Acting on Beliefs:
Feedback and Best-Response Deviations in Theory and Experiment

Alessandro Sontuoso,[a,b] *Cristina Bicchieri*,[b] *Alexander Funcke*,[b] and *Einav Hart*[c]

June 30, 2026

Abstract. We theoretically and experimentally study a multiplayer game with risky prospects to clarify why individuals often fail to act consistently on their stated beliefs. In this game, one action becomes uniquely optimal whenever the expected share of others choosing it falls below a threshold: in that case, a player's best response is well defined regardless of risk attitude, providing a clean test of how beliefs map to actions. We embed this game in a stochastic-choice framework that decomposes best-response failures into three sources of belief–action inconsistency: (i) belief dispersion, (ii) belief location relative to the threshold, and (iii) the scale of behavioral noise (i.e., the disturbance distribution). An experiment with and without feedback on others' choices shows that all three factors predict best-response rates; yet only the noise channel accounts for a meaningful part of the treatment difference. Specifically, feedback reduces "excess switching"—switching more often than would be expected if each choice were an independent random draw from a subject's long-run action frequencies—a direct marker of noise-driven behavior.

KEYWORDS: Belief–action coherence; Behavioral noise; Threshold games; Bayesian games.

JEL Classification Numbers: C72, C92, D83, D91.

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

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

I. Introduction

In strategic settings, rational play requires individuals to form beliefs about others' behavior and to choose actions that are optimal given those beliefs. Yet in practice, belief–action coherence is often imperfect: many individuals fail to choose the action that maximizes their expected utility given their own (stated) beliefs. Rather than treating such inconsistencies as anomalies, recent work has turned to examining the conditions under which individuals are more likely to best respond to their beliefs. This literature asks, both theoretically and empirically, which features of the strategic environment—or which interventions—can promote optimal play.

One such intervention is the provision of feedback about others' past behavior, which has been shown to improve decision quality in some settings. Part of this effect may operate through belief formation: by observing what others have done, individuals can form more accurate expectations. But feedback may also serve a distinct role by enhancing *belief–action coherence*, namely, the likelihood that individuals act on (i.e., best respond to) their beliefs, regardless of whether those beliefs are accurate. The extent of this latter effect, and the mechanisms that drive it, remain open questions. The present paper tackles these questions in a novel strategic environment.

Prior research has documented belief–action inconsistency even in simple settings with incentivized belief elicitation. In repeated 2×2 games where opponents' past actions are observable, belief–action coherence typically reaches 75–90%, whether past actions are observed directly (as in Nyarko and Schotter, 2002) or inferred from realized payoffs (as in Manski and Neri, 2013). Belief–action coherence is lower in *repeated games* with larger action spaces, such as 3×3 or 4×4 games (e.g., Danz, Fehr, and Kübler, 2012; Hyndman, Ozbay, Schotter, and Ehrblatt, 2012) and in *one-shot games* (Rey-Biel, 2009; Ivanov, 2011; Polonio and Coricelli, 2019, among others). Overall, failures to best respond to stated beliefs appear to increase when games are more complex or players lack strategic experience, whether because play is one-shot or because repeated play provides no information about past outcomes or others' behavior. Notably, related suboptimalities have been found in non-strategic tasks as well (e.g., Zizzo, Stolarz-Fantino, Wen, and Fantino, 2000; Charness and Levin, 2005), suggesting that such failures are not merely driven by the interpersonal nature of belief formation, nor by artifacts of the belief elicitation process.[1]

[1] Belief elicitation has become a standard tool in the behavioral toolbox, with simple incentive-compatible methods (e.g., frequency or interval scoring) shown to yield truthful measures of expectations without distorting motivations (Charness, Gneezy, and Rasocha, 2021).

Various accounts have been proposed to explain the apparent disconnect between beliefs and actions. These include limited strategic reasoning, reliance on heuristics, inattention to others' payoffs, and explanations based on other-regarding preferences or risk attitudes, which may appear as deviations from best-response behavior while in fact reflecting different objective functions (see the extensive discussion in Alempaki, Colman, Kölle, Loomes, and Pulford, 2022). Other proposed factors include uncertainty and noise in the evaluation of options, which the stochastic-choice literature ties to the strength of the decision-maker's underlying preferences (e.g., Alós-Ferrer and Garagnani, 2022; Alós-Ferrer, Buckenmaier, Garagnani, and Rustichini, 2026).

While interventions that aim to boost strategic sophistication have yielded mixed results,[2] simple feedback provision has been somewhat more effective at increasing best-response rates.[3] In repeated-game settings, providing information about others' past actions (even without revealing outcomes) can reduce non-strategic choices and bring behavior closer to equilibrium play (Danz et al., 2012). Yet the channels through which these improvements occur remain theoretically and experimentally underexplored, especially when feedback is noisy or imprecise. Put another way, does feedback just help agents form (more accurate) beliefs, or does it also help them act on the beliefs they hold?

This paper sheds light on the latter possibility: that feedback enhances the likelihood of best responding to one's own beliefs, even when belief accuracy is unchanged. Our analysis connects the game-theoretic literature on belief–action coherence with the stochastic-choice literature on noisy evaluation, within a single framework that theoretically isolates and empirically tests *belief-based* and *noise-based* channels (a separation that, to our knowledge, has not been drawn in strategic settings, where the relevant beliefs concern others' play). Specifically, we propose a novel game that permits the identification of best-response failures, while ruling out alternative explanations such as risk attitudes or a preference for the safe option. We then use

[2] For instance, Alempaki et al. (2022) implemented a "structured" deliberation treatment, prompting subjects to explicitly compare payoffs before choosing an action, but found no significant increase in best-response rates relative to an unstructured condition; the findings suggest that directly encouraging more systematic reasoning might be insufficient to enhance strategic sophistication. At the same time, the hypothetical reasoning literature indicates that certain interventions can be effective, particularly when they aid subjects in reasoning through contingencies (Niederle and Vespa, 2023).

[3] Although the evidence for feedback-based interventions is generally more promising than for attempts to train strategic reasoning, it is far from definitive. The literature provides a mixed picture: feedback may overwhelm decision-makers, especially in complex environments, leading to cherry-picking of information and biased or unstable beliefs, so that more information may have no benefit or even be harmful (Wilson, 2014; Hall, Ariss, and Todorov, 2007); yet it may also improve performance, both by supporting learning and by increasing engagement and attention (Compte and Postlewaite, 2004; Fischer and Sliwka, 2018).

an experiment to test how feedback affects the mapping from beliefs to actions. Put simply, we ask: when individuals hold the same belief with and without feedback, are they more likely to choose the optimal action under feedback? And if so, through which mechanism?

To address these questions, we introduce a multiplayer game in which optimal actions generally depend on the individual's intrinsic risk preferences, modeled via a Bayesian framework with uncertainty about opponents' types (specifically, their risk attitudes). A crucial feature of this game is that one action becomes uniquely optimal whenever beliefs about others' behavior fall below a given threshold. As a result, within a given belief range, best responding is well defined regardless of a player's risk preferences: for illustration, if a player believes that fewer than a threshold share (e.g., 40%) of others will choose a specific action, denoted "B", then B is the unique optimal choice.

We first provide a theoretical analysis of this game, and then present a between-subjects design in which participants are randomly assigned to one of two treatments: a *main treatment* with feedback and a *control treatment* without. In both conditions, subjects repeatedly play the same game against a large and anonymous population; after each play, they report their belief about the proportion of other participants who chose B. In both conditions, outcome realizations are withheld until the end of the session, ruling out experiential learning so that choices reflect belief-based reasoning rather than reinforcement from past payoffs. In short, the treatments differ solely in the provision of feedback: in each round of the main treatment, after the choice and belief elicitation stages, participants receive noisy information regarding the actions taken by a random sample of other participants. No such feedback is provided in the control treatment.

Even when feedback is very coarse, it offers a test of whether subjects update their beliefs in response to information and, more importantly, allows us to examine how feedback provision shapes behavior independently of belief accuracy. Our theoretical framework identifies three channels through which belief–action inconsistency may arise: (i) *belief dispersion*, the variability of an individual's beliefs across rounds; (ii) *per-round belief location relative to the threshold*, how far beliefs lie from the cutoff in a given round; and (iii) the *scale of behavioral noise*, random departures from optimal play unrelated to beliefs. The first two channels imply that feedback can promote best-response behavior by stabilizing beliefs (reducing dispersion) or by shifting their average away from the cutoff, thereby making expected-utility differences among actions more clear-cut (regardless of belief accuracy). The third channel captures reductions in erratic play (due to psychological factors such as hesitation, disengagement, or other unobserved influences) that are unrelated to the structure of beliefs. We refer to this as the "noise channel" and show how it can be identified empirically: residually, as the treatment effect that remains once the belief-based channels are controlled for, and directly, through a measure of excess switching.[4]

In the raw data, best-response failures account for 11.64% of choices in the control treatment and 6.49% in the main treatment. The econometric analysis shows that all three factors matter in predicting best-response rates. Greater belief dispersion increases the likelihood of best-response failures, and beliefs farther from the threshold reduce them, indicating that the distribution of beliefs relative to the decision cutoff (rather than their accuracy) helps explain mistakes. Yet although these belief-based factors strongly predict best-response rates, they do not account for the treatment effect: behavioral noise, measured by excess switching, is the only identified channel that carries at least part of it. Specifically, the data reveal that feedback reduces *excess switching*—switching actions more often than would be expected if each choice were an independent random draw from one's long-run action frequencies—a direct marker of noise-driven behavior. So, while belief dispersion and cutoff distance remain important predictors across treatments, feedback specifically improves performance by attenuating the noise component (possibly reflecting hesitation, lack of confidence, or other unobserved psychological factors), thereby helping individuals act more consistently on the beliefs they hold.

The game we introduce is also a contribution in its own right: beyond serving as a test of belief–action coherence, it captures key dynamics of large-population interactions under risk. These dynamics arise from threshold-based incentives. In particular, here players choose among two strategic options, such as contributing or not to a collective endeavor (e.g., vaccinate or not), and one exit option that offers a fixed payoff independent of others' behavior or chance (the preferred choice for risk-averse individuals). When the expected share of free-riders falls below a cutoff, a single action becomes strictly optimal for all preference types, and yet subjects often fail to best respond. Our findings show that feedback, even if low in informational content, can raise best-response rates, suggesting a behaviorally grounded lever for improving decision quality in large-population interactions.

---

[4] Beyond these channels, our framework further allows us to rule out alternative explanations for the treatment effect, such as non-neutral risk attitudes, which are controlled for by the structure of the game, and non-standard preferences (which are unlikely to vary systematically across treatments).

The rest of the article is organized as follows: 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*

We start by defining the strategic environment. Let $N = \{1, \dots, n\}$ denote the set of players; for each player $i \in N$, i's payoff is $m_i(s, \theta)$, where $s = (s_i, s_{-i})$ is an action profile (i.e., an n-tuple of actions) and θ is a move by nature. For each player $i \in N$, let $s_i \in S_i = \{A, B, C\}$ and $s_{-i} \in S_{-i} = \prod_{j \in N,\, j \neq i} S_j$. Nature's move is $\theta \in \Theta = \{heads, tails\}$, representing a fair coin toss with equally likely outcomes; the distribution, but not the realization, is known to all players. Each player *simultaneously* chooses an action in $\{A, B, C\}$ with the goal to maximize her expected utility (before Nature's move is realized). In defining the payoffs $m_i(s, \theta)$, we partition action profiles by whether a threshold d is met with respect to the population-level frequency of action B. Below we consider the specific parameterization used in the experiment and illustrate it in terms of a vaccination problem, though the experiment itself made no reference to vaccination.

(s, θ)	***A*** [*vaccinate*]	***B*** [*don't vaccinate*]	***C*** [*self-isolate*]
$\underline{S} \times \{heads, tails\}$ [scenario x: *herd immunity*]	0.5	3	0.75
$\overline{S} \times \{heads\}$ [scenario y: *epidemic outbreak*]	1	-1.5	0.75
$\overline{S} \times \{tails\}$ [scenario z: *no outbreak*]	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, θ) 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 fewer than a fraction d of all players choose B, whereas $s \in \overline{S}$ if at least d do (with $d = 0.4$).

Let $d = 0.4$ (an arbitrary value, which does not affect our hypotheses), and assume this value is common knowledge among all players. Given this, in what follows $\underline{S}$ and $\overline{S}$ denote the sets of *action profiles with*, respectively, *fewer than* 40% and *at least* 40% of players choosing B.[5] Finally, suppose the coin has been tossed and simultaneously all players have chosen an action. The resulting payoff depends on whether the action profile s lies in $\underline{S}$ or $\overline{S}$, and on the coin toss θ; in each case, i's payoff from A, B or C is defined in the respective column of Table 1.[6] To sum up, there are three actions, ranging from very risky ("B"), to mildly risky ("A"), to riskless ("C"). Conditional on $\overline{S}$, actions A and B are mean-preserving spreads of C. The reader can anticipate that equilibrium predictions depend on players' beliefs about others' behavior *and* on their risk attitudes.

It is now useful to emphasize a key feature of the game. When fewer than 40% choose B, the chance move is immaterial: in that case, one's own risk attitude is irrelevant and one's preferences reduce to $B \succ C \succ A$, as in the first row of Table 1 (risk-free scenario x). Put simply, in that case—from an individual standpoint—*one's best response is well defined and type-invariant*, providing a clean test of how beliefs map to actions. (As we will see, the coplayers' risk preferences still matter to the extent that they determine whether an individual's belief lies above or below the 40% threshold.) By contrast, when 40% or more choose B, payoffs vary with the chance move and thus one's own risk attitude becomes relevant, yielding different preference orderings across types.[7]

This payoff structure illustrates why such a game is well suited to study best-response deviations while ruling out confounds from risk preferences: best responses are well defined and type-invariant within a given belief range. Given this, in II.2–II.3 we model uncertainty about the coplayers' risk preferences, which determine whether an individual's belief lies in that range. In II.4–II.5, we establish the equilibrium benchmark, which characterizes the deterministic mapping between beliefs and optimal actions in the absence of errors (this step is essential for predicting how feedback should shift behavior under different signals). Lastly, in II.6 we introduce stochastic disturbances that relax the perfect-optimization assumption, allowing occasional

[5] Formally, $\underline{S} := \left\{ s \in \prod_{j \in N} S_j : \frac{|\{j \in N:\ s_j = B\}|}{|N|} < d \right\}$, where $|\cdot|$ denotes set cardinality; $\overline{S} := S \setminus \underline{S}$.

[6] As an illustration, consider an individual deciding whether to get vaccinated at the onset of a potential epidemic outbreak. By choosing C (*self-isolate*) one is no longer susceptible to becoming infected with the virus. Instead, by choosing A (*vaccinate*) or B (*don't vaccinate*) one positively or negatively contributes to the herd immunity threshold: in that case, one's payoff depends on the population-level behavior and on a move by nature. That is, when less than 40% of all individuals free ride (i.e., more than 60% of the population get a vaccine or self-isolate), then a risk-free situation occurs ("scenario x": herd immunity). By contrast, when 40% or more free ride, then a negative shock *may* occur ("scenario y": epidemic outbreak) or *may not* occur ("scenario z": no outbreak), each with a 50 percent chance.

[7] The common assumption that every player is *risk-neutral* is behaviorally implausible here. In fact, if that were true, then action B would be weakly dominant for all the players: a fact that is clearly refuted by the experimental data.

deviations from the deterministic equilibrium benchmark: this extension enables us to formalize three sources of belief–action inconsistency and ground our main hypotheses.

2. *Risk preferences*

In what follows, we consider three player types, t_a, t_b, and t_c, differing in their attitudes toward risk. Type t_c is risk-averse, preferring the sure prospect with value $m = 0.75$ to any risky prospect of equal expected value. The remaining types, t_a and t_b, are risk-seeking: they differ in the variance of the risky actions they most prefer. Specifically,

$$A \succ_{t_a} C \succ_{t_a} B, \tag{1}$$

$$B \succ_{t_b} A \succ_{t_b} C, \tag{2}$$

$$C \succ_{t_c} A \succ_{t_c} B, \tag{3}$$

where $\succ_t$ denotes the preference relation for type t. Formally, $t \in T = \{t_a, t_b, t_c\}$; for any $s \in \overline{S}$, preferences over risky prospects (conditional on the threshold being met) are defined as above.

A few comments. Note that *risk-neutral* players are indifferent among the three actions once the threshold is met (i.e., when 40% or more choose B); hence, without loss of generality we focus on cases where individuals behave as either risk-seeking or risk-averse players. Also note that, conditional on the threshold being met, the analysis associated with expression (1) would be qualitatively similar under nearby preference orderings such as $A \succ B \succ C$; similar arguments apply to (2) and (3). For simplicity, we therefore restrict attention to preference orderings (1)–(3).

Below, we will model this problem as a Bayesian game (i.e., a game with incomplete information about the others' preferences). To do so, we must cardinalize preference orderings (1)–(3) so that each type's expected utility from action s_i can be computed conditional on $(s_i, s_{-i}) \in \overline{S}$. Thus, for each $t \in T$ we assume t's preferences (conditional on the threshold being met) admit an expected-utility representation u_t consistent with orderings (1)–(3), where Nature's two outcomes are equally likely.[8] For ease of reference, below we respectively denote by α_t, β_t, γ_t the expected utility values from A, B, C conditional on $s \in \overline{S}$.[9] Given this, preference orderings (1)–(3) are equivalent to inequalities (4)–(6), respectively:

[8] This need for cardinalization does not arise when the threshold is not met (i.e., for any action profile $s \in \underline{S}$). There, for all types $t \in \{t_a, t_b, t_c\}$, i's utility values simply equal the monetary payoffs reported in the first row of Table 1.

[9] More formally, whenever the threshold is met, the expected utility values from A, B, and C are $\alpha_t := \sum_{\theta \in \Theta} \Pr(\theta)\, u_t\big(m_i(A, s_{-i}, \theta)\big)$, $\beta_t := \sum_{\theta \in \Theta} \Pr(\theta)\, u_t\big(m_i(B, s_{-i}, \theta)\big)$, and $\gamma_t := \sum_{\theta \in \Theta} \Pr(\theta)\, u_t\big(m_i(C, s_{-i}, \theta)\big)$, respectively, for each $t \in T$.

$$\alpha_{t_a} > \gamma_{t_a} > \beta_{t_a}, \tag{4}$$

$$\beta_{t_b} > \alpha_{t_b} > \gamma_{t_b}, \tag{5}$$

$$\gamma_{t_c} > \alpha_{t_c} > \beta_{t_c}, \tag{6}$$

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 now lay out a framework to capture i's uncertainty about the coplayers' preference types. The type set $T = \{t_a, t_b, t_c\}$ is common knowledge, and each type $t \in T$ is characterized by conditions (4)–(6). Every player knows her own type but not those of the others. Since B is a dominant action only for type t_b (see condition (2) and Table 1), what matters for any non-t_b player's optimal choice is the share of t_b players in the population. We therefore model the aggregate uncertainty directly, through a binary state of the world $\omega \in \Omega = \{H, L\}$: the population is t_b-rich in the high state H and t_b-poor in the low state L. The state is drawn according to $\Pr[\omega = H] = p$ and $\Pr[\omega = L] = 1 - p$, with $p \in (0,1)$.

Conditional on the state, players' types are independent, and the probability that any given player j is of type t_b is

$$\Pr[t(j) = t_b \mid \omega] = \begin{cases} \pi_H, & \omega = H \\ \pi_L, & \omega = L \end{cases} \tag{7}$$

where $0 < \pi_L < 0.4 < \pi_H < 1$; the residual probability (i.e., $1 - \pi_H$ in state H, and $1 - \pi_L$ in state L) is divided between types t_a and t_c.

We interpret the game as a large-population interaction. By the law of large numbers, conditional on the state, the realized population share of t_b players is deterministic: it equals π_H in state H, and π_L in state L. Because $0 < \pi_L < 0.4 < \pi_H < 1$, the state pins down whether the unconditional B-choosers (namely, the t_b players) alone meet the threshold $d = 0.4$: they fall short of it in state L and exceed it in state H. Hence ω (and not the full profile of individual types) is the payoff-relevant state of the world. The prior over ω therefore informs the optimal behavior of t_a and t_c players. Indeed, the preferences of these two types depend on their expectations about the frequency of B choices in the population, which in turn depends on the share of t_b players (and hence on the state ω), since B is a dominant action only for type t_b.[10]

[10] Out of equilibrium, a player may hold a subjective prior over the state; in equilibrium, instead, all players' beliefs are consistent with the prior over ω and the conditional type distribution in (7), per the standard common-prior assumption.

4. *Feedback*

Here we describe noisy signals about the state ω. Let $W_i \in [0, 1]$ denote a *private signal*, observed by player i, whose distribution depends on ω. We interpret W_i as information about others' past play: concretely, the frequency of B choices among a random sample of coplayers in the previous round. (Under the assumption that preferences are stable over short periods, W_i carries imperfect information about the realized share of t_b players, and thus about the state ω.) The reader can anticipate that i's optimal action depends only on whether the expected share of B choices lies below or above the threshold d, so the key payoff-relevant feature of the signal is its location relative to d. We therefore summarize W_i by a piece of feedback $f \in F = \{low, high\}$, defined as follows: $f = low$ if $W_i < d$, and $f = high$ if $W_i \geq d$, where low and $high$ respectively point to states L and H (with underlying shares π_L and π_H). We assume the signal's distribution places probability $q > 0.5$ on the side of d that matches the true state: $\Pr[W_i < d \mid \omega = L] = \Pr[W_i \geq d \mid \omega = H] = q$.

Two implications follow. First, values below (resp. above) the threshold d are more likely in the low (resp. high) state than in the opposite state, since $\Pr[W_i < d \mid L] = q > 1 - q = \Pr[W_i < d \mid H]$; hence, a low draw is evidence for a t_b-poor population. Second, feedback f is informative but imperfect: it does not perfectly identify whether the share of t_b players equals π_L or π_H, but it falls on the correct side with probability q. For a concrete example, hearing that some individuals' (e.g., the neighbors') past choices were 33% B, and thus under 40%, suggests that the overall share of unconditional B-choosers (i.e., t_b players) in the population might be below the 40% threshold; hearing over 40% suggests the opposite. In brief, signals provide a basis for updating beliefs about the state ω, and hence about the share of t_b players.

5. *Equilibrium analysis*

In a Bayesian Nash equilibrium, each player chooses a best response given her type, signal, and correct beliefs about the state and others' strategies. We define expected utility at the interim stage (i.e., after a signal has been observed). Formally, let $t(i)$ denote i's type, and let $s_{-i}(\omega)$ denote the profile of the coplayers' actions at ω. Then, i's expected utility from $(s_i, s_{-i}(\omega))$ is

$$U_i(s_i, f, t(i)) = \sum_{\omega \in \Omega} \sum_{\theta \in \Theta} \Pr[\omega | f] \cdot \Pr[\theta] \cdot u_{t(i)}(m_i(s_i, s_{-i}(\omega), \theta)), \tag{8}$$

where $\Pr[\omega | f]$ is the posterior probability of state ω given feedback f, computed via Bayes' rule. For brevity, we write expression (8) as $U_i(s_i, \cdot)$, suppressing arguments where no confusion arises.

As noted earlier, the expected share of t_b players in the population is the key variable for the threshold comparison. Accordingly, we denote the *prior expected share* of t_b players by $\bar{\pi}(t_b) = p \cdot \pi_H + (1-p) \cdot \pi_L$.[11] Following a signal, Bayes' rule yields the *posterior expected share* $\bar{\pi}(t_b \mid f)$. For example, conditional on $high$ feedback f (i.e., for any signal $W_i \in [0.4, 1]$), we have

$$\bar{\pi}(t_b \mid f = high) = \frac{pq}{pq+(1-p)(1-q)} \cdot \pi_H + \frac{(1-p)(1-q)}{pq+(1-p)(1-q)} \cdot \pi_L; \tag{9}$$

by contrast, conditional on low feedback f (i.e., for any signal $W_i \in [0, 0.4)$), we have

$$\bar{\pi}(t_b \mid f = low) = \frac{p(1-q)}{p(1-q)+(1-p)q} \cdot \pi_H + \frac{(1-p)q}{p(1-q)+(1-p)q} \cdot \pi_L. \tag{10}$$

In the following, we write the *posterior probability of the high state* as $p_f := \Pr(H|f)$, so that $\bar{\pi}(t_b \mid f) = p_f \pi_H + (1-p_f)\pi_L$. These beliefs may be subjective in general, but in equilibrium they coincide with the type and signal distribution implied by the model. We present the equilibria below.

Proposition 1 – *Bayesian Nash equilibria in pure actions.* In every Bayesian Nash equilibrium of the threshold game with preference types defined by (4)–(6), 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 p_f (i.e., the posterior probability of the high state H given feedback f) is greater than a type-specific cutoff point ψ_t,

$$\text{with } p_f > \psi_{t_a} \equiv \min\left\{\frac{5}{2\alpha_{t_a}-2\beta_{t_a}+5}, \frac{9}{12-4\beta_{t_a}}\right\} \text{ and } p_f > \psi_{t_c} \equiv \frac{9}{12-4\beta_{t_c}},$$

then no one chooses B other than t_b players. Specifically, all t_c players choose C and each t_a player chooses her best non-B action, namely A if $p_f \geq \frac{1}{4\alpha_{t_a}-2}$, and C otherwise.

(ii) If at least one of the cutoff conditions is not satisfied, let $r = d - \pi_L$ be the residual room below the threshold in the low state (L). Let $\omega_L(O)$ denote the low-state population share of players in any arbitrary subset $O \subseteq G_f$, where G_f is the set of players for whom the cutoff condition is not satisfied. Given this, two cases ensue.

- If $\omega_L(G_f) < r$, then B is chosen by all players in G_f without reaching the threshold.
- Otherwise, B is chosen by a fraction of players (i.e., by all players in any maximal feasible subset $O \subseteq G_f$ such that $\omega_L(O) < r$). All remaining non-t_b players take their best non-B action; that is, t_c players choose C, while t_a players choose A if $p_f \geq \frac{1}{4\alpha_{t_a}-2}$ and C otherwise.

---

[11] As noted in II.3: the probability of j being type t_b equals π_H with probability p, and π_L with probability $1-p$, where $0 < \pi_L < 0.4 < \pi_H < 1$. Recall that feedback f points to the true state with probability $q > 0.5$. Finally, since the game is played in a large population, knowledge of one's own type does not affect beliefs about the overall share of t_b players.

Proof. See Appendix A.

A brief interpretation follows. Equilibrium (i) in Proposition 1 corresponds to the case where both t_a and t_c players assign sufficiently high posterior probability to the high (t_b-rich) state H, so that the relevant cutoff conditions are satisfied. Since those cutoffs reflect the expected utility values each type assigns to risky actions, satisfying them means that both t_a and t_c players prefer to avoid B and take their best non-B action: thus, in equilibrium (i) t_c players choose C, while t_a players choose A or C (depending on expected utility values). By contrast, the class of equilibria (ii) arises when t_a or t_c players expect relatively few unconditional B-choosers (i.e., t_b players): in that case, some or all of them may find it optimal to choose B, but only to the extent that their doing so does not push the low-state share of B choices up to the threshold.

6. *Stochastic choice and feedback channels*

To allow for deviations from the benchmark equilibrium predictions—arising from lapses, indecisiveness, or other factors—we now relax the perfect-optimization assumption. In the deterministic framework, each player chooses the action maximizing expected utility given their type and belief; instead, below we allow this mapping from beliefs and types to actions to be imperfect by adding an idiosyncratic random disturbance term in the evaluation process.[12]

Formally, player i with type t and belief $\bar{\pi}$ picks an action s_i to maximize $U_i(s_i,\cdot)$, with actual choice subject to an unobserved stochastic disturbance $\varepsilon_i \sim G_{\eta_{\text{condition}}}$, where $G_{\eta_{\text{condition}}}$ is a generic distribution with finite variance and $\eta_{\text{condition}}$ is an unrestricted parameter vector that may vary across environments (i.e., conditions).[13] In a nutshell, ε_i captures random variation in choice; its distribution $G_{\eta_{\text{condition}}}$ does not depend on the player's type or on the actions being compared. Given this, let $s_i^*(t,\bar{\pi})$ denote the deterministic best response ("BR") for player i. With an individual disturbance ε_i, the probability of a best response is

$$P_{it}(\text{BR} \mid t, \bar{\pi}, \text{condition}) = \Pr\left[\varepsilon_i \leq \underbrace{\min_{s_i' \neq s_i^*}\{U_i(s_i^*,\cdot) - U_i(s_i',\cdot)\}}_{\Delta_i(\bar{\pi},t)}\right]. \tag{11}$$

[12] Players are assumed *not* to anticipate errors by others, placing the analysis within the class of stochastic choice models originating with Fechner (1860/1912). See also Luce (1959) and Becker, DeGroot, and Marschak (1963).

[13] We deliberately leave $G_{\eta_{\text{condition}}}$ unrestricted rather than committing to a specific parametric form. A single-parameter error distribution would attribute all belief–action inconsistency to one precision parameter, with no separate role for the belief-based channels that this paper sets out to disentangle (for instance, the logit quantal-response specification of McKelvey and Palfrey, 1995, is one such case).

Here, $\Delta_i(\bar{\pi}, t)$ is the *smallest utility gap* between the best action and any alternative; so, (11) states that the best response is chosen whenever the disturbance does not overturn the smallest expected-utility advantage of the best action. The interpretation is intuitive: the larger the expected utility gap between alternatives, the clearer the choice, and the less susceptible it is to random noise. (The equilibrium benchmark, analyzed before, is obtained when $\varepsilon_i \equiv 0$.) Given this, we identify three channels as *sources of variability in best-response failures*, each linked to the inequality $\varepsilon_i \leq \Delta_i(\bar{\pi}, t)$. Irrespective of belief accuracy, feedback may affect behavior by operating through one or more of these channels:

1. **Belief dispersion channel** – Each round, a player acts on a single (point) belief about how others will play (i.e., the belief on which their best response is defined), which may be more or less settled across rounds. When unsettled, this belief wobbles from round to round and, through that variation, is more likely to fall near the decision cutoff in some rounds (where the expected utility gap $\Delta_i(\bar{\pi}, t)$ between the best action and its closest alternative is small, and near-indifference makes choices most susceptible to noise). By settling the belief, feedback reduces this round-to-round variability; so long as the settled belief is not itself at the cutoff, near-indifference cases become less frequent, and best-response failures decline.
2. **Belief level channel** – Feedback may move beliefs farther from the cutoff, even if their round-to-round variability is unchanged (this channel concerns where beliefs sit relative to the cutoff in a given round, as opposed to how much they vary across rounds). When i's belief lies farther from the threshold in a given round, $\Delta_i(\bar{\pi}, t)$ is larger (as in the channel above), so the condition $\varepsilon_i \leq \Delta_i(\bar{\pi}, t)$ is met more often, and best-response failures decline.
3. **Behavioral noise channel** – Feedback may attenuate noise from psychological factors such as hesitation, disengagement, or other unobserved influences. In our framework, this is modeled as a change in the distribution of disturbances $G_{\eta_{\text{condition}}}$, with parameters that yield smaller deviations. As a result, ε_i is more likely to fall within the range satisfying the best-response condition, even without changes in $\Delta_i(\bar{\pi}, t)$. This channel thus reflects best-response variability beyond that explained by belief dispersion or belief level.

To recap, the first two channels operate by increasing the utility gap $\Delta_i(\bar{\pi}, t)$, while the noise channel operates by reducing the magnitude of ε_i across plays. Together, our novel game and theoretical framework supply the exact ingredients needed for a clean empirical analysis: a *belief region in which the best response is type-invariant* (so that departures from it are identifiable as errors rather than driven by risk attitudes or a preference for the safe option), and a *decomposition of those errors into three channels* (so that any effect of feedback can be attributed to one or more of them). Sections III and IV develop and test these predictions.

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" (i.e., feedback condition).

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 on the screen, labeled "A", "B", "C". Each subject was told that her earnings, beyond a flat participation fee, depended on her own choice, the choices of all the other participants in the session, and the outcome of a fair computer-generated coin toss. After reading the instructions, subjects were required to answer a set of comprehension questions.

A few comments are due. First, we stress that actions were simply labeled A, B, and C (i.e., no reference was made to vaccinations, outbreaks, etc.). Second, *letter-outcome pairs* (e.g., whether the letter B denotes the threshold-relevant option rather than, say, the safe option) *were randomized across participants*, 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 C.) For ease of exposition, the remainder of the paper will use the exact same letter-outcome pairs as in section II.1.

Play Stage. All plays were conducted using Behavery (https://about.behavery.com/), a cloud-based platform 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 once all participants (in that session) had made their choices, the computer would toss a fair coin to determine which payoff scenario applied for that round (the same for all participants; see Table 1 in section II.1). Note that subjects were *not* informed of the scenario, either before or after making their decisions.

(ii) Each subject was prompted to guess what percentage of participants in the same session chose the threshold-relevant action. For instance, (in the case of the letter-outcome pairs of Table 1) the task read as follows: *"... indicate the percentage of the participants in the entire room that you believe have chosen B…"*. Subjects were informed that they would receive an extra payment of \$0.25 if they provided an accurate estimate within ± 1 percentage point of the realized value, and would receive nothing otherwise; they entered their guesses by moving a slider to the desired percentage.[14] Below we denote the *elicited belief report* by μ_i. (In terms of the model, μ_i is the empirical counterpart of the belief entering the threshold comparison: when i reports $\mu_i < d$, we classify the subject as expecting the threshold not to be reached; in that belief region, the payoff structure underlying Proposition 1 implies that B is the unique best response for every type.)

(iii) "Part 2 instructions." Subjects were told they would play an unspecified number of additional rounds of the same decision task. In each round, the scenario was determined jointly by the (session) population's choices and a fresh coin flip. Earnings were disclosed only at the end of the experiment, but between rounds subjects received noisy information about others' revealed preferences. To provide this feedback about a random sample of others in a way that was consistent with the theory and readily interpretable, we introduced a "neighborhood" device: subjects were told that each participant was randomly connected to some others, called neighbors, and that they would privately receive information about the frequency of B choices among these neighbors in the previous round. Crucially, subjects were never informed of the number or identities of their neighbors, nor of the underlying network structure. The network served only to produce noisy but relatable signals, while payoffs always depended on the choices of the entire group.[15]

---

[14] The slider was initially positioned at 50% but could not be left there, so subjects had to indicate a belief about the frequency of the threshold-relevant action (i.e., action B, in the case of Table 1, section II.1). Regarding belief elicitation incentives, Charness, Gneezy, and Rasocha (2021) show that relatively small payments for beliefs do not typically create a meaningful hedging opportunity. Accordingly, in our design, an accurate guess was defined as a guess within one percentage point above or below the realized value.

[15] The study is not designed to analyze learning in relation to network structures, since subjects did not know the specifics of the network. In practice, our software randomly generated a network for each session, assigning each subject 2 or 3 neighbors besides herself: this ensured enough variability in the feedback while maintaining comparability across sessions. Because neighborhoods were formed at random, the design is consistent with the model's assumption that feedback points to the true state with probability $q > 0.5$, on average (section II.4). Hence, low feedback was more likely than not to indicate correctly that the share of unconditional B-choosers (t_b players) was a value π_L below the threshold. By contrast, $high$ feedback was more likely than not to indicate correctly that the share was a value π_H above the threshold. Lastly, because the population is large and anonymous, any single participant's choice has little influence on the population-level

(iv) Before carrying out the choice task in round 2, each subject was given feedback about the percentage of her neighbors that chose the threshold-relevant 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 threshold-relevant 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 threshold-relevant 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.

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

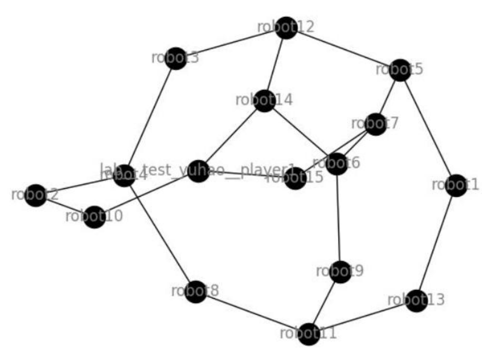

Figure 1 - A random network generated by Behavery (https://about.behavery.com/) as a simulation of the lab environment; each node represents a player. Note: the network structure was used solely as an experimental device to generate noisy, relatable feedback. Subjects never observed the network.

---

frequency that feedback reflects: each signal is thus effectively exogenous to the player's own behavior, consistent with the model.

[16] To comply with Wharton Behavioral Lab guidelines mandating a minimum subject payment, participants were told that if their cumulative payoff across rounds was negative, they would receive only the \$10 participation fee. While such a rule could in principle encourage risk-taking, our study does not aim to measure the extent to which subjects are risk seeking; rather, our focus is on whether subjects commit errors, with and without feedback. In particular, note that the payment mechanism is identical in both the main and control treatments, and hence it *cannot* explain any treatment effects. (For the record, overall only 7% of all the subjects' total payoffs turned out to be negative.)

Lastly, our design includes a (between-subjects) "control treatment" that is the same as the main treatment, except that subjects received *no feedback* about others' choices.

2. ***Hypotheses***

Drawing on the theoretical framework developed earlier, recall that in the equilibrium benchmark players *always* best respond to their beliefs. Section II.6 extends that benchmark by allowing for occasional errors via a stochastic disturbance ε_i: in this specification, deviations occur when the disturbance exceeds the expected utility gap between the best action and its closest alternative (expression (11), section II.6). Table 2 below operationalizes this specification as three testable channels. A treatment effect must operate through *at least one of them.*

Hypothesis	**Intuition**	**Empirical Pattern**
H1: belief dispersion (change in across-round standard deviation of beliefs)	Feedback may settle beliefs, so they fluctuate less across rounds. Formally, a belief that fluctuates more is, through that variation, more likely to fall near the decision cutoff in some rounds, where the expected utility gap $\Delta_i(\bar{\pi}, t)$ between the best action and its nearest alternative is small and choices are most error-prone. Provided beliefs are not centered at the cutoff, less fluctuation makes such rounds rarer.	Best-response failures fall with belief tightening (i.e., when i's across-round belief dispersion falls). Pathway: $belief\ dispersion \downarrow$ $\Rightarrow best\text{-}response\ failures \downarrow$
H2: belief level (change in per-round distance from the threshold)	In a given round, feedback may nudge beliefs away from the decision cutoff, independently of any change in their round-to-round fluctuation. Beliefs that sit farther from the threshold face a wider expected utility gap $\Delta_i(\bar{\pi}, t)$, making choices less error-prone.	Best-response failures decrease as the per-round absolute distance between i's belief and the threshold increases. Pathway: $cutoff\ distance \uparrow$ $\Rightarrow best\text{-}response\ failures \downarrow$
H3: behavioral noise (change in the disturbance)	Any effects of feedback not explained by changes in the expected utility gap $\Delta_i(\bar{\pi}, t)$ via beliefs must operate through the distribution $G_{\eta_{condition}}$ of disturbances ε_i. This reflects reduced random noise in choices, for example due to higher confidence, engagement, or other unobserved psychological factors.	After controlling for belief dispersion and cutoff distance, any remaining treatment effect points to a decrease in behavioral noise not explained by belief structure: this identifies the noise channel residually. Noise can also be measured directly by "excess switching": switching actions more often than would be expected if each choice were an independent random draw from one's long-run action frequencies. Pathway: $excess\ switching \downarrow$ $\Rightarrow best\text{-}response\ failures \downarrow$

Table 2 - Testable mechanisms by which feedback may affect the probability of best responses, derived from the stochastic error specification (section II.6). Each of the effects operates independently of belief accuracy.

Table 2 sums up three channels the model identifies as sources of variability in best-response failures. A treatment effect must operate via at least one of these channels. Feedback can do so by widening the expected utility gap, $\Delta_i(\bar{\pi}, t)$, as in H1 and H2, and/or by altering the distribution of disturbances ε_i (even without any belief change), as in H3. Note that the third mechanism is identified in two complementary ways: residually, as any treatment effect not explained by changes in belief structure; and directly, through the excess-switching measure (switching more often than expected if each choice were an independent random draw from one's long-run action frequencies), as described in IV.3 below. All three mechanisms operate independently of belief accuracy.

IV. Experimental results

1. *Summary statistics and preliminary tests*

	Frequency of choice (%)	Beliefs about population share of B: *Round 1*	Beliefs about population share of B: *Other rounds*	Feedback about neighborhood share of B: *Other rounds*
Main treatment				
A	10.59 (30.79)	54.36 (27.51)	59.04 (23.26)	60.95 (33.37)
B	56.83 (49.55)	59.47 (21.64)	60.62 (23.16)	54.62 (30.66)
C	32.57 (46.88)	56.20 (20.40)	60.23 (21.24)	62.70 (28.39)
Control treatment				
A	9.64 (29.53)	56.33 (18.40)	52.03 (21.50)	
B	64.28 (47.94)	65.51 (20.49)	64.66 (22.32)	
C	26.07 (43.92)	55.04 (20.83)	49.12 (18.15)	

Table 3 - The upper panel reports data from the main treatment (N = 101), and the lower panel from the control treatment (N = 84). Reported values include mean choices, mean beliefs, and—in the main treatment only—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 (all expressed as percentages). Standard deviations are in parentheses. Round-1 beliefs are shown separately, since they were elicited before any feedback was given in the main treatment. No feedback about previous play was provided in round 1 of the main treatment, and none was ever provided in the control treatment.

We begin with summary statistics and overall trends. Sessions were conducted at the Wharton Behavioral Lab: 101 subjects from various academic departments participated in the *main treatment* (the mean age was 24.7 years); 84 subjects took part in the *control treatment* (the mean age was 23.7 years, and other demographics were similar across the two treatments). The control was identical to the main treatment, except that subjects received no feedback about the frequency of B choices among a random sample of other participants. In both treatments, the game was played for ten rounds, with beliefs elicited after each non-final round (before the next play). On average, participants earned \$11.31 in the main treatment and \$6.61 in the control treatment over ten rounds, in addition to the \$10 participation fee.

In the *main treatment*, subjects chose options corresponding to A, B, and C (Table 1, section II.1), respectively, 10.59%, 56.83%, and 32.57% of the time, on average. (Recalling that no one other than the moderately risk-seeking t_a players should choose A, and that their optimal behavior varies with beliefs, the frequency of A choices implies a lower bound of roughly 10% on the share of t_a players, errors aside). Further, a minority of subjects chose the same action across rounds: about 1%, 27%, and 10% always chose A, B, and C, respectively. Errors aside, this means that the share of unconditional B-choosers (i.e., t_b players) in the main treatment was approximately 27% of the population. Turning to the *control treatment*, subjects picked actions A, B, and C, respectively, 9.64%, 64.28%, and 26.07% of the time, on average. Additionally, about 0%, 36%, and 7% of the subjects respectively chose A, B, and C across all rounds, with these figures differing somewhat from the corresponding ones in the main treatment.

To provide context for the observed choice patterns, the middle columns of Table 3 summarize average beliefs about the *population-level* frequency of B choices, held by participants who in a given round chose the option shown on the left-hand side of the table. Also, the last column of Table 3 shows the mean feedback on the *neighborhood-level* frequency of B choices, provided in the main treatment to participants choosing the options on the left-hand side.

Notably, even though average beliefs are relatively similar between conditions, participants in the control treatment (i.e., with no feedback) chose B more frequently than in the main treatment. In the next section we examine how any treatment differences relate to failures to best respond. Before delving into econometric tests of cross-treatment differences, we first verify *within the main treatment* that behavior responds to the feedback signal in the expected direction. Specifically, to connect the data to the model and establish some preliminary patterns of belief updating and behavior under feedback, we check whether receiving *low* versus *high* feedback affects behavior as implied by Proposition 1 (section II.5). That is, in the equilibrium benchmark, players best respond to their beliefs as follows: strongly risk-seeking (t_b) players always choose B (regardless of feedback), whereas risk-averse (t_c) and moderately risk-seeking (t_a) players avoid B when the expected share of B choices exceeds their cutoffs. Intuitively, Proposition 1 implies that feedback below (resp. above) the 40% threshold should increase (resp. reduce) B choices.

Consistent with the model, we refer to feedback as *low* or *high* if (in the previous round) respectively less than or more than 40% of the feedback sample chose B. With this in mind, we present simple pairwise comparisons of choice distributions between low- and high-feedback observations. 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) shows that the risky action B was chosen *more often after low than after high feedback*: 67.73% versus 52.13% of the time, respectively ($z = 2.25$, $p = 0.024$, two-tailed). Put simply, subjects were less likely to choose B after learning that some fellow participants chose B in proportions larger than the threshold. To corroborate these distributional comparisons, individual-level regressions in Appendix B confirm that high feedback significantly reduces the probability of choosing the risky action B; the effect is robust to controls for round-1 priors, round-to-round consistency of feedback, and a time trend, none of which show independent influence. The theoretical rationale is intuitive: high feedback induces an upward revision in the expected share of unconditional B-choosers (i.e., t_b players), pushing risk-averse (t_c) and moderately risk-seeking (t_a) players to move away from B.[17] This diagnostic check allows us to rule out belief-updating frictions as a source of suboptimal play. Note that this concerns how beliefs track the signal's direction within the main treatment; whether feedback shifts the belief-channel variables relative to control is a separate question, taken up below.

2. *Testing belief-based channels of feedback effects on errors (H1 and H2)*

The analysis above shows that, in the main treatment, choices move with the feedback signal (low vs. high) as predicted, indicating that belief updating works as expected. We now

[17] Formally, high (resp. low) feedback implies an upward (resp. downward) updating of the expected share of unconditional B-choosers (i.e., t_b players), so that expression (9) is greater than (10), for any feedback values π_L, π_H with $0 < \pi_L < 0.4 < \pi_H < 1$. Hence, high feedback makes it more likely that risk-averse and moderately risk-seeking players (i.e., t_c and t_a types, respectively) end up with a posterior probability p_f above their respective cutoff points ψ_t. This implies that, on average, high feedback moves some t_a and t_c types toward equilibrium (i) of Proposition 1, while low feedback makes equilibrium (ii) more likely.

turn to tests conditional on stated beliefs: specifically, we examine how often choices deviate from best-response play and whether best-response rates differ between the *main* (feedback) and *control* (no-feedback) treatments. As a reminder, an identifiable best-response failure corresponds to a round in which subject i chooses an action other than B while stating that the overall frequency of B is below the 40% threshold: in any such cases, i commits an error regardless of her preference type.[18]

Empirically, 8.82% of choices belong to the class of errors defined above (averaging per-subject rates across both treatments). Now, to reject the deterministic equilibrium benchmark, it suffices to show that errors occur with strictly positive probability. With nearly 9 percent of choices constituting errors, we can readily conclude that the fully deterministic framework does not fit the data. Indeed, even a conservative Wilcoxon signed-rank test—comparing the median frequency of errors against 5 percent (rather than 0, thus allowing for almost no errors)—is strongly significant (N = 185, z = 2.409, p = 0.016, two-tailed; to satisfy independence of observations, this test is conducted on the sample of per-subject mean choices, where an observation represents the fraction of a participant's choices constituting errors).

We move on to check if there are any between-treatment differences in the distribution of best-response failures. In the main treatment (N = 101), 6.49% of participants' choices consist of errors, compared with 11.64% in the control treatment (N = 84). This pattern is consistent with the prediction that feedback curbs suboptimal behavior. A t-test on per-subject mean choices confirms the difference (N = 185, t = 2.149, p = 0.033, two-tailed). However, these mean comparisons do not account for within-subject variation across rounds. To address this issue, we proceed to test our hypotheses with an econometric analysis of the full sample of observations.

[18] When a subject's elicited belief report $\mu_i < 0.4$, the most natural reading is that she expects the risk-free scenario (x) to obtain, where $B \succ C \succ A$ for every type irrespective of risk attitude. Accordingly, our analysis treats μ_i as a point belief locating the subject's expectation relative to the threshold. For completeness, even under the less natural reading in which μ_i is instead read as the mean of a subjective distribution that places some probability on the threshold binding, the analysis remains valid: in fact, given the monetary payoff ranges in Table 1 in section II.1 (for A, $[0.5, 1]$; for B, $[-1.5, 3]$), it is reasonable to take $\alpha(t)$ and $\beta(t)$ within or near these ranges. Under such conditions, one can verify that if i expects the overall frequency of B choices to be low (i.e., < 0.4), then the best response for i is to choose B, regardless of type.

Non–best-response (error) in round k	[1]	[2]	[3]	[4]
treat=1	-0.640**	-0.706**	-0.555*	-0.656**
	(0.291)	(0.279)	(0.284)	(0.272)
belief SD		3.193*		4.440**
		(1.879)		(2.052)
abs. distance from threshold			-6.535***	-6.672***
			(1.224)	(1.184)
round	-0.005	-0.005	0.001	0.003
	(0.034)	(0.034)	(0.036)	(0.037)
constant	-1.998***	-2.439***	-0.916***	-1.512***
	(0.222)	(0.385)	(0.250)	(0.445)
Pseudo R2	0.013	0.021	0.113	0.125
AIC	986.622	980.940	889.842	879.646
N	1665	1665	1665	1665

Table 4 - Logit coefficients estimating the probability of errors in round k. The dependent variable is an indicator equal to 1 if the chosen action is an identifiable best-response failure (i.e., subject i chose an action other than B while reporting a belief below 40%), and 0 otherwise. In parentheses are robust standard errors clustered at the subject level, accounting for repeated observations within subjects and for the fact that feedback signals, when present, are private and subject-specific. The sample includes 185 participants in the main and control treatments (*, **, and *** respectively indicate $p < 0.10$, $p < 0.05$, and $p < 0.01$; two-tailed tests, z-statistics). Round 10 is excluded because no beliefs were elicited after the final choice task. Note: $treatment$ equals 1 if i is in the main (feedback) condition, and 0 if i is in the control (no-feedback) condition; $belief\ SD$ is the per-subject standard deviation of i's beliefs across rounds; $abs.distance$ is the absolute deviation of i's belief (in round k) from the 40% threshold; $round$ is a linear round-number covariate.

In what follows, we analyze how feedback improves best-response play. To this end, we estimate logit models where the dependent variable is an indicator for an identifiable best-response failure: it equals one when a subject reports a belief below the 40% threshold but chooses an action other than B, and zero otherwise (no identifiable error). The explanatory variables are designed to address the hypothesized mechanisms discussed earlier.

- Treatment: an indicator equal to one for subjects assigned to the feedback condition and zero for subjects assigned to the no-feedback control.

- Belief SD: the per-subject standard deviation of elicited beliefs across rounds; higher values correspond to greater belief dispersion. This variable addresses the belief-tightening hypothesis H1 (*belief dispersion channel*).[19]
- Distance from threshold: the absolute distance between the subject's belief in a given round and the 40% threshold; lower values indicate that the belief in that round is closer to the threshold, where near-indifference might make choice less clear-cut. This variable directly addresses the cutoff-distance hypothesis H2 (*belief level channel*).
- Round: a linear control for round number, to capture any dynamic patterns over time.

We begin by estimating the baseline effect of feedback on best-response failures, and then progressively introduce regressors linked to our hypothesized channels. Throughout, we cluster standard errors at the subject level.

Treatment effect. In model [1], our baseline logit specification includes the treatment dummy and round controls. Specifically, feedback significantly reduces the likelihood of a best-response failure (coef. = −0.640, $p < 0.05$), which confirms the raw pattern that subjects in the feedback condition commit fewer errors. Note that the analysis above retains all observations rather than conditioning on below-threshold beliefs, which feedback itself can shift. In an intention-to-treat spirit, model [1] therefore estimates the overall effect of feedback on identifiable errors while avoiding selection on a post-randomization variable. This is in fact conservative: estimating the same logit model on the restricted sample of observations with below-threshold beliefs (not shown in Table 4) gives a treatment coefficient larger in absolute magnitude, at −0.777 ($p < 0.05$); hence, the full-sample estimate in model [1] of Table 4 actually understates the treatment effect.[20]

[19] In our model, best responses and the expected utility gap Δ_i are defined relative to a single payoff-relevant belief about whether the share of B choices lies below or above the threshold. Empirically, the elicited point belief μ_i locates the subject's expectation about the population share of B choices relative to the threshold. Accordingly, the *round-to-round variation* in belief reports is the relevant notion of belief dispersion here. Alternatively, one might think that within-round belief uncertainty could be elicited directly. Yet that would presuppose a different model, one where each subject holds a full subjective distribution over others' play: a strong assumption that, especially for a repeated game, is psychologically demanding and would risk imposing structure that isn't necessarily there. Identifying such a distribution reliably across ten rounds would also be cumbersome. In brief, the point belief is the object our predictive theory requires for the test of belief-action coherence, not a proxy for it.

[20] As a further robustness check, we address the possibility a subject with a stated belief *just under the threshold* might in principle be pivotal in a small session, so that avoiding B to preserve the risk-free scenario could itself be a best response (a marginal concern given session sizes of about 17). Restricting the sample further to beliefs below 35% addresses this near-threshold argument: despite the smaller number of such observations, relative to model [1] the treatment effect is even larger (coef. = −0.860, $p < 0.05$, $N = 272$ obs. from 100 qualifying subjects, with standard errors clustered by subject).

Belief-tightening hypothesis H1. In model [2], we add the subject-level dispersion of beliefs across rounds to test H1. Dispersion is positively associated with errors (coef. = 3.193, $p < 0.10$), which implies that subjects whose beliefs vary more across rounds are less likely to best respond. Notably, the treatment effect remains negative and significant (coef. = −0.706, $p < 0.05$). This indicates that while belief dispersion helps explain some heterogeneity in errors, the treatment effect cannot be attributed solely to a "volatility" reduction mechanism.

Cutoff-distance hypothesis H2. In model [3], we add the absolute distance between the per-round belief and the 40% threshold. Distance is strongly and negatively associated with errors (coef. = −6.535, $p < 0.01$), confirming that beliefs lying farther from the decision cutoff reduce the chance of near-indifference and hence lower error rates. The treatment effect, though somewhat attenuated, remains negative and significant at the 10% level (coef. = −0.555, $p < 0.10$). Thus, the treatment effect cannot be fully explained by a shift in belief level away from the threshold.

Behavioral noise hypothesis H3. Lastly, in model [4] we include both belief dispersion and cutoff distance simultaneously. Both belief dispersion (coef. = +4.440, $p < 0.05$) and cutoff distance (coef. = −6.672, $p < 0.01$) remain highly significant predictors of error. Still, the treatment effect persists (coef. = −0.656, $p < 0.05$), which suggests that feedback improves best-response play through an additional channel beyond the belief-based mechanisms in H1 and H2. We interpret this residual effect as preliminary evidence of reduced noise due to unobserved psychological factors (e.g., higher confidence, decisiveness, or engagement), consistent with H3.

Interpretation and further tests. Models [2]–[4] of Table 4 establish that our belief variables predict best-response deviations (i.e., errors). To assess whether a channel explains the treatment effect, however, we must ask not only whether the relevant channel variable predicts errors, but also whether feedback affects that variable. For feedback to improve best responses *by way of* the belief variables, it would first have to affect them. Crucially, auxiliary regressions (not shown in Table 4) indicate that feedback has no significant effect on either belief dispersion or cutoff distance: a subject-level OLS of *belief SD* with treatment as the sole regressor shows no detectable effect (coef. = 0.018, $p > 0.10$, $N = 185$). Similarly, an OLS of *distance from threshold* with treatment as the sole regressor shows no effect (coef. = 0.011, $p > 0.10$, $N = 185$).[21] So although both belief variables

[21] For clarity, these auxiliary regressions are run at the subject level (185 observations), since *belief SD* is constant within subjects: note that using round-level data would not add information and could inflate significance. (This caveat does not

predict errors, feedback does not move them. This does not mean that subjects fail to respond to feedback; rather, it means that *the measured treatment difference in errors is not explained by changes in belief dispersion or cutoff distance*. Within the three candidate channels, we therefore find no evidence that feedback improves best responses via belief dispersion or cutoff distance.

In summary, our theoretical framework predicts that best-response failures arise when stochastic disturbances outweigh the expected utility gap between the best action and its closest alternative (section II.6). Feedback can reduce such failures in two broad ways: it may widen the expected utility gap by changing the structure of beliefs, and/or it may reduce the effective magnitude of disturbances through psychological factors beyond belief structure. (Either way, the probability that noise undermines best-response play is lowered, irrespective of belief accuracy.) More precisely, we posit three channels through which feedback may operate within this framework: (H1) it may reduce belief dispersion; (H2) it may shift the per-round belief level away from the decision cutoff; and/or (H3) it may reduce the scale of noise. The econometric analysis above provides residual evidence for H3. We investigate more below.

3. *Direct evidence on the behavioral noise channel (H3)*

We now present tests that more directly evaluate the noise channel (H3). We use action switching as a measure of erratic play: switching is defined as changing one's action relative to the previous round. Switching across repetitions of the same decision problem is commonly used as evidence of behavioral noise in the stochastic-choice literature (e.g., Hey and Orme, 1994; Ballinger and Wilcox, 1997). Our measure refines this approach by comparing observed switching to the switching rate implied by independent draws from the subject's own long-run action frequencies. This benchmark matters because randomization per se need not reflect noise: subjects may deliberately mix across actions. We therefore define "excess switching" as the *difference between a subject's observed switching rate and the rate predicted by independent randomization based on that subject's empirical action shares*.

Implementation. We construct the excess-switching measure in three steps.

First, for each subject i we calculate their empirical action shares over rounds 1–10, denoted $\hat{p}_{iA}, \hat{p}_{iB}, \hat{p}_{iC}$. These correspond to i's fractions of plays of actions A, B, C.

---

affect the round-level regressions in Table 4, where the dependent variable varies across rounds and clustering at the subject level addresses within-subject dependence.)

Second, under the null of independent randomization with these shares, i's (subject-specific) *expected per-round switching rate* is

$$\Pr(switch_i) = 1 - \left(\hat{p}_{iA}^2 + \hat{p}_{iB}^2 + \hat{p}_{iC}^2\right),$$

which we refer to as the i.i.d. benchmark.

Third, at the round level we define an indicator $switch_{ik} = 1$ if the subject's action in round k differs from the previous round (0 otherwise), and compute i's *excess switching at* k as

$$excess\_switch_{ik} = switch_{ik} - \Pr(switch_i).$$

Positive values indicate switching more often than the i.i.d. benchmark predicts (erratic play), while negative values indicate greater persistence than predicted.[22]

Results. A linear regression of excess switching on the treatment dummy and round (model [1] in Table 5, with clustered standard errors at the subject level) shows that the main treatment significantly reduces excess switching (coef. = −0.049, $p < 0.01$), while switching also declines with round number (coef. = −0.014, $p < 0.01$). Hence, feedback curbs erratic noise-like switching between actions, corroborating the behavioral noise channel (H3).

To test whether the baseline result might simply reflect subjects who happened to receive stable signals, in model [2] of Table 5 we include $info\_consistency$ as an additional regressor. This variable tracks whether the feedback signal stayed the same from one round to the next (it is always equal to one in the control group, where subjects consistently received no feedback). Adding this control shows that the treatment effect on excess switching remains negative and statistically significant, while $info\_consistency$ itself is small and not significant. This test confirms that the reduction in erratic play persists even when considering subjects who received inconsistent feedback signals.

[22] For example, for a subject who plays A and B with equal probability, excess switching would be zero if they switched on exactly half of the transitions $K-1$ (with K being the total number of rounds); by contrast, a perfectly alternating sequence such as $ABABAB$ ... would yield the largest positive excess switching. Note that the benchmark used in the body $(1 - \sum_{s_i} \hat{p}_{s_i}^2)$ is slightly downward biased in finite samples as an estimate of the expected switching rate. For robustness, we also used an exact permutation benchmark, which compares the observed sequence with random reorderings of the same realized choices. This benchmark equals $\frac{K}{K-1}(1 - \sum_{s_i} \hat{p}_{s_i}^2)$ and removes any finite-sample bias, subject by subject: under this specification, the treatment effect remains unchanged from the one reported in Table 5 (coef. = −0.049, $p < 0.01$). Finally note that, in principle, switching could reflect deliberate alternation rather than noise; in our setting, however, subjects have no strategic incentive to alternate across actions (unpredictability has no value, since subjects face a large anonymous population and do not observe outcomes). Moreover, in the belief region used to identify best-response deviations, one action is uniquely optimal. We thus interpret excess switching as a proxy for erratic choice, providing a direct behavioral measure of the noise channel beyond the residual evidence from Table 4. (See also Alós-Ferrer et al., 2026, who find deliberate randomization is infrequent and does not account for choice inconsistency in repeated settings.)

Taken together with the results in IV.2, this analysis points to behavioral noise as a channel that carries at least part of the treatment effect. The belief-based variables predict errors but do not account for the treatment difference (as shown in IV.2 under *"Interpretation and further tests"*); by contrast, feedback significantly reduces excess switching, as seen in Table 5.[23]

Excess switching in round k	[1]	[2]
treat=1	-0.049*** (0.015)	-0.049*** (0.016)
round	-0.014*** (0.004)	-0.014*** (0.004)
info_consistency		-0.001 (0.032)
constant	0.138*** (0.027)	0.140*** (0.042)
R2	0.015	0.015
AIC	1276.579	1278.576
N	1665	1665

Table 5 - Ordinary Least Squares (OLS) coefficients. The dependent variable is subject i's excess switching in round k, defined as the difference between i's observed switching indicator at k and the expected switching rate under i.i.d. randomization with subject-specific action frequencies. In parentheses are robust standard errors clustered at the subject level, accounting for repeated observations within subjects and for the fact that feedback signals, when present, are private and subject-specific. The sample includes 185 participants in the main and control treatments (*, **, and *** respectively indicate $p < 0.10$, $p < 0.05$, and $p < 0.01$; two-tailed tests, t-statistics). The first round is excluded since excess switching is undefined in round 1. Note: $treatment$ equals 1 in the feedback condition, and 0 in the control condition; $round$ is a linear round-number covariate; $info\_consistency$ equals 1 when consecutive feedback signals do not differ in their location relative to the threshold (and for all rounds in the control group, which consistently received no feedback), and 0 otherwise.

[23] An alternative reading is that feedback raises coherence not by reducing noise in actions, but by serving as a common anchor for ex-post rationalization. Yet the elicitation order—choice, then stated belief—is identical in both treatments, so it cannot by itself generate a treatment difference; any such difference would require the mechanism to operate differentially through feedback. The data weigh against that. First, the treatment's effect on excess switching persists once we control for the consistency of a subject's feedback stream (Table 5, model [2]); the reduction is therefore not confined to subjects who happened to receive stable signals to track. Second, *feedback has no detectable effect on either of the belief-based variables* (belief dispersion and cutoff distance, as concluded in IV.2), leaving little room for a belief-side anchor to drive the alignment. Hence, the common-anchor account finds no support in our data.

4. ***Linking errors to excess switching: final tests***

The evidence in Table 4 showed that belief dispersion (H1) and cutoff distance (H2) are strong predictors of best-response failures, though auxiliary regressions indicated that feedback has no significant effect on either belief dispersion or cutoff distance, providing indirect support for the noise channel (H3). Table 5 then provided more direct evidence on H3, showing that feedback curbs action switching beyond what would be expected under deliberate mixing. To corroborate this conclusion, we now examine per-subject mean data to link these pieces of evidence explicitly.

Mean error (non–best-response)	[1]	[2]	[3]
treat=1	-0.802*** (0.303)	-0.588* (0.306)	-0.612** (0.278)
mean excess switching		4.285*** (1.383)	2.301* (1.350)
belief SD			9.149*** (2.533)
mean abs. distance from threshold			-11.261*** (2.256)
constant	-1.959*** (0.213)	-2.293*** (0.263)	-1.353*** (0.454)
Pseudo R2	0.021	0.054	0.145
AIC	112.325	110.673	104.512
N	185	185	185

Table 6 - Fractional logit coefficients estimating the determinants of mean best-response failure rates. The dependent variable is subject i's mean error rate, defined as the fraction of rounds in which i chose an action other than B while reporting a belief below 40%. In parentheses are robust standard errors (*, **, and *** respectively indicate $p < 0.10$, $p < 0.05$, and $p < 0.01$; two-tailed tests, z-statistics). Note: $treatment$ equals 1 in the feedback condition, and 0 in the control condition; $mean\ excess\ switching$ is the subject's average deviation between observed switching and the expected switching rate implied by i.i.d. randomization with subject-specific action frequencies; $belief\ SD$ is the subject-level standard deviation of beliefs across rounds; $mean\ abs.distance$ is the subject's average absolute deviation of beliefs from the 40% threshold. All variables are computed per subject over rounds 2–9, since excess switching is undefined in round 1 and beliefs were not elicited in the final round.

Subject-level tests. The regressions in Tables 4 and 5 exploit round-level variation; yet because best-response failures are well defined only when beliefs fall below the 40% threshold, the number of error observations per subject is constrained by the stated beliefs and uneven across rounds. To capture overall individual tendencies, we now collapse the data to the subject level, computing each subject's mean excess switching and mean error rate (i.e., the rate of identifiable best-response failures) across rounds. This provides a final test of whether erratic switching and best-response failures are systematically linked, and whether feedback operates through that channel. More specifically, since average error rates are proportions bounded between 0 and 1, in what follows we estimate fractional logit models (which are designed for fractional outcomes and avoid the limitations of treating proportions as linear): the dependent variable is each subject's average error rate, while regressors include treatment, mean excess switching, and the belief-based variables identified in H1 and H2. Table 6 reports three specifications, summarized below.

- **Treatment and best-response failures.** Model [1] of Table 6 reports the treatment difference in mean best-response failure rates. The treatment coefficient is negative and statistically significant (coef. = −0.802, $p < 0.01$), indicating that participants in the feedback treatment commit fewer best-response failures than those in the control treatment. This provides the per-subject counterpart to the choice-round evidence above (see model [1] of Table 4): feedback is associated with a lower error rate before accounting for excess switching or belief-based predictors.
- **The role of excess switching.** Model [2] adds mean excess switching to the treatment-only specification. The coefficient on excess switching is positive and statistically significant (coef. = 4.285, $p < 0.01$), indicating that subjects who switch more frequently (i.e., more often than is predicted by their own long-run action frequencies) also commit more best-response failures. Adding excess switching leaves the treatment coefficient negative but smaller in absolute magnitude, from −0.802 in model [1] to −0.588 in model [2], where it becomes marginally significant. This pattern is consistent with part of the feedback effect running via reduced switching (the marker for the behavioral-noise channel).
- **Excess switching plus belief controls.** Model [3] adds the two belief-based predictors: belief dispersion and mean distance from the threshold. Both enter with the predicted signs: belief dispersion is positively associated with best-response failures (coef. = 9.149,

$p < 0.01$), while mean cutoff distance is negatively associated with failures (coef. = −11.261, $p < 0.01$). Notably, adding these variables substantially improves model fit (AIC falls from 110.7 in model [2] to 104.5 in model [3]). Also, excess switching remains positive and becomes marginally significant, suggesting that its association with best-response failures partly overlaps with the belief-based predictors. At the same time, the treatment coefficient remains close to its value in model [2] (−0.612 versus −0.588), confirming that belief dispersion and cutoff distance explain heterogeneity in failures but do not account for the treatment difference in our data. This pattern points instead to the noise channel.

Overall interpretation. Taken together, these results sharpen our interpretation of the earlier findings. Belief dispersion and cutoff distance are strong predictors of best-response failures, but they do not account for the feedback effect: adding them improves model fit while leaving the treatment coefficient essentially unchanged (Table 6, model [2] → model [3]). By contrast, excess switching is both affected by feedback (as shown in Table 5) and predicts best-response failures (Table 6, models [2]–[3]). Hence, among the three candidate channels, only excess switching carries part of the treatment difference. In short, the analysis supports the interpretation that feedback reduces errors at least in part by curbing erratic noise-like choice, as opposed to operating primarily through changes in belief dispersion or belief location.

V. Concluding remarks

We asked why individuals often fail to act on their own stated beliefs, and whether feedback about others' behavior can close that gap. To address this question, our theoretical framework develops a novel game in which one action is uniquely optimal whenever beliefs about others fall below a threshold. This feature of the game makes the best response type-invariant within a known belief region, so that departures from it are identifiable as errors rather than driven by risk attitudes or a preference for the safe option. Our experiment can then separate mechanisms that are usually confounded: feedback may improve play by reshaping beliefs, or it may improve play by making the mapping from beliefs to actions less erratic.

The data give a sharp answer. Best-response failures fall from 11.64% without feedback to 6.49% with feedback. Belief structure strongly predicts errors: greater belief dispersion raises best-response failures, while distance from the threshold lowers them. But feedback does not significantly move either belief dispersion or cutoff distance. What feedback does move instead is erratic play: in the feedback treatment, subjects switch actions less than their own long-run choice frequencies would imply, and this reduction in excess switching is the only measured channel that carries part of the treatment difference. Feedback thus enhances belief-action coherence not primarily by reshaping beliefs, but by tempering noise likely arising from psychological factors such as hesitation or disengagement.

This speaks to an open question in the literature on feedback and strategic play: whether information improves decisions by changing beliefs or by helping people act on the beliefs they already hold. By conditioning on stated beliefs (regardless of their accuracy), our setting isolates the latter route and shows that it accounts for part of the feedback effect. While this evidence is not a complete decomposition of the treatment effect, the part it identifies is robust. The noise channel is the only channel that is *both* affected by feedback *and* predicts errors; the belief channels fail the first test (they predict errors but feedback does not move them). Moreover, the noise channel's role holds whether measured residually or through the direct excess-switching marker, so the result does not turn on a single specification or measure.

Finally, the game relates to a broad class of settings in which a safe exit option coexists with a risky collective endeavor governed by a participation threshold. Vaccination is one instance; similar incentives arise whenever agents can take a fixed outside option or instead join a risky initiative whose payoff depends on how many others join. For example, a firm may weigh a known return against a risky investment in a joint venture whose payoff depends on participation levels and on a possible shock. The present paper theoretically characterizes the equilibrium predictions in this class of interactions, and our empirical results point to a practical lever for improving decision quality: even when feedback is not very informative about equilibrium fundamentals, it can make choices more consistent with stated beliefs by damping erratic behavior.

APPENDIX A

Proposition 1 (Bayesian Nash equilibria in pure actions)

Proof. In every equilibrium, all t_b players choose B. This is because B is strictly optimal for type t_b both when the threshold is not reached (where the monetary payoff from B is highest) and when the threshold is reached, where $B \succ_{t_b} A \succ_{t_b} C$.

Consider now a player i of type $t \in \{t_a, t_c\}$. Let $p_f = \Pr(H \mid f)$ denote the posterior probability of the high state after feedback f. Since the game is interpreted as a large-population interaction, the parameters π_H and π_L can be read as the corresponding aggregate shares of t_b players in the high and low states, respectively. Thus, since $\pi_H > d$, the threshold is reached in the high state whenever all t_b players choose B. Similarly, in the low state, if only t_b players choose B, the share of B choices is $\pi_L < d$, so the threshold is not reached.

Suppose first that player i's choice of B does not make the threshold bind in the low state. Then the expected utilities from A, B, and C are

$$U_i(A) = \frac{1}{2}(1 - p_f) + \alpha_t p_f, \qquad U_i(B) = 3(1 - p_f) + \beta_t p_f, \qquad U_i(C) = \frac{3}{4}.$$

Throughout, suppose ties are broken in favor of B. For type t_a, action B is weakly optimal if and only if it weakly dominates both A and C. The comparison with A gives

$$3(1 - p_f) + \beta_{t_a} p_f \geq \frac{1}{2}(1 - p_f) + \alpha_{t_a} p_f,$$

which is equivalent to

$$p_f \leq \frac{5}{2\alpha_{t_a} - 2\beta_{t_a} + 5}.$$

The comparison with C gives

$$3(1 - p_f) + \beta_{t_a} p_f \geq \frac{3}{4},$$

which is equivalent to

$$p_f \leq \frac{9}{12 - 4\beta_{t_a}}.$$

Therefore B is weakly optimal for type t_a whenever

$$p_f \leq \psi_{t_a} \equiv \min\left\{\frac{5}{2\alpha_{t_a} - 2\beta_{t_a} + 5}, \frac{9}{12 - 4\beta_{t_a}}\right\}.$$

For type t_c, action C strictly dominates A among the non-B actions: in the low state, C pays $3/4$ while A pays $1/2$, and in the high state $C \succ_{t_c} A$. Hence the relevant comparison is between B and C:

$$3(1-p_f)+\beta_{t_c}p_f \geq \frac{3}{4},$$

or equivalently

$$p_f \leq \psi_{t_c} \equiv \frac{9}{12-4\beta_{t_c}}.$$

Thus, for each non-t_b type t, the cutoff condition fails exactly when $p_f \leq \psi_t$. Define the set of eligible non-t_b players by

$$G_f = \{i \in N : t(i) \in \{t_a, t_c\} \text{ and } p_f \leq \psi_{t(i)}\}.$$

These are the non-t_b players for whom B is weakly optimal provided that their choosing B does not make the low-state threshold bind. Also let

$$r = d - \pi_L$$

be the residual room below the threshold in the low state. For any subset $O \subseteq G_f$, let $\omega_L(O)$ denote the *share of the total population, in the low state, accounted for by players in* O. If exactly the players in O, in addition to all t_b players, choose B, then the low-state share of B choices is

$$\pi_L + \omega_L(O).$$

Hence the threshold is not reached in the low state if and only if

$$\omega_L(O) < r.$$

We now distinguish two cases.

First, suppose $\omega_L(G_f) < r$. Then all players in G_f can choose B without reaching the threshold in the low state. Each such player is best responding by the cutoff condition. Every non-t_b player outside G_f takes her best non-B action: each t_c player chooses C, while each t_a player chooses A if

$$p_f \geq \frac{1}{4\alpha_{t_a}-2},$$

and C otherwise. This fallback rule for t_a follows from comparing A and C:

$$\frac{1}{2}(1-p_f)+\alpha_{t_a}p_f \geq \frac{3}{4} \quad \Leftrightarrow \quad p_f \geq \frac{1}{4\alpha_{t_a}-2}.$$

No non-t_b player outside G_f has a profitable deviation to B. If the deviation does not make the low-state threshold bind, then B is not optimal because $p_f > \psi_{t(i)}$. If instead the deviation makes the low-state threshold bind, then the threshold is reached in both states and the deviation yields $\beta_{t(i)}$. For both t_a and t_c, $\beta_{t(i)} < 3/4$, while the best non-B action yields at least $3/4$. Hence such a deviation is not profitable.

Second, suppose $\omega_L(G_f) \geq r$. Let $O \subseteq G_f$ be any maximal feasible subset, meaning that

$$\omega_L(O) < r,$$

and, for every player $i \in G_f$ with $i \notin O$,

$$\omega_L(O \cup \{i\}) \geq r.$$

Players in O choose B. Since $\omega_L(O) < r$, the threshold is not reached in the low state, so each player in O is best responding by the cutoff condition. All remaining non-t_b players take their best non-B action: each t_c player chooses C, while each t_a player chooses A if

$$p_f \geq \frac{1}{4\alpha_{t_a} - 2},$$

and C otherwise.

Consider a player $i \in G_f$ with $i \notin O$. If i deviates to B, maximality of O implies that the low-state threshold is reached. The threshold is already reached in the high state, so the deviation yields $\beta_{t(i)}$ for sure. Since $\beta_{t(i)} < 3/4$ for $t(i) \in \{t_a, t_c\}$, while the best non-B action yields at least $3/4$, the deviation is not profitable. Now consider any non-t_b player outside G_f. If a deviation to B does not make the low-state threshold bind, then B is not optimal because $p_f > \psi_{t(i)}$. If the deviation does make the low-state threshold bind, then the deviation yields $\beta_{t(i)} < 3/4$, again below the payoff from the best non-B action. Hence no such player has a profitable deviation.

Thus, in every pure-action equilibrium, all t_b players choose B. If both cutoff conditions are satisfied, no non-t_b player chooses B. If at least one cutoff condition is not satisfied, B is chosen either by all eligible players in G_f, when $\omega_L(G_f) < r$, or by all players in a maximal feasible subset $O \subseteq G_f$, when $\omega_L(G_f) \geq r$. All remaining non-t_b players take their best non-B action. ■

APPENDIX B

Additional data analysis

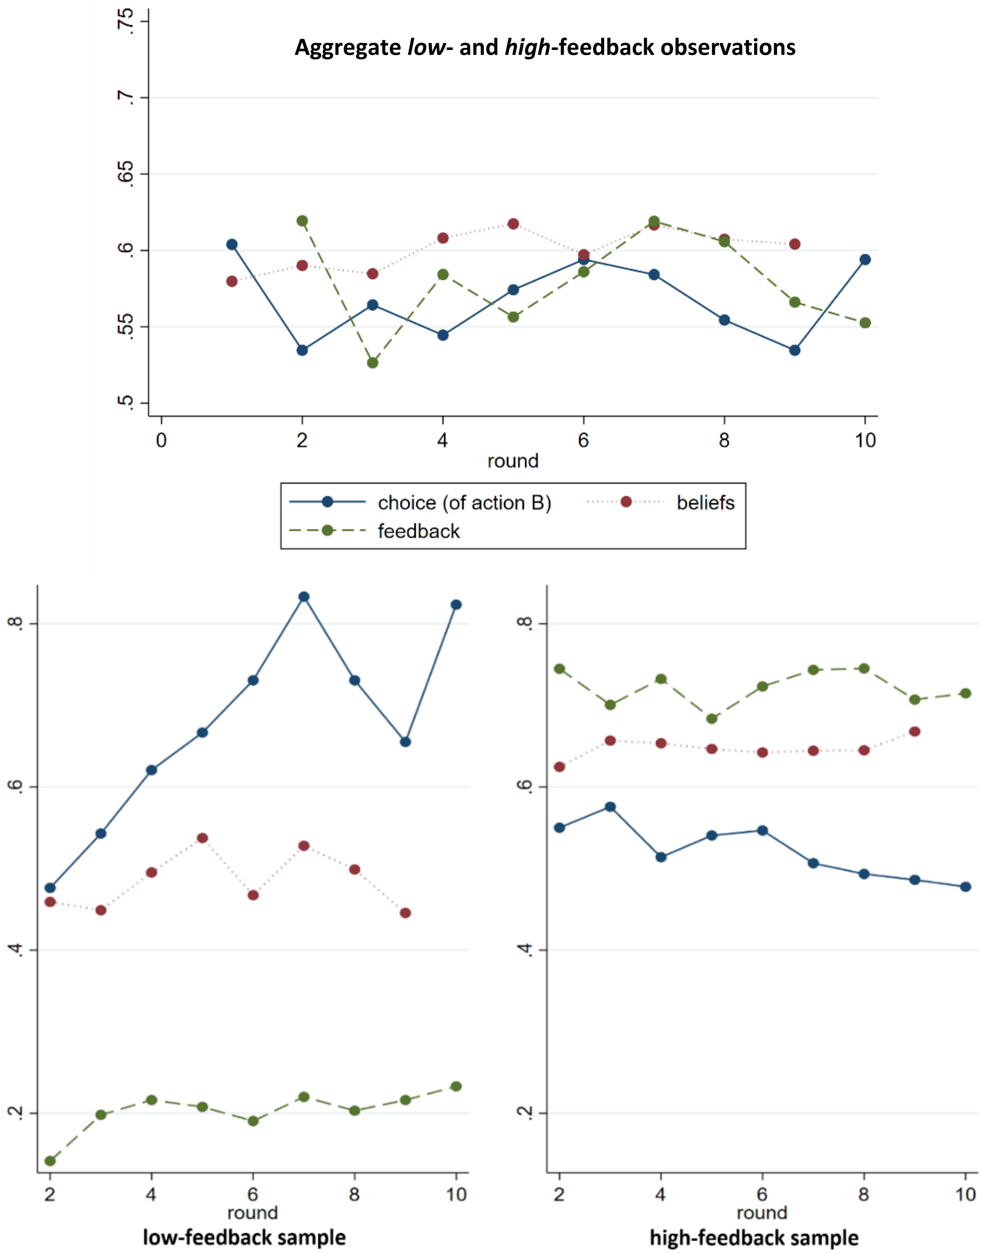

Figure 1B - Main treatment. The upper panel shows line graphs depicting mean values (by round, averaged across all sessions) for: frequency of *B* choices, beliefs about population-level *B* choices, and feedback about neighborhood-level *B* choices. The lower panel breaks down *low*- vs *high*-feedback observations (i.e., all subject–round pairs in which the subject privately observed feedback below vs above the 40% threshold). 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 section III.1 in the main text.

Round-by-round trends: main and control treatments. To provide a more granular view of the *main treatment*, Figure 1B above plots line graphs of round-by-round mean values for: (i) the frequency of B choices; (ii) beliefs about population-level B choices; (iii) feedback about neighborhood-level B choices. The lower panel of Figure 1B contrasts *low*- versus *high*-feedback observations (i.e., rounds in which a subject privately observed feedback below vs. above the 40% threshold) and shows clear behavioral differences, which we analyze econometrically later.

For completeness, Figure 2B below presents analogous line plots for the *control treatment* (no feedback). A quick visual comparison of Figure 2B with the upper panel of Figure 1B shows that, even though average beliefs are relatively similar between conditions, participants in the control treatment chose B more frequently than in the main treatment overall. This contrast is especially striking given that, as shown in the main-text analysis, the control treatment exhibits more best-response failures (i.e., avoiding B, having stated a low belief).

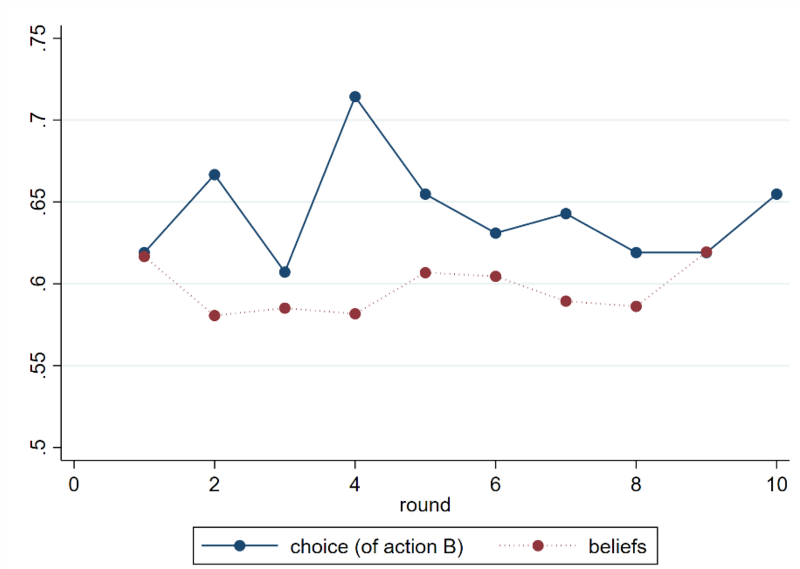

Figure 2B - Control treatment. Line graphs depicting mean values (by round, averaged across all sessions) for frequency of B choices and for beliefs about population-level B choices.

Main treatment (low- vs high-feedback). Having outlined overall trends across treatments, we now turn to the main treatment, which provides further perspective on how feedback shaped beliefs and choices. As noted in the main text, 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) shows that B was chosen more often after low than after high feedback: 67.73% versus 52.13%

of the time, respectively (z = 2.25, p = 0.024, two-tailed). For completeness, here we report the same test for the other actions. For action A, the test shows no meaningful differences in the proportions of choices across samples, which were respectively 10.36% and 10.64%. Lastly, the same test shows that the riskless action C was chosen less often after low than after high feedback: 21.91% versus 37.23% of the time, respectively (z = −2.25, p = 0.024, two-tailed). Together, these patterns provide evidence that exposing subjects to high feedback causes them to shift from a very risky action (B) to a riskless action (C).

choice of action B in round k	[1]	[2]
feedback=high	-0.656*** (0.237)	-0.648*** (0.234)
first-round belief=high		0.046 (0.355)
info_consistency		-0.047 (0.217)
round		0.009 (0.023)
constant	0.741*** (0.225)	0.682* (0.385)
Pseudo R2	0.014	0.015
AIC	1230.695	1236.409
N	909	909

Table 1B - Logit coefficients estimating a subject's choice of B in round k of the main treatment (i.e., feedback condition). In parentheses are robust standard errors clustered on 101 subjects, accounting for repeated observations within subjects and for the fact that feedback signals are private and subject-specific (*, **, and *** respectively indicate $p < 0.10$, $p < 0.05$, and $p < 0.01$; two-tailed tests, z-statistics). The analysis uses all the observations except for round $k = 1$, for which there was no $k - 1$ feedback. Note: $feedback$ equals 1 if the share of neighbors' previous B choices (at $k - 1$) is above the 40% threshold, and 0 otherwise; $first\text{-}round\ belief$ equals 1 if the subject's belief at $k = 1$ (prior to receiving any feedback) is above the threshold, and 0 otherwise; $info\_consistency$ equals 1 when consecutive feedback signals do not differ in their location relative to the threshold, and 0 otherwise; $round$ is a linear round-number covariate.

The tests of proportions discussed above capture average trends across rounds but do not account for individual-specific variation. For robustness, we now conduct a regression analysis of subjects' choice of action B (versus the alternatives combined, i.e., "not B"). As a benchmark, we first report a logit model (see [1] in Table 1B) consisting of the low/high feedback indicator as the sole predictor: unsurprisingly, model [1] corroborates the previous tests, showing a significant negative effect of the high feedback on one's choice of the risky action B (robust standard errors are clustered at the subject level as throughout). Next, to control for any individual-specific differences across rounds, model [2] in Table 1B 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; and a time (i.e., round k) variable. Model [2] confirms the significant impact of the low/high feedback indicator, with no significant effects for the remaining predictors.

In sum, Table 1B confirms that feedback significantly shifts the probability of choosing action B in the expected direction (even after controlling for priors, feedback consistency, and time). This supports the main-text interpretation: behavior moves in the direction predicted by Proposition 1, making belief-updating frictions unlikely to be a key source of suboptimal play. Consistent with this interpretation, an OLS regression of stated beliefs on the continuous private signal—within the main treatment—shows that beliefs move in the direction of the signal (coef. = 0.317, $p < 0.001$). To be clear, this result asks whether, among subjects who receive feedback, stated beliefs move with the signal they receive; they do. This within-treatment responsiveness should not be confused with the main-text mechanism result: the main text asks a different question, namely whether assignment to the feedback treatment changes the belief-channel variables (i.e., belief dispersion and cutoff distance), relative to the control treatment; it does not. This reinforces the main conclusion: *within the feedback treatment*, signals move behavior and beliefs in the expected direction; however, the fall in best-response failures *relative to the control treatment* is linked not to the measured belief channels, but to a reduction in erratic noise-like choice.

Detailed breakdown of private signals

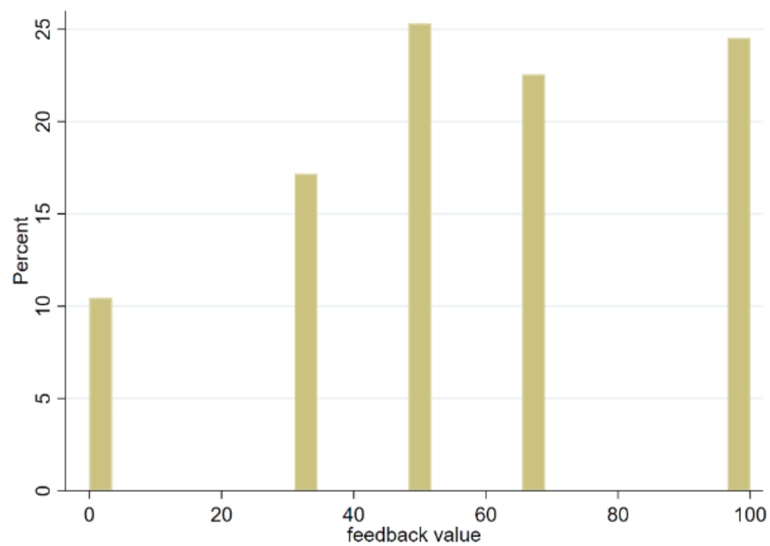

Figure 3B - Histogram of private signals. Bars show the distribution of low (i.e., below 40%) and high (above 40%) feedback values across sessions, computed over all rounds with feedback (rounds 2–10). The figure's specific support arises from the randomly generated neighborhoods (unbeknown to subjects, session networks were randomly generated so that each subject had 2 or 3 neighbors; see section III.1 for details on the experimental design). In practice, 5.94% and 58.42% of subjects received exactly one and two distinct feedback values across rounds, respectively. For illustration, the leftmost bar corresponds to 0%: no neighbor chose B in the previous round.

APPENDIX C

Experimental instructions and screenshots

NOTE: As discussed in section III.1 of the main text, *letter-outcome pairs* (e.g., whether B is associated with the threshold-relevant 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 threshold-relevant option is associated with action A: accordingly, in the below screenshots, 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.

%

0% 100%

Ending Round 1

Waiting for other participants...

Please Wait ...

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.

%

0% 100%

Ending Round 2

Waiting for other participants...

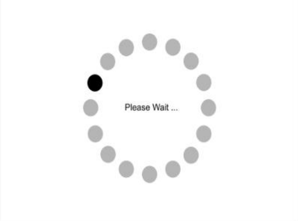

0.0% of your neighbors chose A in the previous round.

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

Alempaki, Despoina, Andrew M. Colman, Felix Kölle, Graham Loomes, and Briony D. Pulford. 2022. "Investigating the Failure to Best Respond in Experimental Games" *Experimental Economics*, 25(2): 656-679.

Alós-Ferrer, Carlos, Johannes T. Buckenmaier, Michele Garagnani, and Aldo Rustichini. 2026. "Why Is Choice Stochastic? Deliberate Randomization vs. Strength of Preference" *Management Science*, Articles in Advance, pp. 1-25.

Alós-Ferrer, Carlos and Michele Garagnani. 2022. "Strength of Preference and Decisions Under Risk" *Journal of Risk and Uncertainty*, 64(3): 309-329.

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.

Charness, Gary, Uri Gneezy, and Vlastimil Rasocha. 2021. "Experimental Methods: Eliciting Beliefs" *Journal of Economic Behavior & Organization*, 189: 234-256.

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.

Compte, Olivier and Andrew Postlewaite. 2004. "Confidence-Enhanced Performance" *American Economic Review*, 94(5): 1536-1557.

Danz, David N., Dietmar Fehr, and Dorothea Kübler. 2012. "Information and Beliefs in a Repeated Normal-Form Game" *Experimental Economics*, 15(4): 622-640.

Fechner, Gustav Theodor. 1860/1912. "Elements of Psychophysics". In B. Rand (Ed.), *The Classical Psychologists: Selections Illustrating Psychology from Anaxagoras to Wundt* (pp. 562-572). Boston, MA: Houghton Mifflin. Original work published 1860.

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.

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.

Hey, John D. and Chris Orme. 1994. "Investigating Generalizations of Expected Utility Theory Using Experimental Data" *Econometrica*, 62(6): 1291-1326.

Hyndman, Kyle, Erkut Y. Ozbay, Andrew Schotter, and Wolf Ze'ev Ehrblatt. 2012. "Convergence: An Experimental Study of Teaching and Learning in Repeated Games" *Journal of the European Economic Association*, 10(3): 573-604.

Ivanov, Asen. 2011. "Attitudes to Ambiguity in One-Shot Normal-Form Games: An Experimental Study" *Games and Economic Behavior*, 71(2): 366-394.

Luce, R. Duncan. 1959. "Individual Choice Behavior: A Theoretical Analysis". New York, NY: Wiley.

Manski, Charles F. and Claudia Neri. 2013. "First- and Second-Order Subjective Expectations in Strategic Decision-Making: Experimental Evidence" *Games and Economic Behavior*, 81: 232-254.

McKelvey, Richard D. and Thomas R. Palfrey. 1995. "Quantal Response Equilibria for Normal Form Games" *Games and Economic Behavior*, 10(1): 6-38.

Niederle, Muriel and Emanuel Vespa. 2023. "Cognitive Limitations: Failures of Contingent Thinking" *Annual Review of Economics*, 15(1): 307-328.

Nyarko, Yaw and Andrew Schotter. 2002. "An Experimental Study of Belief Learning Using Elicited Beliefs" *Econometrica*, 70(3): 971-1005.

Polonio, Luca and Giorgio Coricelli. 2019. "Testing the Level of Consistency Between Choices and Beliefs in Games Using Eye-Tracking" *Games and Economic Behavior*, 113: 566-586.

Rey-Biel, Pedro. 2009. "Equilibrium Play and Best Response to (Stated) Beliefs in Normal Form Games" *Games and Economic Behavior*, 65(2): 572-585.

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.