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

Proof. See Appendix 10.1. □

10.3 X-games

As with QRE and NBE, the QNBE for any X-game (or, more generally, any 2×2 game with a unique mixed strategy NE) is unique for fixed primitives. However, 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 10.1. □

Figure 15 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

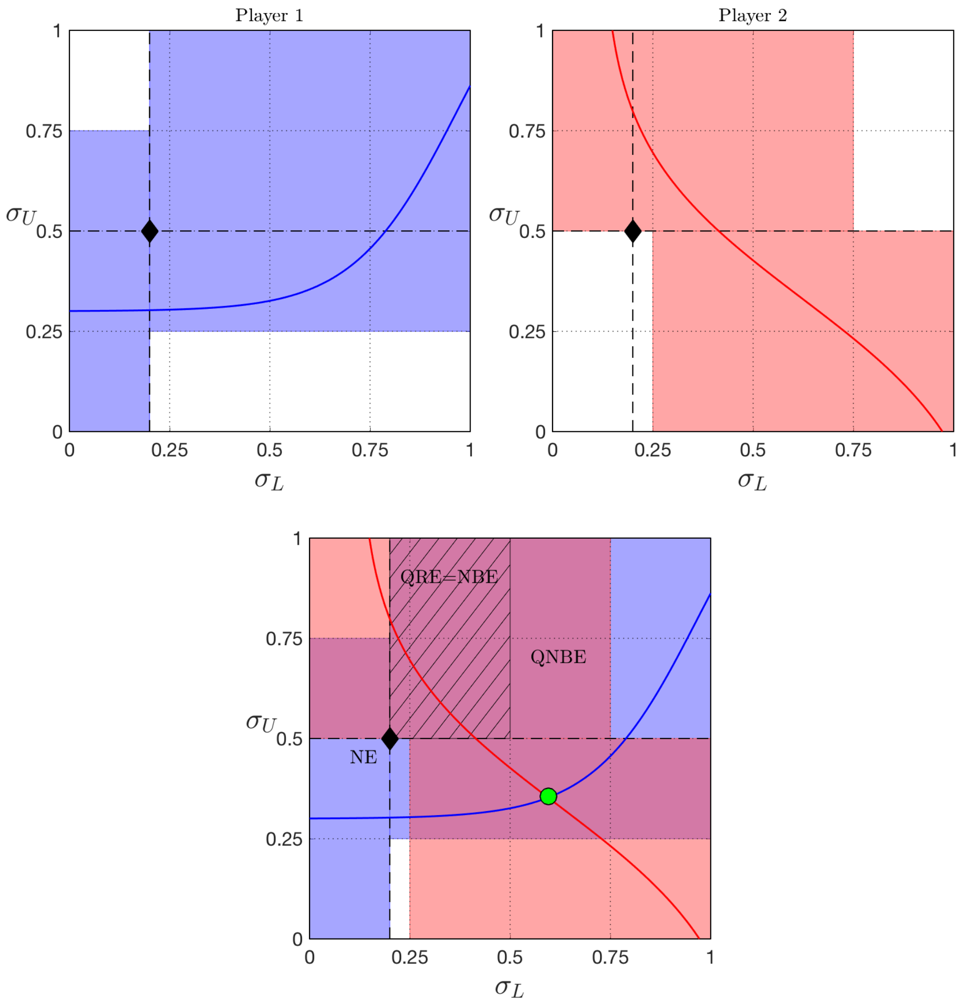

Figure 15: *QNBE in game* $X = 80$. The first panel gives the region in which player 1's QNBE reaction must lie, with an example drawn in blue. The second panel gives the region in which player 2's QNBE reaction must lie, with an example drawn in red. The third panel plots the intersection of the two 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.

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 15, 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.

As previously mentioned, Friedman [2020b] shows that the sets of attainable QRE and NBE mixed action profiles coincide for any X-game. In Figure 15, for $X = 80$, we plot this set as a cross-hatched rectangle, which has a measure of 15% (see Goeree et al. [2005] and Friedman [2020b] 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.

We extend Proposition 1 to a characterization of QNBE for any finite number of X-games.

Proposition 4. *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),

(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 10.1. □

Hence, a dataset of mixed actions can be supported as QNBE outcomes for a single set of primitives if and only if the data from each game can be rationalized as QNBE outcomes individually *and* a particular comparative static in X holds. This comparative static is the same that holds for QRE and NBE, and hence, the experiment is predicted to give the desired variation in actions and beliefs if there is noise in actions, beliefs, or both.

10.4 Games and randomization

In addition to the X-games, we also included the games whose payoffs are in Table 8. $D1$ and $D2$ are dominance solvable games, which are identical up to which player faces which set of payoffs: in Di, it is player i who has a dominant action. $X80s$ ("s" for "scale") is the same as $X80$, except with all payoffs divided by 10. $R1$ and $R2$ are similar to $X5$, except the symmetry of player 2's payoffs have been broken.

Table 9 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. 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.

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.

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 8: *Additional games.*

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

Table 9: *Games by section.*

10.5 Eliciting beliefs using random binary choice

We used the random binary choice (RBC) mechanism (Karni [2009]), also known as the "lottery method", to incentivize subjects to state their beliefs accurately. In an RBC, subjects are asked which option they prefer from a list of 101 binary choices, as in Table 10 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-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 10: 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 10.10 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.

10.6 The effects of belief elicitation

In the top panels of Figure 16, 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.[24] For both players, we observe stastistically significant differences between the two stages. This is confirmed in the first 2 columns of Table 11 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 Figure 16 plot the action frequencies from [A,A] and [A,A], and Columns 3-4 of Appendix Table 11 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.

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

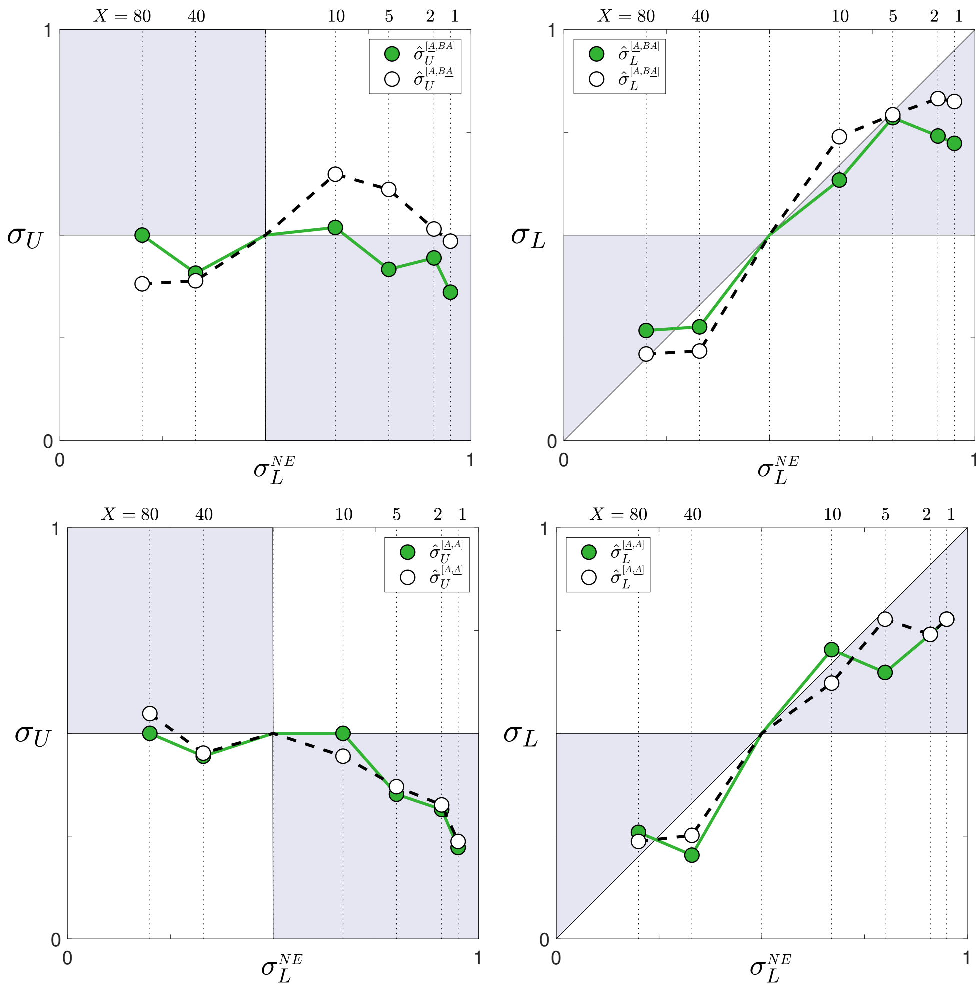

Figure 16: *Effects of belief elicitation.* The top panels plot first-stage and second-stage actions from [A,BA], and shows a systematic difference between the two stages. The bottom panels plot first-stage and second-stage actions from [A,A], and shows no difference between the stages.

	[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 11: *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.

10.7 Random utility estimation

The data of subject s in role i is a set of 30 action-belief pairs $\{\hat{a}^{iX}_{sl}, \hat{b}^{iX}_{sl}\}_{lX}$ 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 by adding a constant $\epsilon > 0$ (arbitrarily pre-set to 0.001) to each payoff. We also normalized utility so that it is between 0 and 1:[25]

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

This utility function induces, for each game X and stated belief $\hat{b}^{iX}_{sl}$, a vector of expected utilities $\bar{w}^X_i(\hat{b}^{iX}_{sl}; \rho) = (\bar{w}^x_{i1}(\hat{b}^{iX}_{sl}; \rho), \bar{w}^X_{i2}(\hat{b}^{iX}_{sl}; \rho))$. 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 $\mu_a > 0$:[26]

$$p^X(\hat{a}^{iX}_{sl} | \hat{b}^{iX}_{sl}; \rho, \mu_a) = \frac{\bar{w}^X_{i1}(\hat{b}^{iX}_{sl}; \rho)^{\frac{1}{\mu_a}}}{\bar{w}^X_{i1}(\hat{b}^{iX}_{sl}; \rho)^{\frac{1}{\mu_a}} + \bar{w}^X_{i2}(\hat{b}^{iX}_{sl}; \rho)^{\frac{1}{\mu_a}}}. \tag{3}$$

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

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

For player 2-subjects, who face symmetric payoffs, ρ and μ_a are not separately identified, and so we set $\rho = 0$ prior to estimation.

10.8 Explaining the failure of monotonicity

We observe, for player 1 only, a systematic failure of *monotonicity*, as shown in Figure 8. Under the maintained assumption of expected utility, utility is linear in matrix payoffs since they are in probability points. *Monotonicity* thus predicts that $\hat{Q}_U$ should cross the one-half line at the indifferent belief plotted in the figure (dashed vertical lines). However, we

[25] By construction, $w(0; \rho) = 0$ and $w(80; \rho) = 1$.

[26] The Luce rule (3) 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.

observe systematic failures: for $X > 20$, $\hat{Q}_U$ crosses to the right of the indifferent belief, and for $X < 20$, $\hat{Q}_U$ crosses to the left of the indifferent belief. If, however, there is non-linearity in the utility function over matrix payoffs, the actual indifferent belief–and thus where *monotonicity* predicts $\hat{Q}_U$ crosses the one-half line–may deviate systematically from that under linear utility. The proposition below states that, with concavity in the utility function, the indifferent belief moves right for $X > 20$ and left for $X < 20$, which is consistent with the observed violations.

Proposition 5. *Let w and v be any strictly increasing Bernoulli utility functions. For player 1 in game X, w induces expected utility vectors $\bar{w}^X = (\bar{w}_U^X, \bar{w}_D^X) : [0,1] \to \mathbb{R}^2$. 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 10.1. □

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 the random utility model with curvature from Section 10.7 to each player 1-subject's actions data given belief statements. 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 likelihood $L(\hat{a}|\hat{b}; \rho, \mu_a) = \sum_s L^s(\hat{a}|\hat{b}; \rho, \mu_a)$. We find the estimate $\hat{\rho} = 0.87$, indicating concavity. In Figure 17, we reproduce Figure 8 for player 1, but we now plot the indifferent beliefs implied by the best-fit utility function as solid vertical lines. Each such line intersects $\hat{Q}_U$ near to where it crosses the horizontal one-half line. Hence, if the subjects admitted a representative agent with this concave utility, nearly all of the *monotonicity* violations would disappear. This also captures the fact that the regions of violation are larger for the more asymmetric games (compare, for example, $X10$ and $X1$ in Figure 17).

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

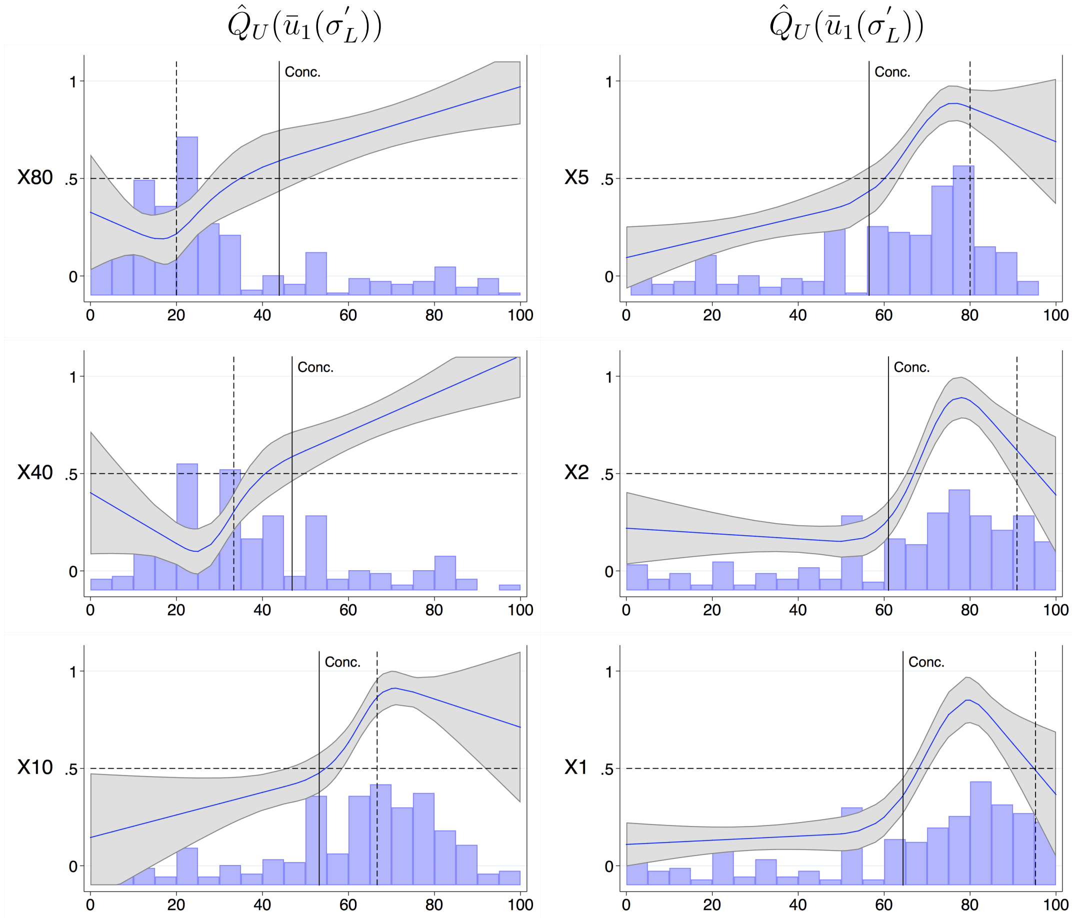

Figure 17: *Concave utility explains monotonicity failures.* For player 1 and each of the X-games, we plot the predicted values (with 90% error bands) from restricted cubic spline regressions of actions on beliefs (4 knots at belief quintiles, standard errors clustered by subject). Belief histograms 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.

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.

10.9 Explaining the failure of 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?

A complete answer is well beyond the scope of this paper, but to get at this question, we report on two accidental discoveries. Taken together, we interpret these as supporting the conclusion that the player role itself causally induces greater sophistication, in the level k sense, in player 1-subjects who face asymmetric payoffs. It is not hard to see how this can generate the observed pattern of bias: 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 that player 2-subjects think this way.

The first discovery is that player 1-subjects have longer response times, as shown in Figure 18, which suggests they think more deeply than player 2-subjects. This suggests that player 1-subjects, by merit of facing asymmetric payoffs, are somehow induced to think more deeply in the X-games than player 2-subjects. This may be related to greater sophistication, but is difficult to interpret.

The second discovery, which can be more directly interpreted in terms of strategic sophistication, is based on stated beliefs in the small number of dominance solvable games (which were included to break up appearance of the X-games; see Section 3.3). The dominance solvable games are reproduced in Figure 19. $D1$ and $D2$ 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 actions is the best response to a uniform distribution and the other is the best response to i's dominant action.

Importantly, by symmetry, player 1-subjects' stated beliefs in $D1$ ($D2$) are fully comparable to player 2-subjects' stated beliefs in $D2$ ($D1$). Also, since all subjects observed exactly the same games throughout the experiment and were randomly assigned to their

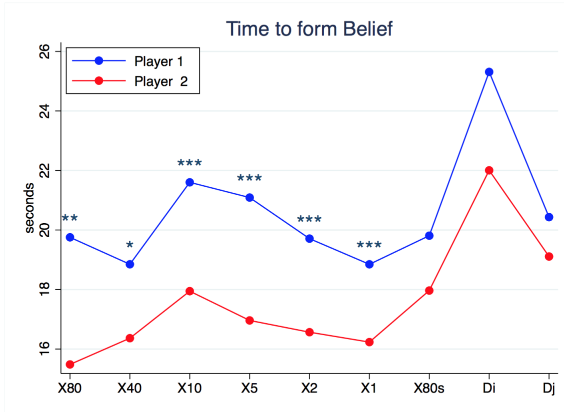

Figure 18: *Average time to form beliefs by game and player.* We plot the average time until stated beliefs are finalized by game and player role.

roles, any differences across player 1- and player 2-subjects in Di or Dj must be caused by their experiences in different roles of the X-games throughout the experiment. In other words, something "spills over" from the asymmetric X-games that is reflected in these neutral environments.[27]

$D1$	$\mathbf{L}_{k\geq 2}$	$\mathbf{R}_{k=1}$
U	6, 0	0, 20
$\mathbf{D}_{k\geq 1}$	8, 20	20, 4

$D2$	**L**	$\mathbf{R}_{k\geq 1}$
$\mathbf{U}_{k=1}$	20, 0	4, 20
$\mathbf{D}_{k\geq 2}$	0, 6	20, 8

Figure 19: *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 ($k \geq 2$).

In the level k framework of strategic sophistication (e.g. Nagel [1995] and Stahl and Wilson [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 game Di, the following characterizes level-types. 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

[27] One concern is that, since experience in the 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.

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.

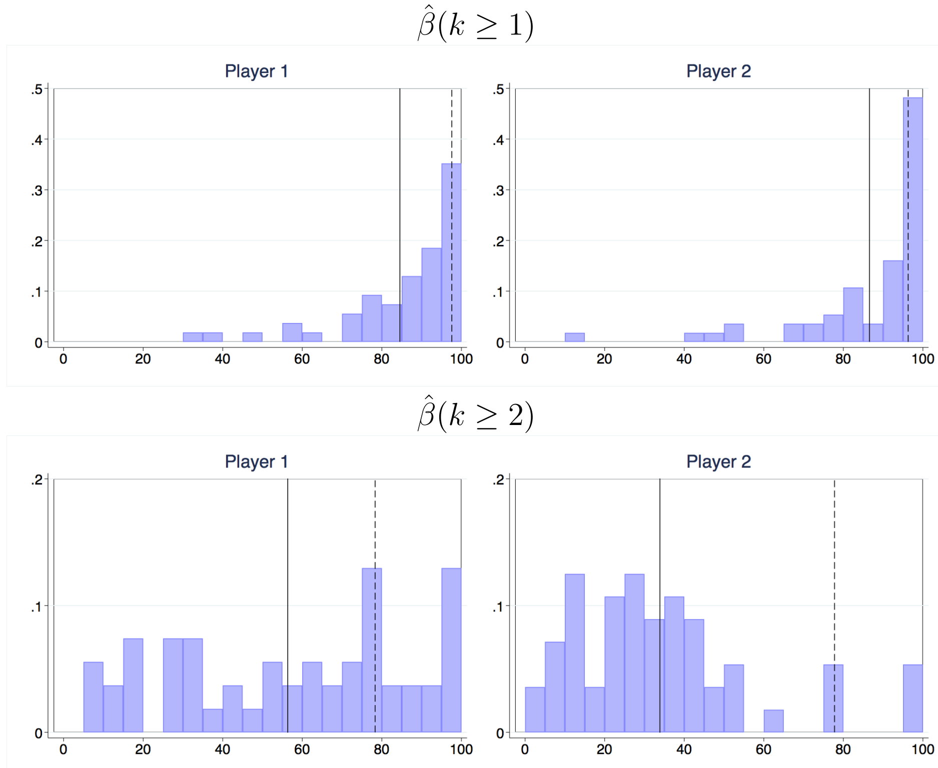

Figure 20: *Sophistication by player.* The top panel gives histograms of $\hat{\beta}(k \geq 1)$, i's belief that j best responds to his dominant action in Dj, across subjects. The bottom panel gives histograms of $\hat{\beta}(k \geq 2)$, i's belief that j best responds to i's dominant action in Di (as 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 $[\underline{A}, \circ]$.

From the top panel of Figure 20, 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). From the bottom panel, however, 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). Hence, by this measure, player 1 appears to be much more sophisticated than player 2.

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, 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.

10.10 Experimental interface

Figure 21 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.

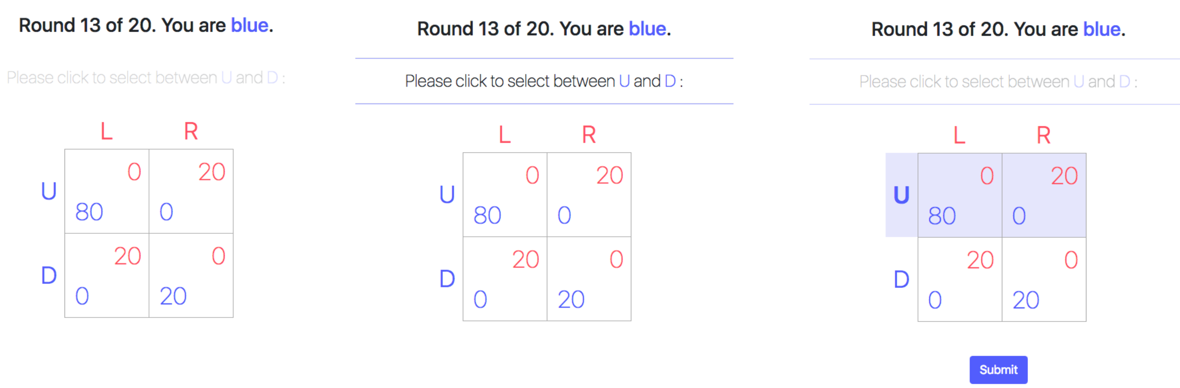

Figure 21: *Screenshots from first stage.*

Figure 23 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

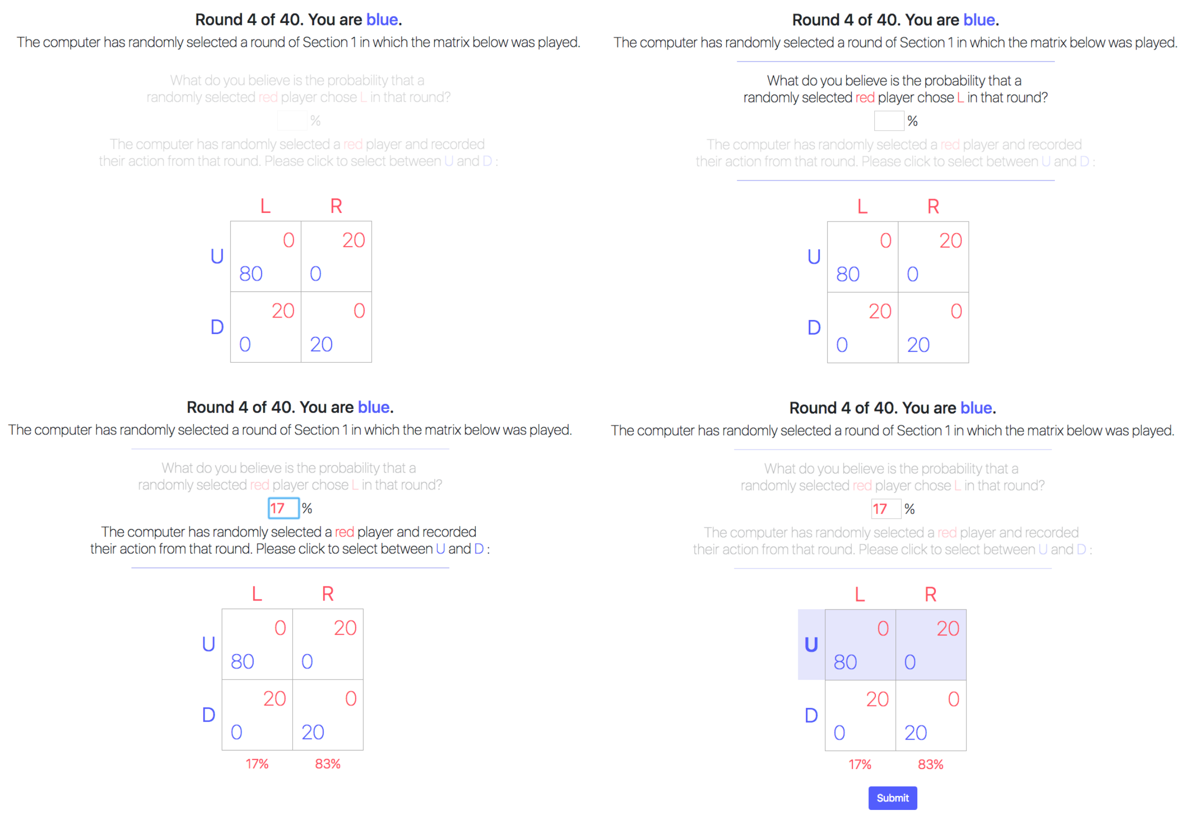

Figure 22: *Screenshots from second stage of [A,BA].*

and action before submitting. After a belief is entered and an action is selected, the 'Submit' button appears (bottom-right screen). Figure 22 shows screenshots for the second stage of [A,A] which is the same as that of [A,BA], except beliefs are not elicited.

10.11 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

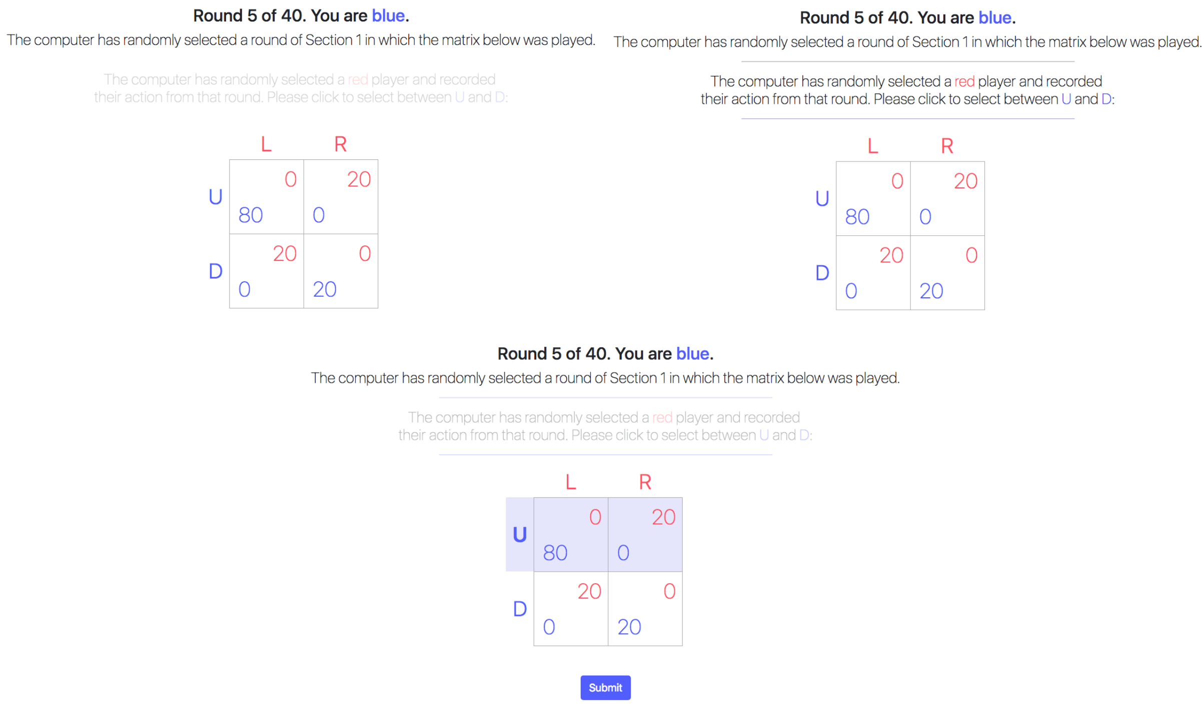

Figure 23: *Screenshots from second stage of [A,A].*

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.

10.12 Additional Figures

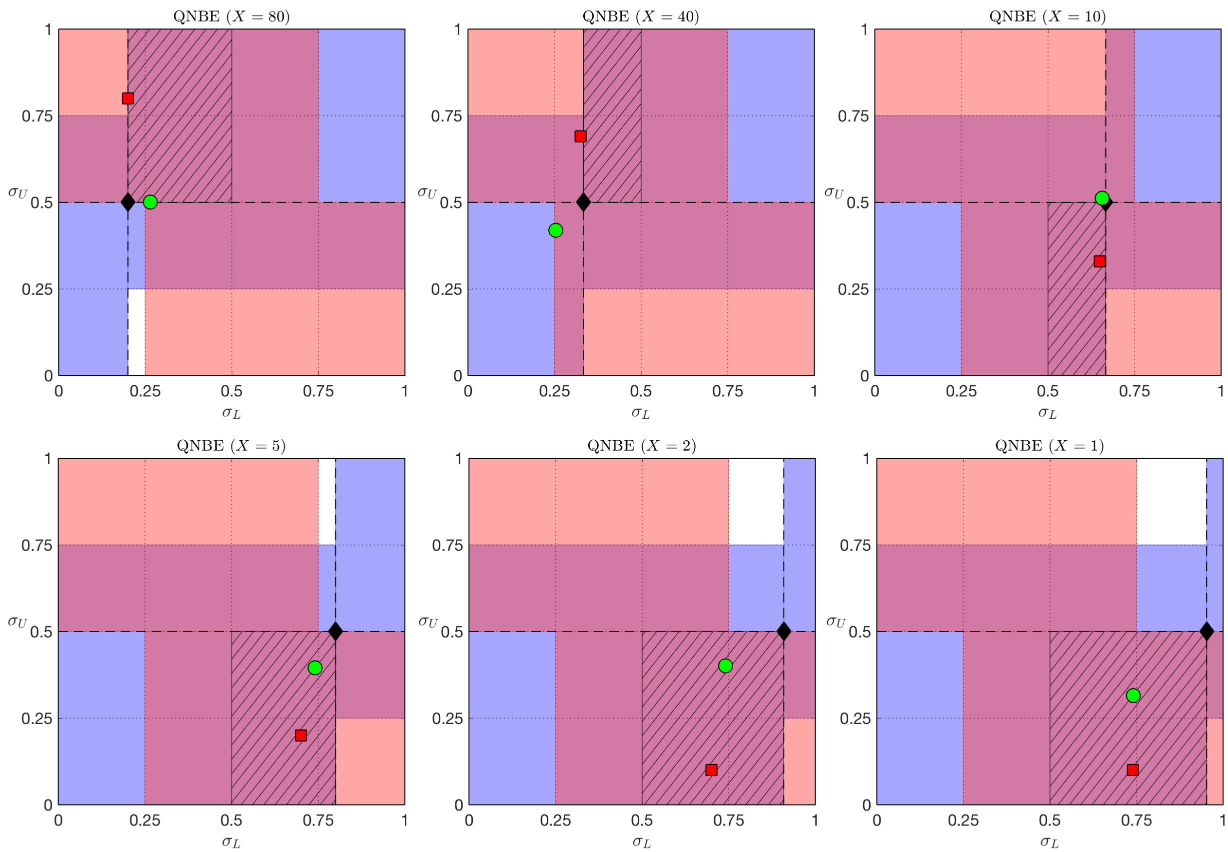

Figure 24: *QNBE and the data.* The green dot gives the empirical action frequencies from [A, ∘], the red square gives the median belief, and the black diamond is the Nash equilibrium.

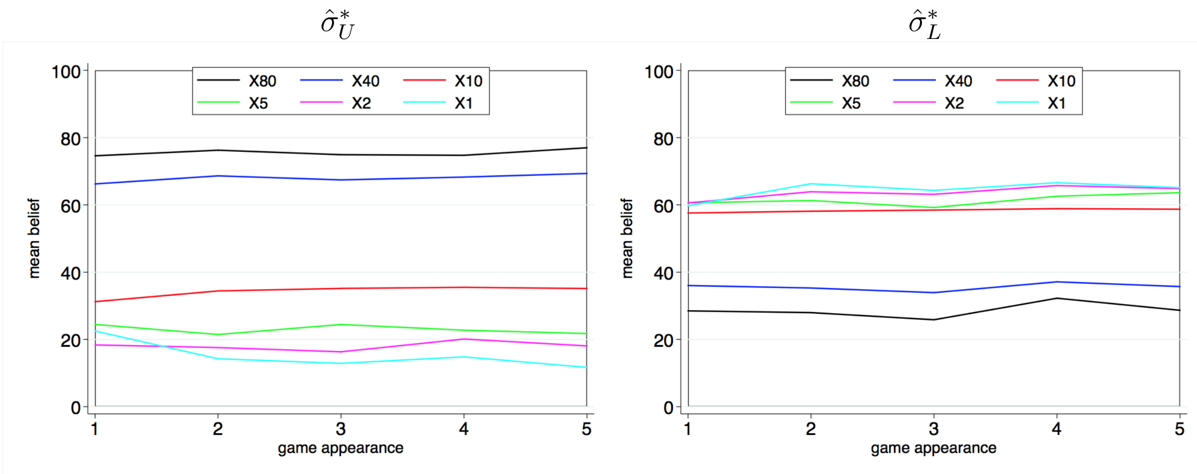

Figure 25: *Stability of average beliefs throughout the experiment.* For each game and player role, this figure plots average beliefs for each of five appearances of the game throughout the experiment.

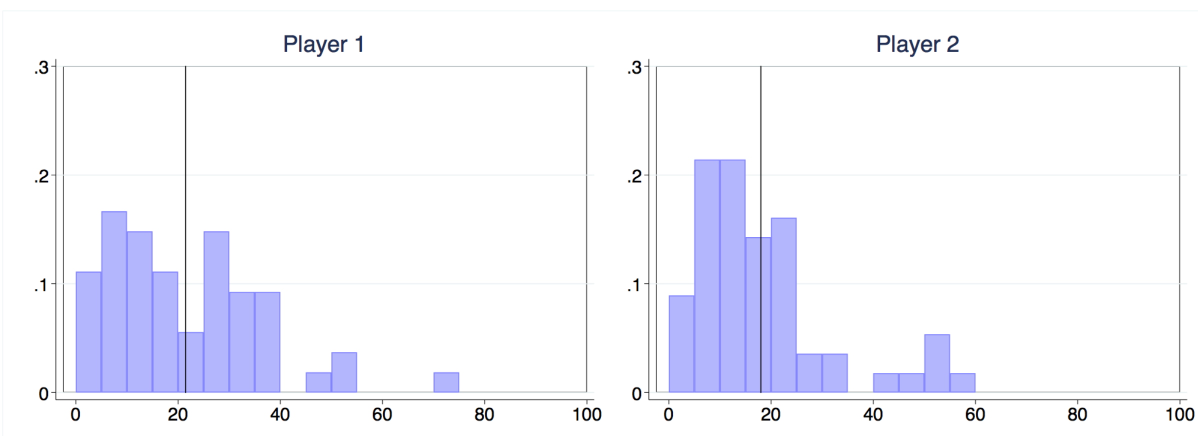

Figure 26: *Subjects' spreads of detrended beliefs.* This figure gives histograms of subjects' spreads of detrended beliefs, averaged across all X-games.

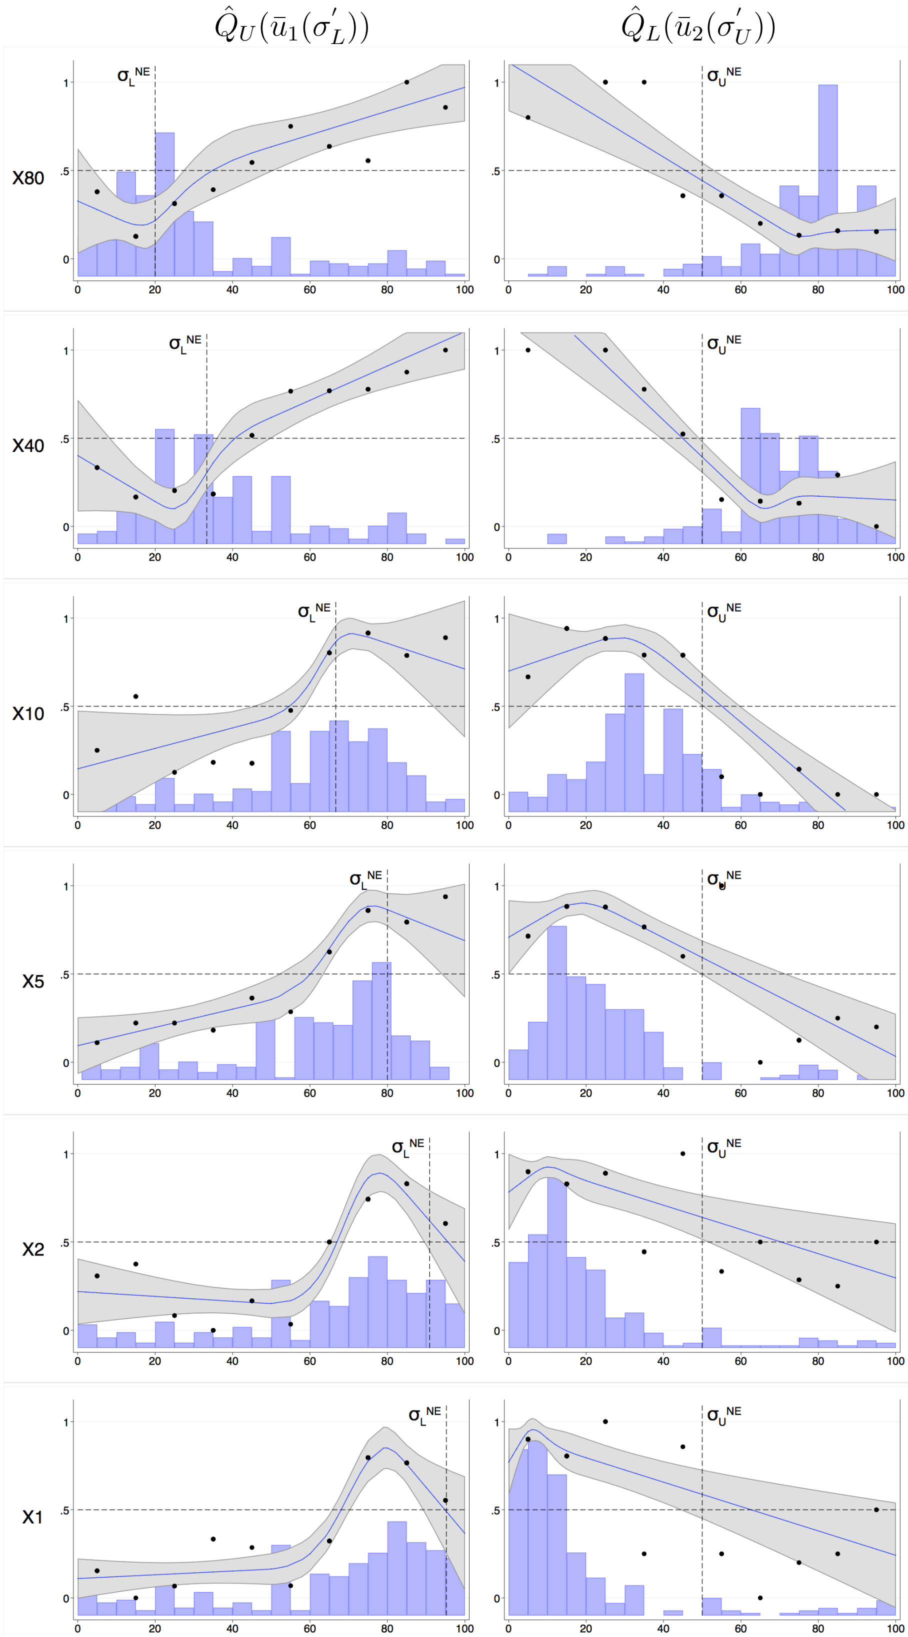

Figure 27: *Action frequencies predicted by beliefs.*

10.13 Additional tables

		$X80$	$X40$	$X10$	$X5$	$X2$	$X1$
actions	$\hat{\sigma}_U^{[\underline{\text{A}},\circ]}$	0.50	0.42	0.51	0.40	0.40	0.31
	$\hat{\sigma}_L^{[\underline{\text{A}},\circ]}$	0.27	0.25	0.66	0.74	0.74	0.74
	$\hat{\sigma}_U^{[\text{A},\text{B}\underline{\text{A}}]}$	0.38	0.39	0.65	0.61	0.51	0.49
	$\hat{\sigma}_L^{[\text{A},\text{B}\underline{\text{A}}]}$	0.21	0.22	0.74	0.78	0.83	0.82
beliefs	$med(\hat{\sigma}_U^*)$	0.80	0.69	0.33	0.20	0.10	0.10
	$med(\hat{\sigma}_L^*)$	0.20	0.33	0.65	0.70	0.70	0.74
	$mean(\hat{\sigma}_U^*)$	0.76	0.68	0.34	0.23	0.18	0.15
	$mean(\hat{\sigma}_L^*)$	0.29	0.36	0.58	0.61	0.64	0.64

Table 12: *Empirical action frequencies, median and mean belief statements.*

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 13: *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.

	(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 14: *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 15: *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.