



WINPEC Working Paper Series No. E2403

May 2024

Why People Seek and Avoid Feedback: A Rank-Based Ego-Utility Model of Payoff-Irrelevant Information

Kazumi Shimizu and Rongyu Hu

Waseda INstitute of Political EConomy
Waseda University
Tokyo, Japan

Why People Seek and Avoid Feedback: A Rank-Based Ego-Utility Model of Payoff-Irrelevant Information

Kazumi Shimizu*

Rongyu Hu[†]

Abstract

This paper studies why individuals sometimes seek feedback and sometimes avoid it, even when feedback is payoff-irrelevant and cannot affect material outcomes. It develops a behavioral explanation for why the same individual may seek feedback in some states but avoid it in others.

The model combines three elements: a private cue that generates a task-specific self-evaluation point, rank-based ego utility defined by the percentile position of that self-evaluation within a subjective Beta belief distribution over fixed latent ability, and outcome-concordance salience in the prospective evaluation of feedback. Realized feedback is processed by Bayes' rule.

The model yields a unique feedback-seeking threshold. For interior self-evaluation points, feedback is sought if and only if the current self-evaluation lies below this threshold. Belief asymmetry determines the threshold's position relative to the mode of the belief distribution, while Bayesian updating shifts the

*Faculty of Political Science and Economics, Waseda University, 1-104 Totsukamachi, Shinjuku-ku, Tokyo 169-8050, Japan. Email: skazumi1961@gmail.com.

[†]Graduate School of Economics, Waseda University, 1-6-1 Nishiwaseda, Shinjuku-ku, Tokyo 169-0051, Japan. Email: rongyuhu@asagi.waseda.jp.

threshold over time. Holding self-evaluation fixed, observed negative feedback can therefore generate a reversal from feedback seeking to feedback avoidance in subsequent decisions.

The paper provides a tractable threshold-based account of payoff-irrelevant information demand that separates task-specific self-evaluation from beliefs about latent ability. It links ego utility, information avoidance, and feedback seeking within a single model and generates testable predictions showing that feedback choice depends not only on current self-evaluation but also on the underlying belief distribution.

JEL codes: C11, D80, D83, D91

Keywords: ego utility, feedback seeking, information avoidance, payoff-irrelevant information, Bayesian updating, rank-based utility

1 Introduction

1.1 Motivation: A Feedback-Seeking Puzzle

Consider a completed task for which viewing the correct answer cannot affect any later payoff-relevant action or performance measure. For example, an e-learning platform may provide a costless “show the correct answer” button after a stand-alone practice question, with no effect on grades, rankings, payments, or future opportunities. Yet users differ sharply in whether they click it. Some immediately seek verification, especially when they feel unsure. Others postpone, ignore, or actively avoid seeing the solution, especially when it may threaten a favorable self-image. Similar behavior appears in everyday low-stakes settings, such as home cooking, small do-it-yourself projects, language practice, or informal skill exercises, when the feedback has no material consequences: people sometimes actively seek feedback but at other times prefer not to know.

The behavioral intuition we seek to capture is simple. A person who has a relatively unfavorable assessment of her current performance may actively seek verification, whereas a person with a more favorable self-evaluation may be more willing to postpone or avoid feedback. Our model does not treat this as a simple confidence effect. Instead, feedback seeking depends on where the current task-specific self-evaluation point lies relative to the feedback-seeking threshold. Thus, the same self-evaluation can lead to different feedback choices for individuals with different subjective belief distributions over latent ability.

The puzzle is not simply why people avoid bad news. Standard accounts of information avoidance emphasize that individuals may shy away from belief-threatening information (Golman et al., 2017). Feedback behavior may nevertheless have a stage-dependent structure: uncertainty about current performance can motivate information seeking, while unfavorable self-relevant information can reduce the appeal of later feedback. A unified behavioral account should therefore allow both initial feedback

seeking and subsequent feedback avoidance within the same framework.

We study payoff-irrelevant feedback: feedback about a completed task that cannot change the submitted answer or affect material payoffs, grades, rankings, or opportunities. It may update beliefs and thereby affect subsequent feedback choices, but it has no instrumental value for material outcomes in the baseline environment. This environment isolates the belief-based and self-image-based value of feedback. Nevertheless, feedback can matter psychologically because it changes the subjective belief distribution relative to which the individual’s current task-specific self-evaluation is assessed.

1.2 Behavioral Mechanism and Contributions

The model distinguishes two objects that are often conflated. The first is a subjective belief distribution over possible values of a fixed latent ability parameter $\alpha \in [0, 1]$, interpreted as the long-run success probability in a homogeneous task class. The latent parameter itself is fixed but unknown to the decision maker; the belief distribution represents her subjective uncertainty about it. The second object is a task-specific self-evaluation point x_t , formed after attempting the current task and based on private cues such as perceived difficulty, fluency, hesitation, familiarity, and confidence. The former summarizes subjective beliefs about latent ability; the latter summarizes how the current answer feels at the moment of the feedback decision.

We represent ego utility as rank-based. A self-evaluation point generates higher ego utility when it occupies a higher percentile position within the individual’s own subjective belief distribution over latent ability. Thus, the same self-evaluation can have different psychological value for two individuals with different belief distributions. As formalized in Section 2.4, this percentile position is represented by the cumulative distribution function of the individual’s belief distribution.

The feedback decision is governed by a simple behavioral mechanism. Before feedback is observed, the individual evaluates the possible ego consequences of two outcomes: positive feedback, indicating that the current answer is correct, and negative

feedback, indicating that the current answer is incorrect. We model this prospective evaluation through an outcome-concordance salience rule: when considering a possible feedback outcome, the individual places greater psychological weight on those possible values of latent ability under which the contemplated outcome is more likely. Once feedback is actually observed, beliefs are updated by Bayes' rule. The departure from standard Bayesian evaluation therefore lies not in updating after feedback is observed, but in the salience-based prospective evaluation of possible feedback before it is observed.

The behavioral primitives are the private-cue-based self-evaluation point, rank-based ego utility, and outcome-concordance salience in the prospective evaluation of feedback; realized feedback is processed by Bayes' rule. From these primitives, the analysis derives three results. First, payoff-irrelevant feedback can have positive or negative subjective value despite having zero instrumental value in the standard benchmark. Second, for interior self-evaluation points, feedback is sought if and only if the task-specific self-evaluation point lies below a unique threshold. Third, Bayesian updating shifts this threshold and, holding the self-evaluation point fixed, can generate history-dependent reversals between feedback seeking and feedback avoidance.

1.3 Related Literature

Ego utility and motivated beliefs. Ego-utility models (Kőszegi, 2006; Kőszegi et al., 2022) formalize the utility from self-image and show that individuals may avoid harmful information. Related self-signaling work treats beliefs about imperfectly known personal characteristics as a source of self-image and diagnostic utility (Prelec and Bodner, 2003), while a recent survey organizes the broader economics of self-image around beliefs about one's core traits, values, goals, and social ties (Bénabou and Henkel, 2025). Recent experimental evidence shows that individuals engage in motivated belief updating and rationalization when processing ego-relevant information (Drobner and Goerg, 2024). Relatedly, Möbius et al. (2022) document systematic asymmetries in belief updating—overweighting positive signals and updating too little overall. These

asymmetries are substantially weaker in an ego-free control condition, suggesting that they are driven in part by preferences over beliefs rather than by a general feature of information processing. Building on this insight, we provide a rank-based formulation of ego utility that yields a feedback-seeking threshold. By studying payoff-irrelevant feedback, we isolate the belief-based channel from instrumental motives for information acquisition. The model also connects to work on reference-dependent preferences and attention allocation (Kovach and Suleymanov, 2023), while shifting the focus to belief-based utility in payoff-irrelevant feedback settings.

Information avoidance. Our work builds on the broader literature showing that individuals may avoid information because of its anticipated psychological consequences (Golman et al., 2017). In our setting, the decision maker chooses whether to observe payoff-irrelevant feedback about a completed task. Our contribution is to show that this choice is governed jointly by the subjective belief distribution over latent ability and the current task-specific self-evaluation point. The model therefore characterizes both feedback seeking and feedback avoidance through a single threshold rule.

Initial feedback seeking and subsequent feedback avoidance. In an incentivized learning environment in which feedback also has instrumental value, low confidence in current performance predicts greater feedback seeking (Cagna et al., 2024). Other evidence documents asymmetric processing of favorable and unfavorable self-relevant information (Eil and Rao, 2011), while information demand can depart systematically from standard informational benchmarks (Hoffman, 2016). More generally, experimental work on learning from outcome feedback in repeated environments (Pittnauer and Hohnisch, 2023) motivates our focus on how the feedback-seeking threshold, and thus information demand, evolves over time. Our model shows how initial feedback seeking and subsequent feedback avoidance following observed negative feedback can arise from a single threshold mechanism that shifts through Bayesian updating.

Optimal optimism and anticipatory utility. Brunnermeier and Parker (2005) study agents who optimally choose subjective beliefs to maximize average felicity, trading off the immediate anticipatory utility of optimism against the deterioration in realized outcomes caused by decisions taken under distorted beliefs. Their framework explains why favorable beliefs may be sustained despite their behavioral costs. By contrast, our model does not study the direct choice of beliefs, but rather feedback seeking conditional on a given belief distribution, and therefore emphasizes how belief asymmetry and current self-evaluation jointly govern information acquisition.

Caplin and Leahy (2001, 2004) analyze anticipatory utility in settings where beliefs generate anxiety, suspense, and other feelings before uncertainty is resolved, which makes the timing and disclosure of information central for welfare. Their focus is therefore on the resolution of uncertainty and, in the 2004 paper, on strategic information provision by an informed sender. Our model instead abstracts from sender strategy and studies when an individual seeks or avoids feedback given an existing subjective belief distribution and a current self-evaluation point.

Bénabou and Tirole (2003, 2009, 2016) develop models in which agents manage beliefs because confidence, self-image, and anticipatory feelings affect motivation, behavior, and strategic interaction. In this literature, maintaining favorable beliefs can be instrumentally valuable because it supports effort, influences self-inference, or shapes bargaining outcomes. By contrast, our framework isolates feedback seeking in a payoff-irrelevant environment, so information acquisition is not driven by the instrumental value of beliefs for future performance or strategic behavior.

Taken together, these literatures show that beliefs may matter because they generate anticipatory utility, sustain motivation, or shape strategic interaction. Our contribution is to show that, abstracting from these channels, payoff-irrelevant feedback seeking admits a sharp threshold characterization: the feedback-seeking threshold depends on belief asymmetry and task-specific self-evaluation, and shifts predictably under Bayesian updating.

1.4 Paper Roadmap

The remainder of the paper proceeds as follows. Section 2 presents the behavioral environment, including the private-cue stage, the latent ability interpretation, the Beta belief specification, and rank-based ego utility. Section 3 analyzes the model, characterizing the feedback-seeking threshold, its comparative statics, and its dynamic evolution through experience. Section 4 discusses alternative mechanisms, states testable predictions, and concludes.

2 Model

2.1 Repeated Tasks, Latent Ability, and Feedback

The decision maker faces a sequence of independent tasks drawn from a homogeneous task class $\mathcal{T}$. She possesses a latent ability parameter $\alpha \in [0, 1]$, representing the probability of correctly answering a randomly drawn task from $\mathcal{T}$. This parameter is fixed but unknown to the decision maker. If she were to solve a large number of tasks from $\mathcal{T}$, the empirical frequency of correct answers would converge to α . Thus, α is interpreted as her long-run success probability within this task class.

Many learning situations naturally fit this structure: a student solving calculus problems of similar difficulty, a language learner answering vocabulary questions, a medical trainee repeatedly diagnosing similar cases, or a chess player solving tactical puzzles. In each case, there is an underlying ability to solve tasks from that class, and each individual problem can provide a noisy signal about persistent ability.

For each task t , conditional on the fixed but unknown latent ability α , the outcome follows a Bernoulli distribution: $s_t \mid \alpha \sim \text{Bernoulli}(\alpha)$. Equivalently,

$$P(s_t = 1 \mid \alpha) = \alpha, \quad P(s_t = 0 \mid \alpha) = 1 - \alpha.$$

Here $s_t = 1$ denotes a correct answer and $s_t = 0$ denotes an incorrect answer. The

decision maker knows the Bernoulli data-generating process but not the fixed value of α . Her uncertainty about α is represented by a subjective belief distribution introduced in Section 2.3, from which her subjective predictive probabilities for the current task are derived.

Once the answer has been submitted, the outcome s_t is realized but is not yet observed by the decision maker. She then chooses whether to observe feedback. Feedback is payoff-irrelevant: it cannot change the submitted answer, material payoffs, or grades. If she observes feedback, the realized s_t is revealed and enters Bayesian updating. If she declines feedback, the realization remains unobserved and does not update her beliefs in the baseline voluntary-feedback environment.

2.2 Private Cues and Task-Specific Self-Evaluation

After submitting an answer to task t , but before deciding whether to view feedback, the decision maker obtains a private cue $r_t \in \mathcal{R}$ about the current task. The cue summarizes task-specific impressions formed during problem solving, such as perceived difficulty, fluency, hesitation, familiarity, and confidence.

Let μ_t denote the decision maker’s current subjective belief distribution over possible values of the fixed latent ability α . The private cue, together with the current belief distribution, generates a task-specific self-evaluation point through a measurable function

$$x_t = \chi(r_t, \mu_t) \in [0, 1].$$

Higher values of x_t indicate a more favorable task-specific self-evaluation of the current answer, whereas lower values indicate a less favorable one.

The variables x_t and α play different roles. The fixed latent parameter α represents the individual’s persistent long-run success probability within the task class, whereas x_t represents her evaluation of the particular answer just submitted. We place x_t on the same normalized performance scale $[0, 1]$ as the possible values of α , so that the position of x_t within the individual’s subjective belief distribution μ_t is well defined.

For the feedback-choice analysis, we restrict attention to interior self-evaluation

points $x_t \in (0, 1)$. The boundary cases are degenerate because the subjective value of feedback satisfies $V(0) = V(1) = 0$, as established in the feedback-seeking threshold analysis in Section 3.2.

The private cue enters the baseline model only through x_t ; it is not an additional likelihood term for updating beliefs over α . This is a modeling restriction, not a claim that task-specific impressions are uninformative about correctness. The analysis conditions on the realized x_t ; Section 4.3 discusses the omitted joint stochastic structure.

2.3 Beliefs over Latent Ability

The decision maker is uncertain about the fixed latent ability α and holds a subjective belief distribution μ_t over its possible values at each task t . We use $\theta \in [0, 1]$ to index a generic possible value of α in the support of this belief distribution. Thus, α denotes the fixed but unknown latent ability, whereas θ is the variable over which the decision maker's subjective belief is defined.

We now specialize μ_t to the Beta family. Specifically, at the moment of the feedback decision for task t , beliefs are represented by $\mu_t = \text{Beta}(a_t, b_t)$. The parameters (a_t, b_t) summarize the current belief, incorporating the initial belief and all feedback about past tasks that has been observed up to that point. For the single-task analysis, we suppress the task index and write $\mu = \text{Beta}(a, b)$.

Formally, the density of the subjective belief μ , written as a function of the possible value θ of the latent ability, is:

$$f(\theta | a, b) = \frac{1}{B(a, b)} \theta^{a-1} (1 - \theta)^{b-1}, \quad a > 1, b > 1, \quad (1)$$

where $B(a, b)$ is the Beta function.

The parameters (a, b) characterize two features of the current subjective belief distribution. Its mean,

$$\mathbb{E}_\mu[\theta] = \frac{a}{a + b},$$

represents the individual's subjective expected long-run success probability. The con-

centration parameter $a + b$ captures the concentration of that belief. Holding the belief mean fixed, larger values of $a + b$ correspond to more concentrated beliefs, while smaller values imply greater uncertainty.

Under the Bernoulli probability model, the subjective predictive probabilities implied by the current belief distribution are

$$\begin{aligned} P^\mu(s = 1) &= \int_0^1 \theta \mu(d\theta) = \frac{a}{a + b}, \\ P^\mu(s = 0) &= \int_0^1 (1 - \theta) \mu(d\theta) = \frac{b}{a + b}. \end{aligned} \tag{2}$$

These probabilities integrate over the decision maker's uncertainty about the fixed latent ability α . Heterogeneity in belief mean and belief concentration therefore enters the model through heterogeneity in (a, b) . Beliefs with $a < b$ are called *pessimistic*: both the mean and mode lie below $1/2$. Beliefs with $a > b$ are called *optimistic*: both the mean and mode lie above $1/2$. A symmetric belief ($a = b$) is centered at $\theta = 0.5$.

The Beta family is the natural conjugate family for the Bernoulli likelihood; its support $[0, 1]$ matches the set of possible values of α , indexed by θ , and the pair (a, b) maps transparently into the belief mean and belief concentration.¹ Figure 1 illustrates three representative belief distributions.

¹By varying (a, b) , the Beta family accommodates left-skewed, right-skewed, and symmetric unimodal beliefs under the maintained restriction $a, b > 1$.




Figure 1: Subjective belief distributions over possible values of the latent ability α , indexed by θ , for three individuals with different belief asymmetries. The solid curve represents an optimistic belief, $\text{Beta}(5, 2)$, with mode at $\theta_{\text{mode}} = 0.80$. The dashed curve represents a pessimistic belief, $\text{Beta}(2, 5)$, with mode at $\theta_{\text{mode}} = 0.20$. The dash-dotted curve represents a symmetric belief, $\text{Beta}(3, 3)$, with mode at $\theta_{\text{mode}} = 0.50$. The horizontal axis represents a possible value θ of the latent ability α , and the vertical axis represents the belief density $f(\theta | a, b)$. *Source: Authors own work.*

2.4 Rank-Based Ego Utility

“Ego utility” is the utility derived directly from self-image—the psychological payoff from perceiving oneself as capable (Kőszegi, 2006). We model ego utility as depending on the relative position of the current task-specific self-evaluation within the individual’s own subjective belief distribution over latent ability.

This specification is motivated by three related ideas in behavioral economics. First, models of ego utility and belief utility emphasize that beliefs about oneself can enter preferences directly, rather than merely serving as inputs into material payoff maximization (Kőszegi, 2006; Bénabou and Tirole, 2016). Second, reference-dependent preference theory models subjective value as depending on comparisons with a reference standard, rather than on absolute outcomes alone (Kőszegi and Rabin, 2006). Third, theories of relative utility and status emphasize that utility or well-being may depend not only on the absolute level of an outcome, but also on its *rank* or relative

position within a comparison distribution (Hopkins and Kornienko, 2004; Brown et al., 2008). The latter literature studies relative position in social comparison environments, where an individual's standing is evaluated relative to others. Relatedly, Decision by Sampling evaluates an attribute by its relative rank among comparison values drawn from the current context and memory (Stewart et al., 2006). Our model adapts this relative-rank logic to the domain of self-image. The relevant comparison distribution is not a social distribution of other individuals' outcomes or a sample of externally encountered attribute values, but the individual's own subjective belief distribution over possible values of the fixed latent ability α .

Recall that x_t and the possible values of α , indexed by θ , are represented on the same normalized scale, although x_t is not itself a belief about α .

We define rank-based ego utility directly as the cumulative belief mass weakly below the current task-specific self-evaluation point:

$$U(x_t, \mu_t) = \mu_t([0, x_t]).$$

Hence, ego utility equals the cumulative subjective belief probability assigned to possible values of α that are no greater than x_t . Equivalently, it is the percentile position of x_t relative to the individual's current subjective belief distribution over the fixed latent ability α . A higher percentile corresponds to a more favorable self-evaluation relative to that belief distribution.

Under the Beta specification $\mu_t = \text{Beta}(a_t, b_t)$, and suppressing the task index as in the single-task analysis, let $F(\cdot \mid a, b)$ denote the cumulative distribution function associated with the belief density $f(\cdot \mid a, b)$. The current pre-feedback CDF is

$$F(z \mid a, b) = \frac{1}{B(a, b)} \int_0^z u^{a-1} (1-u)^{b-1} du = \frac{1}{B(a, b)} B_e(z; a, b), \quad z \in [0, 1], \quad (3)$$

where $B_e(z; a, b)$ is the incomplete Beta function. Rank-based ego utility therefore becomes

$$U(x \mid a, b) = F(x \mid a, b).$$

Thus, changes in the belief distribution can change ego utility even when the task-specific self-evaluation point is held fixed.

This distributional feature is essential to the mechanism. A specification based only on the belief mean $\mathbb{E}_\mu[\theta]$ would discard information about the shape and concentration of the belief distribution and therefore could not capture the percentile position of x within that distribution. Earlier work has also used cumulative distribution functions as utility or relative evaluation functions (Van Praag and Kapteyn, 1973; Gregory, 1980; Chen and Novick, 1982).²

2.5 Belief Mode and the Location of Self-Evaluation

Given the current subjective belief density $f(\theta \mid a, b)$ over possible values of α , let θ_{mode} denote the mode of this belief distribution. For a $\text{Beta}(a, b)$ distribution with $a, b > 1$,

$$\theta_{\text{mode}} = \frac{a - 1}{a + b - 2}.$$

This is the possible value of the latent ability at which the current belief density is maximized. We describe evaluation points satisfying $x < \theta_{\text{mode}}$ as *below-mode* and those satisfying $x > \theta_{\text{mode}}$ as *above-mode*. These labels refer only to the location of x relative to the mode of the individual’s subjective belief; they do not compare x with the true latent ability α .

3 Analysis and Results

3.1 Bayesian Updating and Feedback Evaluation

At the time of the feedback decision, the task outcome has already been realized but is not yet observed by the decision maker. Suppressing the task index, let the current pre-feedback subjective belief be $\mu = \text{Beta}(a, b)$, with density $f(\theta \mid a, b)$ as defined in (1). The corresponding subjective predictive probabilities are given in (2).

²We use these papers as formal precedents for CDF-based relative evaluation, not as direct empirical evidence for the present ego-utility mechanism.

As specified in Section 2.2, the private cue affects x but is not an additional likelihood term. If feedback is chosen, the realized outcome is revealed and the belief over α is updated by Bayes' rule.

3.1.1 Posterior Distribution

By conjugacy of the Beta–Bernoulli model, conditional on observing a correct answer ($s = 1$), the decision maker's posterior belief distribution is $\text{Beta}(a+1, b)$. Conditional on observing an incorrect answer ($s = 0$), her posterior belief distribution is $\text{Beta}(a, b + 1)$. An algebraic derivation is given in Appendix A.

Figure 2 illustrates Bayesian updating.




Figure 2: Bayesian updating after observing feedback. The solid curve represents the symmetric pre-feedback belief distribution $\text{Beta}(a, b)$, where $a = b$. The dashed curve represents the posterior belief distribution after observing that the answer is correct, $\text{Beta}(a + 1, b)$. The dash-dotted curve represents the posterior belief distribution after observing that the answer is incorrect, $\text{Beta}(a, b + 1)$. The horizontal axis represents a possible value θ of the latent ability α , and the vertical axis represents belief density. *Source: Authors own work.*

3.1.2 Outcome-Concordance Salience Rule

Recall from Section 2.4 that rank-based ego utility is $F(x \mid a, b)$. To distinguish the two actions at the feedback decision, define the ego utility from avoiding feedback as

$$F_{\text{avoid}}(x) := F(x \mid a, b).$$

Using the posterior distributions characterized above, the standard Bayesian expected-utility benchmark evaluates prospective feedback by averaging the two posterior ego utilities using the subjective predictive probabilities defined in (2). Hence,

$$\frac{a}{a+b}F(x \mid a+1, b) + \frac{b}{a+b}F(x \mid a, b+1) = F_{\text{avoid}}(x), \quad (4)$$

where the equality follows from $B(a+1, b) = \frac{a}{a+b}B(a, b)$, $B(a, b+1) = \frac{b}{a+b}B(a, b)$, and

$$\theta^a(1-\theta)^{b-1} + \theta^{a-1}(1-\theta)^b = \theta^{a-1}(1-\theta)^{b-1}.$$

Thus, under ordinary subjective predictive averaging, expected ego utility after realized feedback equals ego utility from avoiding feedback, so payoff-irrelevant feedback has zero subjective value. The behavioral model departs from this benchmark only in its prospective evaluation of possible feedback outcomes; Bayesian updating after observed feedback remains standard.

The outcome-concordance salience rule (hereafter, the *salience rule*) is motivated by three related bodies of evidence. First, work on representativeness and judgment under uncertainty suggests that people evaluate uncertain events partly by the states or types for which those events appear representative or diagnostic (Kahneman and Tversky, 1972; Tversky and Kahneman, 1974). In the present setting, a possible positive outcome is more concordant with higher values of the latent ability, whereas a possible negative outcome is more concordant with lower values. Second, salience-based models in behavioral economics show that features that stand out in a given choice context can receive disproportionate decision weight (Bordalo et al., 2012). Third, evidence on self-

relevant information shows that people process and acquire good and bad news about themselves asymmetrically, especially when the information concerns ego-relevant qualities such as ability or attractiveness (Eil and Rao, 2011; Möbius et al., 2022). Related work on information avoidance and psychological expected utility also emphasizes that people may avoid information because of its anticipated psychological consequences (Sweeny et al., 2010; Caplin and Leahy, 2001).

To formalize this idea, consider a possible value θ of the latent ability and a possible outcome $s \in \{0, 1\}$. Define the Bernoulli likelihood

$$L(\theta; s) := P(s \mid \alpha = \theta) = \theta^s (1 - \theta)^{1-s}.$$

Hence,

$$L(\theta; 1) = \theta, \quad L(\theta; 0) = 1 - \theta.$$

We use $L(\theta; 1) = \theta$ and $L(\theta; 0) = 1 - \theta$ as outcome-concordance salience weights. This is a reduced-form representation, not a claim that agents consciously compute the weights.³

Assumption 1 (Outcome-concordance salience). *Before deciding whether to observe feedback, the decision maker prospectively evaluates the ego consequences of possible feedback by placing greater psychological weight on those possible values θ of the fixed latent ability α under which each possible outcome is more likely. In the Beta–Bernoulli environment, this outcome-concordance weighting is represented by the Bernoulli likelihood values: the posterior density following a positive outcome is weighted by $L(\theta; 1) = \theta$, and the posterior density following a negative outcome is weighted by $L(\theta; 0) = 1 - \theta$.*

This assumption changes only the prospective ego evaluation of possible feedback. The subjective predictive probabilities in (2) define the standard benchmark but do not weight the two branches under the salience rule. We therefore define the following

³The cited literatures motivate the idea of outcome-concordance salience but do not uniquely imply the likelihood-based weights used here. We adopt $L(\theta; 1) = \theta$ and $L(\theta; 0) = 1 - \theta$ as a parsimonious specification that introduces no additional salience parameters.

salience-adjusted prospective ego-evaluation kernel:

$$\tilde{f}_{\text{seek}}(\theta) = \theta f(\theta \mid a+1, b) + (1-\theta)f(\theta \mid a, b+1). \quad (5)$$

No additional branch-level parameters are introduced: the positive and negative branches are specified symmetrically, and their relative contributions are implied by the salience weights and the normalization constant.

Because (5) is an ego-evaluation kernel rather than a probability density, we normalize it. Its total mass is

$$\begin{aligned} Z_S(a, b) &:= \int_0^1 \tilde{f}_{\text{seek}}(\theta) d\theta \\ &= \frac{a+1}{a+b+1} + \frac{b+1}{a+b+1} = \frac{a+b+2}{a+b+1}. \end{aligned} \quad (6)$$

The salience-adjusted prospective ego utility from seeking feedback is the normalized cumulative mass of this kernel up to x :

$$F_{\text{seek}}(x) = \frac{1}{Z_S(a, b)} \int_0^x [\theta f(\theta \mid a+1, b) + (1-\theta)f(\theta \mid a, b+1)] d\theta. \quad (7)$$

Equivalently,

$$F_{\text{seek}}(x) = \frac{a+1}{a+b+2} F(x \mid a+2, b) + \frac{b+1}{a+b+2} F(x \mid a, b+2). \quad (8)$$

The mixture representation (8) provides a useful algebraic interpretation of the salience rule. Under the standard benchmark (4), the positive branch uses the posterior rank-based ego utility $F(x \mid a+1, b)$, and the negative branch uses $F(x \mid a, b+1)$. Under the salience rule, multiplying the positive posterior density by θ produces a density proportional to $f(\theta \mid a+2, b)$, while multiplying the negative posterior density by $1-\theta$ produces a density proportional to $f(\theta \mid a, b+2)$:

$$\theta f(\theta \mid a+1, b) \propto f(\theta \mid a+2, b), \quad (1-\theta)f(\theta \mid a, b+1) \propto f(\theta \mid a, b+2).$$

Thus, the salience rule reweights each possible posterior in the direction favored by the corresponding outcome. This is a transformation of prospective ego evaluation, not a second Bayesian update. Unlike under the standard benchmark, payoff-irrelevant feedback can therefore have positive or negative subjective value.

3.2 Feedback-Seeking Threshold

Under Assumption 1, $F_{\text{avoid}}(x) = F(x \mid a, b)$ is the ego utility from avoiding feedback, whereas $F_{\text{seek}}(x)$ in (7) is the salience-adjusted prospective ego utility from seeking feedback. The subjective value of feedback is their difference (see Appendix B):

$$\begin{aligned} V(x) &= F_{\text{seek}}(x) - F_{\text{avoid}}(x) \\ &= \frac{1}{(a+b+2)B(a,b)} x^a (1-x)^b \left\{ \frac{a-b}{ab} + \frac{a+b}{b}(1-x) - \frac{a+b}{a}x \right\}. \end{aligned} \quad (9)$$

The decision maker observes feedback if the salience-adjusted prospective ego utility from seeking strictly exceeds the ego utility from avoidance. Thus, she chooses to observe feedback if $V(x) > 0$, and avoids observing feedback if $V(x) \leq 0$, with indifference resolved in favor of avoidance.

The value function $V(x)$ becomes zero at three points: $x = 0$, $x = 1$, and

$$C = \frac{a - b + a^2 + ab}{(a+b)^2}.$$

The roots at $x = 0$ and $x = 1$ occur because both F_{seek} and F_{avoid} equal 0 at $x = 0$ and 1 at $x = 1$. The threshold C lies in the open interval $(0, 1)$ for $a, b > 1$.

Proposition 1 (Feedback-Seeking Threshold). *There exists a unique threshold $C \in (0, 1)$ given by*

$$C = \frac{a - b + a^2 + ab}{(a+b)^2} = \frac{(a-1)b + a + a^2}{(a+b)^2}$$

such that

- $V(x) > 0$ for $x \in (0, C)$: *feedback seeking*;

- $V(x) \leq 0$ for $x \in [C, 1)$: *feedback avoidance*.

Proof. Since

$$\frac{a-b}{ab} + \frac{a+b}{b}(1-x) - \frac{a+b}{a}x = \frac{(a+b)^2}{ab}(C-x),$$

we can write

$$V(x) = \frac{(a+b)^2}{ab(a+b+2)B(a,b)} x^a(1-x)^b(C-x).$$

For $x \in (0, 1)$, the coefficient multiplying $(C-x)$ is strictly positive. Hence $V(x) > 0$ if and only if $x < C$.

Moreover,

$$a^2 + ab + a - b = a^2 + a + (a-1)b > 0,$$

and

$$(a+b)^2 - (a^2 + ab + a - b) = a(b-1) + b^2 + b > 0.$$

Thus, $C \in (0, 1)$. □

Figure 3 illustrates the subjective value of feedback $V(x)$ for the case $(a, b) = (2, 2)$. The function is positive for $x \in (0, 0.5)$ and negative or zero for $x \in [0.5, 1)$, where $C = 0.5$ is the threshold.




Figure 3: Subjective value of feedback $V(x) = F_{\text{seek}}(x) - F_{\text{avoid}}(x)$ for $(a, b) = (2, 2)$. The function is positive for $x \in (0, C)$, negative for $x \in (C, 1)$, and zero at $x = 0$, $x = C$, and $x = 1$, where $C = 0.5$. For interior self-evaluation points, $(0, C)$ is the feedback-seeking region and $[C, 1)$ is the feedback-avoidance region. The horizontal axis represents the task-specific self-evaluation point x ; the vertical axis represents the change in ego utility from observing feedback. *Source: Authors own work.*

3.3 Comparative Statics: Threshold Position and Belief Structure

For each current Beta belief with parameters (a, b) , feedback is sought if and only if $x < C(a, b)$. The sign of $a - b$ characterizes the belief's asymmetry, using the terminology introduced in Section 2.3.

To understand how belief asymmetry affects feedback-seeking behavior, we compare the threshold C with the belief mode θ_{mode} . The difference between the two is

$$\theta_{\text{mode}} - C = \frac{2(a - b)}{(a + b)^2(a + b - 2)}. \quad (10)$$

Since the denominator is strictly positive, the sign of $\theta_{\text{mode}} - C$ is determined by the sign of $a - b$.

Proposition 2 (Threshold Position and Belief Asymmetry).

1. If $a > b$ (*optimistic belief*), then $C < \theta_{\text{mode}}$.
2. If $a = b$ (*symmetric belief*), then $C = \theta_{\text{mode}}$.
3. If $a < b$ (*pessimistic belief*), then $C > \theta_{\text{mode}}$.

Proof. Direct calculation gives (10). When $a > b$, this expression is positive, implying $C < \theta_{\text{mode}}$. When $a = b$, it is zero, implying $C = \theta_{\text{mode}}$. When $a < b$, it is negative, implying $C > \theta_{\text{mode}}$. $\square$

The belief mode θ_{mode} provides a reference point for interpreting the threshold rule. As defined in Section 2.5, an evaluation point is below the mode when $x < \theta_{\text{mode}}$ and above the mode when $x > \theta_{\text{mode}}$. Comparing C with θ_{mode} clarifies how belief asymmetry shifts the feedback-seeking region. When $a > b$, the threshold lies strictly to the left of the mode, so the feedback-seeking region $\{x < C\}$ is entirely contained in the below-mode region. When $a = b$, the threshold coincides with the mode. When $a < b$, the threshold lies to the right of the mode, so the feedback-seeking region extends beyond the belief mode and includes all below-mode points plus some above-mode points. This is a statement about the position of the seeking region relative to

the mode of the individual’s current belief, not about the absolute size of the seeking region across different belief distributions.

Figure 4 illustrates how the position of the threshold varies with belief asymmetry.




Figure 4: Threshold position $C(a, b)$ under different belief asymmetries. The solid curve represents an optimistic belief, Beta(5, 2), with mode at $\theta_{\text{mode}} = 0.80$. The dashed curve represents a pessimistic belief, Beta(2, 5), with mode at $\theta_{\text{mode}} = 0.20$. The vertical dashed lines mark the belief modes, and the vertical dotted lines mark the corresponding feedback-seeking thresholds. For the optimistic belief, $C \approx 0.776 < \theta_{\text{mode}}$, so the feedback-seeking set $(0, 0.776)$ is contained within the below-mode region. For the pessimistic belief, $C \approx 0.224 > \theta_{\text{mode}}$, so the feedback-seeking set $(0, 0.224)$ extends into the above-mode region. The horizontal axis uses the common $[0, 1]$ scale on which possible values of the latent ability α (indexed by θ), the task-specific self-evaluation point, and the feedback-seeking threshold are represented. *Source: Authors own work.*

The threshold-position result describes the feedback choice for a given self-evaluation point x . By itself, it does not rank how frequently people with different belief distributions seek feedback, because their task-specific self-evaluations may also differ.

It may be natural to suppose that individuals with pessimistic beliefs are more prone to forming task-specific evaluation points below their respective belief modes, whereas individuals with optimistic beliefs are more prone to forming task-specific evaluation points above their respective belief modes.⁴

⁴This auxiliary behavioral hypothesis is not required for the threshold-position result. Relative self-assessments and different forms of under- and overconfidence vary with task difficulty and the comparative structure of the judgment problem (Moore and Cain, 2007; Moore and Healy, 2008). Manipulated prior beliefs about task performance can affect reported confidence without affecting objective performance (Van Marcke et al., 2024), while lower task-specific confidence is associated with greater feedback seeking when feedback also has instrumental learning value (Cagna et al., 2024). These find-

Under this auxiliary behavioral hypothesis, the evaluation point and the feedback-seeking threshold tend to lie on opposite sides of the belief mode. For a pessimistic belief, the evaluation point is more likely to lie below the threshold, favoring feedback seeking. For an optimistic belief, the evaluation point is more likely to lie above the threshold, favoring feedback avoidance.

The model therefore *accommodates* the intuitive pattern that individuals with relatively unfavorable task-specific self-evaluations are more likely to seek feedback, whereas those with relatively favorable self-evaluations are more likely to avoid it. Importantly, this pattern is not driven by self-evaluation alone, but by its position relative to the feedback-seeking threshold. It is a conditional behavioral implication rather than an unconditional consequence of belief asymmetry. Appendix C provides a precise sufficient condition under which the corresponding population-frequency comparison holds.

3.4 Effect of Experience on Subsequent Feedback Choice

Recall that, for interior self-evaluation points, feedback is sought if and only if $x < C(a, b)$. Because the cue rule allows $x_t = \chi(r_t, \mu_t)$, the results below isolate the effect of belief updating on the threshold while holding the self-evaluation point fixed. They are therefore *ceteris paribus* comparative statics, not an unconditional prediction that an individual's realized choice path must reverse after any particular sequence of outcomes.

We now investigate how this threshold evolves in response to accumulated observed feedback. Let m represent the number of observed positive feedback outcomes and n the number of observed negative feedback outcomes. Under the conditionally i.i.d. Bernoulli data-generating process given α , Bayesian updating transforms a current pre-feedback belief $\text{Beta}(a, b)$ into $\text{Beta}(a + m, b + n)$ after these feedback outcomes are

ings motivate the hypothesis but do not establish either the belief mode as the relevant reference point or the proposed mode-relative propensity.

observed. The corresponding threshold becomes

$$C_{m,n} = \frac{(a+m) - (b+n) + (a+m)^2 + (a+m)(b+n)}{(a+b+m+n)^2}.$$

Proposition 3 (Effect of Experience on Threshold). *For any $a, b > 1$ and any $m, n \in \mathbb{Z}_{\geq 0}$:*

1. *For every fixed n , $C_{m,n}$ is strictly increasing in m , and*

$$\lim_{m \rightarrow \infty} C_{m,n} = 1.$$

2. *For every fixed m , $C_{m,n}$ is strictly decreasing in n , and*

$$\lim_{n \rightarrow \infty} C_{m,n} = 0.$$

Proof. See Appendix D. □

Figure 5 visualizes the comparative statics of the threshold under Bayesian updating. Starting from the baseline belief $\text{Beta}(2, 2)$, for which $C_{0,0} = 0.5$, m positive and n negative feedback observations yield $C_{m,n} = C(2+m, 2+n)$. The figure shows the resulting threshold for different combinations of accumulated positive and negative feedback. Additional positive feedback observations raise $C_{m,n}$, whereas additional negative feedback observations lower it. The overlaid paths trace illustrative positive and negative feedback streaks.




Figure 5: Evolution of the feedback-seeking threshold $C_{m,n}$ with experience. Starting from Beta(2, 2), additional positive feedback observations increase the threshold, while additional negative feedback observations decrease it. Darker shading indicates higher values of $C_{m,n}$, whereas lighter shading indicates lower values. Each contour line connects combinations of positive and negative feedback observations that yield the same threshold value. The two marked paths illustrate a positive streak and a negative streak. *Source: Authors own work.*

Observed feedback affects feedback-seeking behavior through its effect on the subjective belief distribution and hence on the threshold. Observed positive feedback increases a relative to b , raises the belief mean, and moves the threshold rightward. This expansion of C enlarges the region $\{x < C_{m,n}\}$ in which feedback seeking is optimal. Conversely, observed negative feedback increases b relative to a , lowers the belief mean, and moves the threshold leftward. This contraction of C reduces the set of evaluation points for which feedback seeking is optimal. At each decision point, the individual seeks feedback if and only if the task-specific self-evaluation point satisfies $x < C_{m,n}$.

Holding the evaluation point fixed, this threshold dynamic can generate two types of feedback-choice reversal. First, consider an individual whose evaluation point satisfies $x < C_{m,n}$, so that she initially seeks feedback. Each additional negative feedback observation lowers the threshold. After sufficiently many such negative feedback ob-

servations, the threshold may fall below the evaluation point, so that $x \geq C_{m,n+k}$ for some $k \geq 1$. At that point, feedback avoidance becomes optimal. Thus, the seeking-to-avoidance reversal is endogenous in the baseline voluntary-feedback environment: an individual initially seeks feedback, observes negative outcomes, updates her beliefs, and may subsequently enter the avoidance region.

The converse avoidance-to-seeking reversal is different. Holding the evaluation point fixed, once an individual declines voluntary feedback, she continues to decline it in the baseline environment unless new information changes the threshold. An avoidance-to-seeking reversal through the threshold can therefore arise when positive performance information about the same task class arrives through an external channel, such as mandatory disclosure, unsolicited feedback, or publicly observable performance outcomes. Under sufficiently favorable externally observed information, the threshold moves rightward and may eventually exceed the fixed evaluation point, making feedback seeking optimal again.

Proposition 4 (Reversals under Observed Feedback). *Fix a common self-evaluation point $x \in (0, 1)$ across decisions, and let $C_{m,n}$ denote the threshold after m positive and n negative feedback observations.*

- (i) *Seeking-to-avoidance reversal after negative feedback.* If $x \in (0, C_{m,n})$, there exists a finite $k^* \geq 1$ such that

$$C_{m,n+k^*} \leq x.$$

Hence, after k^* additional negative feedback observations, feedback avoidance becomes optimal.

- (ii) *Avoidance-to-seeking reversal after external positive information.* If $x \in [C_{m,n}, 1)$, there exists a finite $k^{**} \geq 1$ such that

$$x < C_{m+k^{**},n}.$$

Thus, conditional on observing k^{**} additional positive performance outcomes from the

same task class through an external channel, feedback seeking becomes optimal.

Proof. See Appendix E. □

4 Discussion and Conclusion

4.1 Relation to Asymmetric Updating and Alternative Mechanisms

The threshold mechanism also highlights a distinction between asymmetric belief updating and endogenous feedback avoidance. For example, Möbius et al. (2022) document asymmetric updating of ego-relevant beliefs, with positive feedback receiving greater weight than negative feedback. Our model does not generate such asymmetry conditional on an observed signal: once feedback is observed, beliefs are updated according to Bayes' rule. Instead, the model generates an asymmetry in subsequent feedback seeking. By Proposition 3, negative feedback lowers the threshold $C_{m,n}$. Holding x fixed, repeated negative feedback can therefore move the individual out of the feedback-seeking region, so that subsequent feedback is no longer voluntarily observed. This mechanism, in which negative feedback can reduce subsequent feedback seeking, is complementary to models of asymmetric updating conditional on observed signals.

This perspective distinguishes our framework from mechanisms that relate information demand directly to the expected valence of news. In our model, feedback seeking is not determined by optimism, pessimism, or expected news valence alone. Instead, it is governed by the interaction between two objects: the subjective belief distribution, which determines the position of the feedback-seeking threshold $C(a, b)$, and the task-specific self-evaluation point x , which determines where the individual currently stands relative to that threshold.

Confirmation bias (Nickerson, 1998) and cognitive dissonance (Festinger, 1957) offer mechanisms through which belief-threatening information may be discounted or avoided. Anticipatory-utility models analyze how beliefs about future outcomes and the timing of information affect current utility and information demand (Brunnermeier

and Parker, 2005; Caplin and Leahy, 2001). Our mechanism is distinct: feedback choice depends jointly on the current belief distribution and the location of the task-specific self-evaluation point relative to the threshold. Optimism or expected news valence alone is therefore insufficient to determine feedback choice in the present model.

4.2 Empirical Implications and Testable Predictions

The model can be tested experimentally in controlled feedback environments in which participants complete objectively scored tasks and decide whether to observe payoff-irrelevant feedback. The key empirical objects are the task-specific self-evaluation point x_t , the subjective belief distribution over latent ability, and the subsequent feedback choice. These objects can be elicited separately, allowing the model's central threshold condition for interior self-evaluation points, $x_t < C(a_t, b_t)$, to be tested directly.

The central prediction is not that feedback seeking is determined by self-evaluation alone. Rather, feedback choice depends on the position of the task-specific self-evaluation point relative to the feedback-seeking threshold implied by the individual's subjective belief distribution. This yields the following testable predictions.

1. **Threshold effects beyond self-evaluation.** Conditional on task-specific self-evaluation x_t , feedback choice should vary systematically across subjective belief distributions that imply different thresholds $C(a_t, b_t)$; the estimated threshold should therefore explain feedback choice beyond self-evaluation alone.
2. **Seeking-to-avoidance dynamics after negative feedback.** For comparable values of x_t , voluntarily observed negative feedback should lower the future threshold and increase subsequent feedback avoidance.
3. **External positive information and renewed seeking.** If positive performance information about the same task class arrives through an external channel after voluntary feedback has been declined, the threshold should rise and subsequent feedback seeking may resume.

The first prediction distinguishes the model from a self-evaluation-only account; the latter two test the *ceteris paribus* threshold dynamics.

4.3 Limitations and Extensions

The model makes several simplifying assumptions that future research could relax.

Private cue formation. The present model gives the private cue an explicit role in generating x_t , but it does not impose a full joint stochastic structure linking private cues, current correctness, task difficulty, and ability. Such a structure would be necessary to derive the distribution of x_t across tasks from underlying factors such as task difficulty, latent ability, and the distribution of private cues. This is beyond the scope of the present paper, which conditions on the realized self-evaluation point at the feedback decision.

Bayesian updating. We assume individuals update beliefs according to Bayes' rule after observed feedback. Extensive evidence documents systematic deviations from Bayesian updating, including base-rate neglect (Kahneman and Tversky, 1973), conservatism (Edwards, 1968), and confirmation bias (Nickerson, 1998). Incorporating these biases would generate richer predictions about how the threshold evolves over time.

Heterogeneous ego sensitivity and attention costs. Our baseline model allows individuals to differ in their belief distributions through heterogeneity in (a, b) . Individuals may also differ in the importance they attach to self-image and the subjective cost of attending to feedback. Let $\lambda_i > 0$ denote individual i 's sensitivity to ego utility, and let $c_i \geq 0$ denote the individual-specific attention or psychological cost of viewing feedback. The utilities from avoiding and seeking feedback can be written as

$$U_i^{\text{avoid}}(x) = \lambda_i F_{\text{avoid}}(x),$$

and

$$U_i^{\text{seek}}(x) = \lambda_i F_{\text{seek}}(x) - c_i.$$

Hence, individual i seeks feedback if and only if $\lambda_i V(x) > c_i$.

Holding beliefs and the evaluation point fixed, feedback seeking is more likely when ego sensitivity λ_i is higher or attention cost c_i is lower, provided that $V(x) > 0$. When $c_i = 0$, the baseline decision rule is recovered because $\lambda_i > 0$ does not alter the sign of $V(x)$. When $c_i > 0$, feedback is sought only when its ego-utility gain is sufficiently large to exceed the individual-specific cost.

Instrumental motives. We focus on payoff-irrelevant feedback to isolate ego-utility-based preferences. In many real-world settings, feedback affects subsequent decisions and material outcomes. Integrating ego utility with instrumental information acquisition would allow the model to address environments where both channels operate simultaneously, such as educational settings where grades affect both self-image and future opportunities.

4.4 Conclusion

This paper develops a behavioral model of why individuals seek or avoid payoff-irrelevant feedback. The model has three primitives: a private cue that generates a task-specific self-evaluation point, rank-based ego utility, and outcome-concordance salience in prospective feedback evaluation. From these primitives, the analysis derives a unique feedback-seeking threshold. For interior self-evaluation points, feedback is sought if and only if the current task-specific self-evaluation point lies below the threshold.

The threshold depends on the subjective belief distribution over latent ability. Optimistic beliefs place the threshold below the belief mode, restricting feedback seeking to sufficiently low self-evaluation points. Pessimistic beliefs place the threshold above the belief mode, allowing feedback seeking to extend into the above-mode region. Bayesian updating shifts the threshold over time: observed positive feedback raises

it, while observed negative feedback lowers it. For a fixed evaluation point, this can generate history-dependent transitions between feedback seeking and feedback avoidance.

Taken together, these results imply that feedback seeking depends not on confidence or optimism alone, but on the position of the task-specific self-evaluation point relative to the feedback-seeking threshold implied by the subjective belief distribution. The model therefore offers a tractable behavioral account of why payoff-irrelevant feedback may be sought in some states, avoided in others, and reversed over time through belief-driven threshold movements.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Competing interests

The authors declare no competing interests.

Data Availability

Data sharing is not applicable to this article as no new data were created or analyzed in this study.

Declaration of generative AI use

During the preparation of this work, the authors used ChatGPT Plus (OpenAI) to copy-edit their original text for language and readability and to check terminology, cross-references, and bibliographic formatting. The authors independently reviewed all suggested revisions and verified every citation, cross-reference, and mathematical expression, as well as the manuscript's internal consistency. They made all final decisions and take full responsibility for the submitted manuscript. The tool was not used to generate research data or results and is not listed as an author.

A Derivation of Posterior Distributions

This appendix provides an algebraic derivation of the posterior distributions stated in Section 3.1.

Suppose the decision maker's current pre-feedback subjective belief over possible values of the fixed latent ability α is $\text{Beta}(a, b)$. Using θ to index a generic possible value of the latent ability, this belief has density

$$f(\theta | a, b) = \frac{1}{B(a, b)} \theta^{a-1} (1 - \theta)^{b-1}.$$

Conditional on the fixed latent ability α , the task outcome follows a Bernoulli distribution:

$$s_t | \alpha \sim \text{Bernoulli}(\alpha),$$

or equivalently,

$$P(s_t = 1 | \alpha) = \alpha, \quad P(s_t = 0 | \alpha) = 1 - \alpha.$$

The decision maker knows this Bernoulli probability model but not the fixed value of α . For a fixed observed outcome s_t , these conditional probabilities, viewed as functions of a possible value θ of the latent ability, define the Bernoulli likelihood:

$$L(\theta; s_t) = P(s_t | \alpha = \theta) = \theta^{s_t} (1 - \theta)^{1-s_t}.$$

In particular,

$$L(\theta; 1) = \theta, \quad L(\theta; 0) = 1 - \theta.$$

Correct answer ($s_t = 1$). After a correct outcome is observed, Bayes' rule gives

$$f(\theta | s_t = 1) = \frac{P(s_t = 1 | \alpha = \theta) f(\theta | a, b)}{\int_0^1 P(s_t = 1 | \alpha = \theta) f(\theta | a, b) d\theta}.$$

The denominator is the decision maker's subjective predictive probability of a correct outcome:

$$P^\mu(s_t = 1) = \int_0^1 \theta f(\theta | a, b) d\theta = \mathbb{E}_\mu[\theta] = \frac{a}{a+b}.$$

Substituting gives

$$f(\theta | s_t = 1) = \frac{a+b}{a} \frac{1}{B(a, b)} \theta^a (1-\theta)^{b-1},$$

which is the $\text{Beta}(a+1, b)$ distribution, using the identity $B(a+1, b) = \frac{a}{a+b} B(a, b)$.⁵

Incorrect answer ($s_t = 0$). Similarly, the denominator is the decision maker's subjective predictive probability of an incorrect outcome:

$$P^\mu(s_t = 0) = \int_0^1 (1-\theta) f(\theta | a, b) d\theta = \mathbb{E}_\mu[1-\theta] = \frac{b}{a+b},$$

yielding

$$f(\theta | s_t = 0) = \frac{a+b}{b} \frac{1}{B(a, b)} \theta^{a-1} (1-\theta)^b,$$

which is the $\text{Beta}(a, b+1)$ distribution.

⁵This identity follows from $B(a, b) = \Gamma(a)\Gamma(b)/\Gamma(a+b)$ and $\Gamma(a+1) = a\Gamma(a)$.

B Derivation of the Saliency-Adjusted Prospective Ego Utility and $V(x)$

Notation. Throughout this appendix, x denotes the task-specific self-evaluation point for the current task, while α denotes the fixed latent ability and θ indexes the possible values of the latent ability over which the decision maker's subjective belief is defined.

By integrating the normalized kernel defined in (5), we obtain the saliency-adjusted prospective ego utility from seeking:

$$\frac{a+b+1}{a+b+2} \frac{1}{B(a,b)} \left\{ \frac{a+b}{a} B_e(x; a+2, b) + \frac{a+b}{b} B_e(x; a, b+2) \right\}. \quad (\text{B.1})$$

From the definition of the regularized incomplete Beta function,

$$B_e(x; a, b) = I_x(a, b) B(a, b),$$

where $I_x(a, b)$ is the regularized incomplete Beta function. We can derive the relationship between $B_e(x; a+2, b)$ and $B(a, b)$ as follows:

$$\begin{aligned} B_e(x; a+2, b) &= I_x(a+2, b) B(a+2, b) \\ &= \left[I_x(a+1, b) - \frac{x^{a+1}(1-x)^b}{(a+1)B(a+1, b)} \right] B(a+2, b) \\ &= \left[I_x(a, b) - \frac{x^a(1-x)^b}{aB(a, b)} - \frac{x^{a+1}(1-x)^b}{(a+1)B(a+1, b)} \right] B(a+2, b) \\ &= \left[\frac{B_e(x; a, b)}{B(a, b)} - \frac{x^a(1-x)^b}{aB(a, b)} - \frac{x^{a+1}(1-x)^b}{(a+1)B(a+1, b)} \right] B(a+2, b). \end{aligned}$$

Since

$$B(a+1, b) = \frac{a}{a+b} B(a, b),$$

we have

$$B(a+2, b) = \frac{a+1}{a+b+1} B(a+1, b) = \frac{a(a+1)}{(a+b)(a+b+1)} B(a, b).$$

Thus, $B_e(x; a+2, b)$ can be rearranged as

$$\left[\frac{B_e(x; a, b)}{B(a, b)} - \frac{x^a(1-x)^b}{aB(a, b)} - \frac{(a+b)x^{a+1}(1-x)^b}{a(a+1)B(a, b)} \right] \frac{a(a+1)}{(a+b)(a+b+1)} B(a, b). \quad (\text{B.2})$$

Similarly,

$$\begin{aligned} B_e(x; a, b+2) &= I_x(a, b+2)B(a, b+2) \\ &= \left[I_x(a, b+1) + \frac{x^a(1-x)^{b+1}}{(b+1)B(a, b+1)} \right] B(a, b+2) \\ &= \left[I_x(a, b) + \frac{x^a(1-x)^b}{bB(a, b)} + \frac{x^a(1-x)^{b+1}}{(b+1)B(a, b+1)} \right] B(a, b+2) \\ &= \left[\frac{B_e(x; a, b)}{B(a, b)} + \frac{x^a(1-x)^b}{bB(a, b)} + \frac{x^a(1-x)^{b+1}}{(b+1)B(a, b+1)} \right] B(a, b+2). \end{aligned}$$

With

$$B(a, b+1) = \frac{b}{a+b} B(a, b),$$

and

$$B(a, b+2) = \frac{b+1}{a+b+1} B(a, b+1) = \frac{b(b+1)}{(a+b)(a+b+1)} B(a, b),$$

we obtain

$$\left[\frac{B_e(x; a, b)}{B(a, b)} + \frac{x^a(1-x)^b}{bB(a, b)} + \frac{(a+b)x^a(1-x)^{b+1}}{b(b+1)B(a, b)} \right] \frac{b(b+1)}{(a+b)(a+b+1)} B(a, b). \quad (\text{B.3})$$

According to (B.2) and (B.3), the salience-adjusted expression (B.1) can be rearranged as

$$\begin{aligned} & \frac{1}{a+b+2} \left[\frac{(a+1)B_e(x; a, b)}{B(a, b)} - \frac{(a+1)x^a(1-x)^b}{aB(a, b)} - \frac{(a+b)x^{a+1}(1-x)^b}{aB(a, b)} \right. \\ & \quad \left. + \frac{(b+1)B_e(x; a, b)}{B(a, b)} + \frac{(b+1)x^a(1-x)^b}{bB(a, b)} + \frac{(a+b)x^a(1-x)^{b+1}}{bB(a, b)} \right] \\ &= \frac{B_e(x; a, b)}{B(a, b)} + \frac{x^a(1-x)^b}{(a+b+2)B(a, b)} \left[\frac{b+1}{b} + \frac{a+b}{b}(1-x) - \frac{a+1}{a} - \frac{a+b}{a}x \right] \\ &= \frac{B_e(x; a, b)}{B(a, b)} + \frac{x^a(1-x)^b}{(a+b+2)B(a, b)} \left[\frac{a-b}{ab} + \frac{a+b}{b}(1-x) - \frac{a+b}{a}x \right]. \end{aligned}$$

Therefore, we derive the following expression for $F_{\text{seek}}(x)$:

$$\frac{1}{B(a, b)}B_e(x; a, b) + \frac{1}{(a + b + 2)B(a, b)}x^a(1 - x)^b \left\{ \frac{a - b}{ab} + \frac{a + b}{b}(1 - x) - \frac{a + b}{a}x \right\}. \quad (\text{B.4})$$

Subtracting $F_{\text{avoid}}(x) = F(x \mid a, b) = B_e(x; a, b)/B(a, b)$ from (B.4) gives

$$V(x) = F_{\text{seek}}(x) - F_{\text{avoid}}(x) = \frac{x^a(1 - x)^b}{(a + b + 2)B(a, b)} \left\{ \frac{a - b}{ab} + \frac{a + b}{b}(1 - x) - \frac{a + b}{a}x \right\}. \quad (\text{B.5})$$

Since $x^a(1 - x)^b > 0$ for $x \in (0, 1)$, the sign of $V(x)$ is determined solely by the expression in braces. Setting that expression to zero and multiplying through by $ab/(a + b)$ yields

$$\frac{a - b}{a + b} + a(1 - x) - bx = 0 \implies x = \frac{a - b + a(a + b)}{(a + b)^2} = \frac{a^2 + ab + a - b}{(a + b)^2} \equiv C.$$

Equation (B.5) provides a closed-form expression for $V(x)$, and the unique interior root of $V(x) = 0$ on $(0, 1)$ is $x = C$.

C A Sufficient Condition for Comparing Feedback-Choice Frequencies

This appendix formalizes the conditional frequency comparison discussed in Subsection 3.3, *Comparative Statics: Threshold Position and Belief Structure*. Consider two otherwise comparable groups with mirrored beliefs. Let O denote the optimistic-belief group, whose members hold $\text{Beta}(a, b)$, and let P denote the pessimistic-belief group, whose members hold $\text{Beta}(b, a)$, where $a > b > 1$.

The belief modes and feedback-seeking thresholds for the two groups are defined by

$$\begin{aligned}\theta_{\text{mode},O} &= \frac{a-1}{a+b-2}, & C_O &= C(a, b), \\ \theta_{\text{mode},P} &= \frac{b-1}{a+b-2}, & C_P &= C(b, a).\end{aligned}$$

Define

$$\delta := \frac{2(a-b)}{(a+b)^2(a+b-2)} > 0. \quad (\text{C.1})$$

Applying equation (10) to the two groups gives

$$\begin{aligned}C_O &= \theta_{\text{mode},O} - \delta, \\ C_P &= \theta_{\text{mode},P} + \delta.\end{aligned} \quad (\text{C.2})$$

For each group $j \in \{O, P\}$, define the mode-relative self-evaluation and its cumulative distribution function by

$$Z_j := x_j - \theta_{\text{mode},j}, \quad H_j(z) := \Pr(Z_j \leq z).$$

Assumption 2 (Directional ordering of mode-relative self-evaluation). *The random variables Z_O and Z_P have continuous cumulative distribution functions, and $x_j \in (0, 1)$ almost*

surely for $j \in \{O, P\}$. Moreover,

$$H_P(z) \geq H_O(z) \quad \text{for every } z \in \mathbb{R}, \quad (\text{C.3})$$

$$H_O(-\delta) < H_O(\delta). \quad (\text{C.4})$$

Condition (C.3) states that Z_O weakly first-order stochastically dominates Z_P . Equivalently, the distribution of mode-relative self-evaluation in the pessimistic-belief group is stochastically no higher than that in the optimistic-belief group. Under continuity, condition (C.4) is equivalent to

$$\Pr(-\delta < Z_O < \delta) > 0.$$

Thus, the distribution of mode-relative self-evaluation in the optimistic-belief group assigns positive probability to the interval between the two relevant cutoffs.

The assumption is an explicit sufficient distributional restriction. It does not require every member of the pessimistic-belief group to evaluate herself below the belief mode or every member of the optimistic-belief group to evaluate herself above it.

This stochastic-order restriction is neither implied by nor equivalent to the qualitative behavioral hypothesis stated in the main text. Rather, it provides one sufficient distributional condition consistent with the broader idea that mode-relative self-evaluations tend to be lower under pessimistic beliefs than under optimistic beliefs.

Proposition 5 (Feedback-choice frequencies under directional self-evaluation). *Under Assumption 2 and the baseline feedback-choice rule, the pessimistic-belief group seeks feedback more often than the mirrored optimistic-belief group:*

$$\Pr(\text{seek} \mid P) > \Pr(\text{seek} \mid O). \quad (\text{C.5})$$

Equivalently, the optimistic-belief group avoids feedback more often:

$$\Pr(\text{avoid} \mid O) > \Pr(\text{avoid} \mid P). \quad (\text{C.6})$$

Proof. For the optimistic-belief group, equation (C.2) and the baseline feedback-choice rule imply

$$\begin{aligned} \Pr(\text{seek} \mid O) &= \Pr(x_O < C_O) \\ &= \Pr(x_O - \theta_{\text{mode},O} < -\delta) \\ &= \Pr(Z_O < -\delta) \\ &= H_O(-\delta). \end{aligned}$$

Similarly, for the pessimistic-belief group,

$$\begin{aligned} \Pr(\text{seek} \mid P) &= \Pr(x_P < C_P) \\ &= \Pr(x_P - \theta_{\text{mode},P} < \delta) \\ &= \Pr(Z_P < \delta) \\ &= H_P(\delta). \end{aligned}$$

The final equalities in the two displays use the continuity of H_O and H_P , respectively.

Assumption 2 therefore gives

$$\Pr(\text{seek} \mid P) = H_P(\delta) \geq H_O(\delta) > H_O(-\delta) = \Pr(\text{seek} \mid O).$$

The comparison of avoidance probabilities follows because seeking and avoidance are complementary actions under the baseline feedback-choice rule. $\square$

The proposition is deliberately conditional. Without a restriction on the distributions of x_O and x_P relative to their respective belief modes, the threshold-position result alone does not rank population feedback-seeking frequencies.

D Proof of Proposition 3

This appendix proves Proposition 3. We first derive the updated threshold after accumulated positive and negative feedback observations and then establish its monotonicity and limiting behavior.

After m positive feedback observations and n negative feedback observations, Beta-Bernoulli conjugacy implies that the decision maker's posterior belief distribution is $\text{Beta}(a + m, b + n)$. Since the threshold is a function of the Beta parameters, the updated threshold is obtained by substituting $a + m$ and $b + n$ into the formula for $C(a, b)$:

$$C_{m,n} = C(a + m, b + n) = \frac{(a + m) - (b + n) + (a + m)^2 + (a + m)(b + n)}{(a + b + m + n)^2}. \quad (\text{D.1})$$

For notational convenience, define

$$A := a + m, \quad B := b + n, \quad S := A + B = a + b + m + n.$$

Because $a, b > 1$ and $m, n \geq 0$, we have $A > 1$ and $B > 1$. With this notation,

$$C_{m,n} = \frac{A - B + A^2 + AB}{(A + B)^2} = \frac{A(S + 1) - B}{S^2}. \quad (\text{D.2})$$

Lemma 1 (Monotonicity of the threshold). *For every fixed n , $C_{m,n}$ is strictly increasing in m . For every fixed m , $C_{m,n}$ is strictly decreasing in n .*

Proof. Although m and n are integers in the model, extend $C_{m,n}$ to a differentiable function of real variables $m, n \geq 0$. Since $A = a + m$ and $B = b + n$, differentiating with respect to m is equivalent to differentiating with respect to A , and differentiating with respect to n is equivalent to differentiating with respect to B .

Let

$$C(A, B) = \frac{A - B + A^2 + AB}{(A + B)^2} = \frac{A(S + 1) - B}{S^2}, \quad S = A + B.$$

First consider the effect of increasing m , equivalently A . Let

$$N(A, B) = A - B + A^2 + AB = A(S + 1) - B, \quad D(A, B) = S^2.$$

Then

$$\frac{\partial N}{\partial A} = 2A + B + 1, \quad \frac{\partial D}{\partial A} = 2S.$$

By the quotient rule,

$$\frac{\partial C}{\partial A} = \frac{(2A + B + 1)S^2 - 2SN}{S^4} = \frac{(2A + B + 1)S - 2N}{S^3}.$$

Using $N = A(S + 1) - B$, the numerator becomes

$$(2A + B + 1)S - 2\{A(S + 1) - B\} = (B + 1)S + 2B - 2A.$$

Since $S = A + B$, this can be rewritten as

$$(B + 1)(A + B) + 2B - 2A = A(B - 1) + B(B + 3).$$

Because $A > 1$ and $B > 1$, this expression is strictly positive. Hence

$$\frac{\partial C}{\partial A} > 0.$$

Therefore, for every fixed n , $C_{m,n}$ is strictly increasing in m .

Next consider the effect of increasing n , equivalently B . We have

$$\frac{\partial N}{\partial B} = A - 1, \quad \frac{\partial D}{\partial B} = 2S.$$

Thus,

$$\frac{\partial C}{\partial B} = \frac{(A - 1)S^2 - 2SN}{S^4} = \frac{(A - 1)S - 2N}{S^3}.$$

Again using $N = A(S + 1) - B$, the numerator is

$$(A - 1)S - 2\{A(S + 1) - B\} = -(A + 1)S - 2A + 2B.$$

Since $S = A + B$, this becomes

$$-(A + 1)(A + B) - 2A + 2B = -A^2 - 3A - B(A - 1).$$

Because $A > 1$ and $B > 1$, this expression is strictly negative. Hence

$$\frac{\partial C}{\partial B} < 0.$$

Therefore, for every fixed m , $C_{m,n}$ is strictly decreasing in n . Since the continuous extension is strictly increasing in m and strictly decreasing in n , the same monotonicity holds for integer changes in m and n . $\square$

Lemma 2 (Limits of the threshold). *For every fixed n ,*

$$\lim_{m \rightarrow \infty} C_{m,n} = 1.$$

For every fixed m ,

$$\lim_{n \rightarrow \infty} C_{m,n} = 0.$$

Proof. Using (D.2), write

$$C_{m,n} = \frac{A(S + 1) - B}{S^2} = \frac{A}{S} \left(1 + \frac{1}{S}\right) - \frac{B}{S^2}.$$

First fix n . Then $B = b + n$ is fixed, while $A = a + m \rightarrow \infty$ and $S = A + B \rightarrow \infty$ as $m \rightarrow \infty$. Hence

$$\frac{A}{S} \rightarrow 1, \quad \frac{1}{S} \rightarrow 0, \quad \frac{B}{S^2} \rightarrow 0.$$

Therefore, $\lim_{m \rightarrow \infty} C_{m,n} = 1$.

Next fix m . Then $A = a + m$ is fixed, while $B = b + n \rightarrow \infty$ and $S = A + B \rightarrow \infty$ as

$n \rightarrow \infty$. Hence

$$\frac{A}{S} \rightarrow 0, \quad \frac{1}{S} \rightarrow 0, \quad \frac{B}{S^2} \rightarrow 0.$$

Therefore, $\lim_{n \rightarrow \infty} C_{m,n} = 0$. □

Combining Lemmas 1 and 2 proves Proposition 3.

E Proof of Proposition 4

This appendix proves Proposition 4. The proof uses the monotonicity and limiting properties of the threshold established in Appendix D.

Let $x \in (0, 1)$ be fixed, and let $C_{m,n}$ denote the threshold after m observed positive feedback outcomes and n observed negative feedback outcomes.

Part (i): Observed negative-feedback reversal (seeking $\rightarrow$ avoidance).

Suppose

$$x \in (0, C_{m,n}).$$

Then feedback is initially sought. By Lemma 1, $C_{m,n}$ is strictly decreasing in n , holding m fixed. By Lemma 2,

$$\lim_{k \rightarrow \infty} C_{m,n+k} = 0.$$

Since $x > 0$, there exists a finite integer $k \geq 1$ such that

$$C_{m,n+k} \leq x.$$

Let

$$k^* := \min\{k \geq 1 : C_{m,n+k} \leq x\}.$$

Then $C_{m,n+k^*} \leq x$. Thus, after k^* additional observed negative feedback outcomes, the evaluation point lies weakly above the threshold:

$$x \geq C_{m,n+k^*}.$$

By the feedback-seeking rule, feedback avoidance becomes optimal. This proves Part (i).

Part (ii): Positive-signal reversal under external disclosure (avoidance $\rightarrow$ seeking).

Suppose

$$x \in [C_{m,n}, 1).$$

Then feedback is initially avoided. Holding the evaluation point fixed, the decision maker continues to decline voluntary feedback in the baseline environment unless new information changes the threshold. Therefore, an avoidance-to-seeking reversal through the threshold requires positive performance information about the same task class to arrive through an external channel.

Suppose that the decision maker subsequently observes additional positive performance outcomes from the same task class through such an external channel. By Lemma 1, $C_{m,n}$ is strictly increasing in m , holding n fixed. By Lemma 2,

$$\lim_{k \rightarrow \infty} C_{m+k,n} = 1.$$

Since $x < 1$, there exists a finite integer $k \geq 1$ such that

$$C_{m+k,n} > x.$$

Let

$$k^{**} := \min\{k \geq 1 : C_{m+k,n} > x\}.$$

Then $x < C_{m+k^{**},n}$. Thus, after k^{**} additional positive performance outcomes from the same task class observed through an external channel, the evaluation point lies below the threshold. By the feedback-seeking rule, feedback seeking becomes optimal. This proves Part (ii).

References

- Roland Bénabou and Luca Henkel. Identity as self-image. NBER Working Paper 34297, National Bureau of Economic Research, 2025.
- Roland Bénabou and Jean Tirole. Intrinsic and extrinsic motivation. *Review of Economic Studies*, 70(3):489–520, 2003. doi: 10.1111/1467-937X.00253.
- Roland Bénabou and Jean Tirole. Over my dead body: Bargaining and the price of dignity. *American Economic Review: Papers & Proceedings*, 99(2):459–465, 2009. doi: 10.1257/aer.99.2.459.
- Roland Bénabou and Jean Tirole. Mindful economics: The production, consumption, and value of beliefs. *Journal of Economic Perspectives*, 30(3):141–164, 2016. doi: 10.1257/jep.30.3.141.
- Pedro Bordalo, Nicola Gennaioli, and Andrei Shleifer. Salience theory of choice under risk. *The Quarterly Journal of Economics*, 127(3):1243–1285, 2012. doi: 10.1093/qje/qjs018.
- Gordon D. A. Brown, Jonathan Gardner, Andrew J. Oswald, and Jing Qian. Does wage rank affect employees’ well-being? *Industrial Relations*, 47(3):355–389, 2008. doi: 10.1111/j.1468-232X.2008.00525.x.
- Markus K. Brunnermeier and Jonathan A. Parker. Optimal expectations. *American Economic Review*, 95(4):1092–1118, 2005. doi: 10.1257/0002828054825493.
- Christopher J. Cagna, Jamil P. Bhanji, Da’Quallon Smith, Mauricio R. Delgado, and Elizabeth Tricomi. Decisions to seek cognitive performance feedback: Potential determinants of feedback value and consequences for learning. *Learning and Motivation*, 88:102051, 2024. doi: 10.1016/j.lmot.2024.102051.
- Andrew Caplin and John Leahy. Psychological expected utility theory and anti-

- patory feelings. *Quarterly Journal of Economics*, 116(1):55–79, 2001. doi: 10.1162/003355301556347.
- Andrew Caplin and John Leahy. The supply of information by a concerned expert. *Economic Journal*, 114(497):487–505, 2004. doi: 10.1111/j.0013-0133.2004.0228a.x.
- James J Chen and Melvin R Novick. On the use of a cumulative distribution as a utility function in educational or employment selection. *Journal of Educational Statistics*, 7(1):19–35, 1982. doi: 10.3102/10769986007001019.
- Christoph Drobner and Sebastian J. Goerg. Motivated belief updating and rationalization of information. *Management Science*, 70(7):4583–4592, 2024. doi: 10.1287/mnsc.2023.02537.
- Ward Edwards. Conservatism in human information processing. In Benjamin Kleinmuntz, editor, *Formal Representation of Human Judgment*, pages 17–52. John Wiley & Sons, New York, 1968.
- David Eil and Justin M. Rao. The good news–bad news effect: Asymmetric processing of objective information about yourself. *American Economic Journal: Microeconomics*, 3(2):114–138, 2011. doi: 10.1257/mic.3.2.114.
- Leon Festinger. *A Theory of Cognitive Dissonance*. Stanford University Press, Stanford, CA, 1957.
- Russell Golman, David Hagmann, and George Loewenstein. Information avoidance. *Journal of Economic Literature*, 55(1):96–135, 2017. doi: 10.1257/jel.20151245.
- Nathaniel Gregory. Relative wealth and risk taking: A short note on the Friedman-Savage utility function. *Journal of Political Economy*, 88(6):1226–1230, 1980. doi: 10.1086/260936.
- Mitchell Hoffman. How is information valued? Evidence from framed field experiments. *The Economic Journal*, 126(595):1884–1911, 2016. doi: 10.1111/ecoj.12401.

- Ed Hopkins and Tatiana Kornienko. Running to keep in the same place: Consumer choice as a game of status. *American Economic Review*, 94(4):1085–1107, 2004. doi: 10.1257/0002828042002705.
- Daniel Kahneman and Amos Tversky. Subjective probability: A judgment of representativeness. *Cognitive Psychology*, 3(3):430–454, 1972. doi: 10.1016/0010-0285(72)90016-3.
- Daniel Kahneman and Amos Tversky. On the psychology of prediction. *Psychological Review*, 80(4):237–251, 1973. doi: 10.1037/h0034747.
- Botond Köszegi. Ego utility, overconfidence, and task choice. *Journal of the European Economic Association*, 4(4):673–707, 2006. doi: 10.1162/JEEA.2006.4.4.673.
- Botond Köszegi and Matthew Rabin. A model of reference-dependent preferences. *The Quarterly Journal of Economics*, 121(4):1133–1165, 2006. doi: 10.1093/qje/121.4.1133.
- Botond Köszegi, George Loewenstein, and Takeshi Murooka. Fragile self-esteem. *The Review of Economic Studies*, 89(4):2026–2060, 2022. doi: 10.1093/restud/rdab060.
- Matthew Kovach and Elchin Suleymanov. Reference dependence and random attention. *Journal of Economic Behavior & Organization*, 215:421–441, 2023. doi: 10.1016/j.jebo.2023.09.020.
- Markus M Möbius, Muriel Niederle, Paul Niehaus, and Tanya S Rosenblat. Managing self-confidence: Theory and experimental evidence. *Management Science*, 68(11):7793–7817, 2022. doi: 10.1287/mnsc.2021.4294.
- Don A. Moore and Daylian M. Cain. Overconfidence and underconfidence: When and why people underestimate (and overestimate) the competition. *Organizational Behavior and Human Decision Processes*, 103(2):197–213, July 2007. doi: 10.1016/j.obhdp.2006.09.002.
- Don A Moore and Paul J Healy. The trouble with overconfidence. *Psychological Review*, 115(2):502–517, 2008. doi: 10.1037/0033-295X.115.2.502.

- Raymond S. Nickerson. Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology*, 2(2):175–220, 1998. doi: 10.1037/1089-2680.2.2.175.
- Sabine Pittnauer and Martin Hohnisch. Observing the creation of new knowledge in the economics laboratory—do participants discover how to learn from outcome feedback in a dynamic decision problem? *Journal of Economic Behavior & Organization*, 215:382–405, 2023. doi: 10.1016/j.jebo.2023.09.015.
- Drazen Prelec and Ronit Bodner. Self-signaling and self-control. In George Loewenstein, Daniel Read, and Roy F. Baumeister, editors, *Time and Decision: Economic and Psychological Perspectives on Intertemporal Choice*, pages 277–298. Russell Sage Foundation, New York, 2003.
- Neil Stewart, Nick Chater, and Gordon D. A. Brown. Decision by sampling. *Cognitive Psychology*, 53(1):1–26, 2006. doi: 10.1016/j.cogpsych.2005.10.003.
- Kate Sweeny, David Melnyk, William Miller, and James A. Shepperd. Information avoidance: Who, what, when, and why. *Review of General Psychology*, 14(4):340–353, 2010. doi: 10.1037/a0021288.
- Amos Tversky and Daniel Kahneman. Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157):1124–1131, 1974. doi: 10.1126/science.185.4157.1124.
- Hélène Van Marcke, Pierre Le Denmat, Tom Verguts, and Kobe Desender. Manipulating prior beliefs causally induces under- and overconfidence. *Psychological Science*, 35(4): 398–409, 2024. doi: 10.1177/09567976241231572.
- Bernard MS Van Praag and Arie Kapteyn. Further evidence on the individual welfare function of income: An empirical investigation in the Netherlands. *European Economic Review*, 4(1):33–62, 1973. doi: 10.1016/0014-2921(73)90029-9.