

***Strategic Problems with Risky Prospects:  
Modeling and Testing the Impact of Feedback in Complex Interactions***

*Alessandro Sontuoso,<sup>a,b</sup> Cristina Bicchieri,<sup>b</sup> Alexander Funcke,<sup>b</sup> and Einav Hart<sup>c</sup>*

May 20, 2024

**Abstract.** We study decision-making in interactions where the presence of multiple unknowns makes for complex hypothetical reasoning. Specifically, we propose a threshold game played by a large population of players, where strategic uncertainty is compounded by risky prospects (i.e., chance moves with commonly-known conditional probabilities). In the first half of the paper, we provide a Bayesian model of decision-making that accounts for players' inherent risk preferences. In the second half of the paper, we test the model via an experiment that verifies if decision quality is affected by exposure to noisy feedback. The results confirm that the feedback brings about a belief revision in the expected direction. Yet, contrary to standard predictions, the data indicate that the feedback induces also a decrease in "behavioral errors"; that is, subjects who receive feedback are less likely to fail to respond optimally to their beliefs, regardless of whether their beliefs are correct or not.

**KEYWORDS:** Behavioral errors; Hypothetical reasoning; Strategic uncertainty; Belief revision; Risk; Complexity; Bayesian games.

JEL Classification Numbers: C72, C92, H40, D83.

**Acknowledgments:** We are grateful to Zev Berger, Erik Kimbrough, David Rojo Arjona, and Leeat Yariv for their helpful comments. Also, we thank Jayson Dorsett for running the experimental sessions.

**Contact:** <sup>a</sup>Department of Economics, City University of London, Northampton Sq., London EC1V 0HB. <sup>b</sup>Center for Social Norms and Behavioral Dynamics, University of Pennsylvania, 249 S. 36th St., Philadelphia, PA, 19104. <sup>c</sup>School of Business, George Mason University, 4400 University Dr., Fairfax, VA, 22030. E-mail: a.sontuoso@gmail.com; cb36@sas.upenn.edu; funcke@0z.se; ehart8@gmu.edu.

## I. Introduction

The assumption that an agent is a (subjective) expected utility maximizer underlies most economic models and policy applications. In combination with expected utility theory, a common assumption is that an agent processes new information by consistently revising her beliefs according to some rule (e.g., Bayesian or otherwise). We refer to a theory that incorporates those two broad assumptions as an *updated* expected utility (henceforth UEU) framework. While this is the most common paradigm for examining choice behavior in both non-strategic and strategic (i.e., interactive) domains, violations of the UEU framework are not infrequent: typically such violations are ascribed to two main causes, that is, agents fail to update their beliefs,<sup>1</sup> or exhibit non-standard risk (and ambiguity) attitudes.<sup>2</sup> Still, it has been shown that some UEU violations may arise from the complexity of the problem, in a way that is not attributable solely to failures in belief updating or to non-standard attitudes; for example, data from non-strategic problems suggest that the complexity of a task often causes difficulties in hypothetical reasoning, which result in deviations from expected-value maximization (e.g., Charness and Levin, 2005; Zizzo, Stolarz-Fantino, Wen, and Fantino, 2000). This paper fits in the broad literature on suboptimal choices by studying a novel instance of a complex environment, where outcomes depend on chance and on the behavior of a large population of agents.

Research in cognitive science shows how decision-makers are especially prone to errors in complex environments: indeed, when facing a complex choice task, people have a tendency to idiosyncratically account for some contingencies and ignore others, as a consequence of limited working memory and cognitive load (Legrenzi, Girotto, and Johnson-Laird, 1993). In this regard, it has been shown that cumbersome tasks can deplete a decision-maker's cognitive resources and thus lead to more instinctive, less methodical choices (Kahneman, 2003, 2011). Recent research in behavioral economics has further shown that people's tendency to make suboptimal choices increases with the problem's computational difficulty and – all else equal – with the presence of

---

<sup>1</sup> Agents exhibit *unresponsiveness to information* because they ignore signals due to inattention, or otherwise feature a sticky belief-updating process. Specifically, unresponsiveness may be explained by the fact that some information is more readily “available” than other (Tversky and Kahneman, 1973). Further – when processing signals drawn from a sample – one may misinterpret information (Rabin and Schrag, 1999) or else misunderstand the mechanism that generates signals (Benjamin, Rabin, and Raymond, 2016); alternatively, one's belief might be affected by some form of positive reinforcement (Erev and Roth, 1998). Other biases inducing unresponsiveness to information are surveyed in Rabin (2013).

<sup>2</sup> For an overview of non-standard risk attitudes, see O'Donoghue and Somerville (2018); for an extensive review of models accounting for non-neutral ambiguity attitudes, see Machina and Siniscalchi (2013).

hypothetical contingencies (as contrasted with an equivalent task without hypothetical contingencies; Martínez-Marquina, Niederle, and Vespa, 2018).

Although the social welfare loss resulting from suboptimal choices in complex interactive environments can be substantial (think of the wrong choice of mutual funds or health-care plans), it is unclear how to improve the quality of an individual's decision *without* altering the structure of the choice task or its incentives (e.g., by exogenously simplifying the task, or by adding an incentivized third-party adviser). The question of how to effectively improve decision quality is especially relevant in cases where the outcome of the choice task is not observed immediately, and hence experiential learning is not possible. With this in mind, we set out to investigate if decision quality may be affected by a common feature of our everyday decision-making experience, that is, exposure to noisy feedback about the choices of others.

To assess whether this kind of information leads to better or worse decision-making in complex choice tasks, below we present theoretical and experimental analyses of a game *with* and *without* feedback about (a sample of) fellow participants' choices. Our analyses will address the question of whether feedback impacts the frequency of suboptimal choices, that is, failures to respond optimally to information. In this respect, there are two plausible yet opposing arguments as to why the occurrence of such failures may vary with and without transmission of feedback. One argument implies that receiving feedback – while attending to a complex choice task – may aggravate the agent's information overload, leading to some cherry-picking of data and thus to biased beliefs; accordingly, more information would not improve decisions at best, or could even worsen them (e.g., Wilson, 2014; Hall, Ariss, and Todorov, 2007). Conversely, another argument implies that receiving feedback does not only promote learning, but also enhances the agent's motivation and engagement with the task, which in turn improves the agent's outcomes (e.g., Compte and Postlewaite, 2004; Fischer and Sliwka, 2018). It is important to note that these arguments are usually put forward in specific decision-theoretic contexts, and it is unclear how they might generalize to other contexts. Our study sheds light on these predictions regarding the (positive versus negative) impact of noisy feedback on decision quality: we do so by considering a novel environment that is readily amenable to real-world applications.

To that end, we propose a class of games where the presence of many unknowns (i.e., strategic uncertainty and risky prospects) makes for complex hypothetical reasoning.<sup>3</sup> Specifically, we propose an interactive problem in the form of a threshold game, which for ease of exposition we illustrate in terms of a vaccination decision: in this game, outcomes depend on the population's behavior and on a move by nature, as follows. *If* the frequency of a socially-undesirable action reaches a given threshold (e.g., too many individuals in the population do not vaccinate), then a random shock will occur with a commonly-known conditional probability (an epidemic outbreak may or may not happen). *If* instead that threshold is not reached (e.g., enough individuals choose to vaccinate), outcomes will not depend on any chance element. The game presents players with two strategic options (do or do not vaccinate), in addition to an exit option with a sure payoff that is independent of others or nature (i.e., the action any risk- or ambiguity-averse individual would strictly prefer).

The first part of the paper lays out a Bayesian model of optimal decision-making in the above-described threshold game. In particular, since optimal actions depend on the individual's intrinsic risk preferences, our model formally defines different player types based on alternative risk attitudes. Given this, we provide an equilibrium analysis of the players' choice behavior, conditional on their risk types and beliefs; to account for idiosyncratic irregularities, we then incorporate a stochastic error-term into the players' utility functions. Lastly, the second part of the paper provides an experiment to verify if behavior in complex games can be accurately described by a model of discrete choice with random errors. Of particular interest is the question of how errors are impacted by noisy feedback: does the frequency of suboptimal choices increase or decrease with feedback?

For that purpose, we present a between-subjects design where subjects are randomly assigned to treatments with and without noisy information about fellow participants' choices, respectively referred to as the "main" and "control" treatments. Participants play multiple rounds of the above-described threshold game, against a large population of fellow participants. In both treatments – after each play – we elicit individual beliefs about the population's behavior (the realized outcomes of each play would be revealed only at the end of the entire session).

---

<sup>3</sup> In keeping with conventional practice, we refer to risky prospects as the case in which the objective probability of random events is available. Such a probability is *not* available in the case of uncertain (i.e., ambiguous) prospects: with specific reference to interactive problems, "strategic uncertainty" indicates a lack of knowledge about other players' actions; note that this may arise even in complete-information games, whenever there are multiple equilibria (e.g., think of any coordination games, threshold public-goods games, etc.; Brandenburger, 1996).

Treatments differ solely in the feedback dimension, as follows. After each play, participants in the main treatment are informed about some of the actions that were just chosen by a randomly-generated sample of fellow participants (with this information amounting to a noisy indicator of the opponents' risk preferences). No such feedback is provided in the control treatment.

This design allows us to test if (i) participants in the *main* treatment update their beliefs about the population's behavior based on the feedback received, and (ii) if participants in the *main* or *control* treatment ultimately commit "behavioral errors", which we define as one's failure to best-respond to one's stated belief, regardless of whether the belief is correct or not.

To address point (i), our analysis verifies whether on average participants who receive feedback (in the main treatment) respond to it in a way that is consistent with Bayesian updating. The results show that this is indeed the case, ruling out collective failures in belief updating as a potential violation of the updated expected utility (UEU) paradigm in our data.

Our analysis addresses point (ii) via the following multi-pronged approach. First, we check whether participants in either treatment make behavioral errors (i.e., fail to best-respond to their beliefs about the population's behavior). Two classes of behavioral errors are identified, with one class applicable to all players irrespective of their risk profile: approximately 9 percent of the choice data fall into this category of errors, across treatments.

When comparing treatments against each other, the data reveal that participants in the main treatment are significantly less likely to commit behavioral errors than those in the control treatment. That is, feedback decreases the likelihood that one will fail to respond optimally to one's stated beliefs. Notably, since this pattern occurs regardless of how correct or incorrect beliefs are, it cannot merely be attributed to a belief revision.

Overall, our findings demonstrate a clear improvement in decision quality resulting from the main treatment. As such, the data refute the argument that presenting decision-makers with additional information (during a complex choice task) exacerbates their cognitive load. On the contrary, the results support the opposing perspective, suggesting that noisy feedback not only facilitates learning but also holds motivational value. In fact, the improvement in decision quality occurs despite the fact that our feedback mechanism provides relatively little informational content (after each play, our subjects do not receive any information about payoffs, nor do they receive any third-party advice, unlike previous studies). It is also worth noting that our feedback mechanism does not in itself reduce the task's computational difficulty, and hence the feedback's beneficial

effect on decision quality cannot be fully explained by common characterizations of learning. For these reasons, we ascribe some of the improvement to increased motivation and engagement; our robustness checks substantiate this interpretation, showing that this effect grows with the number of iterations, and can reduce the likelihood of certain errors by almost 50 percent.

Before proceeding, we note that this paper is related to a burgeoning literature indicating that the complexity of a choice problem can cause important deviations from expected-value maximization in non-interactive lottery tasks (e.g., Charness and Levin, 2005; Zizzo et al., 2000). Complementary findings in the domain of non-strategic problems have shown that suboptimal behavior may be due to difficulties with “hypothetical reasoning” (the act of considering alternative strategically-relevant contingencies: Charness and Levin, 2009; Esponda and Vespa, 2014; Levin, Peck, and Ivanov, 2016; Martínez-Marquina, Niederle, and Vespa, 2018). This literature suggests that when problems are complex due to a multiplicity of contingencies, decision-makers tend to overlook the implications of those contingencies.<sup>4</sup> In this respect, our study contributes to the understanding of complex decision-making by investigating an entirely distinct problem, played against multiple participants, with a novel feedback mechanism.

Summing up, here we present theoretical and experimental analyses of a novel game (with and without feedback) that is readily amenable to business and policy applications. The results underscore a beneficial impact of feedback, which causes a significant drop in expected utility maximization failures (with participants in the main treatment earning more money on average, relative to the control). In particular, the data show that subjects who receive noisy feedback update their beliefs in the expected direction; additionally, they are more likely to respond optimally to their beliefs, regardless of whether their beliefs are correct or not. In a nutshell, our results validate the argument that – besides promoting learning – feedback enhances one’s motivation and engagement with complex tasks. In terms of policy implications, these results shed light on the role of feedback in improving decision-makers’ assessment of risk in complex interactions, and thus suggest new strategies for curbing improper risk-taking via targeted information campaigns.

---

<sup>4</sup> Relatedly, earlier experiments on bidding behavior in common-value auctions show that participants are less prone to behave suboptimally, if the auction format allows them to observe others’ bidding (see, among others, Kagel, 1995; Levin, Kagel, and Richard, 1996). Also, evidence from sequential games with perfect information indicates that many decision-makers do not fully grasp the strategic consequences of their own predictions about others’ behavior *until* the outcome is observed; so, people fail to maximize payoffs because they do not account for the implications of events that have not yet been observed (Hedden and Zhang, 2002; see also Weber, 2003).

The rest of the article is organized in this manner: section II lays out the theoretical model; section III introduces the experimental design and hypotheses; section IV presents the data, and section V concludes.

## II. The model

### 1. *Threshold games with risky prospects*

Here we lay out a class of  $n$ -player games with risky prospects, and later provide an application.

Let  $N = \{1, \dots, n\}$  denote the set of players. For each player  $i \in N$ , payoffs depend on  $i$ 's action  $s_i$ , on the co-players' actions  $s_{-i}$ , as well as on a chance move  $\theta$ . Formally,  $i$ 's payoff is given by  $m_i(s, \theta)$ , where  $s = (s_i, s_{-i})$  denotes an action profile (i.e., an  $n$ -tuple of actions) and  $\theta$  denotes a move by nature. For each  $i \in N$ , let  $s_i \in S_i = \{A, B, C\}$  and  $s_{-i} \in S_{-i} = \times_{j \in N, j \neq i} S_j$ . Next, let  $\theta \in \Theta = \{heads, tails\}$ : this may be thought of as a coin toss, where each outcome has a 50 percent probability, which is common knowledge among players.

Each player *simultaneously* chooses an action in  $\{A, B, C\}$  with the goal to maximize her expected utility, prior to observing the outcome of the coin toss. In defining the payoffs  $m_i(s, \theta)$ , we partition the set of action profiles on the basis of whether a threshold  $d\%$  is or is not eventually met: we specify such a threshold with respect to the population-level frequency of choice of one of the actions (namely, action  $B$ ). That is, we write  $s \in \underline{S}$  if *less than*  $d\%$  of all players choose  $B$ ; instead, we write  $s \in \bar{S}$  if  *$d\%$  or more* of all players choose  $B$ .<sup>5</sup> Given the above, we define a generic “threshold game with risky prospects” as a strategic problem where payoffs vary according to the following constraints:

$$m_i(C, s_{-i}, heads) = m_i(C, s'_{-i}, tails), \text{ for any } s_{-i}, s'_{-i} \in S_{-i}; \quad (1)$$

$$m_i(s_i, s_{-i}, heads) = m_i(s_i, s_{-i}, tails), \text{ for any } s \in \underline{S}; \quad (2)$$

$$m_i(s_i, s_{-i}, heads) \neq m_i(s_i, s_{-i}, tails), \text{ for any } s \in \bar{S}; \quad (3)$$

$$\mathbb{E}[m_i(A, s_{-i}, \theta)] = \mathbb{E}[m_i(B, s_{-i}, \theta)] = \mathbb{E}[m_i(C, s_{-i}, \theta)], \text{ for any } s \in \bar{S}. \quad (4)$$

In plain words, condition (1) says that  $i$ 's payoff from  $C$  is independent of both the co-players' actions and the coin toss; (2) says that, for  $s \in \underline{S}$ , the payoff does *not* depend on the coin toss while (3) says that, for  $s \in \bar{S}$ , the payoff does depend on the coin toss; condition (4) says that, conditional on  $s \in \bar{S}$ , the expected value of the payoffs from  $A$ ,  $B$  or  $C$  is the same.

---

<sup>5</sup> Formally,  $\underline{S} := \{s \in S : \exists M \subset N \text{ s.t. } s_j = B \text{ for } \forall j \in M \text{ and } \frac{|M|}{|N|} < d\%\}$ , where  $|\cdot|$  denotes the cardinality of a set. Further,  $\bar{S} := S \setminus \underline{S}$ , where  $\setminus$  denotes set difference.

Below we consider the specific parameterization that we used in the experiment, and illustrate it in terms of a vaccination problem. To that end, we set  $d\% = 0.4$ , with this value being common knowledge among all players (note that this is an arbitrary value, and our experimental hypotheses do not specifically depend on it). Accordingly, in what follows  $\underline{S}$  denotes the set of action profiles where “less than 40% of all players choose  $B$ ”, and  $\bar{S}$  denotes the set of action profiles where “40% or more of all players choose  $B$ ”. Finally, suppose that the coin has been tossed and all the players have simultaneously chosen an action: depending on whether the action profile  $s$  is contained in  $\underline{S}$  or  $\bar{S}$ , and depending on the coin toss  $\theta$ , one of the below contingencies (“scenarios”) obtains; in each case, player  $i$ ’s payoff from  $A$ ,  $B$  or  $C$  is defined in the respective column of Table 1.<sup>6</sup>

	<b>A</b>	<b>B</b>	<b>C</b>
$(s, \theta)$	[ <i>vaccinate</i> ]	[ <i>don't vaccinate</i> ]	[ <i>self-isolate</i> ]
$\underline{S} \times \{\textit{heads, tails}\}$ “scenario $x$ ” [ <i>herd immunity</i> ]	0.5	3	0.75
$\bar{S} \times \{\textit{heads}\}$ “scenario $y$ ” [ <i>epidemic outbreak</i> ]	1	-1.5	0.75
$\bar{S} \times \{\textit{tails}\}$ “scenario $z$ ” [ <i>no outbreak</i> ]	0.5	3	0.75

**Table 1** - A threshold game with risky prospects, illustrated in terms of a vaccination problem. Each row refers to a possible scenario  $(s, \theta)$  while each of the three rightmost columns refers to  $i$ ’s action  $s_i$ : possible payoffs to  $i$  are reported therein. Note:  $s \in \underline{S}$  if less than 40% of all players choose  $B$ , whereas  $s \in \bar{S}$  if 40% or more of all players choose  $B$ .

<sup>6</sup> For an interpretation, consider an individual’s decision of whether to get a vaccine on the verge of a potential epidemic outbreak. In that case, by choosing  $C$  (*self-isolate*) one is no longer susceptible to becoming infected with the virus. On the other hand, by choosing  $A$  (*vaccinate*) one contributes – at a cost to oneself – to the public good, namely herd immunity: here, one’s utility depends on the others’ decisions (which determine if the threshold is met) as well as on the outcome of a move by nature. By choosing  $B$  (*don't vaccinate*) one negatively contributes to the threshold for herd immunity: again, one’s payoff depends on the population-level behavior and on a move by nature. More specifically, when less than 40% of all individuals free ride (that is, when more than 60% of the population get a vaccine or self-isolate), then a risk-free situation occurs (“scenario  $x$ ”, e.g., herd immunity). By contrast, whenever 40% or more free ride, then a negative shock *may* occur (“scenario  $y$ ”, e.g., epidemic outbreak) or *may not* occur (“scenario  $z$ ”, e.g., no outbreak), each with a 50 percent chance.

To sum up, players choose among three actions, ranging from a very risky option “*B*” (featuring the highest and lowest possible payoffs), to a mildly risky option “*A*”, to a riskless option “*C*”. It should be noted that, conditional on  $\bar{S}$ , actions *A* and *B* are mean-preserving spreads of *C*. The reader can anticipate that the equilibria of the game vary with players’ beliefs about the others’ behavior *and* with their risk preferences. Below we define such preferences and formalize an environment where each player is uncertain about the (risk) preferences of others.<sup>7</sup>

## 2. Risk preferences

We now introduce our account of risk preferences. Before doing so, it is worth noting that *no* move by nature is triggered if the threshold is not met (for any action profile  $s \in \underline{S}$ , that is, when less than 40% of all players choose *B*): in that case, risk attitudes do not play any role and all the players’ preferences reflect the ordering  $B > C > A$ , consistent with the payoffs specified in the first row of Table 1 (p. 7) under risk-free scenario  $x$ . By contrast, whenever the threshold is actually met (i.e., for any action profile  $s \in \bar{S}$ ), a chance move is expected to take place and thus risk attitudes come into play: here, different attitudes imply different preference orderings across players. In the following, we consider some common attitudes that have been widely documented in the literature.

One such attitude consists of risk aversion: accordingly, we assume that some players are represented by a *risk-averse* type, denoted  $t_c$ , where this type prefers the certain prospect with value  $m$  over any risky prospect with expected value  $m$  (in our case, 0.75). Further, we assume that the rest of the players exhibit a risk-seeking attitude: specifically, we consider two types of *risk-seeking* individuals, namely  $t_a$  and  $t_b$ , where each such type respectively prefers an action with a lower and higher payoff variance.

Formally, we denote a generic type by  $t$ , with  $t \in T$  and  $T = \{t_a, t_b, t_c\}$ . In each case, preferences over risky prospects conditional on the threshold being met (that is, for any  $s \in \bar{S}$ ) are defined by the following orderings:

$$A \succ_{t_a} B \succ_{t_a} C, \tag{5}$$

$$B \succ_{t_b} C \succ_{t_b} A, \tag{6}$$

$$C \succ_{t_c} A \succ_{t_c} B, \tag{7}$$

---

<sup>7</sup> Note that, given the payoff structure of the game, the common assumption that every player is *risk-neutral* is behaviorally implausible here. Indeed, if such an assumption held, then action *B* would be weakly dominant for all the players (a fact that is clearly confuted by the experimental data).

where  $\succ_t$  refers to the preference relation of type  $t$ .

A few comments. First, note that each  $t$  identifies a player's own (privately-known) preferences, and not her information about others, which will be defined later. Also, note that *risk-neutral* players are indifferent between the three actions (conditional on the threshold being met): for this reason, without loss of generality we focus on cases where individuals behave as either risk-seeking or risk-averse players. Finally, note that since (in a given game) players are prompted to choose one action, and all three actions are available at once, it follows that – conditional on the threshold being met – expression (5) is observationally equivalent to  $A \succ C \succ B$ ; analogous arguments apply to (6) and (7). For simplicity, we therefore restrict attention to preference orderings (5)-(7).

In the next subsection, we will represent this interactive choice problem as a Bayesian game (i.e., a game with incomplete information about the others' preferences). To that end, we must cardinalize preference orderings (5)-(7) so as to reflect how each type computes the expected utility from action  $s_i$  conditional on  $(s_i, s_{-i}) \in \bar{S}$ . Thus, for each  $t \in T$  we assume that there exists a function  $u_t$ , such that iff  $(s_i, s_{-i}) \succ_t (s'_i, s_{-i})$  with  $s \in \bar{S}$ , then  $\sum_{\theta \in \Theta} \Pr(\theta) u_t(m_i(s_i, s_{-i}, \theta)) > \sum_{\theta \in \Theta} \Pr(\theta) u_t(m_i(s'_i, s_{-i}, \theta))$ , where  $\Pr(\theta) = 0.5$  for each  $\theta \in \Theta = \{heads, tails\}$ .<sup>8</sup>

For ease of reference, below we respectively denote by  $\alpha(t)$ ,  $\beta(t)$ ,  $\gamma(t)$  the expected utility values from  $A$ ,  $B$ ,  $C$  conditional on  $s \in \bar{S}$ .<sup>9</sup> Given this, preference orderings (5)-(7) are equivalent to inequalities (8)-(10), respectively:

$$\alpha(t_a) > \beta(t_a) > \gamma(t_a), \tag{8}$$

$$\beta(t_b) > \gamma(t_b) > \alpha(t_b), \tag{9}$$

$$\gamma(t_c) > \alpha(t_c) > \beta(t_c), \tag{10}$$

with  $\alpha(t), \beta(t) \in \mathbb{R}$  and  $\gamma(t) = 0.75$  for each  $t \in T$ .

### 3. Uncertainty about the others' preferences

We can now turn to lay out a framework for capturing  $i$ 's uncertainty about the preference type of the opponents.

---

<sup>8</sup> Of course, this cardinalization need does not arise when the threshold is *not* met (i.e., for any action profile  $s \in \underline{S}$ ). In that case, for all types  $t \in \{t_a, t_b, t_c\}$  it is assumed that player  $i$ 's utility values correspond to the monetary payoffs reported in the first row of Table 1 (p. 7); that is,  $m_i(A, \cdot) = 0.5$ ,  $m_i(B, \cdot) = 3$ ,  $m_i(C, \cdot) = 0.75$ .

<sup>9</sup> More explicitly, when the threshold *is* met (i.e., for any action profiles  $s \in \bar{S}$ ) the type-specific expected utility values from  $A$ ,  $B$ , and  $C$  are given by  $\alpha(t) := \sum_{\theta \in \Theta} \Pr(\theta) u_t(m_i(A, s_{-i}, \theta))$ ,  $\beta(t) := \sum_{\theta \in \Theta} \Pr(\theta) u_t(m_i(B, s_{-i}, \theta))$ , and  $\gamma(t) := \sum_{\theta \in \Theta} \Pr(\theta) u_t(m_i(C, s_{-i}, \theta))$ , respectively, for each  $t \in T$ .

We posit that the set of preference types  $T = \{t_a, t_b, t_c\}$  is commonly known, where each type  $t \in T$  is characterized by (8)-(10). We then denote the set of *states of the world* by  $\Omega$ , with generic member  $\omega \in \Omega = T^n$ ; that is, each state  $\omega$  corresponds to an  $n$ -tuple of types (i.e., one type per player, with  $n$  denoting the cardinality of the set of players  $N$ ). In other words, a state  $\omega$  is a possible profile of risk preferences, such that each player is assigned one of the three types in  $T = \{t_a, t_b, t_c\}$ .

Each player's type is assumed to be drawn independently from the same distribution  $\pi$ . For any player  $j \in N$ , with a slight abuse of notation we write the probability that  $j$  is of type  $t$  as  $\pi(t(j))$ , where  $\pi(t) \in [0, 1]$  for all  $t \in T$ , and  $\pi(t_a) + \pi(t_b) + \pi(t_c) = 1$ .<sup>10</sup> We further assume that the probability of  $j$  being assigned type  $t_b$  is itself a random variable: that is,  $\pi(t_b) \equiv \pi_H$  with probability  $p$  while  $\pi(t_b) \equiv \pi_L$  with probability  $1 - p$ , where  $p \in (0, 1)$  and  $0 < \pi_L < 0.4 < \pi_H < 1$ . Similarly, the probability of  $j$  being assigned either type  $t_a$  or  $t_c$  equals  $1 - \pi_H$  with probability  $p$ , and equals  $1 - \pi_L$  with probability  $1 - p$ .

The reader can anticipate that such priors inform the optimal behavior of  $t_a$  and  $t_c$  players: in fact, the preferences of these two types vary with their expectation about the amount of  $B$  choices at the population level; in turn, this expectation depends on the share of  $t_b$  players in the population (note that  $B$  is a dominant action only for  $t_b$  players).

#### 4. Feedback

Here we describe noisy signals about the realized profile of preference types (at state  $\omega$ ).

Let  $f$  denote a piece of feedback, with  $f \in F$  and  $F = \{low, high\}$ , where  $low \equiv \pi_L$  and  $high \equiv \pi_H$ . Given this, we define a function  $\tau_i: \Omega \rightarrow F$ , where  $f = \tau_i(\omega)$  is interpreted as a signal about the others' preferences, that is received (privately) by  $i$  in the form of noisy feedback about past play. Under the assumption that preferences do not change drastically over a short time, then it follows that past choices can reveal the true profile of preference types, albeit with some noise. That is, we do *not* assume that the private signal conveys perfect certainty about whether the share of  $t_b$  players is exactly  $\pi_L$  or  $\pi_H$ ; rather, we assume that the signal correctly points to the true share of  $t_b$  players with some probability  $q$ , where  $q > 0.5$ .

---

<sup>10</sup> Under the assumption that player  $i$  knows her own type (and accordingly attaches probability zero to states that are inconsistent with it), then it follows that – for any  $i, j \in N$  – player  $i$ 's belief about the profile of preference types at state  $\omega$  is defined by  $\Pr[\omega] := \pi(t(i)) \cdot \prod_{j \in N, j \neq i} \pi(t(j))$ . Note that since  $i$  knows her own type, from  $i$ 's standpoint  $\pi(t(i))$  equals one at any state  $\omega$  that is consistent with  $i$ 's actual type and zero otherwise, for each  $t \in T$ .

In plain words, we interpret  $f = \text{low}$  (or  $\text{high}$ ) as feedback that – more likely than not – points to a state  $\omega$  such that the share of  $t_b$  players in the population is lower (or higher) than the 40% threshold. For a concrete illustration, suppose that you hear from the neighbors that the frequency of  $B$  choices in a previous time period amounted to a value  $< 40\%$ : in that case, the *low* feedback is interpreted as a signal that the share of unconditional  $B$ -choosers (i.e.,  $t_b$  players) in the population must be below the threshold, with probability  $q > 0.5$ . If instead you hear a value  $> 40\%$ , then the *high* feedback is interpreted as a signal that the share of unconditional  $B$ -choosers (i.e.,  $t_b$  players) in the population must be above the threshold, with probability  $q > 0.5$ . We will next see that such signals can be used to update a player’s belief about the realized profile of preference types at state  $\omega$ .

### 5. *Equilibrium analysis*

In a Bayesian Nash equilibrium, each player chooses the best available action, given her beliefs about the possible state *and* given the other players’ actions (according to each state), with a requirement that beliefs are correct.

We now proceed to express a player’s expected utility in the “interim” representation of the game (i.e., after receiving a signal about the state of the world). As per standard accounts of incomplete information, here player  $i$ ’s choice behavior depends on her own preference type and on her belief about the others’ preference types (at each state  $\omega$ ). Formally, let’s consider some generic player of type  $t(i)$ , and denote the other players’ actions at  $\omega$  by  $s_{-i}(\omega)$ ; then, player  $i$ ’s expected utility from  $(s_i, s_{-i}(\omega))$  is given by

$$\sum_{\omega \in \Omega} \sum_{\theta \in \Theta} \Pr[\omega|f] \cdot \Pr[\theta] \cdot u_{t(i)}(m_i(s_i, s_{-i}(\omega), \theta)), \quad (11)$$

where  $\Pr[\omega|f]$  indicates the probability assigned to state  $\omega$  conditional on  $f$ , calculated via Bayes’ rule. For ease of exposition, below we compactly denote expression (11) by  $U_i(s_i)$ .

Even though the set of states of the world  $\Omega$  is very large (with  $3^n$  a-priori possible states), we will see that – for the purposes of calculating  $i$ ’s expected utility – there is a convenient way to group together states  $\omega$  that are strategically similar, in the sense that their profile of risk types is such that the expected frequency of  $B$  choices is *either* above *or* below the 40% threshold. Indeed, as we noted before, the probability of either event is a function of the expected share of  $t_b$  players. In light of that, we shall denote the a-priori expected share of  $t_b$  players (i.e., as computed

before receiving any signal) by  $\bar{\pi}(t_b) = p \cdot \pi_H + (1 - p) \cdot \pi_L$ .<sup>11</sup> In the event of a signal, we can then use Bayes' rule to compute the a-posteriori expected share of  $t_b$  players,  $\bar{\pi}(t_b | f)$ . That is, conditional on a *high* signal, we have

$$\bar{\pi}(t_b | f = \text{high}) = \frac{pq}{pq+(1-p)(1-q)} \cdot \pi_H + \frac{(1-p)(1-q)}{pq+(1-p)(1-q)} \cdot \pi_L; \quad (12)$$

instead, conditional on a *low* signal, we have

$$\bar{\pi}(t_b | f = \text{low}) = \frac{p(1-q)}{p(1-q)+(1-p)q} \cdot \pi_H + \frac{(1-p)q}{p(1-q)+(1-p)q} \cdot \pi_L. \quad (13)$$

For ease of reference, below we generically indicate the expected share of  $t_b$  players (given some signal  $f$ ) by writing  $\bar{\pi}$ , where  $\bar{\pi} \equiv \bar{\pi}(t_b | f)$ . We now proceed to our first result, which concerns the Bayesian Nash equilibria of the game.

**Proposition 1—Bayesian Nash equilibria in pure actions.** In every Bayesian Nash equilibrium of a threshold game (with preference types as defined by (8)-(10) above), all  $t_b$  players choose  $B$ , whereas the actions chosen by  $t_a$  and  $t_c$  players depend on their beliefs, as follows.

- (i) If the expected share of  $t_b$  players is greater than a type-specific cut-off point  $\psi_t$ , with  $\bar{\pi} > \psi_{t_a} \equiv \frac{5}{2\alpha(t_a)-2\beta(t_a)+5}$  and  $\bar{\pi} > \psi_{t_c} \equiv \frac{9}{12-4\beta(t_c)}$ , then all  $t_a$  players choose  $A$  and all  $t_c$  players choose  $C$ .
- (ii) When the first (second) [both] of the inequalities above is/are not satisfied, then a fraction  $g$  of  $t_a$  players (of  $t_c$  players) [of  $t_a$  or  $t_c$  players] will choose  $B$  iff  $\bar{\pi} + (1 - \bar{\pi})g$  is less than  $\psi_{t_a}$  (less than  $\psi_{t_c}$ ) [less than  $\psi_{t_a}$  or  $\psi_{t_c}$ ], with  $0 \leq g \leq 1$ , and will choose  $C$  otherwise.

*Proof.* See Appendix A.

A brief commentary is in order. First, note that equilibrium (i) of Proposition 1 represents the case where  $t_a$  and  $t_c$  players expect that the share of  $t_b$  players is greater than a type-specific cut-off point  $\psi_t$  (each such cut-off intuitively reflects the expected utility values a type attaches to the risky actions). In that case, all  $t_a$  and  $t_c$  players prefer to choose respectively  $A$  and  $C$ , in order to avoid the possibly negative consequences attached to action  $B$ . By contrast, the class of equilibria (ii) represents the case where  $t_a$  or  $t_c$  players expect that there will *not* be that many

---

<sup>11</sup> A few comments. First, recall from section II.3 that the probability of someone being assigned type  $t_b$  is modeled as a random variable, which equals some value  $\pi_H$  (with  $0.4 < \pi_H < 1$ ) with probability  $p$ , or some value  $\pi_L$  (with  $0 < \pi_L < 0.4$ ) with probability  $1 - p$ . Also, recall that  $q$  denotes the probability that the feedback  $f$  correctly points to the true state, with  $q > 0.5$ . Finally, recall that the game is played by a very large population of players: thus, it is reasonable to posit that one's knowledge of one's own preferences will not affect one's expectation as to the share of  $t_b$  players.

$t_b$  players: this allows a fraction  $g$  of them to choose  $B$  themselves without passing their respective cut-off points.

## 6. *Stochastic errors*

In any Bayesian Nash equilibrium, each player chooses a best-response to her beliefs about the other types' behavior, under the assumption that each player's preference relation is rational, and thus stable. Yet experimental studies have long shown inconsistent choice patterns across repetitions in non-strategic problems, such as risky lotteries (e.g., Hey and Orme, 1994; Ballinger and Wilcox, 1997) and uncertain lotteries (e.g., Hey, Lotito, and Maffioletti, 2010). For a model to capture the occasional irregularity, random errors may be specified in such a way that one option is *more likely* (as opposed to *always*) preferred by a type – ceteris paribus – if that option yields a higher utility than another option. This class of models with random errors was originally proposed by Fechner (1860), and later studied by Becker, DeGroot, and Marschak (1963), among others (e.g., Luce's 1959 model is a special case of the class of Fechner models where random errors are drawn from the logistic distribution). In short, Fechner models combine discrete choice with random errors, which may be variously interpreted as arising from complexity, tiredness, or carelessness.

Building on the Fechner approach, here we incorporate errors into a strategic model. Before doing so, let's consider again a (standard) player that is *not* prone to errors: as we noted above, (in any equilibrium) each player  $i$  chooses a best-response to the others' expected actions; compactly, action  $s_i$  is a best-response for player  $i$  if  $U_i(s_i) - U_i(s'_i) \geq 0$  and  $U_i(s_i) - U_i(s''_i) \geq 0$ , ceteris paribus, for any  $s_i, s'_i, s''_i \in S_i = \{A, B, C\}$ , where  $U_i(\cdot)$  is defined as in (11). Now, assuming that  $i$ 's choice is actually affected by occasional errors, then we say that  $s_i$  is a best-response for  $i$  if

$$U_i(s_i) - U_i(s'_i) \geq \varepsilon'_i \text{ and } U_i(s_i) - U_i(s''_i) \geq \varepsilon''_i, \quad (14)$$

with  $\varepsilon'_i, \varepsilon''_i \geq 0$  denoting random errors that distort the comparison between pairs of actions.<sup>12</sup>

The general interpretation is as follows: if an action (for example,  $A$ ) is much more desirable in terms of its expected utility than the alternatives (namely,  $B$  and  $C$ ) ceteris paribus,

---

<sup>12</sup> Since errors occur inadvertently, we assume that  $i$  does *not* anticipate the possibility that players may commit any such errors. Further, we assume that the parameters of the distribution from which random errors are drawn vary with environmental stimuli. From an econometric viewpoint, it is worth noting that whenever random errors are drawn from the normal or logistic distribution, the Fechner model is equivalent to a homoscedastic probit or logit (for an extensive account of decision-theoretic models incorporating random errors in the spirit of Fechner, 1860, see Blavatsky, 2018).

then large random errors  $\varepsilon', \varepsilon''$  are necessary to prevent a player from choosing that action. Overall, our model implies that the players' choice behavior depends on their types, but still exhibits irregularities in the form of random errors. In the next section, we present an experiment that tests the framework above, with a goal to assess the extent to which players commit errors in a complex problem with (and without) feedback.

### III. Experimental design and hypotheses

#### 1. *Design and procedures*

Our experimental sessions were conducted at the University of Pennsylvania's Wharton Behavioral Lab. Upon arrival at the lab subjects were randomly allocated to computer terminals, where they expressed their consent to participate in an interactive decision-making experiment. On average, a session had about 17 subjects and lasted about 50 minutes. Each session consisted of the following stages: Introduction Stage; Play Stage; Payment Stage.

Below we describe the "main treatment".

**Introduction Stage.** After granting consent, subjects were asked to read the on-screen instructions; they were informed that they would go through a set of decision tasks, where each participant would be prompted to choose one of the actions represented by options on the screen, labeled as "A", "B", and "C". Each subject was instructed that any money she would earn (besides a flat participation fee) depended on her choice and on the choices made by all other participants in the lab session, as well as on the outcome of a (fair) coin tossed by the computer. In particular, each subject was informed that the choices made by all other participants – together with the outcome of a coin flip – would determine one of three scenarios. After reading the instructions, subjects were required to answer a set of comprehension questions.

Before moving on to the Play Stage, a few comments are due. First, we stress that participants' actions were simply denoted by "A", "B", "C" (i.e., no reference was made to vaccinations, epidemic outbreaks, etc.). Second, *letter-outcome pairs* (e.g., whether the letter *B* actually denotes the socially-undesirable option rather than, say, the exit option) were *randomized across participants*. This was done in order to control for the fact that letters that come first in the alphabet may be perceived as more prominent. (For an instance of the experimental instructions featuring alternative letter-outcome pairs, please refer to Appendix B.)

For ease of exposition, in addressing the subjects' actions, the remainder of the paper will use the exact same letter-outcome pairs as in section II above.

**Play Stage.** All plays were conducted using Behavory (<https://behavory.com/>), a software for laboratory, online, and field experiments. The order of tasks was as reported below.

- (i) Each subject was asked to choose one of the options “A”, “B”, or “C”. Subjects were told that, after all participants had made their choices, a fair coin would be tossed by the computer and the scenario for the current play would be determined (i.e., the same scenario for all participants). Note that subjects were not informed of the scenario they were in, either before or after making decisions.
- (ii) Each subject was prompted to guess how many participants in the same session chose the option corresponding to the socially-undesirable action. Thus, (in the case of the letter-outcome pairs of Table 1, p. 7) the task read as follows: “... *indicate the percentage of the participants in the entire room that you believe have chosen B...*”. Subjects entered their guesses by positioning a slider to the desired percentage. Upon doing so, they were informed that they would receive an additional payment of \$0.25, if they provided an accurate estimate within  $\pm 1$  percentage point of the realized value (and would receive nothing otherwise).<sup>13</sup> Below we denote such a belief by  $\mu_i$ .<sup>14</sup>
- (iii) “Part 2 instructions.” Subjects were told that they would go through an unspecified number of additional rounds involving the same decision task; in each round, the scenario would be determined by the new round’s population-level behavior and a new coin flip. Subjects were informed that only at the end of the experiment they would learn about the money earned over the rounds. However, in between rounds, subjects would receive noisy information about the others’ (revealed) preferences: to ensure that this information would be *both* consistent with the theoretical model *and* easily relatable

---

<sup>13</sup> The slider was initially positioned at a value of 50%. Subjects could not leave the slider in the initial position; so, they had to take a stance and express their beliefs about the frequency of the action in relation to which the threshold is defined (i.e., action *B*, in the case of Table 1, p. 7). For discussion on the merits of incentivizing the elicitation of beliefs, see Trautmann and van de Kuilen (2015). In particular, it has been shown that relatively small incentives for beliefs do not typically create a meaningful hedging opportunity; accordingly, here the bracket for an accurate guess is 1 percentage point in either direction of the realized value. (See also Rutström and Wilcox, 2009.)

<sup>14</sup> In terms of the model, the belief  $\mu_i$  equals the expected share of unconditional *B*-choosers (i.e., *all*  $t_b$  players), *plus* the conditional *B*-choosers (*some*  $t_a$  or  $t_c$  players, if any). The expected share of  $t_b$  players is measured by  $\bar{\pi}(t_b | f)$ , henceforth denoted simply  $\bar{\pi}$ ; then, the expected fraction of  $t_a$  or  $t_c$  players who might choose *B* is given by  $(1 - \bar{\pi})g$ , with  $0 \leq g \leq 1$ . In summary,  $\mu_i := \bar{\pi} + (1 - \bar{\pi})g$ . See the proof of Proposition 1 (in Appendix A) for the relationship between  $\mu_i$  and *i*'s behavior, given *i*'s type.

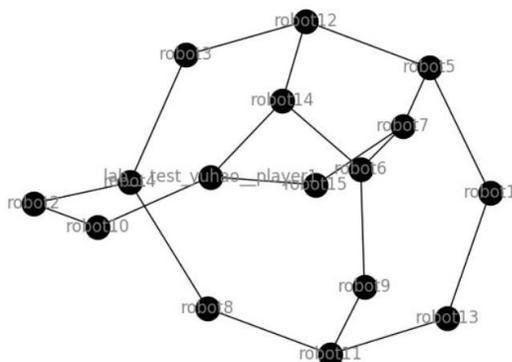
for the subjects, we devised the following mechanism. Subjects were told that they would privately receive information about the frequency of  $B$  choices in their “neighborhood” in the previous round. To be clear, since our model is not concerned with structural properties of networks, such neighborhoods are purely an experimental device for generating feedback that is both noisy and relatable, in the sense that it pertains to some (randomly-selected) fellow participants. Specifically, subjects were told that each participant in the room was connected to some others at random, referred to as neighbors: note that this underlying network determines solely what feedback was passed on to a participant, *not* her incentives (i.e., payoffs depend on the entire room’s behavior, and not just on the neighbors’ behavior). We stress that the network structure as well as the number – and identities – of the neighbors were invisible to the subjects. Such features of the feedback are intended to reflect the everyday experience of receiving noisy information from some people close by.<sup>15</sup>

- (iv) Before carrying out the choice task in round 2, each subject was given feedback about the percentage of her neighbors that chose the socially-undesirable action in round 1; e.g., “0.0% of your neighbors chose  $B$  in the previous round” ...
- (v) Round  $k$  (**choice** task): each subject was asked to choose an option (“A”, “B”, or “C”).
- (vi) Round  $k$  (**belief** elicitation): each subject was prompted to guess the percentage of participants in the entire room that she believed chose the socially-undesirable action in the current round  $k$ .
- (vii) Round  $k + 1$  (**feedback** re. round  $k$ ): each subject was given feedback about the percentage of her *neighbors* that chose the socially-undesirable action.
- (viii) Round  $k + 1$  (**choice** task): each subject was asked to choose an option (“A”, “B”, or “C”).
- (ix) Steps *vi.* to *viii.* were repeated a number of times (subjects played 10 rounds in total).
- (x) Subjects were given a brief demographic questionnaire.

---

<sup>15</sup> Again, the experiment is not designed to analyze learning in relation to structural properties of a network, since subjects did not know the specifics of the network. Unbeknown to subjects, the experiment software was coded to randomly generate a network for each lab session, such that neighborhoods contained 2 or 3 fellow participants (besides the subject herself); this feature of the design guarantees sufficient variability in the feedback passed on to the subjects, while ensuring that subjects are fully comparable across sessions. Note that since neighborhoods are generated at random, the design is consistent with the model’s assumption that the feedback correctly points to the true state with probability  $q > 0.5$ , on average (p. 10). That is, low feedback *more likely than not* implies correctly that the true profile of preference types is such that the share of unconditional  $B$ -choosers (namely,  $t_b$  players) in the room may be a value ( $\pi_L$ ) lower than the 40% threshold. Similarly, high feedback *more likely than not* implies correctly that the share of  $t_b$  players in the room may be a value ( $\pi_H$ ) higher than the 40% threshold.

**Payment Stage.** The payment mechanism consisted of two parts: each subject received a \$10 participation fee, plus any payoffs earned over the ten rounds.<sup>16</sup> Note that subjects did *not* learn about the money earned over the rounds until the end of the experiment.



**Figure 1** - A random network generated by Behavory (<https://behavory.com/>) as a simulation of the lab environment; each node represents a player. Note: the networked structure is a design device for generating feedback that is both noisy and easy to relate to. Subjects did not see the network.

Our experimental design includes a “control treatment” that is the same as the main treatment, except that subjects received *no feedback* about the others’ choices. Note that this is a between-subjects design. (We ran 6 sessions of the main treatment, and 5 sessions of the control treatment; no subject was allowed to participate in more than one session.)

## 2. Hypotheses

The goal of the experiment is to assess whether, in complex problems, behavior can be accurately described by a model of discrete choice with random errors, and how any such errors are impacted by feedback.

We begin by laying out a comparative-statics exercise that is intended to provide an initial check of whether, on average, subjects respond to the feedback consistently with an updated expected utility (UEU) framework. To that end, our first hypothesis verifies whether – in the main treatment – the choice of action *B* varies as a result of subjects updating their beliefs.

---

<sup>16</sup> To comply with the Wharton Behavioral Lab guidelines (which require a minimum compensation for subjects), each subject was informed that if the sum of her payoffs earned across rounds was negative, then she would only receive her \$10 participation fee. Whereas this payment mechanism might be thought to encourage risky choices, note that the goal of the study is *not* to collect data about the extent to which subjects are risk seeking; rather, our focus is on whether subjects commit errors, in a game with and without feedback. In particular, note that the payment mechanism is identical in both the main and control treatments, and therefore it *cannot* explain any treatment effects. (For the record, overall only 7% of all the subjects’ total payoffs turned out to be negative.)

**Hypothesis 1 (H1).** In any given round of the main treatment (regardless of a subject’s prior), receiving *high* feedback decreases the frequency of *B* choices, vis-à-vis *low* feedback.

*Theoretical rationale.* Per the model, high (low) feedback implies an upward (downward) updating of the expected share of  $t_b$  players, denoted  $\bar{\pi}$  (i.e., formally, (12) is greater than (13), for any feedback values  $\pi_L, \pi_H$  with  $0 < \pi_L < 0.4 < \pi_H < 1$ ; p. 10). Thus, high feedback makes it more likely for risk-averse and moderately risk-seeking players (i.e., respectively  $t_c$  and  $t_a$  types) to expect that the share of  $t_b$  players is larger than their respective cut-off points  $\psi_t$ . This implies that on average – regardless of any individual errors – high feedback drives some  $t_a$  and  $t_c$  types toward equilibrium (i) of Proposition 1. Instead, low feedback drives some  $t_a$  and  $t_c$  players toward equilibrium (ii) of Proposition 1.

H1 above allows us to verify if we can rule out collective failures in belief updating as a potential violation of the updated expected utility (UEU) paradigm in our data.<sup>17</sup> Next, we turn to verify if there are any individual-level violations of UEU in the form of suboptimal choices; in other words, we aim to assess whether subjects commit behavioral errors (i.e., expected utility maximization failures) in the sense that they do not best-respond to their stated beliefs, across the main and control treatments.

**Hypothesis 2 (H2).** Subjects do not deterministically best-respond to their beliefs.

*Theoretical rationale.* Per expression (14), each player’s behavior reflects her own preferences, albeit with random errors. (This departs from the standard game-theoretic approach, which deterministically predicts that  $i$  always chooses whatever action yields the highest expected utility, given  $i$ ’s preferences.) Hence, to test H2 we just need to establish if expected utility maximization failures occur with positive probability, regardless of whether beliefs are correct or not. Since a player’s behavior often varies with her type, we will check for behavioral errors by focusing on (a set of) beliefs with respect to which there is a unique, identical best-response for all the preference types. That is, we focus on the case where  $i$  expects that the frequency of *B* choices

---

<sup>17</sup> Note that our design does not aim to disentangle Bayesian and “naïve” belief-updating processes (DeGroot, 1974), since our design was not conceived to investigate learning in relation to structural properties of networks. For recent network-theoretic accounts of learning, we refer the reader to: Gale and Kariv (2003); Golub and Jackson (2010); Acemoglu, Dahleh, Lobel, and Ozdaglar (2011). We further note that models of social learning typically focus on equilibria with risk-neutral players. Thus, our exercise rests on broadly different assumptions than those made by experimental tests of social learning (e.g., Chandrasekhar, Larreguy, and Xandri, 2020; Grimm and Mengel, 2020).

(in the entire population) is low: when that is the case – irrespective of  $i$ 's type –  $i$  commits an error if  $i$  chooses an action other than  $B$ .<sup>18</sup>

Our final hypothesis (H3) is concerned with the observed distribution of errors.

**Hypothesis 3 (H3).** The frequency of expected utility maximization failures differs between the main and control treatments.

*Theoretical rationale.* According to the model, each player's behavior is consistent with her underlying preferences, even though it is prone to occasional errors. Specifically, expression (14) implies that errors affect  $i$ 's best-responses, with random variables  $\varepsilon'_i, \varepsilon''_i$  being drawn from a bivariate distribution, whose parameters vary with environmental stimuli (see footnote 12). Since the main and control treatments feature a different environment (with and without feedback), then violations of expected utility theory should *not* be equally distributed between the two treatments.

In this regard, previous studies have proposed two (opposing) arguments as to how such violations would vary between environments with and without feedback. One argument says that receiving feedback – while attending to a complex choice task – could worsen the agent's cognitive overload, leading to poor choices (e.g., Hall, Ariss, and Todorov, 2007): while several variations of this argument have been put forward, a common prediction is that more information may distract the agent, and will thus *worsen decision quality* (in terms of our model, this means that  $\varepsilon'_i, \varepsilon''_i$  should be larger in the main treatment, relative to the control treatment). By contrast, another argument says that feedback will *improve decision quality*, especially over time (e.g., Fischer and Sliwka, 2018): in this case, feedback is thought to promote learning as well as enhance the agent's motivation and engagement with the task (so,  $\varepsilon'_i, \varepsilon''_i$  should be smaller in the main treatment, relative to the control treatment). It is worth noting that these arguments are usually put forward in specific decision-theoretic contexts, and it is unclear how they would generalize to other contexts. The econometric analysis will test these opposing predictions.

---

<sup>18</sup> The attentive reader will notice that the optimal actions for  $t_a$  and  $t_c$  types formally depend on their cut-off points  $\psi_t$ . As stated in Proposition 1, cut-off points  $\psi_t$  vary with the expected utility values from  $A, B$ , which are respectively denoted by  $\alpha(t), \beta(t)$ . So, one may wonder what values may plausibly be taken on by such variables. In this connection, it is worth noting that, per Table 1 (p. 7), the monetary payoffs from each risky prospect lie in the following range, respectively: for  $A$ ,  $[0.5, 1]$ ; for  $B$ ,  $[-1.5, 3]$ . Consequently, it is reasonable to assume that  $\alpha(t), \beta(t)$  must fall within (or else do not depart substantially from) the range of outcomes above. In any such case, it is easy to see that if  $i$  expects that the overall frequency of  $B$  choices is low (e.g.,  $< 0.4$ ), then  $i$ 's best-response is to choose  $B$ . (Later on, the econometric analysis will also consider an additional class of behavioral errors.)

## IV. Experimental results

### 1. Main treatment: summary statistics and preliminary tests

We begin by addressing our main treatment (comprising the low- and high-feedback samples). A total of 101 subjects from several academic departments took part in our sessions at the Wharton Behavioral Lab; the mean age was 24.7 years. Subjects on average earned a total payoff of \$11.31 (over ten rounds), in addition to the \$10 participation fee.

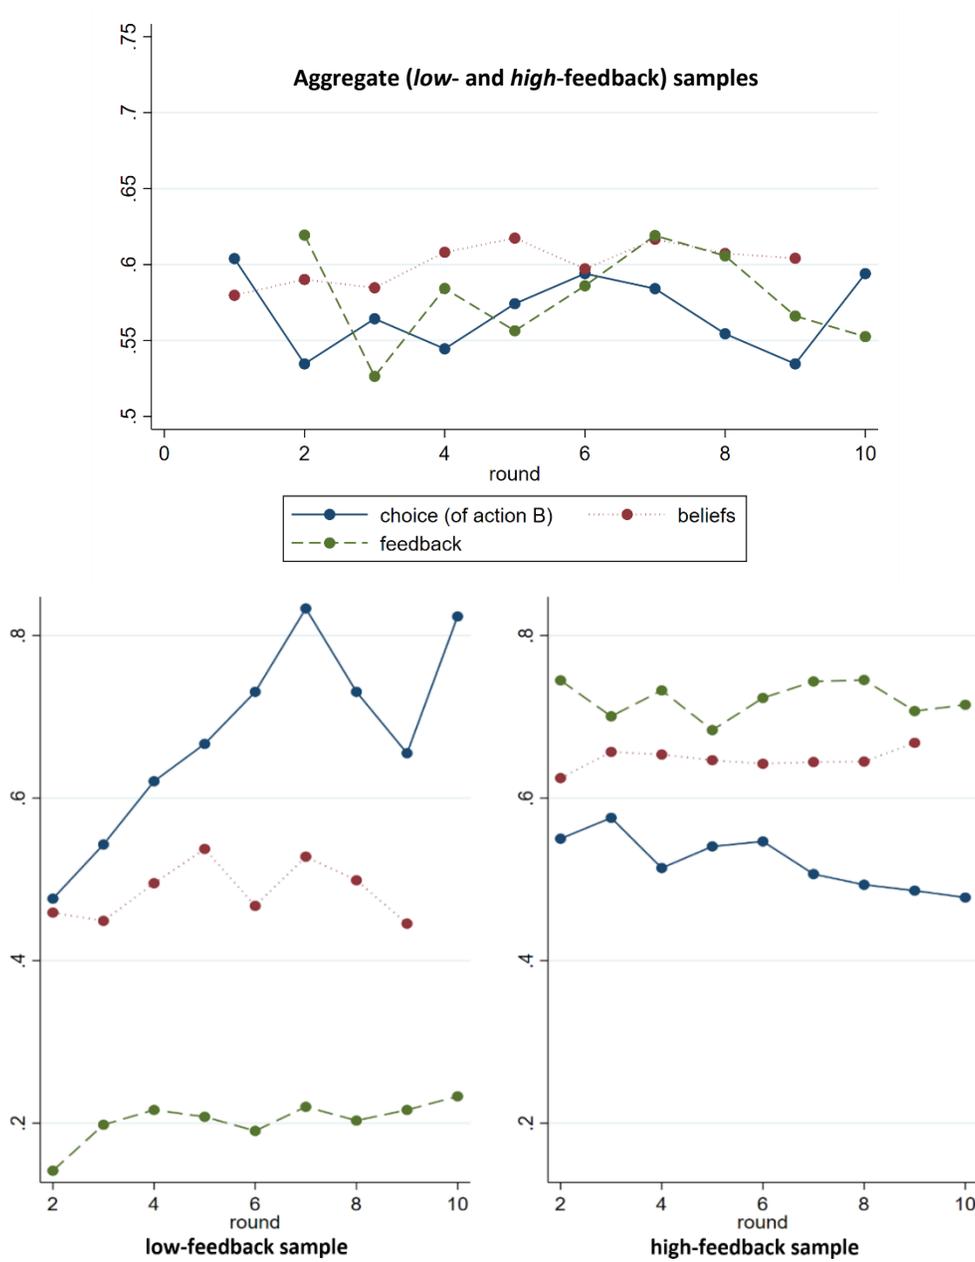
On average subjects chose the option associated with actions  $A$ ,  $B$ , and  $C$  of Table 1 (p. 7) 10.59%, 56.83%, and 32.58% of the time, respectively. A minority of subjects chose the same action across rounds; that is, about 1%, 27%, and 10% of the subjects chose  $A$ ,  $B$ , and  $C$  respectively in each and every round (out of a total of 101 participants in the main treatment). This means that the share of unconditional  $B$ -choosers (i.e.,  $t_b$  players) is 27% of the population; further, since no one other than the moderately risk-seeking  $t_a$  players should ever choose  $A$ , the *overall* distribution of choices implies that at least about one in ten subjects is of type  $t_a$  (i.e., errors aside, this is a lower bound estimate, since the optimal behavior of a  $t_a$  type varies with her beliefs).

	Beliefs about population: <i>Round 1</i>	Beliefs about population: <i>Other rounds</i>	Feedback about neighbors: <i>Other rounds</i>
<b>A</b>	54.36 (27.51)	59.04 (23.26)	60.95 (33.37)
<b>B</b>	59.47 (21.641)	60.62 (23.16)	54.62 (30.66)
<b>C</b>	56.20 (20.40)	60.23 (21.24)	62.70 (28.39)

**Table 2** - Main treatment. Mean beliefs and feedback about  $B$  choices held by the subjects who in a given round chose the option indicated on the left-hand side of the table. In parentheses is the standard deviation. Note: 101 subjects took part in the main treatment. We report the beliefs held in round 1 in a separate column, since subjects stated such beliefs before receiving any feedback; naturally, no feedback about previous play was provided in round 1.

The first two columns of Table 2 above report average beliefs (about the *population-level* frequency of  $B$  choices) held by the subjects who in a given round chose the option indicated on the left-hand side of the table. The last column of Table 2 reports the mean feedback (about the

neighborhood-level frequency of *B* choices) provided to subjects who chose each of the options on the left-hand side. In all cases, beliefs and feedback are above the 40% threshold, on average.



**Figure 2** - Main treatment. The upper panel shows line graphs depicting mean values (by round, averaged across all sessions) for: frequency of *B* choices, beliefs about population-level *B* choices, and feedback about neighborhood-level *B* choices. The lower panel breaks down *low-* vs *high-* feedback samples. Note: no feedback about previous play was provided in round 1; no beliefs were elicited after the last choice task was carried out (in round 10). For the sequence of experimental tasks, see p. 16.

For the purpose of providing a more granular depiction of the data, Figure 2 above breaks down the (mean) values of each of the following variables, by round: (i) frequency of  $B$  choices; (ii) beliefs about population-level  $B$  choices; (iii) feedback about neighborhood-level  $B$  choices. Also, to provide a rough overview of the data, in what follows we report some non-parametric tests and later assess our hypotheses more formally (via a regression analysis).

As a preliminary test of the equality of choice distributions across low- versus high-feedback samples, here we propose some simple pairwise comparisons. In short, a test of proportions (adjusted for clustering on 101 subjects, using data from all the rounds in which feedback was provided; i.e., rounds 2-10) indicates that action  $B$  was chosen more often in the *low*- than in the *high*-feedback sample; that is, 67.73% and 52.13% of the time, respectively ( $z = 2.25$ ,  $p = 0.024$ , two-tailed). In other words, subjects were less likely to choose  $B$  after learning that fellow participants chose  $B$  in proportions larger than the threshold.

Furthermore, the same test shows no meaningful differences in the proportions of choice of action  $A$  across samples, which were respectively 10.36% and 10.64%. By contrast, performing the same test with respect to action  $C$  reveals that the riskless option was chosen less often in the *low*- than in the *high*-feedback sample; that is, 21.91% and 37.23% of the time, respectively ( $z = -2.25$ ,  $p = 0.024$ , two-tailed). Taken together, these patterns provide some preliminary evidence that exposing subjects to high feedback might cause them to shift from a very risky action ( $B$ ) to a riskless action ( $C$ ).

## **2. More on the main treatment: tests of $H1$**

The tests above capture across-rounds average trends, but do not at all account for individual-specific differences across rounds. For this reason, we formally test our first hypothesis via a regression analysis of a subject's choice of action  $B$  (against the other two actions jointly, i.e., "not  $B$ ").

As a benchmark, we first propose a simple logit model consisting of the low/high feedback indicator as the sole predictor (see model [1] in Table 3 below); as usual, robust standard errors are clustered on 101 subjects. Unsurprisingly, model [1] corroborate the insights of the test of proportions above, showing a significant negative effect of the high feedback on one's choice of the risky action  $B$ .

Next, to control for any individual-specific differences across rounds, model [2] in Table 3 includes the following predictors: a dummy variable indicating if one’s stated belief in round 1 (i.e., prior to receiving any feedback at all) is below/above the threshold; a dummy variable indicating if one received the same feedback (either *low* or *high*) across rounds; a time (i.e., round  $k$ ) variable. Again, robust standard errors are adjusted for clustering on the subjects. Given this, model [2] confirms the significant impact of the low/high feedback indicator; moreover, model 2 shows no significant impact of the other predictors.

choice of action $B$ at round $k$	[1]	[2]
low/high ( $k - 1$ ) feedback indicator	-.656*** (.237)	-.566** (.227)
belief round 1		.079 (.362)
same-feedback indicator		.243 (.329)
round		.010 (.023)
<i>constant</i>	.741*** (.225)	.408 (.462)
Pseudo R2	0.014	0.017
AIC	1230.695	1233.763
Obs.	909	909

**Table 3** - Logit coefficients of two models estimating a participant’s choice of  $B$  at round  $k$  of the main treatment. In parentheses are robust standard errors clustered on 101 subjects (\*, \*\*, and \*\*\* respectively indicate  $p < 0.10$ ,  $p < 0.05$  and  $p < 0.01$ , for the relevant Z-statistic, two-tailed tests). The models use all the choice tasks except for round  $k = 1$ , for which there was no  $k - 1$  feedback. Note: the *low/high feedback* indicator takes on value 1 if the frequency of previous  $B$  choices (at  $k - 1$ ) in the subject’s neighborhood is greater than or equal to 40%, it takes on value 0 otherwise. For robustness purposes, model [2] includes a dummy variable indicating if the subject’s belief at  $k = 1$  (i.e., prior to receiving any feedback at all) is below/above the threshold. The *same-feedback* indicator takes on value 1 if the subject received consistent feedback across all rounds, 0 otherwise.

In a nutshell, Table 3 confirms that the feedback significantly affects one’s choice of the risky action  $B$ . This provides evidence in support of H1, which states that “in any given round (regardless of a subject’s prior), transmitting *high* feedback decreases the frequency of  $B$  choices, vis-à-vis *low* feedback”. Generalizing, whenever one learns that some other participants chose

$B$  in proportions larger (smaller) than the threshold, one may infer that the population as a whole may include a large (small) number of risk-seeking individuals; therefore, one is less (more) likely to choose  $B$  oneself.<sup>19</sup>

### 3. *Control treatment: summary statistics*

We turn to our control treatment. This was the same as the main treatment, except that participants received *no feedback* about the others' choices. A total of 84 subjects took part in the control treatment; the mean age was 23.7 years, and other demographics were similar across the two treatments. Participants in the control treatment on average earned a total of \$6.61 (over ten rounds), in addition to the \$10 participation fee. The reader might recall that participants in the main treatment made about twice the money (i.e., \$11.31, plus the participation fee): a two-tailed Wilcoxon-Mann-Whitney test conducted on the entire sample confirms that payoffs were significantly lower in the control than in the main treatment ( $N = 185$  subjects,  $Z = -5.275$ ,  $p = 0.000$ ). In what follows we investigate what may have contributed to such a difference.

We begin by reporting summary statistics relating to the choice data. On average, participants in the control chose the option associated with actions  $A$ ,  $B$ , and  $C$  of Table 1 (p. 7) 9.64%, 64.29%, and 26.07% of the time, respectively. (In the main treatment, the same actions were chosen respectively 10.59%, 56.83%, and 32.58% of the time.) Also, a non-trivial fraction of subjects chose the same action in each and every round: that is, out of a total of 84 participants in the control treatment, about 0%, 36%, and 7% of the subjects respectively chose  $A$ ,  $B$ , and  $C$  in every round. Later we will formally contrast the data across treatments, but for now we note that these figures appear somewhat different than the corresponding figures from the main treatment (1%, 27%, and 10%, respectively).

We further note that, on average, beliefs held by participants in the control treatment are above the 40% threshold, as was the case in the main treatment: see Table 4 and Figure 3 below.

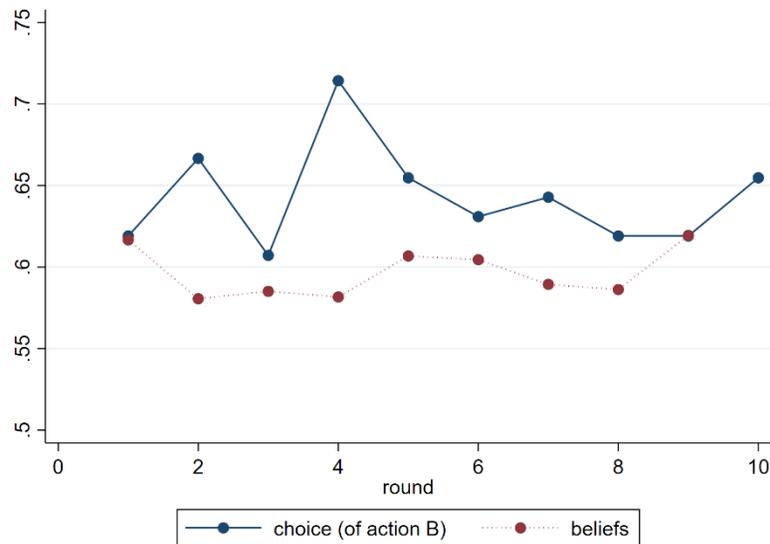
---

<sup>19</sup> This result rejects the *risk contagion* hypothesis in our environment. In this regard, Reiter et al. (2019) reported evidence of risk contagion in non-strategic problems, such that subjects were more likely to choose a risky option – in the form of a lottery – over a sure option, if they observed others' risky choices. For similar results, see also Gioia (2017) and Fatas, Hargreaves Heap, and Rojo Arjona (2018), among others. (A possible explanation for the contagion effect might be conformist preferences, whereby one *likes* to do what one believes others will do; Charness, Naef, and Sontuoso, 2019.)

	Beliefs about population: <i>Round 1</i>	Beliefs about population: <i>Other rounds</i>
<b>A</b>	56.33 (18.40)	52.03 (21.50)
<b>B</b>	65.51 (20.49)	64.66 (22.32)
<b>C</b>	55.04 (20.83)	49.12 (18.15)

**Table 4** - Control treatment. Mean beliefs about *B* choices held by the subjects who in a given round chose the option indicated on the left-hand side of the table. In parentheses is the standard deviation. Note: 84 subjects took part in the control treatment. To be consistent with the style of Table 2 above, we report the beliefs for round 1 in a separate column.

Comparing Figure 3 with the upper panel of Figure 2 (p. 21) reveals that even though beliefs are relatively similar between treatments on average, participants in the control treatment chose *B* in proportions greater than in the main treatment. Put differently, the upper panel of Figure 2 (p. 21) suggests that exposing subjects to feedback – in aggregate – might cause a decrease in the choice of the riskiest action (*B*). The analysis below will illuminate these patterns by formally contrasting behavior across the main and control treatments.



**Figure 3** - Control treatment. Line graphs depicting mean values (by round, averaged across all sessions) for: frequency of *B* choices and beliefs about population-level *B* choices.

#### 4. *Analyzing behavioral errors across treatments*

The data from the main treatment (and the tests of H1, pp. 22-23) allow us to rule out collective failures in belief updating as an explanation for potential violations of the updated expected utility (UEU) paradigm. Moving on, we shall check if individual-level violations of UEU may arise from the fact that some subjects actually fail to best-respond to their stated beliefs. Below we refer to the choice of any suboptimal actions (i.e., expected utility maximization failures) as “behavioral errors” or simply “errors”.

***Tests of H2.*** In order to check for any such errors, for the time being we focus on observations where  $i$  states that the expected population-level frequency of  $B$  choices is low, in a given round (by “low” we consider values below the 40% threshold). When that is the case – regardless of  $i$ ’s preference type –  $i$  commits an error if  $i$  chooses an action other than  $B$ , in the same round (see footnote 18 for theoretical details). Given this, the data reveal that on average 8.82% of each participant’s choices constitute errors, and hence violate the predictions of the UEU framework (this result accounts for observations from both the main and control treatments). Now, to reject a fully deterministic model we just need to show that errors occur with strictly positive probability: with nearly 9 percent of the participants’ choices constituting errors, we can readily affirm that the standard framework does not fit the data. Indeed, even a conservative Wilcoxon signed-rank test for whether the median frequency of errors differs from the value of 5 percent (i.e., instead of 0, thus allowing for “almost no errors”) is strongly significant:  $N = 185$ ,  $z = 2.409$ ,  $p = 0.016$ , two-tailed; note that the test is conducted on the sample of per-subject mean choices, to satisfy the assumption of independence of observations (i.e., the test uses one observation for each participant in the main and control treatments, where an observation represents the fraction of a participant’s choices constituting errors). We conclude that the data support H2, which states that “subjects do not deterministically best-respond to their beliefs”.

***Tests of H3.*** We turn to check if there are any between-treatment differences in the distribution of the above-discussed errors (breaking down the data by treatment). In this respect, the data from the main treatment ( $N = 101$ ) show that 6.49% of participants’ choices consist of errors; by contrast, in the control treatment ( $N = 84$ ) 11.64% of participants’ choices constitute errors. This appears

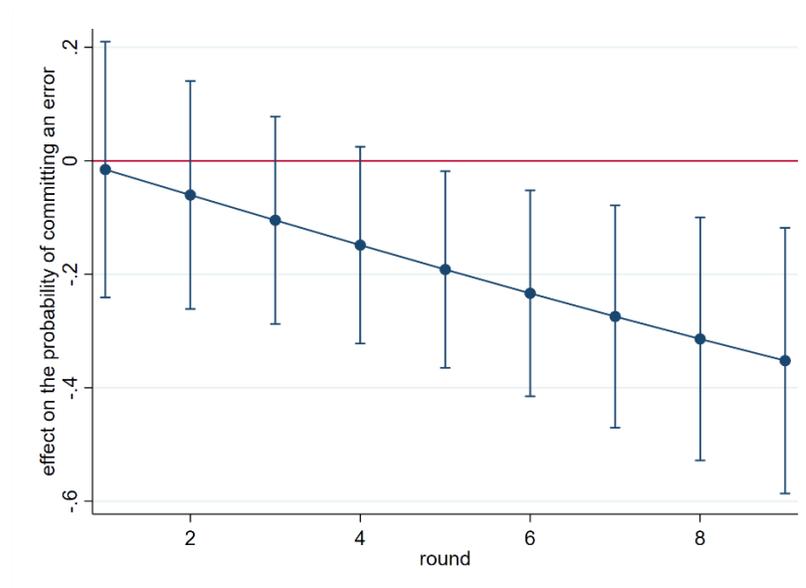
to support our prediction that the frequency of errors should differ between the main and control treatments: in fact, a t-test confirms our intuition ( $N = 185$ ,  $t = 2.149$ ,  $p = 0.033$ , two-tailed). Yet, the mean data above do not account for individual-specific differences across rounds; for this reason, we shall test our third hypothesis by conducting a regression analysis on the entire sample of observations (while adjusting standard errors for clustering).

To that end, we propose a pair of simple logit models: in both cases, the dependent variable is whether an observation at round  $k$  does or does not constitute an error, as previously described. The list of predictors is as follows: model [1] in Table 5 uses a control/main treatment indicator, plus a time (i.e., round  $k$ ) variable; in addition to those two predictors, model [2] includes an interaction variable. (In both models, robust standard errors are clustered on 185 subjects; i.e., 101 and 84 participants in the main and control treatment, respectively.) That said, model [1] shows a significant negative effect of the main treatment on the probability of committing an error. In accounting for the (*treatment indicator \* round*) interaction, model [2] reveals that the beneficial effect of the main treatment increases with time.

choice of a suboptimal action at $k$	[1]	[2]
treatment indicator	-.640** (.291)	.181 (.406)
round	-.005 (.034)	.063* (.037)
treatment indicator * round		-.168** (.068)
<i>constant</i>	-1.998*** (.222)	-2.357*** (.257)
Pseudo R2	0.013	0.019
AIC	986.622	982.591
Obs.	1,665	1,665

**Table 5** - Logit coefficients of two models estimating a participant's choice of a suboptimal action (i.e.,  $A$  or  $C$ , versus  $B$ , given low beliefs) at round  $k$ . In parentheses are robust standard errors clustered on 185 subjects from two treatments (\*, \*\*, and \*\*\* respectively indicate  $p < 0.10$ ,  $p < 0.05$  and  $p < 0.01$ , for the relevant Z-statistic, two-tailed tests). The model uses all the observations from the main and control treatments, except for round  $k = 10$  (no beliefs were elicited after the last choice task). Note: the treatment indicator takes on value 1 if a subject is in the main treatment, value 0 if a subject is in the control treatment.

To better illuminate the results from the latter logit model, Figure 4 graphs the change in the predicted probability of committing an error in each round of either treatment.



**Figure 4** - The effect of the feedback (i.e., main treatment) on the probability of committing an error, relative to the control treatment, with 95% confidence intervals. Specifically, the y-axis measures the discrete change in the predicted probability of choosing a suboptimal action (i.e.,  $A$  or  $C$ , versus  $B$ , given low beliefs) in a certain round. Note: the horizontal line at the 0 value of the y-axis represents no change between the two treatments; all the predicted probabilities are computed from model [2] in Table 5.

Taken together, Table 5 and Figure 4 provide strong evidence in support of H3, which states that “the frequency of errors differs between the main and control treatments”. In fact, the data conclusively imply that random variables  $\varepsilon'_i, \varepsilon''_i$  in expression (14) are smaller in the main treatment, on average. This means that the main treatment does *improve decision quality*, especially over time. Put differently, the data do not back the argument that decision-makers’ cognitive load is exacerbated (and they are distracted) if presented with additional – noisy – information during a complex task. On the contrary, the results suggest that (besides promoting learning) noisy feedback enhances decision-makers’ motivation and engagement with the task. As we can see from Figure 4, this effect increases with the number of iterations.

We stress that the effect above occurs despite the fact that our feedback mechanism provides relatively little informational content (after each play, our subjects do not receive any information about payoffs, nor do they receive any third-party advice, unlike previous studies: e.g., Esponda and Vespa, 2014; Martínez-Marquina, Niederle, and Vespa, 2018). In fact, our

feedback mechanism does not in itself reduce the task’s computational difficulty, and hence the feedback’s effect on decision quality cannot be fully explained by common characterizations of learning. For these reasons, we ascribe the improvement in decision quality to increased motivation and engagement.

***Robustness checks (I).*** The data patterns above show that, even if noisy, feedback does have a beneficial effect in complex interactive environments. In this regard, we note that previous research has shown that noisy information can improve performance in simple decision-theoretic tasks, thanks to some form of confidence boost (Chen and Schildberg-Hörisch, 2019; Fischer and Sliwka, 2018). Indeed – also in our case – it is plausible that some participants in the main treatment may have felt more engaged with the task due to an uplift in confidence: more specifically, some subjects may have experienced an increased trust in their abilities to predict others’ behavior on the basis of the feedback received; in turn, this confidence boost may have strengthened one’s inclination to act on one’s stated beliefs.

In the following, we delve into the claim that a higher confidence may be a reason why some subjects are more motivated to respond optimally to their beliefs. To that end, we shall use the dispersion of a subject’s stated beliefs across rounds as a proxy measure for confidence: under the assumption that some subjects may feel *more or less confident* (about their beliefs) *based on the feedback received*, then it follows that – in the main treatment – confidence will correlate negatively with belief dispersion; therefore, we should find that belief dispersion is a predictor of behavioral errors in the main treatment.

Given the above, we denote by  $sd_{\mu_i}$  the standard deviation of subject  $i$ ’s stated beliefs across rounds, with  $sd_{\mu_i}$  being normalized to take values between 0 and 1 (the non-normalized standard deviation ranges from 0 to 35.65, with an average value of 14.17). Next, we define  $be_i$  to be the frequency with which – across rounds – subject  $i$  commits behavioral errors of the following form: a choice of  $A$  or  $C$  at round  $k$ , having stated a low belief for round  $k$ ; that is,  $be_i$  is defined as a continuous variable ranging from 0 to 1, so that a value of 0 indicates that subject  $i$  has not committed a behavioral error in any round  $k$ , whereas a value of 1 indicates that subject  $i$  has committed a behavioral error in each and every round  $k$ . (Our data show that 36.76 percent of the subjects across treatments have at least once committed an error, as described above.)

We now proceed to set up a linear regression model, with  $be_i$  as the dependent variable and  $sd_{\mu_i}$  as the sole predictor, using one observation for each participant in the main treatment. Such a model reveals that  $sd_{\mu_i}$  has a significant positive effect on  $be_i$  ( $N = 101$ , coef. = .122,  $t = 2.44$ ,  $p = 0.016$ , two-tailed OLS with robust standard errors). This means that a participant in the main treatment is likely to commit more errors the larger is the variability in her stated beliefs.

As an additional test, it is worth considering the same model above, this time using participants from both the main and control treatments (with one observation per subject; i.e.,  $N = 185$ ), while adding a control/main treatment indicator. The results confirm our previous intuitions, in the sense that the main treatment significantly decreases the occurrence of behavioral errors (coef. = -.056,  $t = -2.33$ ,  $p = 0.021$ , two-tailed), whereas belief dispersion increases it (coef. = .090,  $t = 1.74$ ,  $p = 0.084$ , two-tailed). In assessing these results, we note that here the effect of belief dispersion on behavioral errors is still positive and fairly large (though mildly significant): this intuitively suggests that confidence plays a role also independently of the feedback.

In summary, this analysis provides further context to the hypothesis that noisy feedback can enhance one's performance in complex tasks; in particular, the analysis suggests that the motivational value of feedback – in our environment – may be linked to an increased trust in one's ability to predict others' behavior.

**Robustness checks (II).** As we noted before, given that a player's behavior often varies with her type, we checked for errors by focusing on a specific set of (low) beliefs with respect to which there is a unique, identical best-response for all the preference types. In other words, the analysis so far has revolved around a class of suboptimal actions that does not depend on the preference type (but it is conditional on holding a low belief about population-level  $B$  choices). In what follows, instead, we shall investigate the relationship between behavior and beliefs in more general terms. That is, under the common assumption that the subjects' underlying preferences must be similar – on average – across treatment samples, our goal is to check if subjects who state the same beliefs (including low *and* high beliefs) behave similarly across treatments *with* and *without* feedback, as is predicted by a standard model without errors.

To that end, Table 6 presents the results of a logit model consisting of one's choice of action  $B$  as the binary dependent variable, and of the following predictors: a control/main treatment indicator; a time (i.e., round  $k$ ) variable; one's stated belief at round  $k$ ; an interaction

variable for each pair of the above predictors. (Robust standard errors are clustered on 185 subjects.) Given this, the results in Table 6 reveal a significant impact of each of the (main-effects) variables. Moreover, Table 6 shows a significant negative effect of both the (*treatment \* belief*) and (*round \* belief*) interactions: this means that participants who state higher beliefs are less likely to choose *B* in the main treatment *and* in later rounds. For example, simple calculations show that the predicted probability of choosing *B* goes down, on average, by 2.39 percentage points in the main treatment (relative to the control) when beliefs are at 50%; instead, the predicted probability of choosing *B* on average drops by 25.10 percentage points in the main treatment (relative to the control) when beliefs are at 90%.

choice of action <i>B</i> at <i>k</i>	
treatment indicator	1.336** (.519)
round	.113** (.056)
stated belief (at <i>k</i> )	.043*** (.007)
treatment indicator * round	.011 (.029)
treatment indicator * stated belief (at <i>k</i> )	-.029*** (.008)
round * stated belief (at <i>k</i> )	-.002** (.000)
<i>constant</i>	-1.827*** (.462)
Pseudo R2	0.040
Obs.	1,665

**Table 6** - Logit coefficients estimating a participant's choice of action *B* at round *k*. In parentheses are robust standard errors clustered on 185 subjects from the two treatments (\*, \*\*, and \*\*\* respectively indicate  $p < 0.10$ ,  $p < 0.05$  and  $p < 0.01$ , for the relevant Z-statistic, two-tailed tests). The model uses all observations from the main and control treatments except for round  $k = 10$  (since no beliefs were elicited after the last choice task). Note: the treatment indicator takes on value 1 if a subject is in the main treatment, value 0 if a subject is in the control treatment.

In interpreting these results, it is worth noting that our theoretical model makes the following prediction: risk-averse *and* moderately risk-seeking players (i.e.,  $t_c$  and  $t_a$  types,

respectively) should be less likely to choose  $B$ , when beliefs about (population-level)  $B$  choices get higher (for details, see the proof of Proposition 1). Now, the fact that there is a difference in  $B$  choices across treatments (when controlling for any stated beliefs) provides indirect evidence of the following class of suboptimal choices: that is, the case where *some*  $t_c$  or  $t_a$  types choose  $B$  *mistakenly, having stated high beliefs*. More explicitly, the (*treatment \* belief*) interaction in Table 6 indicates that, as beliefs get higher, participants in the main treatment are collectively less likely to choose  $B$ , relative to the control. So, these results suggest that the feedback causes  $t_c$  or  $t_a$  types to make fewer suboptimal choices, given the same stated beliefs.<sup>20</sup>

## V. Concluding remarks

This paper has investigated complex problems where contingencies depend on multiple unknowns, and thus hypothetical reasoning is difficult. In particular, we have formalized a novel multiplayer game that captures some interesting features of real-world risky interactions, and is therefore amenable to business and policy applications. After modeling a Bayesian environment, we have proposed a tractable way to formally incorporate random errors into a model of choice; we have then designed an experiment with the goal to study if noisy feedback has an impact on the likelihood of suboptimal choices (that is, failures to respond optimally to information).

To that end, we have considered two common arguments as to why the occurrence of errors may vary with and without feedback. One argument implies that receiving feedback – while attending to a complex choice task – may exacerbate the agent’s information overload; so, more information would weakly worsen decisions (e.g., Wilson, 2014; Hall, Ariss, and Todorov, 2007). Instead, another argument implies that receiving feedback does not only promote learning, but also enhances motivation and engagement, which in turn improves the agent’s outcomes (e.g., Compte and Postlewaite, 2004; Fischer and Sliwka, 2018).

Our data indicate a beneficial effect of the feedback: besides inducing a belief revision, feedback contributes to reducing the occurrence of two classes of behavioral errors. The first class (i.e., choice of  $A$  or  $C$ , given *low* beliefs) applies to participants regardless of their risk profile:

---

<sup>20</sup> Recall that, by definition, risk-averse and moderately risk-seeking players (i.e.,  $t_c$  and  $t_a$  types, respectively) are the only types whose behavior depends on their beliefs (by contrast, the behavior of  $t_b$  players is not a function of their beliefs, since  $B$  is a dominant action for  $t_b$  players). That said, note that the randomized assignment of subjects to either treatment ensures *a-priori similar* risk profiles across treatments. So, if the underlying preferences are similarly distributed across treatments, then subjects who state the same beliefs should behave similarly, on average, *unless* the tendency to commit errors varies across treatments. (Indeed, our previous analysis – including Table 5 and Figure 4 – confirmed that such a tendency clearly varies across treatments.)

in this respect, we found that nearly 9 percent of the choices across treatments are suboptimal (specifically, 6.49% and 11.64% of the choices in the main and control treatment, respectively). By comparing choices across treatment samples – with and without feedback – we have also isolated a second class of errors (i.e., the case in which risk-averse and moderately risk-seeking types choose *B* mistakenly, having stated *high* beliefs): in this regard, our analysis reveals that participants collectively take fewer risks in the main treatment (relative to the control), given the same stated beliefs. Such a trend increases in later rounds.

It is important to stress that since the feedback's beneficial effect occurs regardless of the beliefs' correctness, then it cannot be fully explained by a belief revision. In fact, our analysis corroborates the argument that noisy feedback not only facilitates learning but also holds motivational value (in that it appears to increase engagement with the task). In this respect, the data suggest that such a motivational value is linked to an increased confidence in one's ability to predict others' behavior on the basis of the information received.

Finally, our findings suggest practical strategies for curbing improper risk-taking via targeted information campaigns. In this connection, we note that numerous everyday life interactions present strategic features that, to some extent, are consistent with our game. Indeed, in addition to vaccination decisions, other applications include a broad range of interactions wherein risky prospects may depend on some threshold. For example, think of the case of private companies deciding whether or not to fund research on new materials in anticipation of a possible supply shock. Also of interest is the case of regional governments deciding whether or not to invest in new sources of energy in readiness for a possible disruption in the energy market. In short, we believe that the present paper provides intriguing insights in relation to a class of risky interactions that may be worth investigating further.

## APPENDIX A

### Proposition 1 (Bayesian Nash equilibria in pure actions)

**Proof.** In every equilibrium, all  $t_b$  players choose  $B$  as it is a dominant action (only) for  $t_b$  players; by contrast, the other players will have to consider the probability of being in a risk-free ( $x$ ) versus risky ( $y$  or  $z$ ) scenario. This depends on the expected frequency of  $B$  choices,  $\mu_i$ ; in turn,  $\mu_i$  equals the expected share of  $t_b$  players (i.e., which is denoted by  $\bar{\pi}$ ), plus the expected fraction of  $t_a$  or  $t_c$  players who might choose  $B$ , which we denote by  $(1 - \bar{\pi})g$ , with  $0 \leq g \leq 1$ . Compactly, the expected frequency of  $B$  choices is given by  $\mu_i := \bar{\pi} + (1 - \bar{\pi})g$ .

It follows that player  $i$ 's expected utility from  $A$ ,  $B$ , and  $C$  is respectively given by

$$U_i(A) = 0.5 \cdot (1 - \mu_i) + \alpha(t) \cdot \mu_i,$$

$$U_i(B) = 3 \cdot (1 - \mu_i) + \beta(t) \cdot \mu_i,$$

$$U_i(C) = 0.75 \cdot (1 - \mu_i) + \gamma(t) \cdot \mu_i,$$

where for each type  $t \in T$ ,  $\alpha(t), \beta(t), \gamma(t)$  are consistent with (8)-(10) in the main text.

$$(i) \text{ For } t = t_a: \text{ if } \bar{\pi} > \psi_{t_a} \equiv \frac{5}{2\alpha(t) - 2\beta(t) + 5}, \quad (1A)$$

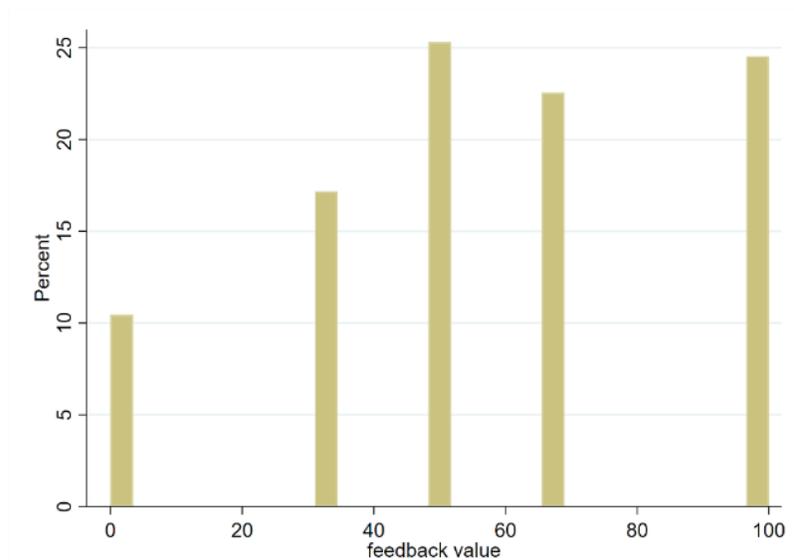
then  $U_i(A) > U_i(B)$  and  $U_i(A) > U_i(C)$ ; hence, each  $t_a$  player prefers to choose  $A$ ;

$$\text{For } t = t_c: \text{ if } \bar{\pi} > \psi_{t_c} \equiv \frac{9}{12 - 4\beta(t)}, \quad (2A)$$

then  $U_i(C) > U_i(B)$  and  $U_i(C) > U_i(A)$ ; hence, each  $t_c$  player prefers to choose  $C$ .

- (ii) When 1A (2A) [both 1A and 2A] is/are not satisfied, then a fraction  $g$  of  $t_a$  players (of  $t_c$  players) [of  $t_a$  or  $t_c$  players] will prefer to choose  $B$ , by the same reasoning as in equilibrium (i), iff  $\bar{\pi} + (1 - \bar{\pi})g$  is less than  $\psi_{t_a}$  (less than  $\psi_{t_c}$ ) [less than  $\psi_{t_a}$  or  $\psi_{t_c}$ ], with  $0 \leq g \leq 1$ . Instead, a fraction  $1 - g$  of such players will prefer to choose  $C$  to avoid passing their respective cut-off points.

## Additional data



**Figure A1** - Histograms for the feedback that was passed on to subjects. As an example, note that the first bar indicates that about 10 percent of the time subjects were informed that none (0%) of their neighbors chose action *B* in the preceding round. The relative frequency of each feedback value is calculated with respect to all the rounds in which feedback was provided, namely, rounds 2-10. Note: the experimental feedback consisted of five possible values because, unbeknown to subjects, the experiment software was coded to randomly generate a network per session such that each node had 2 or 3 neighbors (for details, see section III of the main text).

## APPENDIX B

### *Experimental instructions and screen shots*

**NOTE:** as we discuss in section III of the main text (p. 14), *letter-outcome pairs* (e.g., whether *B* is associated with the socially-undesirable option rather than, say, the exit option) *were randomized across participants*. This was done in order to control for the fact that letters that come first in the alphabet may be perceived as more prominent. Below is an instance of the experimental instructions (for the main treatment) where the socially-undesirable option is associated with action A: accordingly, in the below screen shots, the threshold is defined in relation to action A; hence, the belief elicitation task and the feedback refer to action A. Finally, note that instructions for the control treatment are the same as the main treatment, except that there is no feedback.

#### [Welcome screen]

At the beginning of this study, you will receive instructions on what to do and how your decisions can affect your earnings. Your participation in the study is voluntary. You may end your participation at any point, without loss of any benefits to which you are entitled.

The main purpose of the study is to explore people's decision making in different contexts. The study involves monetary decisions, that can only add to the \$10 (show up fee) you receive for your participation. The duration of the study will be about 50 minutes.

Your final earnings depend on the decisions you and other participants make.

Please click the box if you agree to participate in the study.

## Instructions (1/3)

You will receive a show up fee, and can earn additional money. The additional payment will be determined by your own choices and those made by the other participants, according to rules described below. Your final earnings will be added to your show up fee if positive.

In each round, each participant will be asked to choose one of the actions represented by options on the screen, namely "A", "B", and "C". Please note that the information about the amount of money earned over each round will be provided only at the end of the experiment.

Next

## Instructions (2/3)

The money you will earn in each round depends on your choice, as well as on the choices made by all other participants, and on the outcome of a coin tossed by the computer in each round. The coin may result in either of two outcomes, HEADS or TAILS, each with a 50% chance. Depending on the conditions described above, you will end up in ONE of three alternative *scenarios*:

**If less than 40% of all participants chose A, then *regardless of the coin outcome*:**

- Your earnings for the round will be \$3.0 if you chose A, \$0.5 if you chose B, and \$0.75 if you chose C.

A	B	C
\$3.0	\$0.5	\$0.75

**If 40% or more of all participants chose A, then:**

- If the coin outcome is HEADS  
Your earnings for the round will be \$-1.5 if you chose A, \$1.0 if you chose B, and \$0.75 if you chose C.

A	B	C
\$-1.5	\$1.0	\$0.75

- If the coin outcome is TAILS  
Your earnings for the round will be \$3.0 if you chose A, \$0.5 if you chose B, and \$0.75 if you chose C.

A	B	C
\$3.0	\$0.5	\$0.75

Next

## Instructions (3/3)

After all participants have made their choice, the coin is tossed by the computer, and the scenario for the round is determined.

(Participants will *not* be informed of the scenario they are in before making decisions.)

At any point during the experiment, if you have any questions please raise your hand and an experimenter will approach you.

Next

## Control Questions

If more than 40% of all participants chose A, the coin outcome is HEADS, and you chose C, how much will you earn?

1.0 ▾

If less than 40% of all participants chose A, the coin outcome is TAILS, and you chose C, how much will you earn?

-1.5 ▾

If less than 40% of all participants chose A, the coin outcome is TAILS, and you chose B, how much will you earn?

-1.5 ▾

If more than 40% of all participants chose A, the coin outcome is HEADS, and you chose A, how much will you earn?

-1.5 ▾

If less than 40% of all participants chose A, the coin outcome is TAILS, and you chose A, how much will you earn?

-1.5 ▾

**Hover (using mouse) and Scroll (using arrow keys) to review previous instructions**

Next

## You are currently in round 1 .

Choose an action from below

C  A  B

Hover (using mouse) and Scroll (using arrow keys) to review previous instructions

Next

Move the slider below to indicate the percentage of the participants in **the entire room** that you believe have chosen A in this round.

You will earn \$0.25 if you guess within 2 percentage points (1 point in either direction) of the actual percentage.



## Ending Round 1

Waiting for other participants...



## Instructions part 2 (1/2)

**In the following rounds** you will face the same decision task as before.

Each participant in the room is connected to some others at random, such that everyone is either directly or indirectly connected to everyone else.

Participants who are directly connected to one another are “neighbors” (your neighbors are most likely not the participants sitting next to you).

Those who are indirectly connected to you are your neighbors' neighbors, the neighbors of your neighbors' neighbors, and so on.

Next

## Instructions part 2 (2/2)

**All connections (direct and indirect) remain constant across rounds.** That is, if you are connected to specific participants in round 1, they will be your neighbors in all rounds.

Your neighbors may or may not have the same number of neighbors as you do. That is, each participant may have a different number of connections.

If you have any questions, please raise your hand and an experimenter will approach you.

Next

50.0% of your neighbors chose A in the previous round.

Press next to continue.

Next

## You are currently in round 2 .

Choose an action from below

C  A  B

Hover (using mouse) and Scroll (using arrow keys) to review previous instructions

Next

Move the slider below to indicate the percentage of the participants in **the entire room** that you believe have chosen A in this round.

You will earn \$0.25 if you guess within 2 percentage points (1 point in either direction) of the actual percentage.



## Ending Round 2

Waiting for other participants...



0.0% of your neighbors chose A in the previous round.

Press next to continue.

Next

[...]

## Demographic Survey

Please enter your age in the box below

Please select your gender from below

Continue without responding ▾

Please select your race/ethnicity from below

Continue without responding ▾

Please select your education level from below

Continue without responding ▾

Next

## **Thank You for Participating**

You have successfully completed the experiment.

You began the experiment with a show-up pay of \$10.0. Your earnings at the end of the experiment were \$4.5. Your final pay amounts to \$14.5.

Please wait for your number to be called by the experimenter.

## References

- Acemoglu, Daron, Munther A. Dahleh, Ilan Lobel and Asuman Ozdaglar.** 2011. "Bayesian Learning in Social Networks" *Review of Economic Studies*, 78(4): 1201-1236.
- Ballinger, T. Parker and Nathaniel T. Wilcox.** 1997. "Decisions, Error and Heterogeneity" *Economic Journal*, 107(443): 1090-1105.
- Becker, Gordon M., Morris H. DeGroot and Jacob Marschak.** 1963. "Stochastic Models of Choice Behavior" *Behavioral Science*, 8(1): 41-55.
- Benjamin, Daniel J., Matthew Rabin and Collin Raymond.** 2016. "A Model of Nonbelief in the Law of Large Numbers" *Journal of the European Economic Association*, 14(2): 515-544.
- Blavatsky, Pavlo.** 2018. "Fechner's Strong Utility Model for Choice Among  $n > 2$  Alternatives: Risky Lotteries, Savage Acts, and Intertemporal Payoffs" *Journal of Mathematical Economics*, 79: 75-82.
- Brandenburger, Adam.** 1996. "Strategic and Structural Uncertainty in Games" in *Wise Choices: Games, Decisions, and Negotiations*, ed. R. Zeckhauser, R. Keeney and J. Sibenius. Boston, MA: Harvard Business School Press.
- Chandrasekhar, Arun G., Horacio Larreguy and Juan Pablo Xandri.** 2020. "Testing Models of Social Learning on Networks: Evidence from Two Experiments" *Econometrica*, 88(1): 1-32.
- Charness, Gary and Dan Levin.** 2005. "When Optimal Choices Feel Wrong: A Laboratory Study of Bayesian Updating, Complexity, and Affect" *American Economic Review*, 95(4): 1300-1309.
- Charness, Gary and Dan Levin.** 2009. "The Origin of the Winner's Curse: A Laboratory Study" *American Economic Journal: Microeconomics*, 1(1): 207-236.
- Charness, Gary, Michael Naef and Alessandro Sontuoso.** 2019. "Opportunistic Conformism" *Journal of Economic Theory*, 180: 100-134.
- Chen, Si and Hannah Schildberg-Hörisch.** 2019. "Looking at the Bright Side: The Motivational Value of Confidence" *European Economic Review*, 120: 103302.
- Compte, Olivier and Andrew Postlewaite.** 2004. "Confidence-Enhanced Performance" *American Economic Review*, 94(5): 1536-1557.
- DeGroot, Morris H.** 1974. "Reaching a Consensus" *Journal of the American Statistical Association*, 69(345): 118-121.
- Erev, Ido and Alvin E. Roth.** 1998. "Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria" *American Economic Review*, 88(4): 848-881.
- Esponda, Ignacio and Emanuel Vespa.** 2014. "Hypothetical Thinking and Information Extraction in the Laboratory" *American Economic Journal: Microeconomics*, 6(4): 180-202.
- Fatas, Enrique, Shaun P. Hargreaves Heap and David Rojo Arjona.** 2018. "Preference Conformism: An Experiment" *European Economic Review*, 105: 71-82.
- Fechner, Gustav.** 1860. "Elements of Psychophysics". New York, NY: Holt, Rinehart, and Winston Inc.
- Fischer, Mira and Dirk Sliwka.** 2018. "Confidence in Knowledge or Confidence in the Ability to Learn: An Experiment on the Causal Effects of Beliefs on Motivation" *Games and Economic Behavior*, 111: 122-142.
- Gale, Douglas and Shachar Kariv.** 2003. "Bayesian Learning in Social Networks" *Games and Economic Behavior*, 45(2): 329-346.
- Gioia, Francesca.** 2017. "Peer Effects on Risk Behaviour: The Importance of Group Identity" *Experimental Economics*, 20(1): 100-129.
- Golub, Benjamin and Matthew O. Jackson.** 2010. "Naive Learning in Social Networks and the Wisdom of Crowds" *American Economic Journal: Microeconomics*, 2(1): 112-149.
- Grimm, Veronika and Friederike Mengel.** 2020. "Experiments on Belief Formation in Networks" *Journal of the European Economic Association*, 18(1): 49-82.

- Hall, Crystal C., Lynn Ariss and Alexander Todorov.** 2007. "The Illusion of Knowledge: When More Information Reduces Accuracy and Increases Confidence" *Organizational Behavior and Human Decision Processes*, 103(2): 277-290.
- Hedden, Trey and Jun Zhang.** 2002. "What Do You Think I Think You Think?: Strategic Reasoning in Matrix Games" *Cognition*, 85(1): 1-36.
- Hey, John D., Gianna Lotito and Anna Maffioletti.** 2010. "The Descriptive and Predictive Adequacy of Theories of Decision Making Under Uncertainty/Ambiguity" *Journal of Risk and Uncertainty*, 41(2): 81-111.
- Hey, John D. and Chris Orme.** 1994. "Investigating Generalisations of Expected Utility Theory Using Experimental Data" *Econometrica*, 62(6): 1291-1326.
- Kagel, John H.** 1995. "Cross-Game Learning: Experimental Evidence from First-Price and English Common Value Auctions" *Economics Letters*, 49(2): 163-170.
- Kahneman, Daniel.** 2003. "A Perspective on Judgment and Choice: Mapping Bounded Rationality" *American Psychologist*, 58(9): 697-720.
- Kahneman, Daniel.** 2011. *Thinking, Fast and Slow*. Farrar, Straus and Giroux.
- Legrenzi, Paolo, Vittorio Girotto and Philip N. Johnson-Laird.** 1993. "Focussing in Reasoning and Decision Making" *Cognition*, 49(1-2): 37-66.
- Levin, Dan, John H. Kagel and Jean-Francois Richard.** 1996. "Revenue Effects and Information Processing in English Common Value Auctions" *American Economic Review*, 86(3): 442-460.
- Levin, Dan, James Peck and Asen Ivanov.** 2016. "Separating Bayesian Updating from Non-Probabilistic Reasoning: An Experimental Investigation" *American Economic Journal: Microeconomics*, 8(2): 39-60.
- Luce, R. Duncan.** 1959. "Individual Choice Behavior: A Theoretical Analysis". New York, NY: Wiley.
- Machina, Mark J. and Marciano Siniscalchi.** 2014. "Ambiguity and Ambiguity Aversion" in *Handbook of the Economics of Risk and Uncertainty* (Vol. 1), ed. M. Machina and K. Viscusi. North-Holland.
- Martínez-Marquina, Alejandro, Muriel Niederle and Emanuel Vespa.** 2018. "Failures in Contingent Reasoning: The Role of Uncertainty" *American Economic Review*, 109(10): 3437-3474.
- O'Donoghue, Ted and Jason Somerville.** 2018. "Modeling Risk Aversion in Economics" *Journal of Economic Perspectives*, 32(2): 91-114.
- Rabin, Matthew.** 2013. "Incorporating Limited Rationality into Economics" *Journal of Economic Literature*, 51(2): 528-543.
- Rabin, Matthew and Joel L. Schrag.** 1999. "First Impressions Matter: A Model of Confirmatory Bias" *Quarterly Journal of Economics*, 114(1): 37-82.
- Reiter, Andrea M. F., Shinsuke Suzuki, John P. O'Doherty, Shu-Chen Li and Ben Eppinger.** 2019. "Risk Contagion by Peers Affects Learning and Decision-Making in Adolescents" *Journal of Experimental Psychology: General*, 148(9): 1494-1504.
- Rutström, E. Elisabet and Nathaniel T. Wilcox.** 2009. "Stated Beliefs Versus Inferred Beliefs: A Methodological Inquiry and Experimental Test" *Games and Economic Behavior*, 67(2): 616-632.
- Trautmann, Stefan T. and Gijs van de Kuilen.** 2015. "Belief Elicitation: A Horse Race Among Truth Serums" *Economic Journal*, 125(589): 2116-2135.
- Tversky, Amos and Daniel Kahneman.** 1973. "Availability: A Heuristic for Judging Frequency and Probability" *Cognitive Psychology*, 5(2): 207-232.
- Weber, Roberto A.** 2003. "'Learning' with No Feedback in a Competitive Guessing Game" *Games and Economic Behavior*, 44(1): 134-144.
- Wilson, Andrea.** 2014. "Bounded Memory and Biases in Information Processing" *Econometrica*, 82(6): 2257-2294.
- Zizzo, Daniel John, Stephanie Stolarz-Fantino, Julie Wen and Edmund Fantino.** 2000. "A Violation of the Monotonicity Axiom: Experimental Evidence on the Conjunction Fallacy" *Journal of Economic Behavior & Organization*, 41(3): 263-276.