Stochastic Choice and Noisy Beliefs in Games[*]

Evan Friedman[†]and Jeremy Ward

First version: 15th November 2019
Current version: 29th September 2024

Abstract

We conduct an experiment in which we elicit subjects' beliefs over opponents' behavior multiple times for a given game without feedback. A large majority of subjects have stochasticity in their belief reports, which we argue cannot be explained by learning or measurement error, suggesting significant noise in subjects' unobserved "true" beliefs. Using a structural model applied to actions and beliefs data jointly, we find that such "noisy beliefs" are equally important for explaining our data as "noisy actions"—the sort of stochastic choice given fixed beliefs that is commonly assumed in empirical research. We then test the axioms underlying equilibrium models with noisy actions (quantal response equilibrium) and noisy beliefs (noisy belief equilibrium). We find support for both sets of axioms, except for those that assume beliefs are perfectly unbiased. To fully explain our data, we argue that beliefs and belief-noise are driven by the payoff salience of actions.

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

[*]We thank Alessandra Casella, Yeon-Koo Che, Mark Dean, Ambuj Dewan, Teresa Esteban-Casanelles, Duarte Gonçalves, Navin Kartik, RC Lim, Suanna Oh, Tom Palfrey, Jacopo Perego, Mike Woodford, and seminar participants at Columbia, Hebrew, Purdue, University College London, Queen Mary, Essex, CERGE-EI, Queensland, Paris School of Economics, Econometric Society-NAWM 2021, the 2021 Korean Economic Review International conference, and the Arne Ryde Workshop on Attention in Decision Making for helpful comments. This material is based upon work supported by the National Science Foundation under Grant Number 1730399.

[†]Email: evan.friedman@psemail.eu

A large literature has documented a pattern of stochastic choice in individual decision making. In games, in which payoffs depend on beliefs over opponents' behavior, a potentially important source of stochasticity is in the beliefs themselves. This suggests a decomposition of stochastic choice into (1) "noisy beliefs" and (2) "noisy actions" *conditional on* beliefs. A priori, it is unclear which type of noise is the primary driver of stochastic choice, and yet, little is known of the empirical relevance of noisy beliefs.

Modeling noise appropriately is crucial for making good predictions, particularly out-of sample, as well as for estimating structural parameters. This may be especially true in games, where the interaction of noisy players can lead to equilibrium effects. In particular, noisy actions and noisy beliefs are fundamentally different in their equilibrium implications (Friedman [2022]), so it is important to understand the nature and prevalence of the two sources of noise.

In this paper, we report on the results of a laboratory experiment designed to answer three related questions. First, are beliefs noisy? Second, what is the relative importance of noisy actions and noisy beliefs for explaining data? Third, how valid are the assumptions underlying existing models of stochastic choice in games? These questions concern the quantitative importance and qualitative nature of the two types of noise.

To guide the experimental design, we formulate as a benchmark model a generalization of Nash equilibrium (NE) that allows for both stochastic choice given beliefs and randomness in the beliefs themselves.[1] 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 beliefs as a function of the opponents' mixed actions. Without parametric assumptions, the action- and belief-maps are only restricted to satisfy the axioms of (regular) quantal response equilibrium (QRE) (Goeree et al. [2005]) and noisy belief equilibrium (NBE) (Friedman [2022]), respectively. The model is thus a hybrid model, borrowing from existing models of stochastic choice.

Guided by the hybrid model, the idea behind our experiment is to make observable the empirical action- and belief-maps. These mappings are of interest because they summarize the empirical relationships between actions and beliefs, allowing us to test basic postulates about the nature and importance of the two sources of noise. Additionally, as model primitives, they have a close connection to existing theories.

[1]The model will therefore not be trivially rejected with a single failure of best response or instance of incorrect beliefs.

Noting that both mappings depend on both actions *and* beliefs, identification requires that we identify beliefs in addition to observing actions. In our experiment, we achieve this by *eliciting beliefs directly*.

Identification also requires variation in both actions and beliefs. We achieve this through the family of 2×2 games whose payoffs are in Table 1. Indexed by player 1's payoff parameter $X > 0$, these X*-games* have unique, mixed strategy NE. By varying X, the hybrid model predicts systematic variation in action- and belief-distributions for both players. Hence, by having subjects state beliefs and take actions for different values of X, the model predicts that we will observe multiple points on the empirical action- and belief-maps. In other words, by varying X, we "trace out" these mappings.

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

Table 1: *Game* X. Player 1's payoff parameter $X > 0$ controls the asymmetry of payoffs and is our instrument for varying actions and beliefs.

Prior to previewing our results, we make two essential remarks. First, by "noisy 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 so that there is no scope for learning.[2]

Second, we view *stated beliefs* simply as measures of the unobservable *true beliefs* that subjects hold in their minds and guide their actions. The issue we face is that while we are primarily interested in true beliefs, we only observe stated beliefs that may be subject to random measurement error. We proceed by first taking stated beliefs as equal to true beliefs for our main analysis. Then, we show that our qualitative results—including the existence of noisy beliefs—hold for true beliefs under the weaker (non-parametric) assumption that a stated belief is an unbiased signal of the underlying 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 structural

[2]The selection of games too is such that we would not expect the phenomenon of *no-feedback learning*. Weber [2003] documents, using dominance solvable beauty contest games, that subjects' behavior drifts toward the NE, even without feedback. The X-games, on the other hand, are fully mixed and comparatively simple in structure. In our data, we also find no trends in beliefs (Section 4).

model, that belief-noise is equally important as action-noise for explaining our data. Third, we study the validity of assumptions underlying models of stochastic choice by way of statistically testing the axioms of the hybrid model. We find support for most, but not all, of its axioms; and we offer possible explanations for those that are rejected.

The conclusion that beliefs are noisy is based on within-subject 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. This cannot be driven by learning since there is no feedback, most subjects have no trends in their beliefs, and the 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 subjects' actions. This implies that a high stated belief signals a high true belief, and thus noise in stated beliefs reflects noise in true beliefs.

To conclude that belief-noise is as important as action-noise for explaining the data, we formulate a simple structural model and 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 for that subject if turning it off leads to large prediction errors relative to her observed action frequency. We find that the two counterfactuals perform equally well on average. This suggests that belief-noise, while rarely considered, is just as important as the sort of action-noise commonly assumed in the analysis of experimental data (e.g. in QRE).

In testing the axioms of the hybrid model, we find broad support for all axioms, except for those that assume beliefs are perfectly unbiased. In particular, we find that: (1) for any given belief, actions with higher expected utility are played more often (consistent with *monotonicity*); (2) if the expected utility to some action increases (while keeping other expected utilities fixed), that action will be played with a higher probability (*responsiveness*); and (3) the distribution of beliefs shifts in the same direction as changes in the opponent's behavior (*belief-responsiveness*). However, we also find that (4) the central tendency of the belief-distribution is biased relative to the opponents' behavior (a failure of *unbiasedness*). Hence, while subjects correctly anticipate how behavior changes across games, beliefs are not well-calibrated.

In addition to the failure of *unbiasedness*, we also document that belief-noise is systematically higher for player 1-subjects and decreasing in X for both players. To explain these findings, we propose that the differential *salience* of actions across players *causally induces greater strategic sophistication in player 1-subjects*, i.e. those who have a salient action. This is both capable of explaining the observed patterns and consistent with auxiliary data. In particular, we find that player 1-subjects exhibit much higher levels of strategic sophistication based on a measure that is collected identically for all subjects toward the end of the experiment. Moreover, since subjects were randomized into player roles (that they maintain throughout the experiment), we establish causality: experience in different roles of the X-games has a causal effect on measures of strategic sophistication.

Our results have important implications for modeling stochastic choice in games. Broadly, while beliefs tend to be biased and therefore cannot be fully explained by reduced-form equilibrium models, our results confirm the basic premise of models with noisy actions (e.g., quantal response equilibrium) and noisy beliefs (e.g., noisy belief equilibrium). In particular, we establish through a structural model that belief-noise is indispensable. This suggests that ignoring belief-noise, as is done in a majority of model-fitting applications, can lead to biased parameter estimates and poor out-of-sample predictions (Section 5 provides further discussion). However, we show that allowing for both belief- and action-noise can lead to models that are very permissive in their predictions over standard actions data. Beyond suggesting the value of belief elicitation, this calls for additional research into the factors that drive the two types of noise. In our data, one prominent factor emerges. The payoff salience of actions, intermediated by a causal effect of salience on strategic sophistication, is a primary driver of noisy beliefs.

The remainder of this section reviews related literature. Section 1 presents the theory; Section 2 gives the experimental design; Section 3 provides an overview of the data; Section 4 presents evidence that beliefs are noisy; Section 5 quantifies the relative importance of action- and belief-noise for explaining the data; Section 6 presents the results from testing the axioms; Section 8 establishes the role of payoff silence in driving belief-noise; Section 7 shows that our results, including the existence of noisy beliefs, are robust to measurement error; and Section 9 provides a discussion.

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 [2022]), which are closely related to other models found in the literature.[3] We believe we are

[3]Models with noisy beliefs include those of Friedman and Mezzetti [2005], Rubinstein and Osborne [2003], and Goncalves [2020], among others. 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

novel in studying a hybrid model that allows for noise in both actions and beliefs, but 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 [2021], and Friedman and Mauersberger [2023]). While the focus of our paper is testing axioms using actions and beliefs data jointly, we do show that the hybrid model, despite being fairly permissive, is falsifiable in standard actions data.[4]

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. Agranov and Ortoleva [2017], Agranov et al. [2020], and Tversky [1969]). As we discuss in Section 9.1, our focus on documenting noise in beliefs leads us to a novel experimental design. We do not attempt to distinguish between different theories of noisy beliefs, but we note that the noise may arise from sequential evidence accumulation (e.g. Fudenberg et al. [2018]), sampling from memory (e.g. Goncalves [2020]), cognitive uncertainty (Enke and Graeber [2020]), or as a reaction to strategic uncertainty (e.g. Wolff and Bauer [2018]). Hence, we provide evidence consistent with these theories.

We also contribute to the general understanding of elicited beliefs, their interpretation, and their relationship to subjects' behavior.[5] Our key innovation is to collect multiple elicitations per subject without feedback, which is necessary to study noisy beliefs. This distinguishes us from experiments that elicit beliefs once for each game (e.g. Costa-Gomes and Weizsacker [2008], Rey-Biel [2009], and Ivanov [2011]) 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 the belief elicitation literature has focused on documenting rates of best response, the axioms we consider lead us to study how rates of best response vary across every neighborhood of stated beliefs. In particular, we show that, within-subject, the probability of taking an action varies with changes in stated beliefs—even when restricting attention to neighborhoods of beliefs that imply the same best response. Finally, our non-parametric treatment of measurement error may be broadly applicable to other settings where the assumption that "stated beliefs equal true beliefs" is problematic.

testable restrictions.

[4] Haile et al. [2008] showed that *structural* QRE can rationalize the data from any one game without strong restrictions on the error distributions, raising concerns that QRE models may be non-falsifiable. However, *regular* QRE (Goeree et al. [2005]), as well as the hybrid model that generalizes it, *is* falsifiable.

[5] Schotter and Trevino [2014] and Schlag et al. [2015] provide excellent review articles of belief elicitation methods and related applications.

1 Theory to guide the experimental design

Our primary goal is not to test any particular model, but to study noisy actions, noisy beliefs, and the contribution of each to stochastic choice. Nevertheless, we introduce a benchmark model to ground our work and guide the experimental design.

The benchmark model is a stochastic generalization of Nash equilibrium (NE). It is a hybrid model, defined by an action-map satisfying the axioms of regular quantal response equilibrium (QRE) (Goeree et al. [2005]) and a belief-map satisfying the axioms of noisy belief equilibrium (NBE) (Friedman [2022]). Our main takeaway: as long as there is noise in actions, beliefs, or both, the hybrid model predicts that the X-games that we use in the experiment will give rise to the variation in actions and beliefs required to "trace out" the empirical mappings.

Anticipating the experiment, we present the case of 2×2 games in which there are two players with two actions each, but the model generalizes 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 utility 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.

1.1 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 subjective expected utility 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 utility to action k. We use $v_i = (v_{i1}, v_{i2}) \in \mathbb{R}^2$ as shorthand for an arbitrary such vector: v_i is understood to satisfy $v_i = \bar{u}_i(\sigma_j')$ for some belief σ_j'.

As in QRE, the *action-map* (also known as a "quantal response function") is a function $Q_i : \mathbb{R}^2 \to [0, 1]$ that maps any vector of subjective expected utility to a mixed action. We assume that it satisfies 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.

1.2 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 [2022], we assume the belief-map satisfies 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**: If $\sigma_j < \sigma_j^{'} \in [0, 1]$, $F_i(\bar{\sigma}_j|\sigma_j^{'}) < F_i(\bar{\sigma}_j|\sigma_j)$ 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 $\sigma_j^*(1) = 1$ w.p. 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, belief-distributions shift up in a strict sense of stochastic dominance. (B4) imposes that belief-distributions are unbiased *on median*.[6]

[6] Both median- and mean-unbiasedness can be microfounded via a model of sampling (Friedman [2022]), and the technical axioms allow for either or both. We use median-unbiasedness to derive theoretical results because it turns out to be simpler in our setting, but we test for both types of unbiasedness in our data.

1.3 The hybrid model

The hybrid model generalizes QRE and NBE, defined by an action-map $Q = (Q_1, Q_2)$ satisfying (A1)-(A4) and a belief-map $\sigma^* = (\sigma_2^*, \sigma_1^*)$ satisfying (B1)-(B4).

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 the action-map $Q_i(\bar{u}_i(\sigma_j^{'})) \in [0,1]$. Player i's *expected action-map* is found by integrating over belief realizations: $\Psi_i(\sigma_j; Q_i, \sigma_j^*) := \int_{[0,1]} Q_i(\bar{u}_i(\sigma_j^{'}))dF_i(\sigma_j^{'}|\sigma_j)$. Since $Q_i : \mathbb{R}^2 \to [0,1]$ is single-valued, $\Psi_i : \mathbb{R}^2 \to [0,1]$ is also single-valued. A hybrid equilibrium is defined as a fixed point of $\Psi = (\Psi_1, \Psi_2) : [0,1]^2 \to [0,1]^2$, along with the supporting belief-distributions.

Definition 1. Fix $(\Gamma^{2\times 2}, Q, \sigma^*)$. A hybrid equilibrium is any pair $(\sigma, \sigma^*(\sigma))$ where $\sigma = \Psi(\sigma; Q, \sigma^*)$, Q satisfies (A1)-(A4), and σ^* satisfies (B1)-(B4).

We specialize theory for the family of X-games whose payoffs are in Table 1. These 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 $\sigma_U^{NE,X} = \frac{1}{2}$. 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 freely go between X and $\sigma_L^{NE,X}$ as convenient.

Within any one X-game, the hybrid equilibrium is unique, but flexibility in the primitives gives rise to set-valued predictions, which we completely characterize in Appendix 10.1. For our purposes, however, the following comparative static is more important. By varying X, the hybrid model predicts the variation in actions and beliefs required to "trace out" the empirical action- and belief-maps for both players, justifying the use of X-games in our experiment.

Proposition 1. *Let $\{\sigma_U^X, \sigma_L^X, \sigma_U^{*X}, \sigma_L^{*X}\}_X$ be a dataset of mixed actions and belief-distributions for any finite number of X-games. To be consistent with the hybrid model for some action- and belief-map (Q, σ^*) (held fixed across games), it must be that:*

(i) σ_U^X is strictly decreasing in $\sigma_L^{NE,X}$;

(ii) σ_L^X is strictly increasing in $\sigma_L^{NE,X}$;

(iii) σ_U^{*X} is strictly decreasing in the sense of stochastic dominance in $\sigma_L^{NE,X}$; and

(iv) σ_L^{*X} is strictly increasing in the sense of stochastic dominance in $\sigma_L^{NE,X}$.

Proof. See Appendix 10.3. □

Remark 1. One might be concerned that the fully-mixed X-games are problematic for studying noise as NE predicts that players should be indifferent and therefore *cannot* make mistakes or display "noisy actions." However, our analysis conditions on elicited beliefs, and we find that stated beliefs rarely imply indifference. Moreover, unlike NE, the hybrid model predicts that indifference occurs with probability zero,[7] and so our empirical exercise is internally consistent with our benchmark model.

1.4 QRE and NBE

While Proposition 1 is enough to establish the comparative static prediction that is at the core of our design, it is useful to also consider the special cases of QRE and NBE. QRE is defined by an action-map satisfying (A1)-(A4) and a belief-map that is the identity map imposing "correct beliefs." NBE is defined by a belief-map satisfying (B1)-(B4) and the perfect action-map of "best response." We define these concepts formally in Appendix 10.2.

Friedman [2022] shows that, for any generalized matching pennies game, the set of attainable QRE is equal to the set of mixed action profiles attainable in NBE. It is easy to show that this equivalence extends to any family of X-games.

Proposition 2. *Let $\{\sigma_U^X, \sigma_L^X\}_X$ 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) $\sigma_U^X \in (\frac{1}{2}, 1)$ for $\sigma_L^{NE,X} < \frac{1}{2}$, $\sigma_U^X \in (0, \frac{1}{2})$ for $\sigma_L^{NE,X} > \frac{1}{2}$;
(ii) $\sigma_L^X \in (\sigma_L^{NE,X}, \frac{1}{2})$ for $\sigma_L^{NE,X} < \frac{1}{2}$, $\sigma_L^X \in (\frac{1}{2}, \sigma_L^{NE,X})$ for $\sigma_L^{NE,X} > \frac{1}{2}$;
(iii) σ_U^X is strictly decreasing in $\sigma_L^{NE,X}$; and
(iv) σ_L^X is strictly increasing in $\sigma_L^{NE,X}$.

Proof. See Appendix 10.3. □

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 $\{\sigma_U^X, \sigma_L^X\}_X$ as green dots: the left panel plots σ_U^X and the right panel plots σ_L^X. 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.

[7] This is obviously true by (B1) as any given belief occurs with probability zero. It is also the case that, since hybrid equilibrium action profiles are *not* NE, even a player who randomly realizes "correct beliefs" also strictly prefers one of her actions.

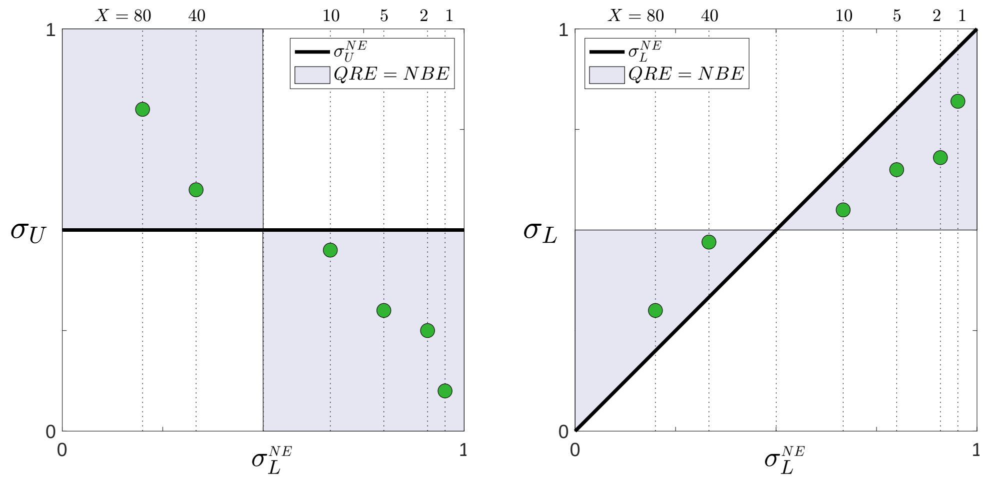

Figure 1: *QRE and NBE in the X-games as a function of* σ_L^{NE}. The left panel gives σ_U and the right panel gives σ_L. The vertical dotted lines correspond to the values of X considered in the experiment (marked at the top). A hypothetical dataset $\{\sigma_U^X, \sigma_L^X\}_X$ is given as green dots. A dataset can be rationalized by QRE or NBE if and only if the datapoints are in the gray regions, decreasing in the left panel, and increasing in the right.

Remark 2. Proposition 2 suggests a convenient way of organizing the data as it gives predicted mixed action profiles as a function of a single parameter; hence, we make use of plots similar to Figure 1 throughout the paper. An analogous plot would not be possible for the hybrid model as the set of predictions for a given value of X is not convex (see Appendix 10.1).

Remark 3. In terms of mixed action profiles, the set of hybrid equilibria is much larger than that of either QRE or NBE.[8] Importantly, however, the hybrid model is rejected if and only if any of the axioms—which are in terms of both actions *and* beliefs—are rejected. Hence, the model is much easier to falsify by augmenting actions data with beliefs data, as we do in our experiment.

2 Experimental design

Recall that our aim is to make observable the empirical action- and belief-maps, which we pursue by collecting actions and beliefs data for a family of X-games.

Overall structure. Our sessions were run in the Columbia Experimental Laboratory in the Social Sciences (CELSS). Subjects were undergraduate students at Columbia and Barnard Colleges.

[8]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. See Appendix 10.1 for details.

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. Subjects maintain their roles throughout the experiment.

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".

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*

In each of the 20 rounds of the first stage, each subject was anonymously and randomly paired with another subject in the opposite role ("random rematching"), 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, rather than being paired against a subject acting simultaneously, subjects played against a randomly selected first-stage subject whose action had already been taken (but was not observed by others). Second, prior to taking an action, each subject was asked for her belief—the probability that a particular action was chosen by a randomly selected first-stage subject in the opposing role.[9] 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 Online Appendix 11.8.

Before the start of each stage, instructions for that stage were read aloud and there were a small number of unpaid practice rounds. Importantly, only after stage 1 were subjects introduced to the notion of a belief and the elicitation mechanism described. Table 2 summarizes the number of subjects who participated in the experiment by

[9] 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.

treatment and player role.[10]

Special procedures. We are interested in observing the stochasticity inherent in beliefs. Hence, we wished to eliminate predictable sources of variation in stated beliefs due to new information or *learning*. For the same reason, we also wished to minimize variation in stated beliefs due to *measurement error*, which we think of as any random misreporting, no matter the cause, of a given underlying belief.

To avoid learning, at no point during the experiment (including the unpaid practice rounds) were subjects provided any feedback. Only at the end of the experiment did subjects learn about the outcomes of the games and belief elicitations that were selected for payment. 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, which we expect is less error prone than using a slider. Of course, we acknowledge that measurement error can never be fully eliminated. Indeed, measurement error in reporting beliefs can come from within the mind of subjects who noisily introspect about their true beliefs. It is for this reason that we weaken the assumption that "stated beliefs equal true beliefs" in Section 7.

Because we wish to analyze stochasticity and patterns in individual subjects' belief data, each game was played multiple times. 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 randomized order, with the restriction that the same game did not appear twice within 3 consecutive rounds.

Incentives. 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 was selected for a subject's action payment, the

[10] 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.

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.

To incentivize subjects to accurately report their beliefs, we used the random binary choice mechanism (Karni [2009]). The important feature is that reporting truthfully is incentive compatible, independent of risk attitude. 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 (randomly selected) 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.

The games. As discussed in Section 1, the X-games take center-stage since they are predicted to give rise to systematic variation in actions and beliefs. Henceforth, we sometimes refer to the game with $X = 80$ as "$X80$" and similarly for the other games.

Another important aspect of the X-games is that they are very simple and fully-mixed. Hence, 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.

For the experiment, we chose the six values of X given in Table 3. These correspond to the vertical dotted 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.

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

Table 3: *Selection of X-games.* In all games, $\sigma_U^{NE} = 0.5$.

Each of the X-games was played 2 times in the first stage and 5 times in the second stage. Since there are six X-games, they appeared a total of 12 times in the first stage and 30 times in the second stage.

We also included some additional 2×2 games that we do not analyze. These were included primarily to break up the appearance of the X-games, similar to what has been done in non-strategic 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 "Incentives" above) and that the X-games take up a similar fraction of total games in each stage.

The games appeared 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. Online Appendix 11.9 shows all of the games and explains the precise randomization.

Alternative designs and generalizability. Many alternative designs are possible and would yield complementary insights. We discuss several, as well as related questions concerning the generalizability of our findings in Section 9.1.

3 Overview of the data

3.1 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], and "[A,B$\underline{\text{A}}$]" refers to second-stage actions from [A,BA].

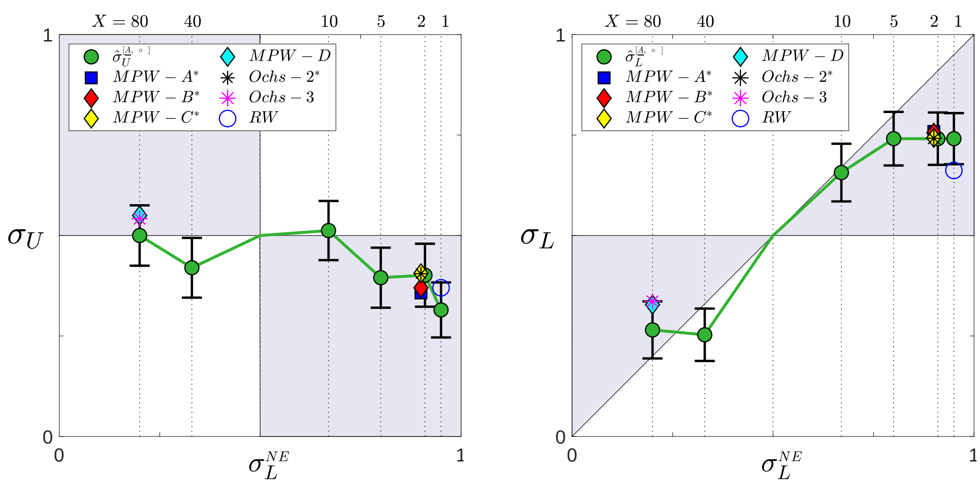

Figure 2: *Actions data.* This figure plots the first-stage empirical action frequencies from [$\underline{\text{A}}$, $\circ$] with 90% confidence bands (clustered by subject), superimposed with the empirical frequencies from other studies.

We are interested in first-stage action data 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 [$\underline{\text{A}}$, $\circ$]. Figure 2 plots the action frequencies from [$\underline{\text{A}}$, $\circ$], which are also given in Appendix Table 11.

We see that the data largely falls within the QRE-NBE region, and F-tests strongly reject that the data generating process is NE. 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 21.

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. We find, however, that our actions data is very similar to that collected under more standard experimental conditions for similar games. Figure 2 plots our action frequencies from [$\underline{\text{A}}$, $\circ$], superimposed with those from three studies (Ochs [1995], McKelvey et al. [2000], and Rutstrom and Wilcox [2009]).[11] We find that our data is statistically indistinguishable from theirs.

We also make use of the second-stage actions data from [A,B$\underline{\text{A}}$] because, in testing axioms on the action-map, we must associate to each belief statement the corresponding action. As shown in Appendix Table 11, there are some differences between the first- and second-stage action frequencies, which we discuss further in Section 3.4.

3.2 Beliefs

Figure 3 plots individual belief statements along with the median and quartiles of beliefs for each game (Appendix Table 11 reports median and mean beliefs). The left panel gives player 2's beliefs over σ_U, and the right panel gives player 1's beliefs over σ_L.

Figure 3 leads to three important observations. First, the central tendency of beliefs is consistent with QRE and NBE. Player 2 tends to believe player 1 will favor U when $X > 20$ and D when $X < 20$. Player 1 tends to believe that player 2 will mostly respond to this, favoring R when $X > 20$ and L when $X < 20$. Second, beliefs clearly respond to changes in X, with medians and quartiles of beliefs varying monotonically in the direction predicted by the hybrid model. Third, player 2's beliefs have much less dispersion than player 1's beliefs, which have average interquartile ranges of 15 and 29 percentage points, respectively.

[11] For inclusion, we sought studies that played games with "sparse" payoffs (in the sense of having zero payoff outcomes) and $\sigma_U^{NE} = \frac{1}{2}$ (for some relabelling of players and actions). 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. We cannot find precedents in the literature for games closely matching our more symmetric games—those with σ_L^{NE} relatively close to $\frac{1}{2}$.

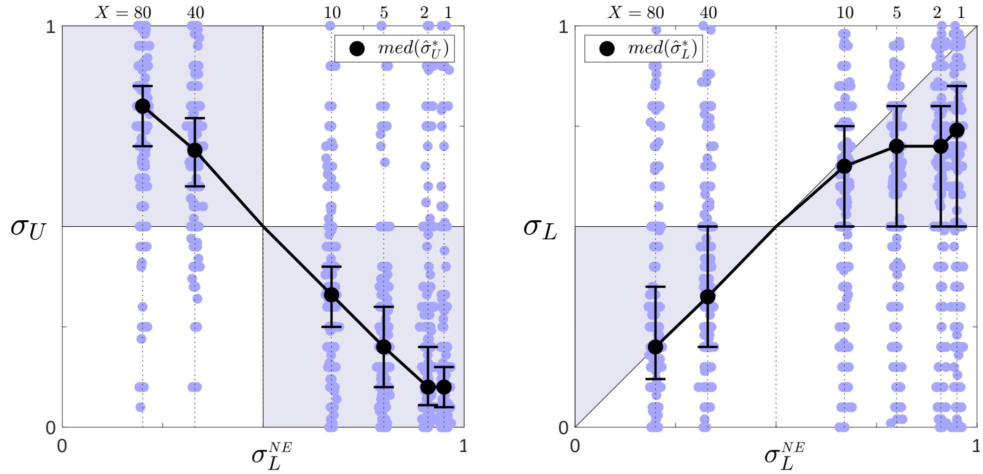

Figure 3: *Beliefs data.* This figure plots individual belief statements along with the median and quartiles of beliefs. The left panel gives player 2's beliefs over σ_U, and the right panel gives player 1's beliefs over σ_L.

To better visualize the distributions of beliefs, Figure 4 plots the empirical CDFs of beliefs for each player and game, with player 1's beliefs in the left panel and player 2's in the right. The figure makes clear that the entire belief-distributions shift monotonically in X in the sense of stochastic dominance. It is also clear there is much more dispersion in player 1's beliefs. In Appendix Figure 22, we display the same information using histograms, which some may find more intuitive.

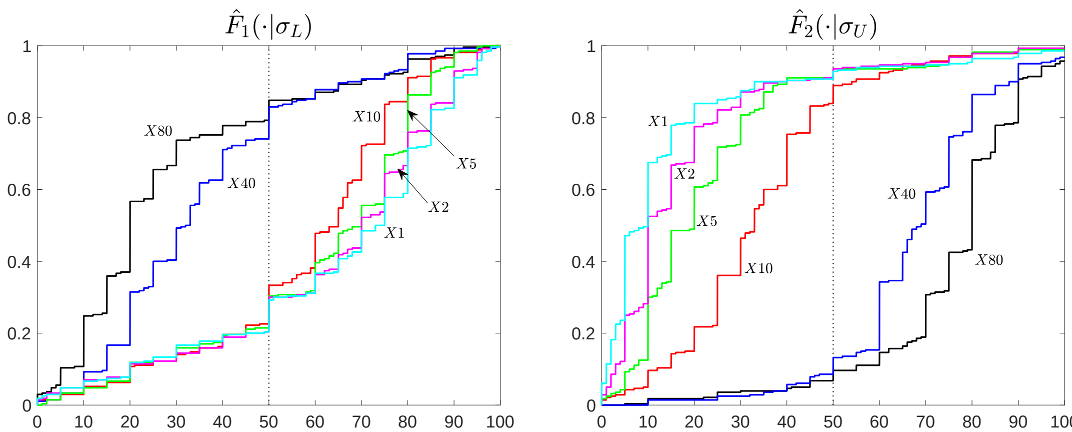

Figure 4: *CDFs of beliefs.* We plot the empirical CDFs of beliefs for each player and game. The left panel is for player 1, and the right panel is for player 2.

Heterogeneity. The focus of this paper is on noise, and therefore on within-subject variation in beliefs. However, it is interesting to also consider heterogeneity across subjects, which a priori may be just as empirically relevant. The fact that we have multiple belief statements for each subject allows us to perform analysis of variance (ANOVA) tests. In Appendix Table 14, we report the results of such tests for each player and game. This reveals that there is actually more between- than within-subject variation in beliefs, and that player 1-subjects' beliefs have more of both types of variation.

3.3 Rates of best response

The *rate of best response*, i.e. the percentage of actions that are a best response to stated beliefs, has been suggested as a metric for validating elicited beliefs (Schotter and Trevino [2014]).

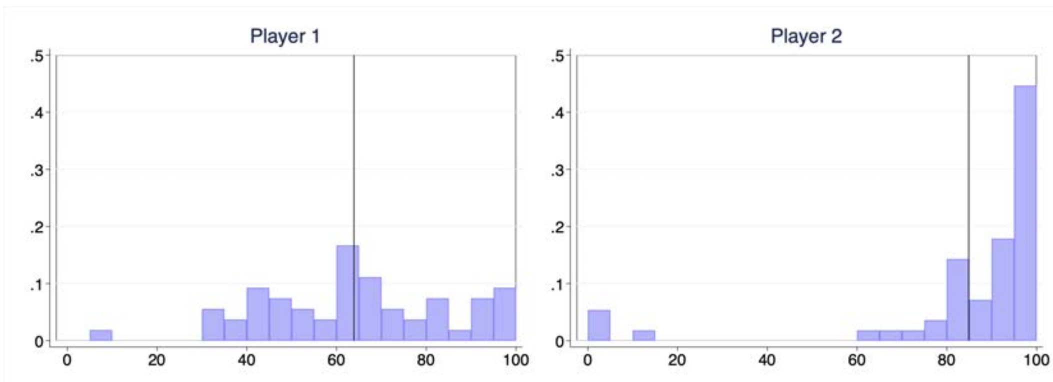

Figure 5: *Subjects' rates of best response.* This figure gives histograms of subjects' rates of best response across all X-games. The average bet response rates are given as vertical lines.

Figure 5 plots histograms of individual subjects' rates of best response, calculated from all six X-games.[12] We find considerable heterogeneity in rates of best response, with averages of of 64% for player 1 and 85% for player 2. 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. Appendix Table 11.7, which gives the average rates for each game, shows that our relatively low rates for player 1 are driven by the very asymmetric games with low values of X. The rates we find are similar to those reported in other studies.[13]

In subsequent sections, the axioms we consider lead us to study not just rates of best response, but how these rates vary across every neighborhood of stated beliefs.

[12]For this exercise, we assume linear utility, which we relax in subsequent sections. Note that, since player 2 faces symmetric payoffs, best response is invariant to non-linearities in the utility function, so this is without loss for player 2.

[13]Nyarko and Schotter [2002] find an average rate of 75% for an asymmetric matching pennies game played many times with feedback. For intermediate values of X, games that resemble the one in Nyarko and Schotter [2002] more closely, we find very similar rates for player 1.

3.4 The effects of belief elicitation

It has been suggested that the very act of eliciting beliefs may change subjects' behavior (Schotter and Trevino [2014]). In general, such effects may have important implications for interpreting results, so while there is little direct evidence for such effects in previous work, we test for it in our data.

In Online Appendix 11.1, we compare the empirical action frequencies from [$\underline{\text{A}}$,BA] and [A,B$\underline{\text{A}}$]. That is, we compare first-stage actions, without belief elicitation, to second-stage actions, each of which was preceded by belief elicitation. For player 2-subjects, we find only minor differences that are not statistically significant. For player 1-subjects, however, we find a sizeable and systematic difference that is statistically significant.

Our hypothesis is that these differences are *caused by* belief elicitation. However, the two stages of [A,BA] 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—identical to [A,BA], except that beliefs are not elicited. In Online Appendix 11.1, we show that the actions data is statistically indistinguishable between the two stages of the [A,A] treatment for both players. Since there is no difference across the two stages in the absence of belief elicitation, we conclude that it was the belief elicitation itself in the [A,BA] sessions that affected player 1-subjects' actions.

Our best guess is that, by focusing their attention on beliefs, the elicitation increased player 1-subjects' degree of strategic sophistication. For player 1, choosing U when $X > 20$ and D when $X < 20$ corresponds to level 1 behavior (Nagel [1995]) if level 0 is assumed to uniformly randomize. After belief elicitation, player 1-subjects are less likely to take these actions, consistent with a shift toward level 3 behavior. We observe no analogous effect for player 2-subjects, perhaps because the commonly held belief that player 1 will tend toward U when X is large and toward D when X is small is very salient. This finding—that only the player facing asymmetric payoffs is affected by belief elicitation—is also consistent with the results of Rutstrom and Wilcox [2009].[14]

Implications for interpreting our results. The fact that player 1's actions are affected by belief elicitation suggests that, for some player 1-subjects, their stated beliefs in the second stage may not be a good indication of the beliefs they formed during the first stage. Hence, for player 1-subjects, the observed relationship between second-stage

[14] Rutstrom and Wilcox [2009] play, with feedback, a game similar to $X1$. Using a structural model, they argue that belief elicitation only affects the path of play for the player facing asymmetric payoffs.

stated beliefs and the first-stage actions that they refer to may give a somewhat biased estimate of the belief-map. On the other hand, the action-map, which summarizes the actions taken *conditional on* beliefs, should not be affected. Finally, while we do not think there is any reason to suppose that belief elicitation systematically affects the variance of belief-distributions, it is a possibility that should be kept in mind while interpreting our results quantifying noise in beliefs.

4 Are beliefs noisy?

An important feature of our data is that we have multiple belief elicitations for each subject and game without feedback. This allows us to answer the basic question: are beliefs noisy?

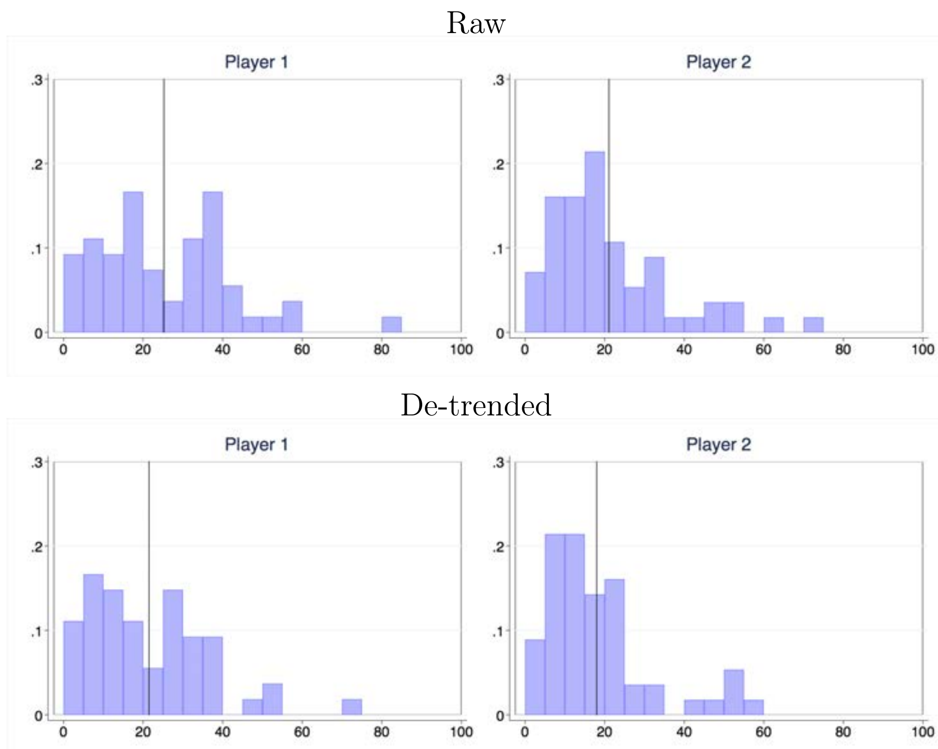

Figure 6: *Subjects' spreads of stated beliefs.* The top panel plots histograms of subjects' spreads of stated beliefs by player role. The spread for a given subject and game is calculated as the highest stated belief minus the lowest, and here we plot each subjects' average spreads, i.e. averaged across all X-games. The bottom panel plots the same after linearly detrending beliefs separately for each subject and game. In both panels, averages are given as vertical lines.

We first derive a simple measure of within-subject variation in stated beliefs. For each subject and X-game, we calculate the spread of her stated beliefs—the highest belief minus the lowest belief—across the five belief statements. We average this across the six X-games to get an average spread measure for each subject. The top panel of Figure 6 plots histograms of subjects' spreads by player role. We see that there is considerable heterogeneity across subjects, as well as a right tail of subjects with very high spreads. The average spreads are also high: 25 and 21 belief-points for player 1- and player 2-subjects, respectively. Is this evidence for noise in beliefs? Or does it simply reflect learning or measurement error?

Our procedures, in particular the lack of feedback, were designed to minimize conventional learning due to new information. 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]). Figure 7, however, shows that average beliefs are very stable throughout the experiment: there is no overall trend in beliefs. To account for subject-specific trends, we linearly detrend beliefs for each subject and game and recalculate for each subject an average detrended spread. The bottom panel of Figure 6 replicates the top panel 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.

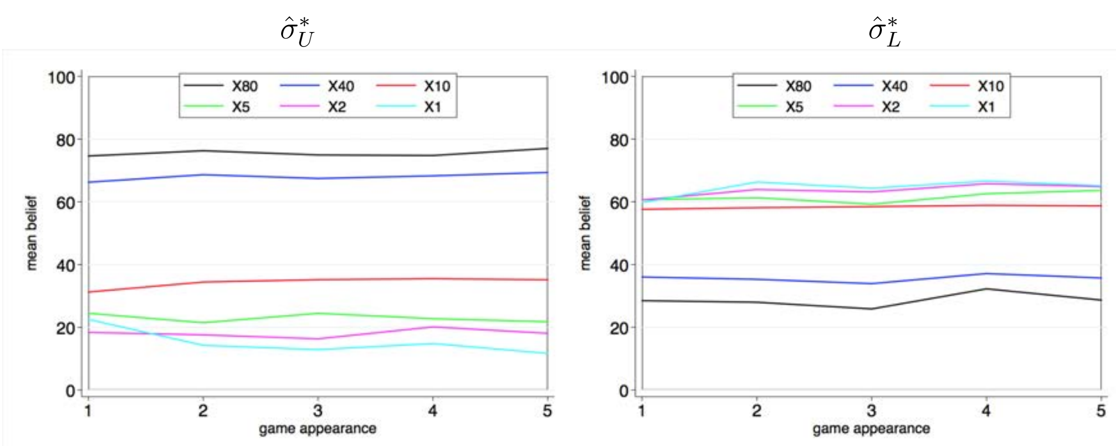

Figure 7: *Stability of average beliefs throughout the experiment.* For each game and player role, this figure plots average stated beliefs for each of five appearances of the game throughout the experiment.

In Section 7, we argue formally that variation in stated beliefs cannot be due entirely to random measurement error. 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 reflects variation in true beliefs.

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.

In Figure 8, we plot the average belief-spreads separately for each game and player role, which shows that belief-noise tends to be higher for player 1-subjects and decreasing in X for both players. In Section 8, we offer an explanation for what we believe is driving these and other patterns in the data.

Result 2. *Variation in belief-noise across game and player role.* The average spread in beliefs is higher for player 1-subjects and decreasing in X for both players.

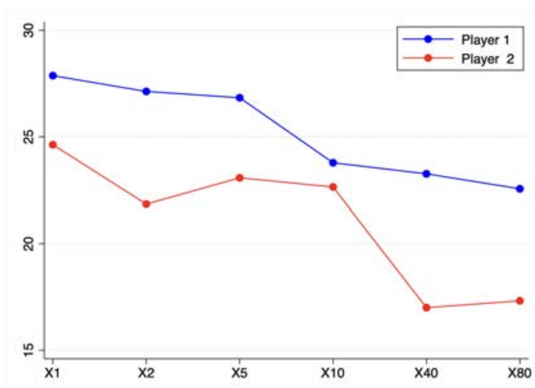

Figure 8: *Average spread of beliefs by game and player role.* We plot the average spread in subjects' beliefs for each game and player role.

5 Action-noise or belief-noise?

Given a standard dataset consisting only of actions data, it is common practice to fit parametric models with action-noise and deterministic beliefs. For example, it is particularly common to fit logit QRE or level k-type models after being augmented with an error structure. However, action-noise and belief-noise can be quite different in their behavioral implications (Friedman [2022]). Hence, as we discuss at the end of this section, if there is considerable unobserved belief-noise, a fitted model that ignores this may have poor performance, especially out-of-sample, or result in biased estimates of structural parameters (e.g. risk aversion).

Following this motivation, we make use of both actions and beliefs data jointly to directly determine if belief-noise is quantitatively important. We ask: which of action- or belief-noise is more important for explaining the data?

We answer the above 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 relative to the data. 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 subject-specific best-fit random utility model. For player 1-subjects, based on evidence from Section 6.5, we also allow for curvature in the utility function.[15] We leave the details of the random utility estimation to Online Appendix 11.5, but emphasize that the estimation is done for each subject separately based on her data from all six X-games. We denote the counterfactual action frequency by $p_s^{iX}(b_{s,med}^{iX};\hat{\rho}_s^i,\hat{\mu}_s^i)$. This is the predicted probability for subject s in role i of taking U if $i=1$ or L if $i=2$ in game X, where $\hat{\rho}_s^i$ is the estimated curvature, $\hat{\mu}_s^i$ is the estimated noise parameter, and $b_{s,med}^{iX}$ is the median belief statement. For player $i=2$, $\hat{\rho}_s^i$ is set to 0, corresponding to linear utility.

To turn off action-noise, we assume subjects best respond to every belief realization. Thus, the counterfactual action frequency is given by $q_s^{iX}(\{b_{sl}^{iX}\}_l;\hat{\rho}_s^i) = \frac{1}{5}\sum_l BR^{iX}(b_{sl}^{iX};\hat{\rho}_s^i)$, where $\{b_{sl}^{iX}\}_l$ are the five belief statements (indexed by $l=1,..,5$) for subject s in role i of game X and BR^{iX} is player i's best response correspondence. For $i=1$, BR^{iX} equals 1 if U is a best response to the given belief and 0 otherwise. For $i=2$, BR^{iX} equals 1 if L is the best response to the given belief and 0 otherwise. Note that BR^{iX} depends on estimated curvature $\hat{\rho}_s^i$ if $i=1$.

For each subject, game, and counterfactual, we calculate the absolute difference between the empirical action frequency $\hat{\sigma}_s^{iX}$ and the counterfactual frequency. Averaging across all six games gives the subject's average counterfactual prediction error:

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

Hence, each subject s is associated with a pair $\varepsilon_s^i = \{\varepsilon_{s,\text{belief-noise}}^i, \varepsilon_{s,\text{action-noise}}^i\}$ that gives

[15] For player 2-subjects, due to symmetry of payoffs, best response is invariant to curvature in the utility function and any curvature parameter cannot be separately identified from the noise parameter.

the prediction errors that result from turning off action- and belief-noise, respectively. In other words, $\varepsilon^i_{s,\text{belief-noise}}$ is the error of the counterfactual with only belief-noise, and $\varepsilon^i_{s,\text{action-noise}}$ 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.

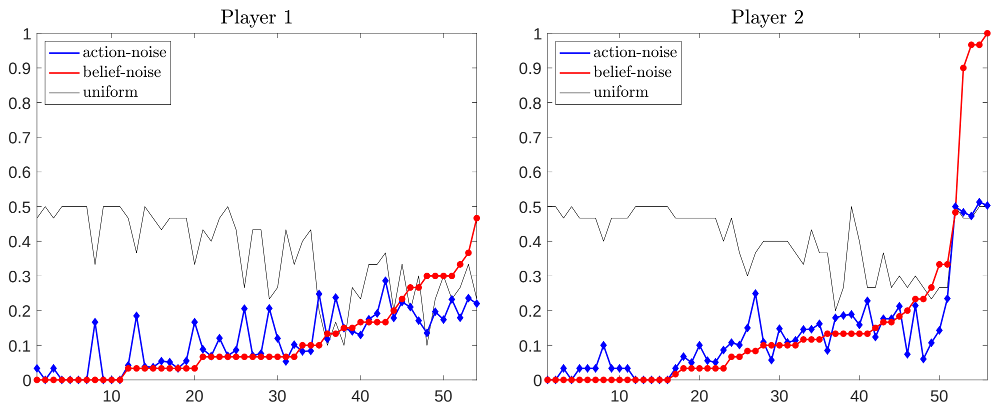

Figure 9: *Counterfactual prediction errors (individual subjects).*

Figure 9 plots all subjects' counterfactual prediction errors. Within each player role i, subjects are sorted so that $\varepsilon^i_{s,\text{belief-noise}}$ is increasing. In addition to plotting $\varepsilon^i_{s,\text{belief-noise}}$ and $\varepsilon^i_{s,\text{action-noise}}$, we also give, as a benchmark, the errors that would result from predicting uniformly random behavior. Table 4 summarizes the average errors by player role ("pooled"), as well as for each player and game.

		$X80$	$X40$	$X10$	$X5$	$X2$	$X1$	pooled
action-noise	Player 1	0.12	0.11	0.16	0.10	0.10	0.09	0.11
	Player 2	0.14	0.16	0.18	0.11	0.10	0.10	0.13
belief-noise	Player 1	0.10	0.11	0.17	0.11	0.10	0.07	0.11
	Player 2	0.17	0.16	0.16	0.16	0.14	0.15	0.16
uniform random	Player 1	0.41	0.37	0.33	0.35	0.34	0.37	0.36
	Player 2	0.40	0.41	0.39	0.39	0.42	0.42	0.41

Table 4: *Mean counterfactual prediction errors.*

Noting that prediction errors can range from 0 to 1, Table 4 suggests that both counterfactuals perform rather well in absolute terms and much better than the random benchmark. The two types of noise do equally well on average for player 1, and action-noise does slightly better on average for player 2.[16] We also find that errors for both

[16]As shown in the right panel of Figure 9, 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

counterfactuals tend to be higher for games with intermediate values of X—games for which stated beliefs tend to be closer to the indifferent belief.

Result 3. *Action-noise and belief-noise are equally important.* Counterfactual action frequencies that result from "turning off" either action-noise or belief-noise lead to similar average prediction errors.

The pitfalls of ignoring belief-noise. Having established that belief-noise is quantitatively important, we discuss how ignoring belief-noise can result in misleading conclusions. Since it is clear that not all behaviors can be explained by noise, we focus on two types of situations for which action-noise is capable of explaining the data in-sample, but ignoring belief-noise leads to poor out-of-sample predictions.

A

	L (40%)	**R** (60%)
U (60%)	0 80	20 0
D (40%)	20 0	0 20

B

	L (5%)	**R** (95%)
U (95%)	8 32	12 32
D (5%)	8 12	12 12

Table 5: *Hypothetical dataset featuring actions with varying sensitivity to belief-noise.*

The first type of situation involves actions with varying degrees of sensitivity to belief-noise. To take a stylized example, consider the hypothetical dataset given in Table 5. Game A is an asymmetric matching pennies game, and the data falls in the interior of the QRE-NBE region with each player taking her action that yields a higher expected payoff slightly more often than not, with probability 60%. Game B is a version of game A for which beliefs play no role. The payoff to each action is given, independently of the opponent's action, by the corresponding empirical expected payoff observed in game A. By construction, each player has a strictly dominant action in B, so we suppose the data from B involves each player taking her dominant action an overwhelming 95% of the time. Now, suppose that a QRE model fitted to the data from game A matches the game A-data perfectly. In this case, the out-of-sample prediction for game B would exactly equal the in-sample prediction for game A, but this would be a very poor prediction. One interpretation is that, in game A, action- and belief-noise are substitutes in the sense

trivially maximizes prediction error. For these subjects, the action-noise counterfactual fairs much better because it predicts uniformly random behavior. Since best responding for player 2-subjects should be easy given their symmetric payoffs, these subjects may simply be confused. After dropping these subjects, the belief-noise counterfactual slightly outperforms the action-noise counterfactual, with pooled prediction errors of 0.09 and 0.10, respectively.

that behavior can be explained equally well by either; therefore, by ignoring belief-noise, the fitted QRE model overstates the degree of action-noise in A. The out-of-sample prediction for game B is poor because the fitted model continues to overstate the degree of action-noise, and, since beliefs play no role in B, the model overstates the *overall* degree of noise as well. Instead, a model fitted to game A that allowed for belief-noise would not overstate the degree of action-noise and therefore result in a better out-of-sample prediction for game B.

The second type of situation involves making inference across environments defined by large differences in payoff magnitude. As is well-known, logit QRE, for a given value of parameter λ, makes predictions that are sensitive to scaling payoffs by positive constants, and yet, such predictions are often rejected (see, for example, McKelvey et al. [2000]). Friedman [2022] also establishes a sense in which this "scaling issue" is considerably more general, applying to many QRE models beyond logit.[17] In contrast, models with belief-noise, such as NBE, random belief equilibrium (Friedman and Mezzetti [2005]), and sampling equilibrium (Rubinstein and Osborne [2003]), make predictions that are invariant to affine transformations of payoffs. Hence, ignoring belief-noise in fitting models can lead to overstating sensitivity to changes in payoff magnitude, leading to poor out-of-sample predictions.

6 Testing the axioms

With a view toward informing models of stochastic choice, we now attempt to reject the four behavioral axioms of the hybrid model: *responsiveness*, *monotonicity*, *belief-responsiveness*, and *unbiasedness*.

In this section, we focus on the data of *player 2*-subjects. The reason is three-fold. First, in light of the finding that player 1's behavior is affected by the belief elicitation itself (Section 3.4), some of the results for player 1 would be harder to interpret. Second, since player 2's payoffs are held fixed as X varies, we are able to pool data across games, which allows for more powerful tests. Third, some of our tests require the identification of preferences, and since player 2 faces symmetric payoffs, our tests are robust to non-linearities in the utility function.[18]

[17] Lemma 3 and Corollary 1 of Friedman [2022] establish that all QRE models satisfying (R1)-(R4) are sensitive to scaling and/or translating payoffs.

[18] In particular, in all games and independent of any curvature, player 2 is made indifferent when she believes player 1 is uniformly mixing.

In Online Appendix 11.4, we conduct analogous tests for player 1-subjects. Broadly, we reach similar conclusions, which we summarize in Section 6.5.

Remark 4. Whereas our previous analysis in Sections 4 and 5 was based on individual subject-level data, our focus is now on aggregate data, which is necessary for statistical power. However, we emphasize that a population of heterogeneous noisy agents, each of whom satisfies the axioms, admits a representative agent who also satisfies the axioms. Hence, the focus on aggregates does not preclude, and is logically consistent with, heterogeneity across individuals. To the extent that we cannot reject an axiom, we think of it as a plausible description of the aggregate data.

6.1 Responsiveness

Responsiveness states that an all-else-equal increase in the subjective expected utility to some action increases the probability that action is played. To test this, we must associate actions and beliefs, and so we use the data from the second stage of [A,BA].

Since subjective expected payoffs are one-to-one with beliefs within a game, *responsiveness* is easily translated in terms of beliefs. Focusing on player 2, *responsiveness* holds if and only if Q_L, the probability she takes action L, is *everywhere strictly decreasing* in belief σ'_U. Since player 2 faces the same payoff matrix across all X-games, we pool data across all games.

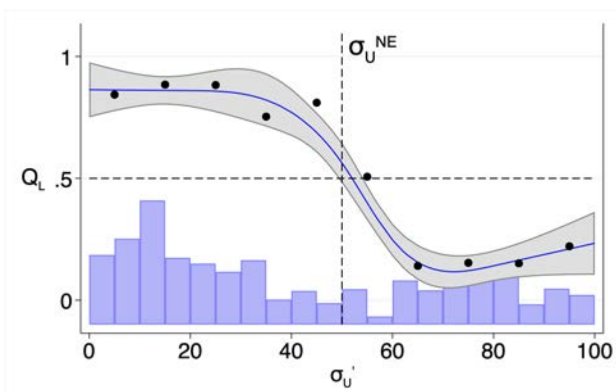

Figure 10: *Action frequencies predicted by beliefs (player 2)*. Using all player 2-subjects and pooling across all X-games, we plot $\hat{Q}_L$, the predicted probability of choosing L (with 90% error bands) as a function of stated beliefs based on a restricted cubic spline regression (5 knots at 20, 40, 50, 60, and 80%, standard errors clustered by subject). We also show a histogram of beliefs and the average action within each of ten evenly spaced bins of beliefs (black dots).

We first visualize the aggregate data in Figure 10, which plots the estimated $\hat{Q}_L$. This is simply the predicted probability of choosing L from regressing actions on beliefs using a flexible specification (see figure caption for details; see Appendix Figure 23 for similar plots for individual X-games). The vertical dashed line gives the *indifferent belief* $\sigma_U' = \sigma_U^{NE} = 50\%$, and the horizontal dashed line is set to one-half. The plot also includes a histogram of stated beliefs, which shows that the belief data is fairly dense throughout the entire space of possible beliefs.

Responsiveness is equivalent to an everywhere strictly decreasing slope for player 2. Inspecting Figure 10, it appears there may be violations. However, since different subjects form different beliefs, $\hat{Q}_L$ is 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 a mechanical issue related to incomplete data—subjects that tend to hold higher beliefs and favor taking L. 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 circumvent this issue, we run fixed-effect regressions. Let (a_{sl}^X, b_{sl}^X) be the lth action-belief pair of subject s in game X. Letting $\bar{a}_s^X \equiv \frac{1}{5}\sum_l a_{sl}^X$ and $\bar{b}_s^X \equiv \frac{1}{5}\sum_l b_{sl}^X$ be the subject-level averages in game X, we run regressions of the following form:

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

Since there is no difference across subjects in the averages of their demeaned variables by construction, the coefficient estimate $\hat{\beta}$ primarily reflects within-subject variation and is similar to an average of individual slopes. Since *responsiveness* concerns the slope at every point, we run separate regressions for different neighborhoods of stated beliefs. Specifically, we first demean the variables. Then, pooling data across all six X-games, we run the regression separately for each quintile of (non-demeaned) belief statements, which we label as "very low," "low," "medium," "high," and "very high" beliefs. The results are displayed in the first column of Table 6. In the second column, as a robustness check, we run the same regressions, except the five belief groups are evenly spaced bins of twenty percentage points.

Consistent with *responsiveness*, we find that every slope is negative and highly statistically significant. Furthermore, the magnitudes are large: all slopes have an absolute value ranging between 0.004 and 0.010, indicating that a 1 percentage point change in belief is associated with a 0.4-1 percentage point change in the probability of taking an action. Since the slopes all have the sign predicted by *responsiveness*, this suggests

	(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 6: *Fixed-effect regressions of actions on beliefs (player 2).* We demean player-2 subjects' belief and action data by subtracting subject-game specific averages. Pooling data across all six X-games, we then regress demeaned actions on demeaned beliefs for each quintile of (non-demeaned) belief statements—very low, low, medium, high, and very high beliefs. These results are in column 1. In column 2, we run the same regressions, except the five belief groups are evenly spaced bins of twenty percentage points. Standard errors are clustered by subject.

broad support for responsiveness.

To help visualize some of the heterogeneity that is hidden in the regressions, Figure 11 plots the data, pooled across all six X-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 very noisy, but still plausibly responsive.

An important question is whether within-subject variations in beliefs have predictive power only insofar as beliefs go on one side or the other of the indifferent belief. Inspecting the second column of Table 6, the answer is definitive. Restricting attention to beliefs that are in the lowest or highest bins—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. Hence, even for player 2, whose indifferent belief is salient, constant across games, and invariant to curvature in the utility function,

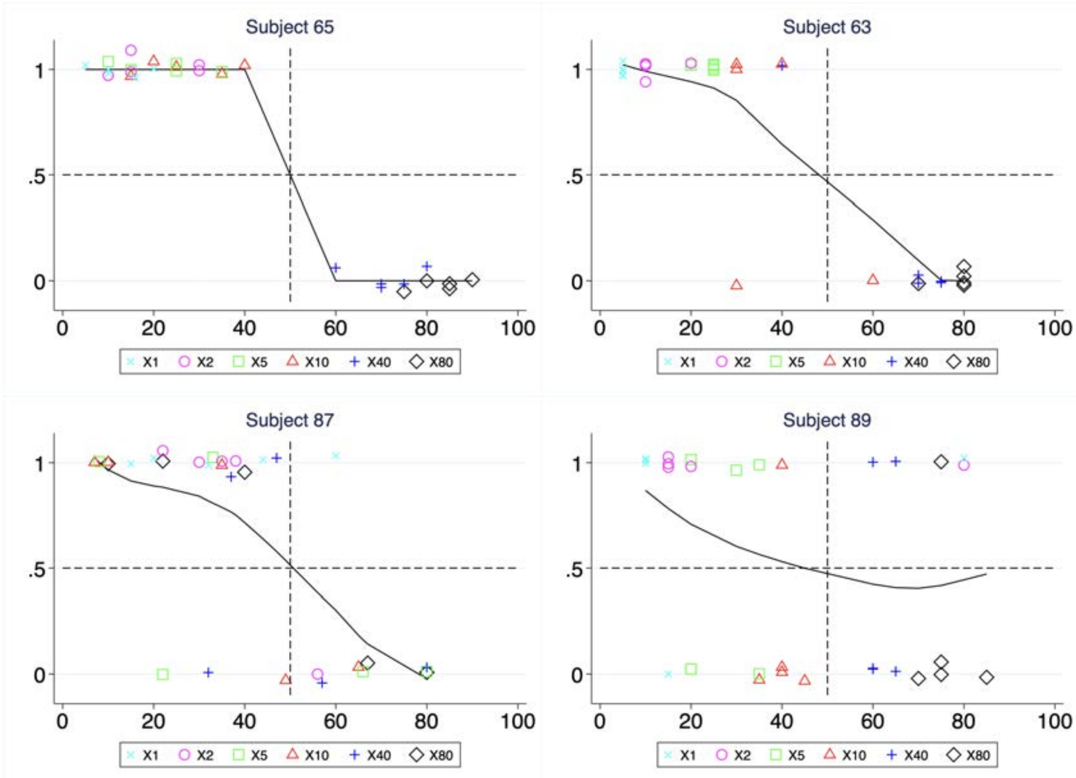

Figure 11: *Individual subjects' actions and beliefs (player 2).* All plots involve individual player 2-subjects whose data is pooled across all games. Action L is coded as 1, and action R is coded as 0. Solid black curves are estimates from local linear regressions. All datapoints involve a value of 1 or 0 on the vertical axis, but are (vertically) jittered for clarity.

all variation in beliefs is highly predictive.

Result 4. *Responsiveness cannot be rejected.* An increase in a subject's stated beliefs is associated with an increase in the probability of taking the action whose payoff is increasing in beliefs. This is true even when restricting attention to subsets of beliefs that imply the same best response.

6.2 Monotonicity

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

For the X-games, since players are indifferent when their beliefs equal the opponent's NE strategy, *monotonicity* takes a particularly simple form. Focusing on player

2, *monotonicity* requires that Q_L, the probability of taking L, is greater than $\frac{1}{2}$ if and only if belief $\sigma_U^{'}$ is less than $\sigma_U^{NE} = \frac{1}{2}$: $Q_L \gtreqless \frac{1}{2} \iff \sigma_U^{'} \lesseqgtr \frac{1}{2}$. Since player 2 faces the same payoff matrix in all X-games, this condition is the same in all games and so we pool the data from all games.

In order to visualize potential *monotonicity* violations, we appeal once again to Figure 10, which plots the estimated $\hat{Q}_L$ using the aggregate data. The vertical dashed line gives the *indifferent belief* $\sigma_U^{'} = \sigma_U^{NE} = 50\%$ 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 2, $\hat{Q}_L$ 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.

Our test for *monotonicity* is the natural one suggested by eyeballing Figure 10. After running flexible regressions of actions on beliefs, we calculate the standard error of the prediction (clustering by subject), which we use to calculate pointwise error bands for the estimated $\hat{Q}_L$. From the figure, one can observe any rejections of the null at the given level of significance. For instance, it is never the case in the figure that $\hat{Q}_L$ is significantly below $\frac{1}{2}$ for beliefs less than 50% or significantly above $\frac{1}{2}$ for beliefs above 50%. Since it is the 90% error band that is plotted, *monotonicity* cannot be rejected with 10% significance, i.e. the p-value is at least 0.1. Hence, we cannot reject *monotonicity.*

Result 5. *Monotonicity cannot be rejected.* For every neighborhood of stated beliefs, subjects tend to take the action which yields the higher payoff more often than not.

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 [$\underline{\text{A}}$, $\circ$].

Focusing on player 2, the right panel of Figure 4 plots the empirical CDFs of player 2's beliefs for all six X-games. Visually, it appears that the belief-distributions are ordered by stochastic dominance, which is confirmed by statistical tests (described below). Furthermore, the belief-distributions 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.

A violation of *belief-responsiveness* occurs whenever, across two games x and y,

$\sigma_U^x > \sigma_U^y$ and $F_2(\cdot|\sigma_U^x) \not\succ_{FOSD} F_2(\cdot|\sigma_U^y)$, meaning beliefs do not go in the same direction as the corresponding action frequencies. Hence, we perform one-sided tests of the null hypotheses $H_0 : \sigma_U^x > \sigma_U^y$ and $H_0 : F_2(\cdot|\sigma_U^y) \succ_{FOSD} F_2(\cdot|\sigma_U^x)$ for all games $x \neq y$. We say that *belief-responsiveness* is rejected whenever we reject *both* $\sigma_U^x > \sigma_U^y$ and $F_2(\cdot|\sigma_U^y) \succ_{FOSD} F_2(\cdot|\sigma_U^x)$.

	Player 1's actions (p-values)					
	$X80$	$X40$	$X10$	$X5$	$X2$	$X1$
$X80$	–	0.89	0.42	0.95	0.93	1.00
$X40$	0.11	–	**0.07***	0.65	0.61	0.96
$X10$	0.58	0.93	–	0.97	0.95	1.00
$X5$	0.05*	0.35	0.03**	–	0.46	0.90
$X2$	0.07*	0.39	0.05**	0.54	–	0.91
$X1$	0.00***	0.04**	0.00***	0.10*	0.09*	–

	Player 2's beliefs (p-values)					
	$X80$	$X40$	$X10$	$X5$	$X2$	$X1$
$X80$	–	0.87	1.00	1.00	1.00	1.00
$X40$	0.00***	–	1.00	1.00	1.00	1.00
$X10$	0.00***	**0.00*****	–	0.84	0.98	0.73
$X5$	0.00***	0.00***	0.00***	–	0.95	0.77
$X2$	0.00***	0.00***	0.00***	0.00***	–	0.78
$X1$	0.00***	0.00***	0.00***	0.00***	0.00***	–

Table 7: *Testing belief-responsiveness (player 2).* The left panel reports p-values from tests of $H_0 : \sigma_U^x > \sigma_U^y$ across game x (row) and game y (column). This is from standard t-tests, clustering by subject. The right panel reports p-values from tests of $H_0 : F_2(\cdot|\sigma_U^x) \succ_{FOSD} F_2(\cdot|\sigma_U^y)$ across game x (row) and game y (column). This is from non-parametric Kolmogorov-Smirnov-type tests in which the test statistic is bootstrapped following Abadie [2002]. We say that a rejection of *belief-responsiveness* occurs whenever we reject *both* $\sigma_U^x > \sigma_U^y$ and $F_2(\cdot|\sigma_U^y) \succ_{FOSD} F_2(\cdot|\sigma_U^x)$. The entries in bold correspond to the only rejection, i.e. rejections of both $\sigma_U^{X40} > \sigma_U^{X10}$ and $F_2(\cdot|\sigma_U^{X10}) \succ_{FOSD} F_2(\cdot|\sigma_U^{X40})$.

Table 7 reports the p-values of these tests for pairs of games x and y in matrix form, with entries in row x and column y (see table caption for details). We find only one significant violation across the many comparisons. This can be seen from the p-values in bold, indicating rejections of both $\sigma_U^{X40} > \sigma_U^{X10}$ and $F_2(\cdot|\sigma_U^{X10}) \succ_{FOSD} F_2(\cdot|\sigma_U^{X40})$. We conclude that *belief-responsiveness* cannot be rejected in our data.

Result 6. *Belief-Responsiveness cannot be rejected.* Beliefs-distributions are ordered across games by stochastic dominance as predicted by the hybrid model. Empirical action frequencies are typically ordered in the same way.

6.4 Unbiasedness

Unbiasedness states that beliefs are unbiased on median. For player 2 forming beliefs over player 1's behavior, this is: $med(\sigma_U^{*,X}) = \sigma_U^X$ for all X. Once again, we use the beliefs data from [A,BA] and the first-stage actions data from [$\underline{\text{A}}$, $\circ$].

Unbiasedness requires that beliefs are unbiased on median, so we plot the aggregate action frequencies and median beliefs, as well as the individual belief statements, in

Figure 12. Appendix Table 13 reports the bias in both median- and mean-beliefs with p-values of the hypothesis that beliefs are unbiased (see caption for details).

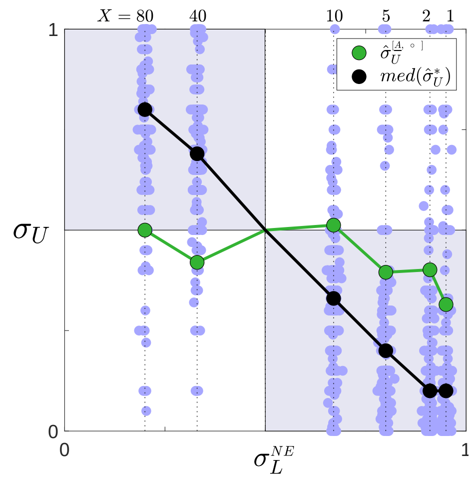

Figure 12: *Bias in beliefs (player 2).* We plot $\hat{\sigma}_U$ from $[\underline{A},\ \circ\ ]$ and the median of player 2-subjects' beliefs over σ_U. Blue circles are individual belief statements.

We find that player 2's beliefs about σ_U are very "extreme": 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 D when X is small. As such, we strongly reject *unbiasedness* for all games (and similarly for mean-unbiasedness). This finding is in contrast with the "conservative" bias documented by Huck and Weizsacker [2002] and Costa-Gomes and Weizsacker [2008] in other settings. In Section 8, we offer an explanation for what we believe is driving this pattern of bias.

Result 7. *Unbiasedness is rejected.* Player 2-subjects have "extreme" beliefs, exaggerating the direction of deviations from uniform play.

6.5 Player 1

In testing the axioms, we focused on the data from player 2-subjects, as discussed at the start of Section 6. In Online Appendix 11.4, we report analogous tests for player 1-subjects, which we summarize below.

Responsiveness. As with player 2, we cannot reject *responsiveness*: fixed-effect regressions similar to those reported in Section 6.1 show that an increase in a subject's stated beliefs is associated with taking the action whose expected payoff is increasing in beliefs more often.

Monotonicity. We find that there are, in fact, intervals of stated beliefs for which subjects fail to take the action that yields the higher expected payoff more often than not. However, we argue that *monotonicity* cannot be rejected. This is because the observed pattern is consistent with plausible *concavity* in the utility function.[19] To illustrate, we take the example of $X80$, which is representative of all X-games. If utility is linear, *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 the indifferent belief $\sigma_L^{NE} = \frac{1}{5}$: $Q_U \gtreqless \frac{1}{2} \iff \sigma_L' \gtreqless \frac{1}{5}$. In other words, Q_U must cross the one-half line from below at belief 20%. Figure 13, analogous to Figure 6 but drawn for player 1-subjects in game $X80$, shows that the estimated $\hat{Q}_U$ crosses the one-half line at 35%, a seeming violation. However, we show that concavity can explain this as it implies an increase in the indifferent belief. Moreover, with a single curvature parameter fitted to the aggregate data from all six X-games, the "violations" in all games disappear.

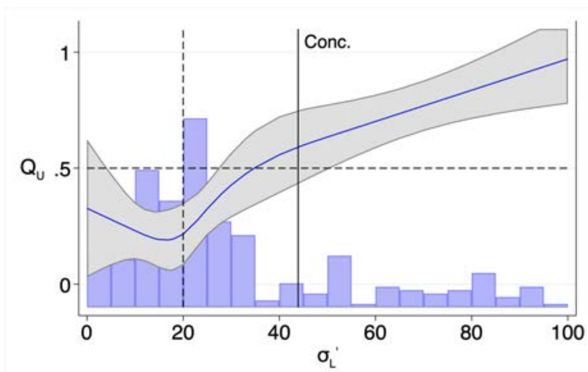

Figure 13: *Action frequencies predicted by beliefs (player 1).* For all player 1-subjects in game $X80$, we plot $\hat{Q}_U$, the predicted probability of choosing U (with 90% error bands) as a function of stated beliefs. The dashed vertical line gives the indifferent belief under linear utility, and the solid vertical line gives the indifferent belief under the estimated concave utility function.

Belief-responsiveness. As with player 2, we cannot reject *belief-responsiveness*: in an exercise similar to that of Section 6.3, we find that player 1-subjects' stated belief-distributions are ordered (by stochastic dominance) in the same way across games as player 2-subjects' empirical action frequencies.

Unbiasedness. Unlike for player 2, we cannot reject *unbiasedness* for player 1. Comparing player 1-subjects' stated beliefs to player 2-subjects' actions, we find that the central tendency of beliefs provides a remarkably close match to player 2's behavior.

[19] Recall that payoffs are in probability points, so curvature should not be interpreted as risk aversion.

7 Measurement error

Throughout the paper, we have treated stated beliefs as equal to the latent or "true" beliefs that subjects hold in their minds and guide their actions. More generally, a stated belief is simply a measure of the underlying true belief, and therefore subject to random *measurement error*—which we define as any random misreporting, no matter the cause, of a given true belief.[20] In this case, does our data allow us to conclude that the unobserved true beliefs are noisy? Can we still reject *unbiasedness* with respect to true beliefs? In this section, we argue that the answer to both questions is "yes" under weak assumptions on the relationship between stated and true beliefs.

We begin by introducing a simple framework that we maintain throughout this section. To this end, suppose that, for a given game, b_0^* and b_s^* are true and stated beliefs, respectively. These are (possibly degenerate) random variables whose supports are contained in $[0, 100]$ and $\{0, ..., 100\}$, respectively. Note that we restrict stated beliefs to be integer-valued so as to respect the experimental elicitation procedure.

We make two key assumptions. First, we assume that, exactly once within each round, a true belief is drawn, and then, conditional on the true belief realization, a stated belief is drawn. We use b_0 to denote an arbitrary realization of true beliefs, and we define $b_s^*(b_0)$ to be random stated beliefs conditional on b_0. Second, we assume that actions depend only on true belief realizations through the function $Q_i(\bar{u}_i(b_0))$. Hence, both stated and true beliefs may be random, the distribution of stated beliefs depends on the realization of true beliefs, and behavior depends *only* on true beliefs.

7.1 Are true beliefs noisy?

Consider some fixed game and a particular subject in that game. If the subject's true belief was fixed and her stated beliefs were simply measures of the underlying belief, then variation in her stated beliefs would not be predictive of her actions. If this were the case, we would see coefficients of 0 in Table 6, but this is strongly rejected. Hence, we conclude that true beliefs are noisy.

7.2 Are true beliefs biased?

We rejected *unbiasedness* for player 2 with respect to stated beliefs. To determine if the axiom can also be rejected with respect to true beliefs, we require additional structure.

[20] A stated belief *is* an action. Hence, given a true belief, noise in stated beliefs—what we call measurement error—is a form of noise in actions.

To this end, we make the following assumption:

Assumption 1. *Stated beliefs are an unbiased signal of the underlying true belief: for all $b_0 \in [0, 100]$, $\mathbb{P}(b_s^*(b_0) > b_0) = \mathbb{P}(b_s^*(b_0) < b_0)$ and $b_s^*(b_0) = b_0$ w.p. 1 if $b_0 \in \{0, 100\}$.*

In words, we assume that the stated belief-distribution is centered, in the sense of median, around the true belief realization. We feel that this formulation is natural in that it closely mirrors the *unbiasedness* axiom, and it makes our analysis very tractable.[21]

We found that player 2-subjects form very biased stated beliefs over σ_U. For instance, consider $X80$, which is representative of all games. Statistical tests support that $med(b_s^*) > \sigma_U$ (see Appendix Table 13). Does this imply that player 2's true beliefs are also biased? To answer this, we first suppose that there is no bias in true beliefs, i.e. $med(b_0^*) = \sigma_U$. This does not imply unbiased stated beliefs, i.e. $med(b_s^*) = \sigma_U$, but this and Assumption 1 jointly imply that $\mathbb{P}(b_s^* > \sigma_U) \leq \frac{3}{4}$.[22] However, we observe that $\hat{\mathbb{P}}(b_s^* > \hat{\sigma}_U)$ is much greater than three-fourths in the data (Figure 12), and so it is very unlikely that *unbiasedness* holds with respect to true beliefs.

Result 8. *The results are robust to measurement error.* Assuming that actions depend on unobservable true belief realizations of which stated beliefs are unbiased signals, we conclude that true beliefs are noisy and that *unbiasedness* is rejected (for player 2) with respect to true beliefs.

8 Toward a theory of noise and bias: the role of salience

We have found that the hybrid model provides a fairly good description of the joint distribution of actions and belief statements. The only exception concerns the rejection of the *unbiasedness* axiom. That player 2's beliefs are biased implies a violation of the hybrid model, as well as any reduced-form equilibrium model that assumes correct beliefs. We take this as the first puzzle to be explained. A second puzzle, embodied in Result 2, is that belief-noise is higher for player 1-subjects and decreasing in X for both

[21] An alternative assumption would be that the distribution of stated beliefs, as a function of the true belief, is governed by a parametric quantal response function (e.g. logit) with payoffs induced by the random binary-choice mechanism.

[22] That $med(b_0^*) = \sigma_U \in (0, 1)$ implies that $\mathbb{P}(b_0^* > \sigma_U) = \mathbb{P}(b_0^* < \sigma_U) = \frac{1}{2}$. Assumption 1 then implies that $\mathbb{P}(b_s^* > \hat{\sigma}_U)$ is maximized when $\mathbb{P}(b_s^*(b_0) > \sigma_U | b_0 > \sigma_U) = 1$ and $\mathbb{P}(b_s^*(b_0) > \sigma_U | b_0 < \sigma_U) = \frac{1}{2}$, which implies that $\mathbb{P}(b_s^* > \sigma_U) = \frac{3}{4}$.

players.[23]

To explain these puzzles, we propose that the differential *salience* of actions across players *causally induces greater strategic sophistication in player 1-subjects*. In forming beliefs, we conjecture that both player 1 and player 2 are instinctively drawn to first consider player 1's payoffs as only she has a salient action—U when $X > 20$ and D when $X < 20$. Naturally, player 2 tends to form *instinctive* beliefs that put significant weight on that action. However, since player 2 does not have a salient action, player 1 must then consider how player 2 reacts to her own salient action. In other words, unlike player 2, player 1 is induced to consider the effect of her own payoffs on her opponent's behavior. As a result, player 1 forms more *contemplative* beliefs that better anticipate her opponent's behavior.[24]

Such a theory can easily explain the observed puzzles: (1) player 2's beliefs are biased toward player 1's salient action because player 1 is "one step ahead"; (2) player 1's beliefs are noisier because they are based on more sophisticated, higher-ordered reasoning; and (3) beliefs become less noisy as X increases because one of player 1's actions becomes even more salient, pushing both players toward more instinctive belief formation.

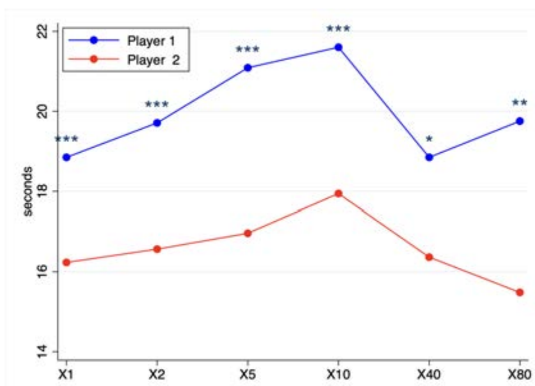

Figure 14: *Average response times by game and player role.* Response times are defined as the time between the start of the round and when the belief statement is finalized.

The salience-induced sophistication hypothesis is consistent with the observed joint distribution of actions and beliefs. We also show that it is consistent with two additional pieces of auxiliary data. First, as shown in Figure 14, player 1-subjects have significantly

[23]Because the hybrid model allows each player to have a different belief-map, this is consistent with the hybrid model. However, since the belief-map is exogenous, the hybrid model does not predict this in particular, nor does it give insight into why this may occur.

[24]Rubinstein [2016] introduces the idea of a typology between instinctive and contemplative players. Unlike Rubinstein [2016], however, we propose that whether a player is instinctive or contemplative is endogenously determined based on her role within the game.

longer response times across all X-games, consistent with more "steps" of reasoning, as in level k (Nagel [1995]). Second, we derive a measure of strategic sophistication (based on level k) from the small number of dominance-solvable games that were included toward the end of the experiment to break up the appearance of X-games (see Online Appendix 11.9). All subjects played in both roles of these games, and so the measure we derive is collected identically for all subjects. By this measure, we find that player 1-subjects are much more sophisticated. Moreover, since all subjects were randomly assigned to their roles at the beginning of the experiment, we are able to establish causality: experience in different roles of the X-games throughout the experiment has a causal effect on independent measures of strategic sophistication. In Online Appendix 11.3, we present this analysis in detail and provide additional discussion.

In our data, that different player roles seem to induce different degrees of strategic sophistication is not easily explained by existing models. In particular, it is hard to reconcile this with the endogenous depth of reasoning theory (Alaoui and Penta [2015]). This is because, for values of $X < 20$, player 1 faces lower expected payoffs against a random opponent, so it is unlikely that she subjectively has more to gain from higher-ordered reasoning; and yet, we find that player 1-subjects have longer response times for all values of X. Rather, we think that our results are better explained by a direct effect of "bottom-up" salience on strategic sophistication, which in turn drives differences in noise and bias across players.[25]

9 Discussion

9.1 Alternative designs and generalizability

We discuss several alternative designs as well as related questions concerning the generalizability of our findings.

Other games. In principle, we could have chosen any set of games. Even restricting attention to 2×2, there are many possible games that differ in terms of dominance-solvability, multiplicity and fragility of NE, etc. So why not run an experiment with many types of games to determine the relationship between noise and game *classes*? Certainly, this would be very interesting. However, as discussed in Section 2, we would be concerned about no-feedback learning effects in other game classes that would lead us to conflate noise and learning. We would also be concerned about spillover effects

[25] Li and Camerer [2022] show that bottom-up visual salience predicts behavior in hide-and-seek games in which choices are locations on an image.

across game classes, so we feel the cleanest design is based on a single class.

To what extent do insights from the X-games generalize to other game classes? By varying a single parameter, our experiment gives a very clean test of how one important dimension of games, i.e. the salience of actions, drives noise. This allows us to speculate about other game classes. For instance, we have found that when X is very large and a particular action of player 1 is very salient, there is significantly less belief-noise for both players. And for any given value of X, player 1 has noisier beliefs than player 2. This allows us to speculate that, in a game where only player 1 has a dominant action, belief-noise would be low for both players, and especially for player 2.

Feedback and learning. A crucial element of our design is to not provide subjects any feedback so as not to conflate learning and noise. An interesting question, however, is whether noise diminishes with learning. To answer this question, one could design an experiment in which feedback is provided for T_1 rounds, followed by T_2 rounds without feedback. By varying T_1 across treatments, and applying our analysis to the last T_2 no-feedback rounds, one could document how noise varies with experience. We speculate that noise would indeed decrease with learning, but would still be significant after any reasonable number of rounds.[26]

In addition to belief-noise, we would also expect belief-bias to decrease with learning. Hence, our rejection of the *unbiasedness* axiom might not hold with sufficient opportunity to learn. However, we would expect subjects to be slow to correct their belief-bias in practice. One indication is that the empirical action frequency of our no-feedback games is statistically indistinguishable from that collected for similar games with feedback in other studies (Section 3.1). Hence, we expect our results to have implications for games played with feedback as well.

Alternative procedures. A common paradigm in the decision-theoretic stochastic choice literature is to give subjects the same choice repeatedly in consecutive rounds (Agranov and Ortoleva [2017]). It is documented that, even here, choices are stochastic, a finding usually interpreted as indicating a preference for randomization. We could have done something similar by playing the same game multiple times in a row. We chose not to do this, however, because we are not interested in preference for randomization, but stochasticity that results from freshly thinking about a game.

Another approach for studying action-noise, used in Agranov et al. [2020], would

[26]One indication comes from Selten and Chmura [2008], who document that, even after 200 rounds of feedback, observed behavior in fully-mixed 2×2 games is still very far from NE and better explained by QRE.

have been to “turn off” belief-noise by first eliciting beliefs and then giving subjects a choice between the objective lotteries implied by their stated beliefs. However, we are interested in stochastic choice given subjective beliefs, and there is no a priori reason to suppose that stochastic choice is the same under subjective and objective beliefs.

9.2 Conclusions

Game theoretic models with stochastic elements have had considerable 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, a phenomenon that cannot be explained by learning or measurement error. Moreover, belief-noise matters: it is equally important as action-noise for explaining our data. As discussed in Section 5, this suggests that ignoring belief-noise, as is done in the large majority of model-fitting applications, could lead to biased parameter estimates and poor performance, especially in making out-of-sample predictions.

In addition to documenting belief-noise, we have found a number of regularities in the relationships between beliefs and actions. Broadly, while beliefs tend to be biased and therefore cannot be fully explained by reduced-form equilibrium models, our results confirm the basic premise of models with noisy actions (e.g., quantal response equilibrium) and noisy beliefs (e.g., noisy belief equilibrium).

Our results, however, highlight some difficulties in formulating parsimonious models of stochastic choice in games. Both action-noise and belief-noise are important for explaining our data, but models incorporating both types of noise as independent factors may be overly flexible. This makes us speculate that the two types of noise should not be modeled as independent, but rather, as determined jointly as part of an optimized response to the environment. In any case, our results call for more research into the factors that drive the relative importance of the two types of noise.

Our data does suggest a number of important factors. In particular, we have argued in Section 8 that the *salience* of payoffs plays a critical role in driving noise, both across games and across players within a game. More generally, it is natural to consider complexity as a potentially important driver of noise, with payoff salience being one of its many aspects. While it is not well understood what makes a game complex, future work should continue to study the relationship between noise, game features, and independent measures of complexity, e.g. response times.

References

Alberto Abadie. Bootstrap tests for distributional treatment effects in instrumental variable models. *Journal of the American Statistical Association*, 2002. 7, 10

Marina Agranov and Pietro Ortoleva. Stochastic choice and preferences for randomization. *Journal of Political Economy*, 2017. (document), 9.1

Marina Agranov, Paul Healy, and Kirby Nielsen. Stable randomization. *Working paper*, 2020. (document), 9.1

Larbi Alaoui and Antonio Penta. Endogenous depth of reasoning. *Review of Economic Studies*, 2015. 7.2

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

Benjamin Enke and Thomas Graeber. Cognitive uncertainty. *Working paper*, 2020. (document)

Evan Friedman. Stochastic equilibria: Noise in actions or beliefs? *American Economic Journal: Microeconomics*, 2022. (document), 1, 1.2, 6, 1.4, 5, 5, 17, 10.3

Evan Friedman and Felix Mauersberger. (a)symmetric quantal response equilibrium: Representation and applications. *Working paper*, 2023. (document)

James Friedman and Claudio Mezzetti. Random belief equilibrium in normal form games. *Games and Economic Behavior*, 2005. 3, 5

Drew Fudenberg, Philipp Strack, and Tomasz Strzalecki. Speed, accuracy, and the optimal timing of choices. *American Economic Review*, 2018. (document)

Jacob Goeree and Philippos Louis. M equilibrium: A theory of beliefs and choices in games. *American Economic Review*, 2021. (document)

Jacob Goeree, Charles Holt, and Thomas Palfrey. Regular quantal response equilibrium. *Experimental Economics*, 2005. (document), 4, 1, 1.1, 10.3

Duarte Goncalves. Sequential sampling and equilibrium. *Working paper*, 2020. 3, (document)

Philip Haile, Ali Hortacsu, and Grigory Kosenok. On the empirical content of quantal response equilibrium. *American Economic Review*, 2008. (document), 4

Steffen Huck and Georg Weizsacker. Do players correctly estimate what others do?: Evidence of conservatism in beliefs. *Journal of Economic Behavior and Organization*, 2002. 6.4

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

Edi Karni. A mechanism for eliciting probabilities. *Econometrica*, 2009. 2

Xiaomin Li and Colin Camerer. Predictable effects of visual salience in experimental decisions and games. *Quarterly Journal of Economics*, 2022. 25

Richard McKelvey and Thomas Palfrey. Quantal response equilibria for normal form games. *Games and Economic Behavior*, 1995. (document)

Richard McKelvey, Thomas Palfrey, and Roberto Weber. The effects of payoff magnitude and heterogeneity on behavior in 2x2 games with unique mixed strategy equilibria. *Journal of Economic Behavior and Organization*, 2000. 3.1, 5

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

Rosemarie Nagel. Unraveling in guessing games: an experimental study. *American Economic Review*, 1995. 3.4, 7.2, 11.3

Yaw Nyarko and Andrew Schotter. An experimental study of belief learning using elicited beliefs. *Econometrica*, 2002. (document), 13

Jack Ochs. An experimental study of games with unique mixed strategy equilibria. *Games and Economic Behavior*, 1995. 3.1

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

Ariel Rubinstein. A typology of players: Between instinctive and contemplative. *Quarterly Journal of Economics*, 2016. 24

Ariel Rubinstein and Martin Osborne. Sampling equilibrium, with an application to strategic voting. *Games and Economic Behavior*, 2003. 3, 5

Elisabet Rutstrom and Nathaniel Wilcox. Stated beliefs versus inferred beliefs: A methodological inquiry and experimental test. *Games and Economic Behavior*, 2009. (document), 3.1, 3.4, 14

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. 5

Andrew Schotter and Isabel Trevino. Belief elicitation in the laboratory. *Annual Review of Economics*, 2014. 5, 3.3, 3.4

Reinhard Selten and Thorsten Chmura. Stationary concepts for experimental 2x2 games. *American Economic Review*, 2008. 26

Amos Tversky. Physical review. *Intransitivity of Preferences*, 1969. (document), 2

Roberto Weber. 'learning' with no feedback in a competitive guessing game. *Games and Economic Behavior*, 2003. 2, 2, 4

Irenaeus Wolff and Dominik Bauer. Why is belief-action consistency so low? the role of belief uncertainty. *Working paper*, 2018. (document)

10 Appendix

10.1 Hybrid equilibrium in the X-games

For any given action- and belief-map (Q, σ^*), the hybrid equilibrium is unique in any X-game. However, since the primitives are only restricted to satisfy axioms, we characterize the set of equilibria that can be attained for *some* primitives. Below, we characterize this set for a given value of X.

Proposition 3. *Fix $X > 0$. An action profile $\sigma = (\sigma_U, \sigma_L)$ can be supported as a hybrid equilibrium for some (Q, σ^*) satisfying (A1)-(A4) and (B1)-(B4) if and only if $\sigma_U \in \Phi_U^X(\sigma_L)$ and $\sigma_L \in \Phi_L^X(\sigma_U)$ where*

$$\Phi_U^X(\sigma_L) = \begin{cases} (0, 3/4) & \sigma_L < \sigma_L^{NE,X} \\ (1/4, 3/4) & \sigma_L = \sigma_L^{NE,X} \\ (1/4, 1) & \sigma_L > \sigma_L^{NE,X}, \end{cases} \quad \text{and} \quad \Phi_L^X(\sigma_U) = \begin{cases} (1/4, 1) & \sigma_U < \frac{1}{2} \\ (1/4, 3/4) & \sigma_U = \frac{1}{2} \\ (0, 3/4) & \sigma_U > \frac{1}{2}. \end{cases}$$

Proof. See Online Appendix 11.2. ☐

Figure 15 illustrates the proposition for $X = 80$, in which case $\sigma_L^{NE,X} = 1/5$. The first panel plots $\Phi_U^X(\sigma_L)$ with a representative expected action-map for player 1, and the second panel plots $\Phi_L^X(\sigma_U)$ with a representative expected action-map for player 2. Where these two regions intersect (third panel) is the set of hybrid equilibrium mixed action profiles. We also give a representative hybrid equilibrium (green dot) and the sets of QRE and NBE, which coincide (cross-hatched region).

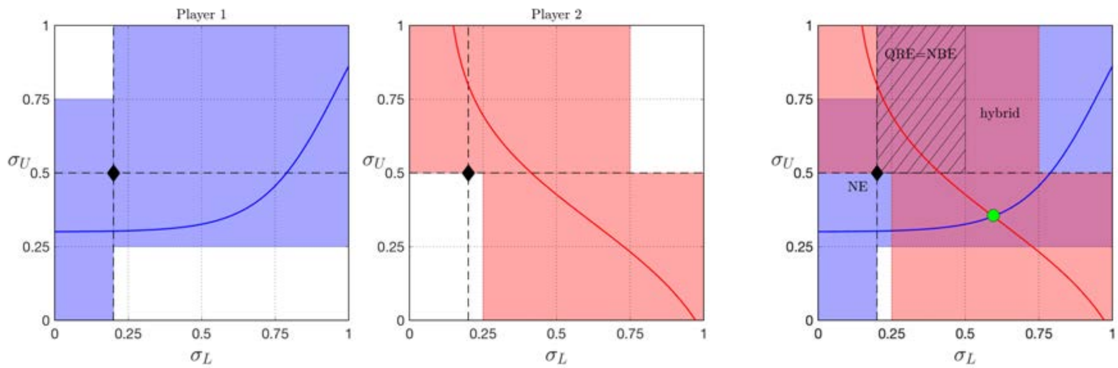

Figure 15: *Hybrid equilibrium in game* $X = 80$.

10.2 Formal definition of QRE and NBE

Definition 2. Fix $(\Gamma^{2\times 2}, Q)$. A quantal response equilibrium any $\sigma \in [0,1]^2$ such that $\sigma = Q(\bar{u}(\sigma))$, where Q satisfies (A1)-(A4).

Definition 3. Fix $(\Gamma^{2\times 2}, \sigma^*)$. A noisy belief equilibrium is any pair $(\sigma, \sigma^*(\sigma))$ with $\sigma \in \psi(\sigma; \sigma^*)$ where $\psi_i(\sigma_j; \sigma_j^*) \equiv \int_{[0,1]} BR_i(\bar{u}_i(\sigma_j'))dF_i(\sigma_j'|\sigma_j)$ is the expected best response correspondence and σ^* satisfies (B1)-(B4).

10.3 Proofs

Proof of Proposition 1. Fix (Q, σ^*) satisfying (A1)-(A4) and (B1)-(B4). Let $\bar{u}^X(\sigma_L') = (\bar{u}_U^X(\sigma_L'), \bar{u}_D^X(\sigma_L'))$ be the expected utilities for player 1 given belief σ_L' in game X. For any given σ_L', $\bar{u}_U^X(\sigma_L')$ is increasing and $\bar{u}_D^X(\sigma_L')$ is decreasing in X, and therefore, by (A3), $Q_U(\bar{u}^X(\sigma_L'))$ is also increasing in X. Hence, $\Psi_U(\sigma_L)$ "shifts up" as X increases. It is easy to show that $\Psi_U(\sigma_L)$ is strictly increasing and $\Psi_L(\sigma_U)$ is strictly decreasing. Thus, it must be that σ_U strictly increases and σ_L strictly decreases as X increases. Noting that X increases if and only if $\sigma_L^{NE,X} = \frac{20}{20+X}$ decreases, we have shown *(i)* and *(ii)*. Parts *(iii)* and *(iv)* follow directly from (B3). ☐

Proof of Proposition 2. The *only if* direction follows for QRE from the results of Goeree et al. [2005] and for NBE from the results of Friedman [2022]. The *if* direction is novel. For QRE, let $\{\sigma_U^X, \sigma_L^X\}_X$ satisfy *(i)-(iv)*. We first construct $Q_1 = (Q_U, Q_D)$ for player 1. Letting $\bar{u}^X(\cdot) = (\bar{u}_U^X(\cdot), \bar{u}_D^X(\cdot))$ be the expected utilities for player 1 in game X, it is easy to show that $\delta(X) := \bar{u}_U^X(\sigma_L^X) - \bar{u}_D^X(\sigma_L^X)$ is strictly increasing in X and $\delta(X) \neq 0$ for all X. Let $\tilde{Q} : \mathbb{R} \to (0,1)$ be any function that is continuous, strictly increasing, and satisfies $\tilde{Q}(0) = \frac{1}{2}$ and $\tilde{Q}(\delta(X)) = \sigma_U^X$ for all X, which exists because $\delta(X)$ is strictly increasing and $\delta(X) \neq 0$ for all X. Now let the quantal response function $Q_1 = (Q_U, Q_D) : \mathbb{R}^2 \to \Delta$ be defined by $Q_U(v_U, v_D) = \tilde{Q}(v_U - v_D)$ and $Q_D(v_U, v_D) = 1 - Q_U(v_U, v_D)$. It is easy to check that this satisfies (A1)-(A4). A similar construction gives $Q_2 = (Q_L, Q_R)$ for player 2, and by construction $Q = (Q_1, Q_2)$ rationalizes the data. For NBE, let $\{\sigma_U^X, \sigma_L^X\}_X$ satisfy *(i)-(iv)*. For player 1, with an arbitrary belief-map $\{F_1(\cdot|\sigma_L)\}_{\sigma_L \in [0,1]}$ satisfying (B1)-(B2), the expected best response function is $\psi_U^X(\sigma_L) = 1 - F_1(\sigma_L^{NE,X}|\sigma_L)$. Hence, we must construct a family of CDFs $\{F_1(\bar{\sigma}_L|\sigma_L)\}_{\bar{\sigma}_L \in [0,1], \sigma_L \in [0,1]}$ such that $1 - F_1(\sigma_L^{NE,X}|\sigma_L^X) = \sigma_U^X$ for all X and that satisfies (B1)-(B4). Given that $\{\sigma_U^X, \sigma_L^X\}_X$ satisfies *(i)-(iv)*, we have that (1) $\sigma_L^X < \sigma_L^{X'}$ whenever $X > X'$, (2) $\frac{1}{2} > \sigma_L^X > \sigma_L^{NE,X}$ if $X > 20$ and $\frac{1}{2} < \sigma_L^X < \sigma_L^{NE,X}$ if $X < 20$, and (3) $\sigma_U^X > \frac{1}{2}$ if $X > 20$ and $\sigma_U^X < \frac{1}{2}$ if $X < 20$. This implies the existence of a *finite* family of CDFs $\{F_1(\cdot|\sigma_L^X)\}_X$ such that, for all X, $1 - F_1(\sigma_L^{NE,X}|\sigma_L^X) = \sigma_U^X$, meaning it can rationalize the data, $F_1(\bar{\sigma}_L|\sigma_L^X)$ is strictly increasing and continuous in $\bar{\sigma}_L \in [0,1]$, $F_1(0|\sigma_L^X) = 0$, $F_1(1|\sigma_L^X) = 1$, and $F_1(\sigma_L^X|\sigma_L^X) = \frac{1}{2}$; and such that $F_1(\bar{\sigma}_L|\sigma_L^{X'}) < F_1(\bar{\sigma}_L|\sigma_L^X)$ for all $\bar{\sigma}_L \in (0,1)$ if $\sigma_L^X < \sigma_L^{X'}$. Hence, the constructed $\{F_1(\cdot|\sigma_L^X)\}_X$ is consistent with (B1)-(B4), and it remains to extend this to $\{F_1(\cdot|\sigma_L)\}_{\sigma_L \in (0,1)}$. For the extension, order the values of X in the dataset: $X^1 > X^2 > ... > X^n$ and $\sigma_L^{X^1} < \sigma_L^{X^2} < ... < \sigma_L^{X^n}$. For $\sigma_L \in (\sigma_L^{X^i}, \sigma_L^{X^{i+1}})$ set $F_1(\bar{\sigma}_L|\sigma_L) = \alpha(\sigma_L)F_1(\bar{\sigma}_L|\sigma_L^{X^i}) + (1 - \alpha(\sigma_L))F_1(\bar{\sigma}_L|\sigma_L^{X^{i+1}})$ for all $\bar{\sigma}_L \in (0,1)$, where $\alpha(\sigma_L)$ is such that $F_1(\sigma_L|\sigma_L) = \frac{1}{2}$, which is uniquely defined. To finish the extension for $\sigma_L < \sigma_L^{X^1}$ and $\sigma_L > \sigma_L^{X^n}$, one can use the construction given in "step 2" of the proof of Theorem 3 of Friedman [2022]. It is easy to check that the constructed belief-map for player 1 satisfies (B1)-(B4). A similar construction gives the belief-map for player 2, and by construction the pair of belief-maps rationalizes the data. □

11 Online Appendix

For Online Publication

11.1 The effects of belief elicitation

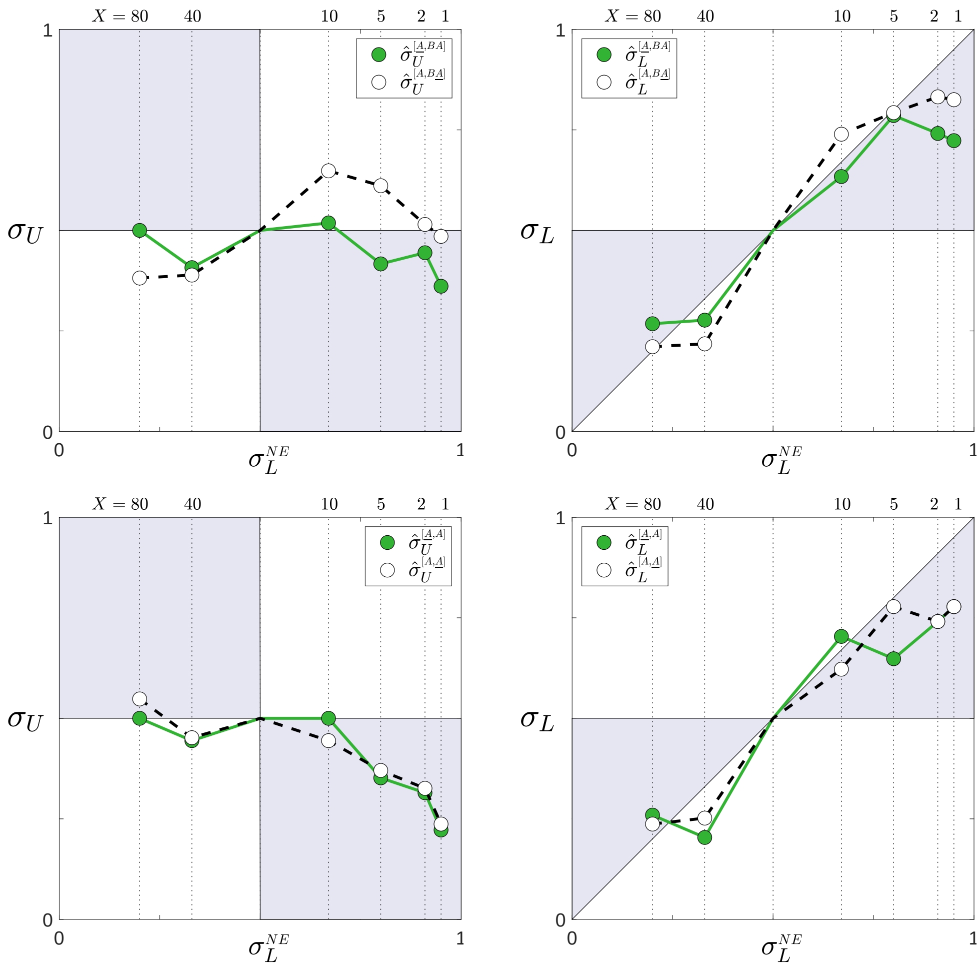

Figure 16: *The effects of belief elicitation.* The top panels plot first- and second-stage action frequencies from [A,BA], and shows a systematic difference between the two stages for player 1 (top-left panel). The bottom panels plot first-stage and second-stage actions from [A,A], i.e. without belief elicitation, and shows no difference between the stages for both players. This suggests that the act of belief itself has a causal effect on behavior for a significant number of player 1-subjects.

In the top panels of Figure 16, we plot the action frequencies from [A̲,BA] and [A,BA̲]. That is, we are comparing first-stage actions, without belief elicitation, to second-stage actions, each of which was preceded by belief elicitation.[27] For player 1 (left), we observe a fairly large difference, both quantitative and qualitative, between first- and second-stage actions. For player 2 (right), there are only some minor quantitative differences. This is confirmed by the F-tests reported in Columns 1-2 of Table 8: we strongly reject that the action frequencies are the same across stages for player 1 ($p < 0.000$), but not for player 2 ($p = 0.064$).

	[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	4.70***	2.11*	0.22	1.03
	(0.000)	(0.066)	(0.967)	(0.423)
[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 8: *The effects of belief elicitation: comparing first- and second-stage actions.* 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 (standard errors clustered by subject). 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. Rejection of this hypothesis means that there is a difference between first- and second-stage actions.

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

[27] The results are similar if, instead, we compare the data from [$\underline{\text{A}}$, ∘] and [A,B$\underline{\text{A}}$], but this would be somewhat confounded by composition effects.

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.

The bottom panels of Figure 16 plot the action frequencies from [A̲,A] and [A,A̲], and Columns 3-4 of Table 8 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 for both players. Since there is no difference across the two stages in the absence of belief elicitation, we conclude that it was the belief elicitation itself in the [A,BA] sessions that affected player 1-subjects' actions.

11.2 Proof of Proposition 3

Proof of Proposition 3. Fix $X > 0$. *Only if*: Let (σ_U, σ_L) be a hybrid equilibrium. Suppose $\sigma_L < \sigma_L^{NE}$. By (B4), player 1's belief-map must satisfy $F_1(\sigma_L|\sigma_L) = \frac{1}{2}$, 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), player 1's action-map must satisfy $Q_U \circ \bar{u}(\sigma_L^{'}) \in (0, \frac{1}{2})$ 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 $\sigma_U \in \Phi_U^X(\sigma_L)$ for all σ_L and $\sigma_L \in \Phi_L^X(\sigma_U)$ for all σ_U. *If*: Let $\sigma_L < \sigma_L^{NE}$ and $\sigma_U \in (0, \frac{3}{4})$. The expected action-map for player 1, Ψ_U, can be made to satisfy $\Psi_U(\sigma_L) = \sigma_U$ by setting the belief-distribution evaluated at σ_L to be $\sigma_L^*(\sigma_L) = \begin{cases} \epsilon & w.p.\ \frac{1}{2} \\ 1-\epsilon & w.p.\ \frac{1}{2} \end{cases}$ for very small $\epsilon > 0$ and the action-map $Q_1 = (Q_U, Q_D) : \mathbb{R}^2 \to \Delta$ to be any that satisfies (A1)-(A4) and $\frac{1}{2}Q_U \circ (\bar{u}_1(\epsilon)) + \frac{1}{2}Q_U \circ (\bar{u}_1(1-\epsilon)) = \sigma_U$. The only restrictions imposed by this construction on the action-map are $Q_U \circ (\bar{u}_1(\epsilon)) \in (0, \frac{1}{2})$ and $Q_U \circ (\bar{u}_1(1-\epsilon)) \in (\frac{1}{2}, 1)$, which is consistent with (A4) and therefore feasible. The constructed belief-distribution, because it is discrete, is not consistent with (B1) and (B2), but $\sigma_L^*(\sigma_L)$ can be modified to be consistent with these axioms and $\Psi_U(\sigma_L) = \sigma_U$ by smoothing out the distribution of $\sigma_L^*(\sigma_L)$ arbitrarily little (along with a corresponding modification of Q_1). Hence, there exists a belief CDF $F_1(\cdot|\sigma_L)$ that rationalizes $\Psi_U(\sigma_L) = \sigma_U$ with $F_1(\bar{\sigma}_L|\sigma_L)$ strictly increasing and continuous in $\bar{\sigma}_L \in [0, 1]$, $F_1(0|\sigma_L) = 0$, $F_1(1|\sigma_L) = 1$, and $F_1(\sigma_L|\sigma_L) = \frac{1}{2}$. All that remains is to extend $F_1(\cdot|\sigma_L)$ to a belief-map, i.e. to a family of CDFs $\{F_1(\cdot|\sigma_L^{'})\}_{\sigma_L^{'} \in [0,1]}$ satisfying (B1)-(B4). This can be done exactly as in the proof of Proposition 2. Using similar arguments, for any (σ_U, σ_L) satisfying $\sigma_U \in \Phi_U^X(\sigma_L)$ and $\sigma_L \in \Phi_L^X(\sigma_U)$, we can construct (Q, σ^*) satisfying (A1)-(A4) and (B1)-(B4) such that the induced expected action-map satisfies $\Psi_U(\sigma_L) = \sigma_U$ and $\Psi_L(\sigma_U) = \sigma_L$. □

11.3 A causal effect of payoff salience on strategic sophistication

In Section 8, we discussed two puzzles in the data: the rejection of the *unbiasedness* axiom for player 2 and the fact that belief-noise is higher for player 1-subjects and decreasing in X for both players. To explain these puzzles, we proposed that the differential *salience* of actions across players *causally induces greater strategic sophistication in player 1-subjects*. Consistent with this hypothesis, we showed in Figure 14 that player 1-subjects have longer response times. In this section, we present another piece of supporting evidence—that player 1-subjects exhibit higher measures of strategic sophistication, which can be interpreted in a causal fashion due to the randomization of subjects into player roles.

The measure of strategic sophistication is based on stated beliefs in the small number of dominance solvable games that were included to break up appearance of the X-games (see Online Appendix 11.9). The dominance solvable games are reproduced in Figure 17. $D1$ 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. In the second stage of the experiment, Di appeared in rounds 7, 21, and 35, and Dj appeared in rounds 14 and 28.

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 17: *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 ($k \geq 2$).

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 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, we are looking at "spillover" effects from the X-games to the dominance solvable games and how these effects vary by player role in the X-games.[28]

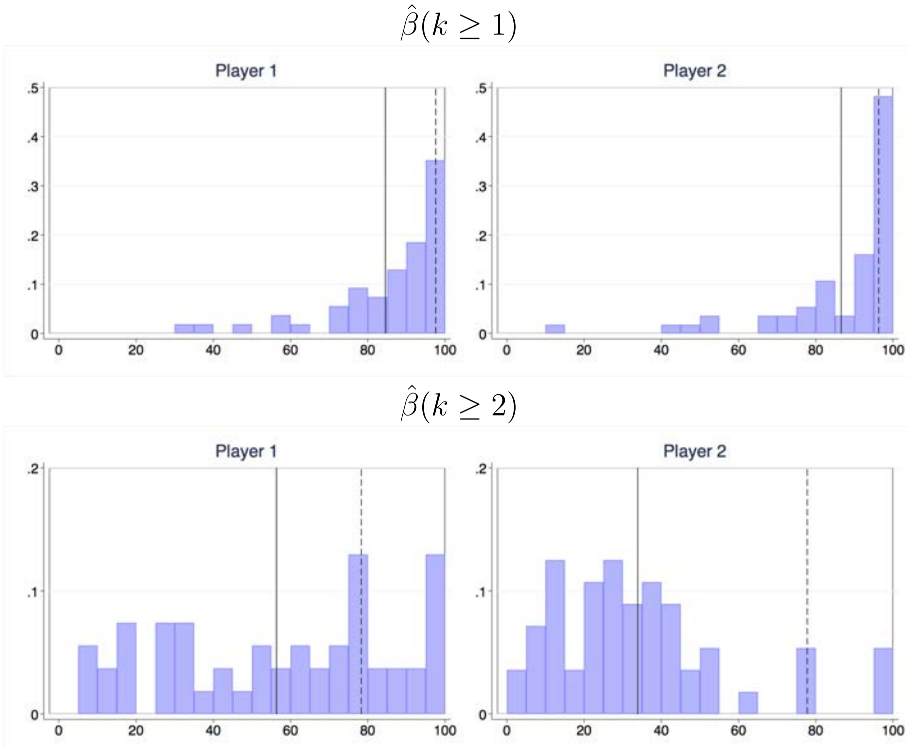

Figure 18: *Sophistication by player.* The top panel gives histograms of $\hat{\beta}(k \geq 1)$, i's belief that j best responds to i's dominant action in Dj, across subjects. The bottom panel gives histograms of $\hat{\beta}(k \geq 2)$, i's belief that j best responds to i's dominant action in Di (as 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 [$\underline{\text{A}}$, $\circ$].

In the level k framework of strategic sophistication (Nagel [1995]), level 0 is assumed to uniformly randomize, level 1 best responds to level 0, and so on, with level k best responding to level $k-1$. In game Di, the following characterizes level-types $k \geq 1$. Player i: levels $k \geq 1$ take the dominant action. Player j: level 1 best responds to a uniform distribution and levels $k \geq 2$ best respond to i's dominant action.

This suggests two benchmark beliefs: (1) i's belief that j takes her dominant action in Dj and (2) i's belief that j best responds to i's dominant action in Di. Assuming i believes j is drawn from a distribution of level types, for any fixed probability that i believes j is level 0, (1) is an increasing function of i's belief that j is any level $k \geq 1$ and (2) is an increasing function of i's belief that j is any level $k \geq 2$, respectively.

[28] One concern is that, since experience in the X-games affects behavior in $D1$ and $D2$, these latter games may also have an effect on behavior in the former. However, we find this implausible since the X-games take up a large majority of the experiment.

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.

Each player states beliefs three times in Di and two times in Dj, so we average beliefs for each subject within-game to yield $\hat{\beta}(k \geq 1)$ and $\hat{\beta}(k \geq 2)$. We plot histograms of these measures in Figure 18. From the top panel, 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, with respective means of 56 and 33—a 23 percentage point difference. Hence, while both players overwhelmingly believe their opponent will take a dominant action, player 1 is much more likely to believe her opponent goes one level deeper. In other words, player 1 is more sophisticated than player 2.

To summarize, since $D1$ and $D2$ are exactly the same up to which player faces which payoffs, the sophistication measure $\hat{\beta}(k \geq 2)$ 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, and we find that player 1-subjects are much more sophisticated by this measure.

11.4 Testing the axioms (player 1)

We replicate the tests of Section 6 for player 1-subjects by attempting to reject the four behavioral axioms (*responsiveness*, *monotonicity*, *belief-responsiveness*, and *unbiasedness*). Since player 1's payoffs vary with X, we cannot pool data across games. For this reason, some of the tests are modified and less powerful, but we reach similar conclusions as we did for player 2-subjects.

11.4.1 Responsiveness

For player 1 and game X, *responsiveness* states that Q_U, the probability of taking action U, is *everywhere strictly increasing* in belief σ_L'.

We first visualize the aggregate data in the left panels of Figure 23, which plots the estimated $\hat{Q}_U$ for each game. As was the case for player 2, we see that the slopes are not everywhere strictly monotonic. However, since different subjects form different beliefs, $\hat{Q}_U$ is 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 a mechanical issue related to incomplete data—subjects that 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.

As before, we get around this issue by running fixed-effect regressions. Let (a_{sl}^X, b_{sl}^X) be the lth action-belief pair of subject s in game X. Letting $\bar{a}_s^X \equiv \frac{1}{5}\sum_l a_{sl}^X$ and $\bar{b}_s^X \equiv \frac{1}{5}\sum_l b_{sl}^X$ be the subject-level averages, we run regressions of the following form for each game X:

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

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.

Since *responsiveness* concerns the slope at every point, we run separate regressions for different neighborhoods of stated beliefs. Specifically, we first demean the variables, and then run the regression separately for each tercile of (non-demeaned) belief statements, which we label as "low", "medium", and "high" beliefs. Since we cannot pool data across games, we use terciles instead of quintiles for more power.

	(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

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

Table 9: *Fixed effect regressions of actions on beliefs (player 1).* 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.

The results are displayed in Table 9. Consistent with *responsiveness*, we find that every slope is positive, with many of these being highly statistically significant. Furthermore, the magnitudes are large: the average slope is 0.065, indicating that a 1 percentage point change in belief is associated with a 0.65 percentage point change in the probability of taking an action. Since the slopes all have the sign predicted by *responsiveness*, this suggests broad support for responsiveness.

11.4.2 Monotonicity

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 the *indifferent belief*. If utility is linear, the indifferent belief coincides with the (linear utility) Nash equilibrium, and so *monotonicity* is the following: $Q_U \gtreqless \frac{1}{2} \iff \sigma_L' \gtreqless \sigma_L^{NE} = \frac{20}{20+X}$. With non-linear utility, the indifferent belief may diverge from σ_L^{NE}.

In order to visualize potential *monotonicity* violations, we plot the estimated $\hat{Q}_U$ in Figure 19 for each game. The vertical dashed line gives σ_L^{NE}. For each game, if utility is linear, $\hat{Q}_U$ should be less than $\frac{1}{2}$ to the left of the vertical dashed line and greater than $\frac{1}{2}$ to the right of the vertical dashed line. From Figure 19, we see what appears to be *monotonicity* violations whereby the estimated $\hat{Q}_U$ is significantly below $\frac{1}{2}$ for some beliefs greater than σ_L^{NE} and significantly above $\frac{1}{2}$ for some beliefs less than σ_L^{NE}. For $X > 20$, this occurs over an interval of beliefs just "right of" σ_L^{NE} , and for $X < 20$, the violations are over an interval of beliefs just "left of" σ_L^{NE} .

The proposition below states that, with concavity in the utility function, the indifferent belief moves right for $X > 20$ and left for $X < 20$. Hence, what we observe may not be violations at all if utility is concave.

Proposition 4. *Let w and v be any strictly increasing (Bernoulli) utility functions over matrix payoffs. For player 1 in game X, w and v induce expected payoff vectors $\bar{w}^X = (\bar{w}_U^X, \bar{w}_D^X) : [0,1] \to \mathbb{R}^2$ and $\bar{v}^X = (\bar{v}_U^X, \bar{v}_D^X) : [0,1] \to \mathbb{R}^2$, respectively. Let $\tilde{\sigma}_L^{w,X}$ and $\tilde{\sigma}_L^{v,X}$ be the unique indifferent beliefs such that $\bar{w}_U^X(\tilde{\sigma}_L^{w,X}) = \bar{w}_D^X(\tilde{\sigma}_L^{w,X})$ and $\bar{v}_U^X(\tilde{\sigma}_L^{v,X}) = \bar{v}_D^X(\tilde{\sigma}_L^{v,X})$, respectively. (i) If w is strictly more concave than v ($w = f(v)$ for strictly concave f), 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 strictly 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. (i): Let w and v be any strictly increasing (Bernoulli) utility functions with $w = f(v)$ for some strictly 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

$\hat{Q}_U$ $\hat{Q}_U$

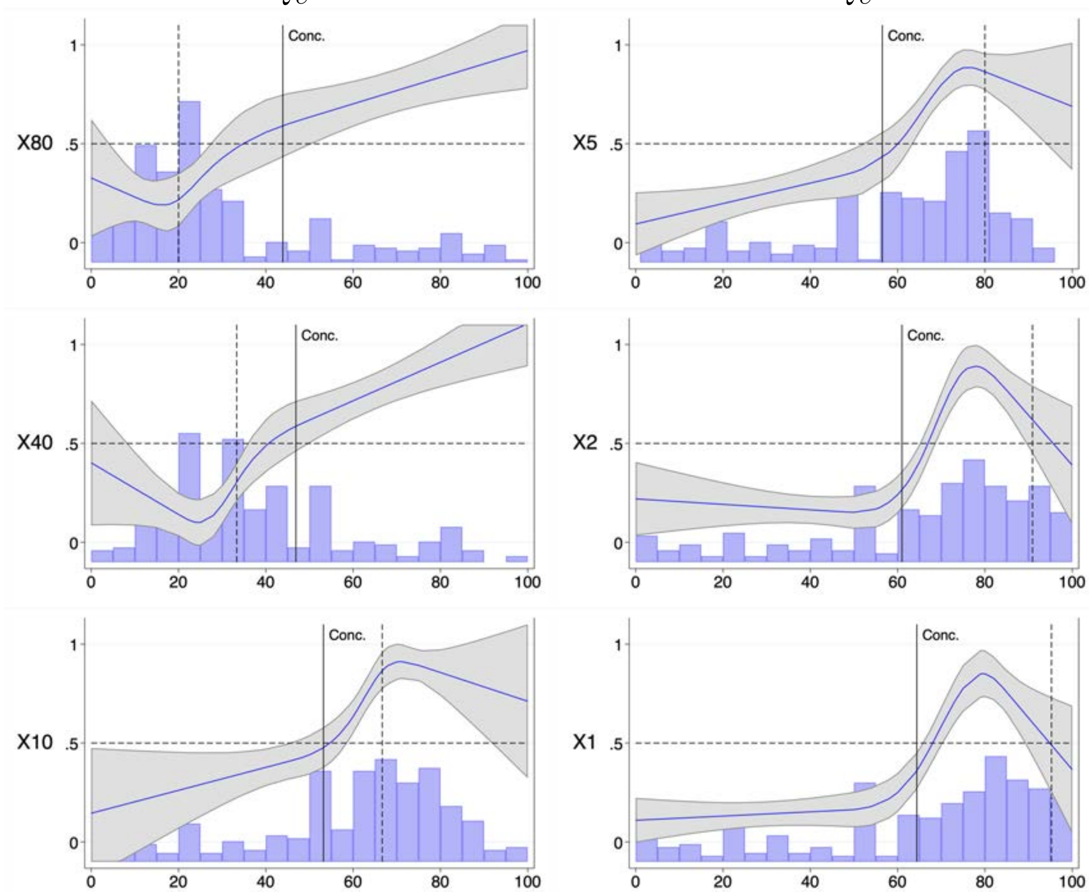

Figure 19: *Concave utility explains monotonicity "violations" (player 1).* For player 1 and each of the X-games, we plot $\hat{Q}_U$, the predicted probability of choosing U (with 90% error bands) as a function of stated beliefs based on restricted cubic spline regressions (4 knots at belief quintiles, standard errors clustered by subject). Belief histograms appear in gray, the vertical dashed line is the indifferent belief under linear utility $\sigma_L' = \sigma_L^{NE}$, 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.

show that the indifferent belief is $\tilde{\sigma}_L^{u,X} = \frac{u(20)}{u(20)+1}$. Since w is strictly more concave than v, $w(20) > v(20)$ and thus $\tilde{\sigma}_L^{w,X} > \tilde{\sigma}_L^{v,X}$. Similarly, if $X < 20$, re-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 strictly more concave than v, $w(X) > v(X)$ and thus $\tilde{\sigma}_L^{w,X} < \tilde{\sigma}_L^{v,X}$. Part *(ii)* is the same, except with $v(z) = z$, which implies $\tilde{\sigma}_L^{v,X} = \frac{20}{20+X} = \sigma_L^{NE,X}$. □

To test for concavity, we fit the random utility model with curvature from Online Appendix 11.5 to each player 1-subject's data, pooled across all X-games. The parameters ρ and μ give curvature and noise, respectively. 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 to the aggregate player 1 data—pooled across all subjects and games. We find the estimate $\hat{\rho} = 0.87$, indicating concavity. In Figure 19, we plot the indifferent beliefs implied by this 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 19). In light of this evidence, we conclude that *monotonicity* cannot be rejected.

11.4.3 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.

We plot the empirical CDFs of player 1's beliefs for all six X-games in the left panel of Figure 4. Visually, it appears that the belief-distributions are ordered by stochastic dominance, which is confirmed by statistical tests (described below). Furthermore, the belief-distributions shift monotonically in X in the direction predicted by the hybrid model: as X increases, player 1 believes that player 2 will play L less often.

A violation of *belief-responsiveness* occurs whenever, across two games x and y, $\sigma_L^x > \sigma_L^y$ and $F_1(\cdot|\sigma_L^x) \not\succ_{FOSD} F_1(\cdot|\sigma_L^y)$, meaning beliefs do not go in the same direction as the corresponding action frequencies. Hence, we perform one-sided tests of the null hypotheses $H_0 : \sigma_L^x > \sigma_L^y$ and $H_0 : F_1(\cdot|\sigma_L^y) \succ_{FOSD} F_1(\cdot|\sigma_L^x)$ for all games $x \neq y$. We say that *belief-responsiveness* is rejected whenever we reject *both* $\sigma_L^x > \sigma_L^y$ and

	Player 2's actions (p-values)							Player 1's beliefs (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***		–	0.00***	0.00***	0.00***	0.00***	0.00***
$X40$	0.42	–	0.00***	0.00***	0.00***	0.00***		0.80	–	0.00***	0.00***	0.00***	0.00***
$X10$	1.00	1.00	–	0.08*	0.08*	0.07*		0.96	0.93	–	0.00***	0.00***	0.00***
$X5$	1.00	1.00	0.92	–	0.50	0.50		0.97	0.92	0.76	–	0.00***	0.00***
$X2$	1.00	1.00	0.92	0.50	–	0.50		1.00	0.72	0.79	0.68	–	0.01**
$X1$	1.00	1.00	0.93	0.50	0.50	–		1.00	0.83	0.66	0.72	0.73	–

Table 10: *Testing belief-responsiveness (player 1).* The left panel reports p-values from tests of $H_0 : \sigma_L^x > \sigma_L^y$ across game x (row) and game y (column). This is from standard t-tests, clustering by subject. The right panel reports p-values from tests of $H_0 : F_1(\cdot|\sigma_L^x) \succ_{FOSD} F_1(\cdot|\sigma_L^y)$ across game x (row) and game y (column). This is from non-parametric Kolmogorov-Smirnov-type tests in which the test statistic is bootstrapped following Abadie [2002]. We say that a rejection of *belief-responsiveness* occurs whenever we reject *both* $\sigma_L^x > \sigma_L^y$ and $F_1(\cdot|\sigma_L^y) \succ_{FOSD} F_1(\cdot|\sigma_L^x)$.

$F_1(\cdot|\sigma_L^y) \succ_{FOSD} F_1(\cdot|\sigma_L^x)$.

Table 10 reports the p-values of these tests for pairs of games x and y in matrix form, with entries in row x and column y (see table caption for details). We find no significant violation across the many comparisons, and so we cannot reject *belief-responsiveness*.

11.4.4 Unbiasedness

Unbiasedness states that beliefs are unbiased on median. For player 1 forming beliefs over player 1's behavior, this is: $med(\sigma_L^{*,X}) = \sigma_L^X$ for all X. Figure 20 plots the aggregate action frequencies and median beliefs, as well as the individual belief statements. Table 13 reports the bias in both median- and mean-beliefs with p-values of the hypothesis that beliefs are unbiased (see caption for details).

Visually, it appears player 1's beliefs are unbiased, and indeed we cannot reject *unbiasedness* for player 1: for no games can we reject *unbiasedness* at conventional levels of significance. In terms of the mean, we find that there is a slight conservative bias: relative to player 2's behavior, average beliefs are closer to the uniform distribution.

11.5 Random utility estimation

The data of subject s is a set of 30 action-belief pairs $(a_{sl}^X, b_{sl}^X)_l$ where $l \in \{1, ..., 5\}$ indexes each elicitation and X indexes the game. We assume that the utility function over matrix payoffs is the constant relative risk aversion (CRRA) utility function with curvature parameter ρ, which has been modified to allow for 0 payoffs by adding a

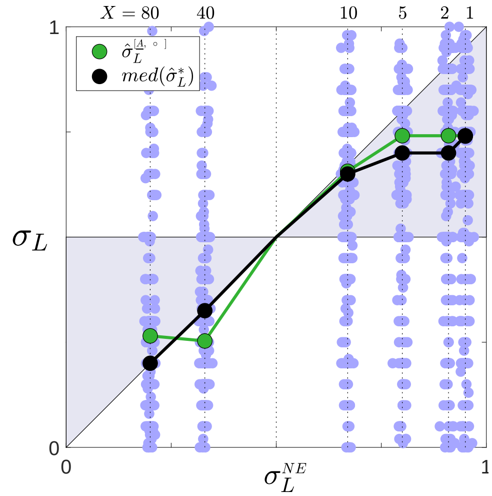

Figure 20: *Bias in beliefs (player 1).* We plot $\hat{\sigma}_L$ from $[\underline{A}, \circ]$ and the median of player 1-subjects' beliefs over σ_L. Blue circles are individual belief statements.

constant $\epsilon > 0$ (arbitrarily pre-set to 0.001) to each payoff. We also normalized utility so that it is between 0 and 1:[29]

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

This utility function induces, for each game X and stated belief b_{sl}^X, a vector of subjective expected utilities $\bar{w}_i^X(b_{sl}^X;\rho) = (\bar{w}_{i1}^x(b_{sl}^X;\rho), \bar{w}_{i2}^X(b_{sl}^X;\rho))$. Coding actions U and L as 1 and D and R as 0, we assume that the probability of taking the action 1 depends only on this vector, based on the Luce quantal response function with sensitivity parameter $\mu > 0$:[30]

$$p_i^X(b_{sl}^X;\rho,\mu) = \frac{\bar{w}_{i1}^X(b_{sl}^X;\rho)^{\frac{1}{\mu}}}{\bar{w}_{i1}^X(b_{sl}^X;\rho)^{\frac{1}{\mu}} + \bar{w}_{i2}^X(b_{sl}^X;\rho)^{\frac{1}{\mu}}}. \tag{1}$$

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

$$L_{is}(\{a_{sl}^X\}_{Xl}|\{b_{sl}^X\}_{Xl};\rho,\mu) = \sum_X \sum_{l=1}^{5} \left[a_{sl}^X \cdot ln(p_i^X(b_{sl}^X;\rho,\mu)) + (1-a_{sl}^X)\cdot ln(1-p_i^X(b_{sl}^X;\rho,\mu))\right].$$

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

[30] The Luce rule (1) 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.

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

11.6 Additional Figures

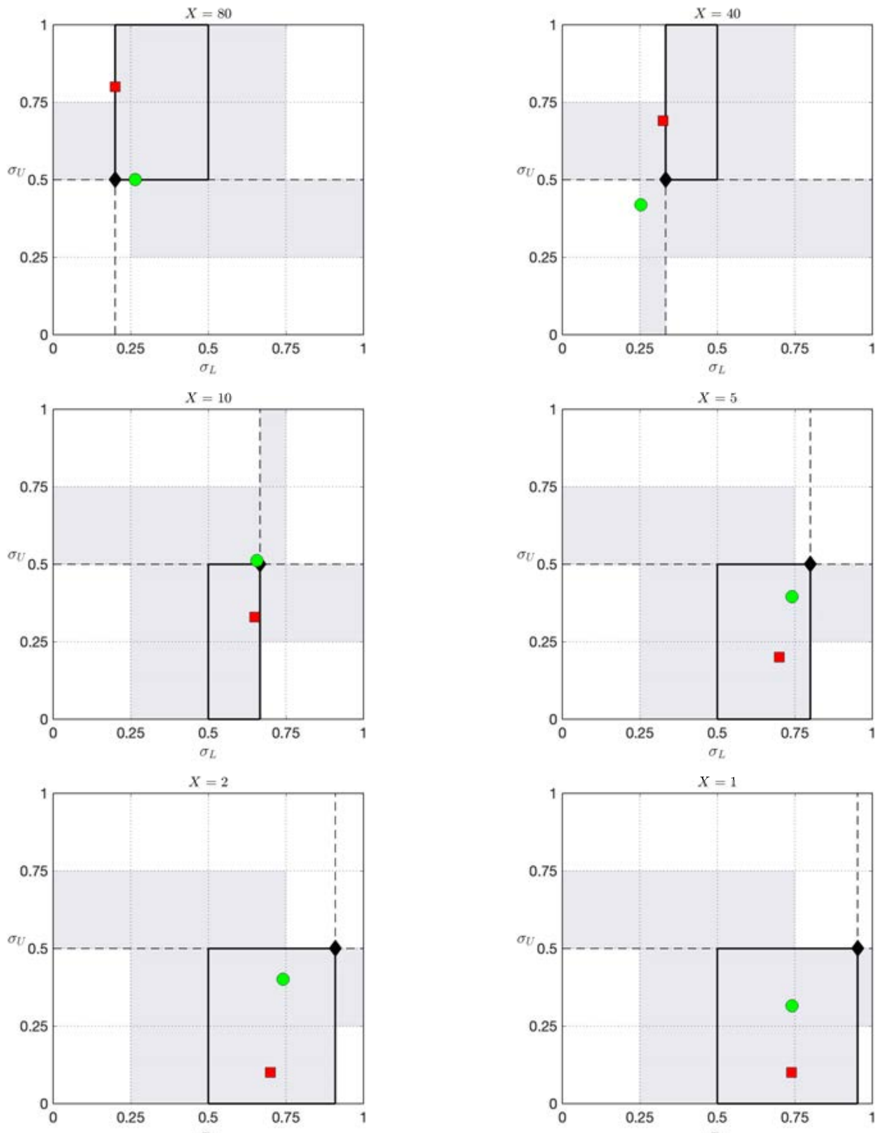

Figure 21: *Hybrid equilibrium and the data.* For each X-game, we plot the set of hybrid equilibria (gray), the set of QRE and NBE (black outline), the empirical action frequencies from [A, ∘] (green circle), median belief (red square), and Nash equilibrium (black diamond).

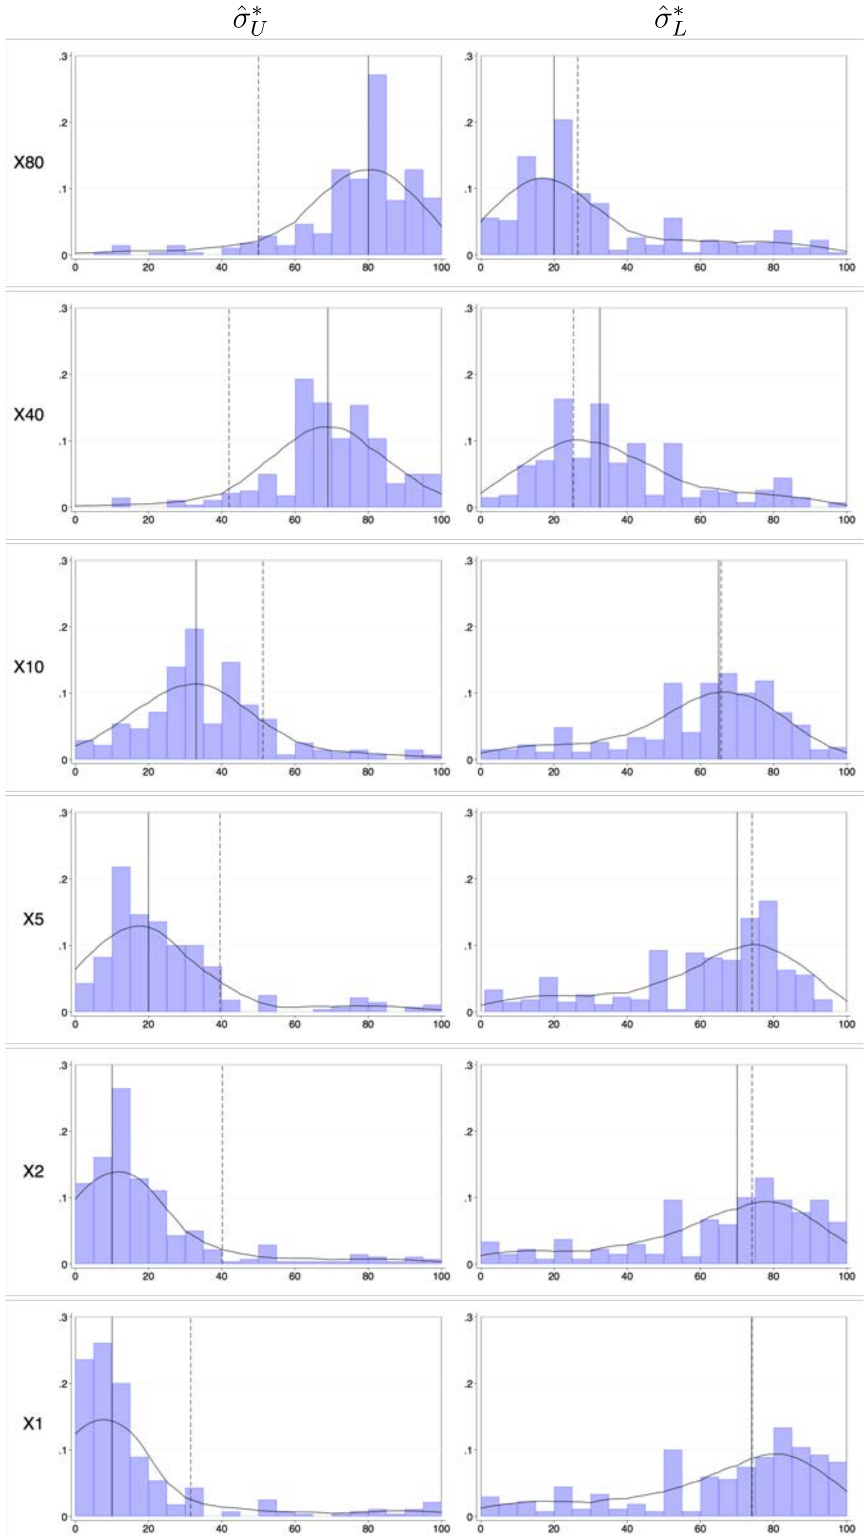

Figure 22: *belief-distributions*. The left panels are for player 2's beliefs over σ_U, and the right panels are for player 1's beliefs over σ_L. The solid lines mark the median of i's beliefs and the dashed line marks the empirical frequency of j's actions.

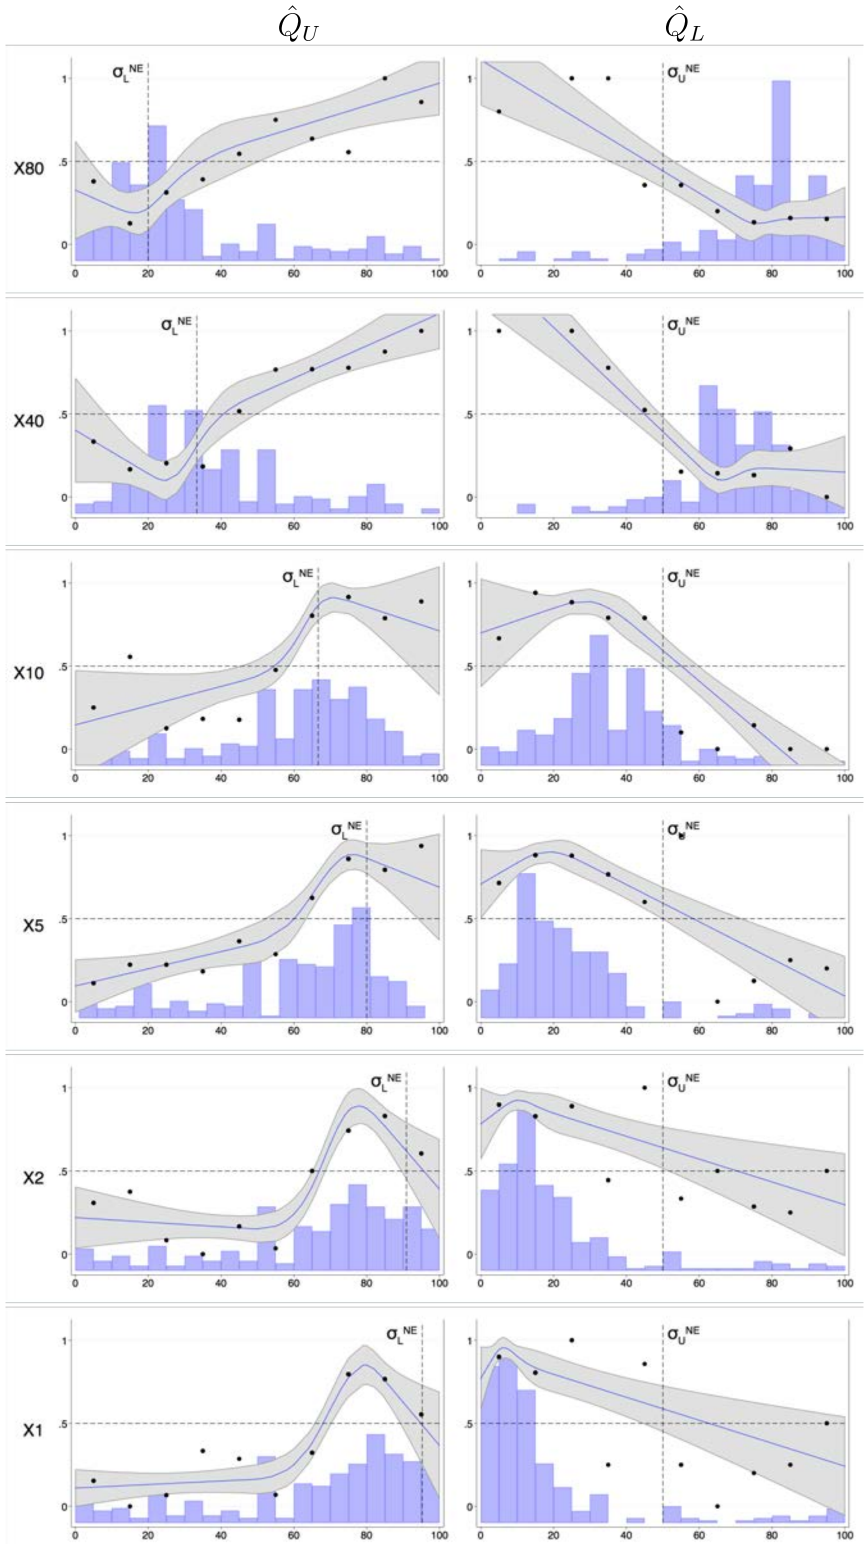

Figure 23: *Action frequencies predicted by beliefs.* For each player and X-game, we plot the predicted action probability as a function of stated beliefs based on restricted cubic spline regressions (4 knots at belief-quintiles, standard errors clustered by subject).

11.7 Additional tables

		$X80$	$X40$	$X10$	$X5$	$X2$	$X1$
actions	$\hat{\sigma}_U^{[\underline{\mathrm{A}},\circ]}$	0.50	0.42	0.51	0.40	0.40	0.31
	$\hat{\sigma}_L^{[\underline{\mathrm{A}},\circ]}$	0.27	0.25	0.66	0.74	0.74	0.74
	$\hat{\sigma}_U^{[\mathrm{A},\mathrm{B}\underline{\mathrm{A}}]}$	0.38	0.39	0.65	0.61	0.51	0.49
	$\hat{\sigma}_L^{[\mathrm{A},\mathrm{B}\underline{\mathrm{A}}]}$	0.21	0.22	0.74	0.79	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 11: *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 12: *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) 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 13: *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.

	Player 1					
	$X80$	$X40$	$X10$	$X5$	$X2$	$X1$
Between MS	2,047.8	1,553.1	1,681.1	1,816.0	2,300.6	2,403.2
Within MS	180.1	155.7	171.2	239.7	236.3	273.7
F	11.4***	10.0***	9.8***	7.6***	9.7***	8.8***
χ^2	227.7***	145.9***	161.4***	202.5***	222.2***	207.1***
	Player 2					
	$X80$	$X40$	$X10$	$X5$	$X2$	$X1$
Between MS	982.8	948.9	974.7	1,057.0	1,105.5	1,009.5
Within MS	121.6	82.7	159.1	181.7	188.3	321.6
F	8.0***	11.5***	6.1***	5.8***	5.9***	3.1***
χ^2	189.9***	123.8***	183.6***	238.8***	327.5***	405.6***

Table 14: *Analysis of Variance of stated beliefs.* For each player role and X-game, we report the results of ANOVA tests: estimated variance of subjects' average beliefs (Between MS), estimated average variance an an individual subject's beliefs (Within MS). This shows that, for every game, there is more between- than within-subject variation and that both types of variation are higher for player 1 than player 2. For each player and game, we also report the results of F-tests of the null hypothesis that the means of all subjects' belief-distributions are equal (F) and of Bartlett's χ^2-tests of equal variances (χ^2). All tests strongly reject the null: means and variances are not the same for all subjects.

11.8 Experimental interface

Figure 24 shows an example round from the perspective of a player 1-subject (blue) in the first stage. 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 her answer before submitting. The subject may undo her action choice by clicking again on the highlighted row.

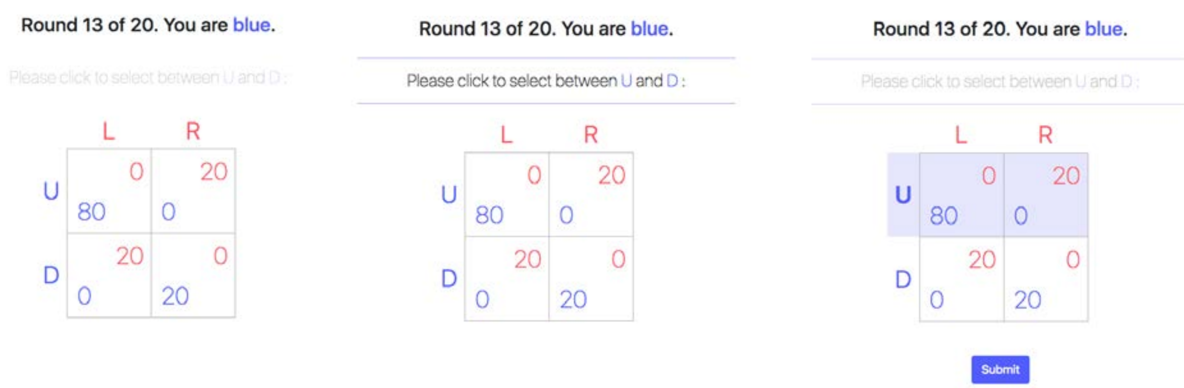

Figure 24: *Screenshots from first stage.*

Figure 25 shows an example round from the perspective of a player 1-subject (blue) in the second stage of [A,BA]. 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 her belief and action before submitting. After a belief is entered and an action is selected, the 'Submit' button

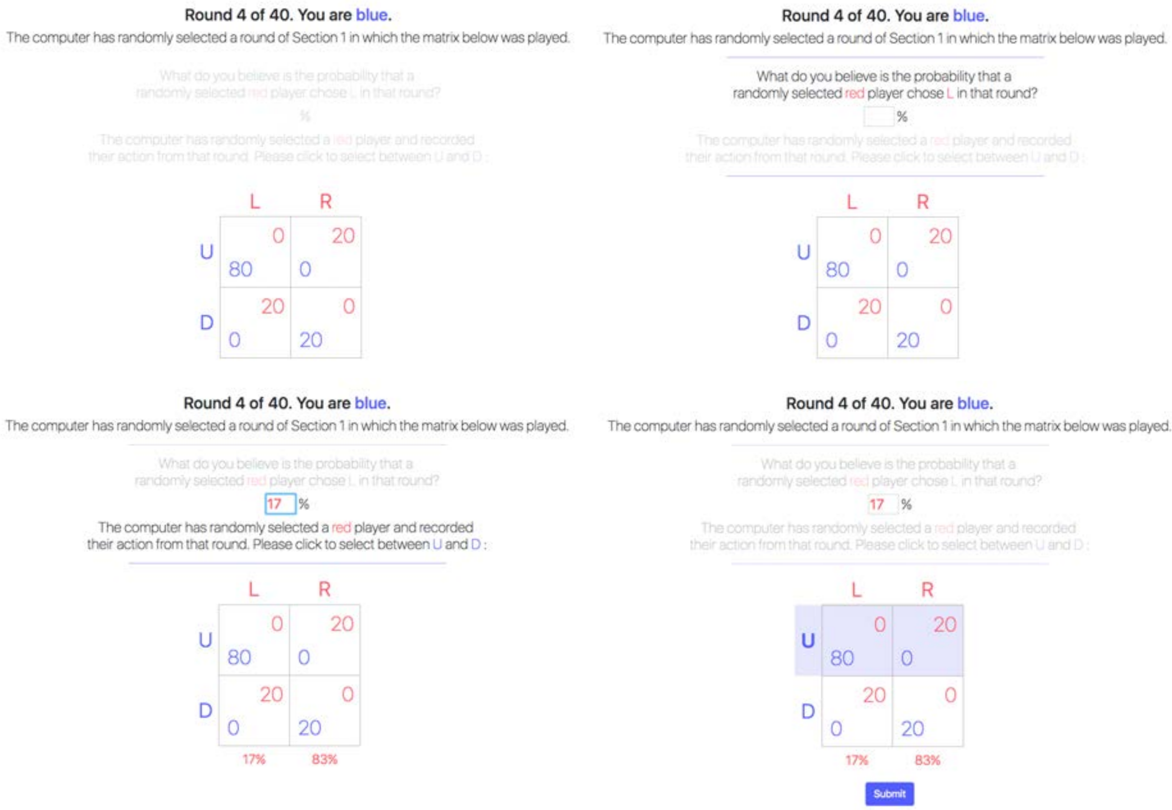

Figure 25: *Screenshots from second stage of [A,BA].*

appears (bottom-right screen). Figure 26 shows screenshots for the second stage of [A,A] which is the same as that of [A,BA], except beliefs are not elicited.

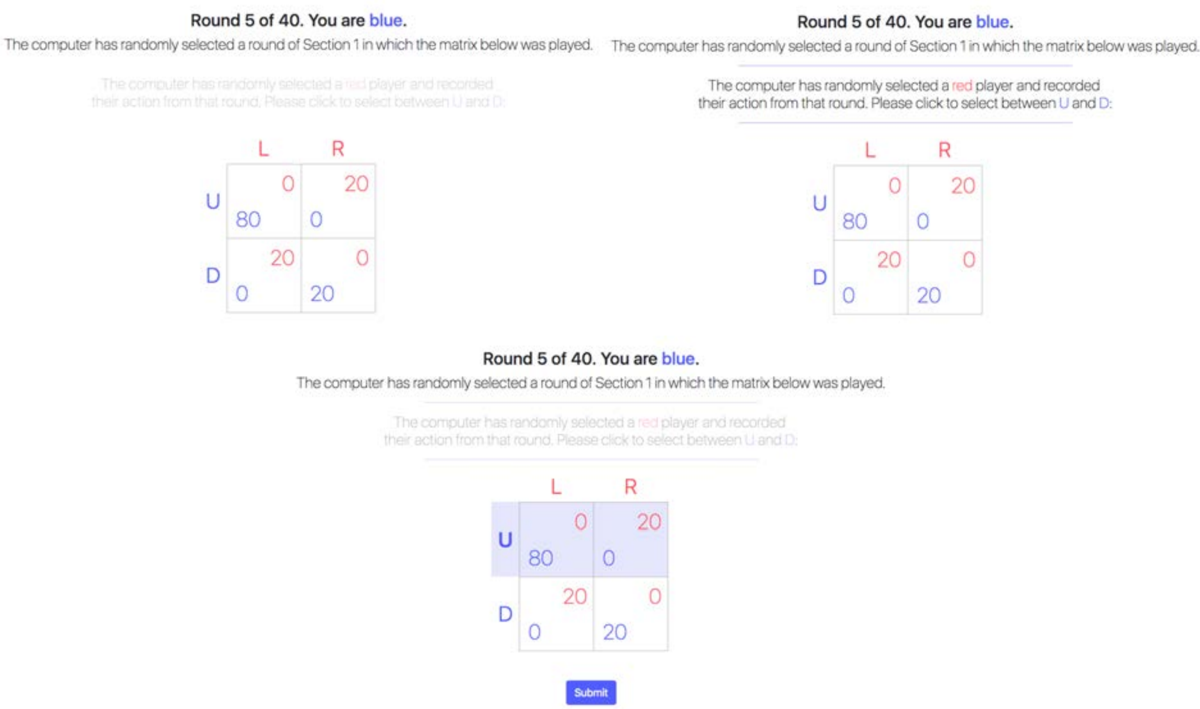

Figure 26: *Screenshots from second stage of [A,A].*

11.9 Games and randomization

In addition to the X-games, we also included the games whose payoffs are in Table 15. $D1$ 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.

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 15: *Additional games.*

Table 16 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.

Stage	Games	Rounds of each	Total 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 16: *Games by section.*