



Article scientifique

Article

2026

Published version

Open Access

This is the published version of the publication, made available in accordance with the publisher's policy.

Feedback-induced attitudinal changes in risk preferences

Nasioulas, Antonios; Potier, Elise; Cerrotti, Fabien; Lebreton, Mael; Palminteri, Stefano

How to cite

NASIOULAS, Antonios et al. Feedback-induced attitudinal changes in risk preferences. In: Nature communications, 2026, vol. 17, n° 1. doi: 10.1038/s41467-025-67729-x

This publication URL: <https://archive-ouverte.unige.ch/unige:191387>

Publication DOI: [10.1038/s41467-025-67729-x](https://doi.org/10.1038/s41467-025-67729-x)

Feedback-induced attitudinal changes in risk preferences

Received: 7 March 2024

Accepted: 5 December 2025

Cite this article as: Nasioulas, A., Potier, E., Cerrotti, F. *et al.* Feedback-induced attitudinal changes in risk preferences. *Nat Commun* (2026). <https://doi.org/10.1038/s41467-025-67729-x>

Antonios Nasioulas, Elise Potier, Fabien Cerrotti, Maël Lebreton & Stefano Palminteri

We are providing an unedited version of this manuscript to give early access to its findings. Before final publication, the manuscript will undergo further editing. Please note there may be errors present which affect the content, and all legal disclaimers apply.

If this paper is publishing under a Transparent Peer Review model then Peer Review reports will publish with the final article.

Title

Feedback-induced attitudinal changes in risk preferences

Author list

Antonios Nasioulas(1,2,*), Elise Potier(3), Fabien Cerrotti(1,4), Maël Lebreton(2,5) & Stefano Palminteri(1,4,*).

Affiliations

¹ Département d'Études Cognitives, Ecole Normale Supérieure, Université de Recherche Paris Sciences et Lettres, 24 Rue Lhomond, 75005, Paris, France

² Paris-Jourdan Sciences Économiques UMR8545, CNRS & Paris School of Economics, 48 Bd Jourdan, 75014 Paris, France

³ Direction interministérielle de la transformation publique, 20 avenue de Segur, 75007 Paris, France

⁴ Laboratoire de Neurosciences Cognitives et Computationnelles, Institut National de la Santé et Recherche Médicale, 24 Rue Lhomond, 75005, Paris, France

⁵ Swiss Center for Affective Science, University of Geneva, Campus Biotech, Chem. des Mines 9, 1202 Geneva, Switzerland

***Corresponding authors antonios.nasioulas@ens.fr, stefano.palminteri@ens.fr**

Keywords

decision-making, risk, feedback, experience description gap

Abstract

Decision-making under risk is often studied with fully described lotteries, where normative theory predicts that post-choice outcome disclosure (feedback) should not influence preferences. However, previous empirical work has generally shown that feedback does affect risk-taking, yet, without reaching a consensus on the consequences of feedback or the underlying cognitive mechanisms. Here, across seven behavioral experiments, we disentangle two competing accounts: the learning hypothesis, where feedback alters subjective values through experience, and the attitudinal hypothesis, where feedback changes preferences in anticipation of outcomes. We find that feedback does not improve maximization but consistently increases risk-taking. Fine-grained temporal analyses reveal that this effect emerges before any outcomes are experienced, ruling out learning as the primary driver. Moreover, the increase of risk-taking in partial feedback seems to be driven by curiosity, while in complete feedback by anticipated regret. Our results indicate that feedback can bias decision-making primarily through attitudinal rather than learning mechanisms.

Introduction

Traditionally, empirical investigations of decision-making under risk have mostly been carried out in behavioral setups limited to one-shot description-based choice problems^{1,2}: unique binary choices between mutually exclusive probabilistic options (lotteries), where relevant information (i.e., prospective outcomes and probabilities) is explicitly displayed, and considered known to the decision maker. This experimental setup matches the scope and limits of both normative and descriptive decision theories, which are generally silent about the effects of feedback and choice repetition³⁻⁵. Arguably, although both theoretically and empirically convenient, this one-shot description-based framework is not representative of the vast majority of decision situations that one faces every day. Most decision problems are recurrent, and, very often, one gets to know the outcome of one's choice (partial feedback) – and sometimes also the outcome of the forgone option (complete feedback)⁶⁻⁹. To address those shortcomings, repetitions and feedback have been gradually incorporated into the study of human decision-making under risk over the last couple of decades¹⁰⁻¹³. This experimental innovation revealed that, in contrast to the normative dictate, human choices and risk preferences elicited in repeated decisions under risk do appear to change depending on the presence versus absence of feedback.

A widespread and intuitive hypothesis concerning the effect of feedback on risk preferences proposes that outcome information modifies the decision-maker's subjective representation of probabilities. Indeed, from studies involving one-shot decisions, it appears clearly that individuals behave as if their subjective representation of probabilities is distorted (overweighting of rare events, underweighting of common events^{2,14}). In the presence of feedback, the realized frequency of the outcomes received in the context of repeated decisions can be used to update (if not correct) the subjective beliefs concerning their probabilities, ultimately affecting one's preferences and choices¹⁵. We shall refer to this category of accounts as the learning hypothesis. Because the integration of feedback in future choices is supposed to correct originally distorted subjective probabilities, the learning hypothesis often assumes that the presence of feedback should correct representational biases and, as a consequence, promote optimal (i.e., expected value maximizing) choices^{7,16}. Admittedly, though, these simple predictions from the learning hypothesis can be challenged, e.g., by the presence of learning biases^{17,18}, or when the probabilities of the outcomes

are extreme and the options are not sufficiently sampled: two conditions that lead the experienced and the actual frequency of the outcomes to diverge¹⁹.

The notion that feedback should enhance optimal decision-making goes beyond academic circles and has been widely proposed as a strategy to debias individuals and improve their decisions in significant applied contexts, such as finance and healthcare^{20–22}. Feedback is, in fact, a fundamental element of “boosting”—an approach distinct from “nudging”—which aims to create a more enduring positive impact on decision-making performance by incorporating a learning component^{23–25}.

A relatively recent instantiation of the learning hypothesis is embedded in the BEAST model (Best Estimate And Sampling Tool), which becomes particularly influential after emerging victorious from a choice prediction competition, notably featuring choices between fully described lotteries followed by complete feedback²⁶. The model, rooted in the decision-by-sampling tradition²⁷, postulates that as soon as feedback is available, the subjective estimations of decision variables (outcomes are probabilities) are based on samples drawn from the past experienced outcomes. The BEAST is currently considered the reference against which new (and old) models of decision-making under risk should be compared, as testified by numerous follow-up studies^{28–31}.

An alternative category of accounts, which we shall refer to as the attitudinal hypothesis, conjectures that the mere presence of feedback changes the decision-maker’s preferences, independently of any learning process. We identify in the literature two main candidate cognitive processes for the attitudinal hypothesis: epistemic curiosity and anticipated regret.

Epistemic curiosity pertains to the idea of gaining utility from the resolution of uncertainty^{32–36}. Indeed, when feedback is available, some options acquire different informational values with respect to the resolution of uncertainty. For instance, when only the outcome of the chosen option is revealed (partial feedback) and the decision features a riskier (high variance) versus a safer (low variance) lottery, choosing the safer lottery resolves less uncertainty about the final state of the world. In other words, there is an extra informational incentive to choose the riskier option if partial feedback is provided. This informational asymmetry explains how curiosity –or an uncertainty minimization drive – may shift choices in favor of the risky options if the participant anticipates that the decision will be followed by feedback^{37,38}.

Regret can also cause attitudinal effects of feedback when it is expressed as an anticipated emotion during decision-making³⁹. The rationale is that, when considering choosing a safe lottery over a risky one, one might forecast the regret elicited by observing a positive resolution from the best alternative outcome (the unchosen, riskier lottery)⁴⁰, and therefore make more risk-seeking choices to avert this possibility. As opposed to epistemic curiosity, anticipated regret should notably emerge when the outcome of the unchosen option is also available (complete feedback). In support of the prominent role of counterfactual emotions in decision-making, regret has recently been invoked to explain certain aspects of choices elicited in paradigms that couple descriptions with feedback^{8,9,26,41}.

Critically, in contrast to the learning hypothesis, which supposes that the effect of the feedback should emerge gradually, the attitudinal hypothesis conceives that it can emerge even before any feedback is experienced – i.e., by anticipation. Because the two hypotheses differ in the temporal relation with respect to the time of choice and outcomes, fine-grained temporal dynamics can dissociate the two accounts: attitudinal effects precede choices while learning effects follow outcomes. In addition, the two hypotheses are also different in their relation to choice optimality. Learning mechanisms are generally supposed to correct representational biases and, accordingly, should increase expected value maximization. Attitudinal ones are more silent in that respect.

The goal of the present study is to evaluate the relative merit of these alternative hypotheses in seven newly conducted experiments (N=540) and a reanalysis of a previous dataset (N=446). The new experiments have been designed to systematically investigate the role of feedback in decision-making under risk across different experimental manipulations that allow discriminating competing accounts.

Here, we show that the presence of feedback reliably increases risk-taking without improving expected value maximization. Temporal analyses reveal that this effect is attitudinal, emerging before any outcomes are experienced, and is therefore not driven by learning mechanisms. Moreover, the underlying psychological drivers differ across feedback regimens: curiosity appears to dominate under partial feedback, whereas anticipated regret plays a key role under complete feedback. Finally, we demonstrate that these behavioral patterns are not well captured by the influential BEAST model, challenging its validity as a general account of decision-making under risk.

Results

Experimental design

To address our research questions, we ran a series of seven incentivized experiments (six online - $N=100$ for each experiment before the application of strict exclusion criteria – and one in the laboratory - $N=30$; see Methods). The six online experiments were variants of the experimental paradigm that we will describe below.

Our experimental design allowed us to overcome several methodological shortcomings identified in previous literature, which overall delivered an unclear picture concerning the directionality of the effect of feedback in decision-making under risk (see Supplementary Note 1 for a quantitative survey of the literature). One specificity of our experimental design is using a factorial design with a within-subject manipulation of post-choice feedback (present or absent); feedback was treated as a between-subjects factor in most of the previous studies, with only one exception²⁶ - our design improves over Erev et al., (2017) by equalizing the number of repetitions for the two conditions (their design featured five repetitions for no-feedback and twenty for feedback) and by counterbalancing the order of the conditions. Furthermore, unlike most of the previous experiments, we included the same number of trials (10) in both feedback conditions, to disentangle between the effect of mere repetition and the effect of feedback itself⁴². Trials featuring the same decision problem were clustered in blocks of 10 trials. Feedback and no-feedback blocks were randomly interspersed (Figure 1A).

We used a binary choice task featuring a sure and a risky option in each trial. The risky option had the form of $(m, p; 0, 1-p)$, namely giving m points with probability p and zero points otherwise. In addition to feedback (present or absent), we also factorially manipulated the option optimality, i.e., whether the risky or the sure option has a higher expected value (EV): in one condition, the risky option maximizes the EV; in the other condition, the sure option maximizes EV (Figure 1B). This allowed us to orthogonalize risk preference and decision optimality –two features that have often been confounded in the literature. Finally, we manipulated within-subject the probability of (positive) gain associated with the risky option (three levels, namely 10%, 50%, and 90%), and the magnitude of the risky option (two levels, 40 and 60 points), thus leading to a decision space of 12 unique decision problems. Together with our feedback/no feedback manipulation and the

repetitions ($\times 10$ per decision context), our final factorial design comprised 240 choices per participant, which is larger than that commonly found in the literature.

The role of feedback and instructions on risk preferences

The first experiment (Exp.1) featured partial feedback, i.e., we revealed only the outcome associated with the chosen option. Participants were not informed about the presence or absence of the feedback before starting a given block. Dependent variables, i.e., the propensity to choose the risky option (R-rate) or the optimal – EV maximizing – option (O-rate) were analyzed using a generalized linear mixed model (GLMM) with the task factors (including presence of feedback and option optimality) as independent variables (see Methods for more details).

Our analyses of the R-rate identified a significant main effect of feedback (logistic GLMM: $z = 5.13$, $P < 0.001$, OR = 1.86, 95% CI [1.47, 2.35], see Table 1), which was characterized by an increased propensity to choose the risky option when trial-by-trial feedback was present (Figure 2A; Figure 1C, Exp1). Interestingly, this increase was of the same size and direction both when the risky option was better and when the sure option was better (with respect to EV) (Bayesian paired t-test $BF_{01} = 8.0$, Bayesian sign test $BF_{01} = 5.3$, indicating moderate evidence in favor of the null; Figure 1D, Exp1), such that there was no effect of feedback on the optimal choice rate (logistic regression: $P = 0.78$, see Table 1; Bayesian paired t-test $BF_{01} = 8.0$; Bayesian sign test $BF_{01} = 5.7$; Figure 1C, Figure 2B). In other terms, there was a significant interaction between feedback and option optimality: feedback decreased the O-rate when the sure option was the EV-maximizing one ($z = -6.45$, $P < 0.001$, OR = 0.5, 95% CI [0.4, 0.62]), but increased it otherwise ($z = 5.01$, $P < 0.001$, OR = 1.67, 95% CI [1.37, 2.05], see Table 1).

Finally, we examined the trial-by-trial unfolding of the main effect of feedback on R-rate. The learning hypothesis predicts that the effect of feedback should be absent at the first trial, then gradually emerge and increase after repeated experience with feedback. Our analyses revealed a slightly different pattern: while indeed no significant effect of feedback could be detected in the first trial, R-rate in feedback and no-feedback conditions abruptly diverged in the second trial and the difference remained significant until the end of the block (two-tailed Wilcoxon Signed-Rank Tests (Benjamini–Hochberg adjusted): Trial 1: $z = 1.54$, $P = 0.124$, rank-biserial $r = 0.23$, 95 % CI [-0.06, 0.51]; Trials 2–10: all pairwise tests are significant, with the most conservative

comparison (Trial 10) yielding $z = 3.16$, $P = 0.002$, $r = 0.47$, 95 % CI [0.20, 0.70], $n = 80$; Trial 1: $BF_{01} \approx 2-4$; Figure 2C).

Overall, the results of this first experiment appeared, at the macroscopic level, in line with most of the existing literature, generally showing an increase in R-rate in the presence of feedback. At a finer grain level, because the effects develop after the first feedback, they seem overall consistent with a learning effect. The fact that the effect is abrupt rather than gradual could be understood as one-shot learning. However, because participants in Exp.1 started each block without knowing whether they would receive feedback or not, the first trial also implicitly but unambiguously informed them about the presence of feedback in the ongoing block (feedback or no-feedback), which may have altered their attitude toward risky options. Thus, although the separation of R-rates in the second trial can be a result of (one-shot) learning, it can also reflect the triggering of an attitudinal change. Of note, also against the learning hypothesis is the fact that the presence of feedback did not improve the optimal choice rate.

To disentangle these two possibilities, we ran a second experiment (Exp.2) in which, at the beginning of each block, participants received explicit instructions (hence block instructions) mentioning whether they would receive post-choice feedback in the upcoming block or not. Everything else was kept the same as in Exp.1. At the aggregate level, Exp.2 replicated Exp.1 in all respects (logistic GLMM for R-rate: $z = 7.88$, $P < 0.001$, OR = 2.06, 95% CI [1.72, 2.46], see Table 1; O-rate, F vs nF: Bayesian paired t-test $BF_{01} = 7.1$; sign test $BF_{01} = 7.2$; Figure 1C and 1D; Figure 2D and 2E). Yet, the between-experiment manipulation of the block instructions produced a significant difference in the effect of feedback on the first-trial R-rates (two-tailed Mann-Whitney U test: $U = 2880.5$, $z = -2.77$, $P = 0.0056$, $r = -0.24$, 95% CI [-0.40, -0.07]; inset of Figure 2F). Actually, the difference in R-rates between feedback and no-feedback blocks in Exp.2 arose from the very first trial -and remained significant for the rest of the block (two-tailed WSRTs (BH-adjusted): Trials 1-10, most conservative at Trial 8: $z = 3.33$, $P < 0.001$, $r = 0.45$, 95 % CI [0.20, 0.67], $n = 95$; Figure 2F). Thus, it seems that the mere anticipation of feedback information induced by the block instructions was enough to change risk preferences before any feedback was actually experienced.

In summary, results from Exp.1 and Exp.2 clearly revealed that the presence of feedback about the outcome of the chosen lottery increased risk propensity but not choice optimality. Besides, while

Exp.1's results only superficially supported the learning hypothesis, the results following the introduction of explicit block instructions in Exp.2 favor the attitudinal hypothesis. Exp.1 and Exp.2 featured a partial feedback regimen and since the result of the sure option is always known by definition and the result of the unchosen option is not disclosed, choosing the risky option provides the participant with more information about the current state of the world. Thus, this pattern of results is consistent with an attitudinal effect created by epistemic curiosity, where the demand for uncertainty resolution increases risk propensity because of the informational asymmetry between the risky and sure lotteries.

Regret as an additional determinant of the effect of feedback on risk preferences

If epistemic curiosity is the only driver of the observed effect, the informational asymmetry between the sure option (whose result can be inferred with certainty even when its outcome is not revealed) and the risky one (whose result cannot be inferred when its outcome is not revealed) causes the increased risk-taking propensity in the presence of feedback. Thus, according to the epistemic curiosity account, the effect should vanish (or, at least, decrease) under a complete (or 'full') feedback regimen, i.e., when the forgone outcome of the unchosen option is additionally revealed. To test this hypothesis, we ran Exp.3 and Exp.4, which were analogous to Exp. 1 (without block instructions) and Exp.2 (with block instructions) except for the fact that they both featured complete feedback.

The complete-feedback experiments replicated the main effects observed in their partial feedback counterparts (Exp.1 and Exp.2). Most importantly, the presence of complete feedback still increased the R-rate (logistic GLMM: $z = 3.86$, $P < 0.001$, OR = 1.42, 95% CI [1.19, 1.70] (Exp3); $z = 6.34$, $P < 0.001$, OR = 1.97, 95% CI [1.60, 2.43] (Exp4), see Table 1; Figure 1C, Exp.3 and Exp.4). Concerning the O-rate, in Exp.3 we found evidence for no effect (Bayesian paired t-test $BF_{01} = 3.7$; Bayesian sign test $BF_{01} = 6.8$) while in Exp.4 results were mixed, with the available evidence, if anything, pointing towards a slight decrease rather than an increase in performance ($z = -1.49$, $P = 0.136$, OR = 0.92, 95% CI [0.82, 1.03], see Table 1; two-tailed WSRT: $z = -2.14$, $P = 0.032$, $r = -0.27$, 95% CI [-0.51, -0.02], $n = 85$).

The pattern of results at the level of the trial-by-trial dynamic was also replicated. In Exp.3 (without block instructions), the divergence induced by the presence versus absence of feedback was detectable from the third trial and remains significant for the rest of the block (two-tailed WSRTs

(BH-adjusted): Trial 1: $z = -1.69$, $P = 0.101$, $r = -0.24$, 95 % CI $[-0.53, 0.04]$; Trial 2: $z = 1.25$, $P = 0.210$, $r = 0.16$, 95 % CI $[-0.08, 0.41]$; Trials 3–10: all pairwise tests are significant with the most conservative comparison (Trial 7) yielding $z = 2.69$, $P = 0.009$, $r = 0.36$, 95 % CI $[0.11, 0.60]$, $n = 86$; Figure 3A). In Exp.4, (the one with block instructions), contrary to the idea of epistemic curiosity being the sole determinant of the change in risk propensity between feedback and no feedback conditions, we found an effect from the first trial. The R-rates in the feedback blocks are significantly higher than the no-feedback ones starting from the first trial and throughout the block (two-tailed WSRTs (BH-adjusted): Trials 1–10, most conservative at Trial 4: $z = 2.62$, $P = 0.009$, $r = 0.35$, 95 % CI $[0.10, 0.59]$; $n = 85$; Figure 3B). As in Exp.1 versus Exp.2, the between-experiment manipulation of the block instructions produced a significant difference in the effect of feedback on first-trial R-rates in Exp.3 versus Exp.4 (two-tailed MWU test: $U = 2661$, $z = -3.11$, $P = 0.0019$, $r = -0.27$, 95% CI $[-0.44, -0.11]$; inset of Figure 3B).

While these results are once again consistent with a feedback-induced attitudinal change in risk preference (the effect arose before any feedback was received), they are not easily accommodated by the epistemic curiosity account, because, under the complete feedback regimen, there is no uncertainty resolution utility bonus attributable to choosing the risky option. An alternative psychological mechanism for this effect that is compatible with the complete feedback scenario is the anticipated regret. To understand why anticipated regret could represent a possible explanation for this effect, it should be first noted that in many economic decision-making settings, regret is generally thought to be dependent on a comparison between the obtained and the forgone outcome and that this comparison is made explicit only in the complete feedback condition, where regret is experienced whenever the forgone outcome is higher than the obtained one. We propose that at the decision stage, the option that gives the best outcome most of the time gets a regret “premium”. Specifically, an option's regret premium depends on the probability that its outcome is greater than that of the alternative. Note that, according to this definition and given the decision problems in our design, the regret premium of the risky option should be monotonically increasing with its probability (of giving its positive outcome).

Thus, to assess the anticipated regret hypothesis, we evaluated the effect of feedback on R-rates as a function of the probability of the best-risky outcome (10%, 50% and 90%) and of the type of feedback (partial and complete). This analysis revealed a clear interaction (linear regression on the differences of participant-level average R-rates between F and nF conditions: significant

interaction between probability and feedback type $t(1034) = 3.27$, $P = 0.0011$, $\beta = 0.048$, 95% CI [0.019, 0.076]; Figure 3E), which was driven by the effect of feedback increasing as a function of the probability of the best-risky outcome in the complete feedback experiments ($t(1034) = 5.63$, $P < 0.001$, $\beta = 0.058$, 95% CI [0.038, 0.079]; see Table 1 for logistic regressions) but being stable in the partial feedback experiments ($t(1034) = 1.04$, $P = 0.3$, $\beta = 0.011$, 95% CI [-0.009, 0.031], $BF_{01} = 13$; see Table 1 for logistic regressions).

Extending the results to moderate risk options and losses

Next, we attempted to further clarify the psychological mechanisms involved in this effect. The fact that, in all experiments, the sure option is systematically a certain prospect leaves open the possibility that the effect of feedback is idiosyncratic to this framing. Indeed, certainty effects are known to heavily weigh decisions and to create robust paradoxes⁴³. In the next two experiments, we therefore assessed the robustness of our results to variation in outcome probabilities, specifically in contexts where the non-risky option is not certain. To do so, we designed Exp.5 and Exp. 6, where we substituted the sure option (which gives a specific amount with certainty) with a 50%-50% low variance lottery with EV equal to the EV of the sure option (Figure 1B). This new option remains relatively safe (given its low variance), yet now features an uncertain outcome. We shall refer to this option as the safe option, to differentiate it from both the sure and the risky. All other things considered (i.e., except for the sure options being substituted with the corresponding safe ones), Exp.5 and Exp.6 were respectively indistinguishable from Exp.2 (partial feedback) and Exp.4 (complete feedback) (Figure 1A). Consolidating our conclusions, all the main results identified in Exp.1-4 were replicated in this modified setup (Figure S1). Notably, the presence of feedback significantly increased risk-taking from the first trial in the partial feedback Exp. 5 and from trial 2 in the complete feedback Exp.6 (two-tailed WSRTs (BH-adjusted within-experiment): Exp. 5: Trials 1–10, all pairwise tests significant, with the most conservative comparison at Trial 10 yielding $z = 3.65$, $P < 0.001$, $r = 0.52$, 95 % CI [0.28, 0.72], $n = 80$; Exp. 6: Trial 1, $z = 1.22$, $P = 0.221$, $r = 0.17$, 95 % CI [-0.11, 0.44]; Trials 2–10, all pairwise tests significant, with the most conservative comparison at Trial 4 yielding $z = 2.82$, $P = 0.005$, $r = 0.39$, 95 % CI [0.14, 0.62], $n = 84$).

This attitudinal effect interacted with the probability of the risky option in the complete feedback condition but not in the partial one. To statistically substantiate this we run a linear regression on

the differences of participant-level average R-rates between F and nF conditions (significant interaction between probability and feedback type: $t(482) = 2.47$, $P = 0.014$, $\beta = 0.052$, 95% CI [0.011, 0.093]; slope in the complete condition: $t(482) = 4.73$, $P < 0.001$, $\beta = 0.070$, 95% CI [0.041, 0.099]; slope in the partial condition: $t(482) = 1.19$, $P = 0.23$, $\beta = 0.018$, 95% CI [-0.012, 0.047], $BF_{01} = 6.8$; see also Table 1). This result illustrates that our key findings are not idiosyncratic to some design choices and might therefore reflect a generalizable psychological effect.

Leveraging this robustness, we completed our demonstration by a comprehensive assessment of our main claims, evaluated over our six experiments. This analysis indicated that the attitudinal effect induced at the first trial was robustly elicited in the experiments featuring block instructions (two-tailed WSRT: $z = 6.42$, $P < 0.001$, $r = 0.45$, 95% CI [0.32, 0.56], $n = 344$; Experiments 2,4,5,6; Figure 4B), vanished in the absence of the said instructions (Bayesian paired t -test $BF_{01} = 11.4$; Bayesian sign test $BF_{01} = 8.4$, $n = 166$; Experiments 1,3; Figure 4A) and was significantly different in the two conditions (two-tailed MWU test: $U = 22108$, $z = -4.08$, $P < 0.001$, $r = -0.22$, 95% CI [-0.32, -0.12]). Consistent with different psychological mechanisms operating under partial or complete feedback regimens, the effect of feedback was identical across all levels of the probability of the risky option in the partial feedback experiments, while significantly modulated by this factor in the complete feedback experiments (significant interaction between probability and feedback type: $t(1520) = 4.11$, $P < 0.001$, $\beta = 0.049$, 95 % CI [0.026, 0.073]; slope in the complete condition ($t(1520) = 7.32$, $P < 0.001$, $\beta = 0.062$, 95 % CI [0.045, 0.079]); slope in the partial condition ($t(1520) = 1.53$, $p = 0.127$, $\beta = 0.013$, 95 % CI [-0.004, 0.029], $BF_{01} = 7.85$); Figures 4C-E).

Concerning O-rates, the partial-feedback experiments exhibited strong evidence for a null effect (Bayesian t -test: $BF_{01} = 14.13$, Bayesian sign test: $BF_{01} = 12.45$, $n = 255$), whereas aggregating the complete-feedback experiments indicated a small decrease in O-rate in the presence of feedback (two-tailed WSRT: $z = -2.85$, $P = 0.004$, $r = -0.21$, 95% CI [-0.36, -0.07], $n = 255$; mean $\Delta \pm \text{std} = -0.020 \pm 0.007$).

Finally, in a last experiment performed in the lab, we also showed that our main findings hold when incentives are increased and monetary losses are introduced (see Figure 1C, Exp7 and Supplementary Note 2).

Trial-by-trial analysis

A follow-up question is whether the abrupt emergence of the feedback effect on the second trial in the experiments without block instructions was dependent on the choice made on the first trial. We observe that the main effect of feedback is observed regardless of the previous choice type (sure/safe versus risky), and additively to the main effect of the previous choice type on R-rates (two-tailed WSRTs F vs. nF : Choice_{t=1} safe: $z = 3.83$, $P < 0.001$, $r = 0.39$, 95% CI [0.20, 0.57], Choice_{t=1} risky: $z = 2.83$, $P = 0.005$, $r = 0.32$, 95% CI [0.11, 0.52], $n = 166$; Figure 5A). Moreover, the magnitude of the feedback effect did not differ between the two cases (Bayesian paired t-test $BF_{01} = 10.9$, Bayesian sign test $BF_{01} = 8.1$; Figure 5C). The fact that, in experiments without block instructions, the second-trial effect is present regardless of the choice made on the first trial is again consistent with the idea that the discovery of the presence of feedback—and the consequent attitudinal change—explains the effect.

The existence of feedback-induced attitudinal effects on risk preferences does not rule out that additional feedback-induced learning processes co-exist. However, feedback-induced learning processes may not be apparent when looking at the average risky choice rate, because their effect depends on the previous trial choice and outcome. To investigate possible learning effects in trial-by-trial dynamics, we analyzed the probability of repeating a risky choice as a function of the outcome received in the previous trial. The logic of this analysis is that virtually any instantiation of a learning process would induce a “positive recency” effect, meaning that the probability of repeating a risky choice should increase after receiving the best possible (positive) outcome, compared to receiving the worst possible (zero) outcome (Figure 5E, left). We tested this hypothesis by analyzing this behavioral variable (probability of repeating a risky choice, $p(R_t | R_{t-1})$) in the feedback condition across all datasets. The results are in sharp contrast with the predictions of the learning hypothesis (Figure 5F). In fact, the probability of repeating a risky choice was lower after receiving positive (i.e., the “max”) feedback as opposite to zero (i.e., the “min”) (logistic GLMM: $z = -11.67$, $P < 0.001$, OR = 0.38, 95% CI [0.32, 0.45], see Table 1). The analysis of trial-by-trial dynamics thus shows no support for any form of feedback-induced learning process, and rather falsifies it. The observed behavioral pattern exhibits in fact negative recency, which is better understood as a manifestation of the gambler’s fallacy (in the laboratory⁴⁴⁻⁴⁶ and in an ecological (real-life) setting⁴⁷), according to which participants would move away

from a recently rewarded risky choice because they (wrongly) assume that the subsequent likelihood of positive feedback will be lower (Figure 5F and Figure 5G). This gambler's fallacy interpretation is further supported by conditioning this analysis on the probability of the risky outcome (recall, in our task we featured three probability levels: 10%, 50%, and 90%). This actually reveals that the effect is modulated by the underlying outcome probability and is maximal when outcomes are rare (difference between $p=10\%$ and $p=50\%$: $z = 6.35$, $P < 0.001$, OR = 2.19, 95% CI [1.72, 2.79]; difference between $p=10\%$ and $p=90\%$: $z = 11.11$, $P < 0.001$, OR = 5.31, 95% CI [3.96, 7.13]) and absent when the outcomes are common (effect of previous risky outcome in $p=90\%$ condition: $z = -0.64$, $P = 0.52$).

Next, we wanted to check whether the abrupt effect of feedback between the first and the second trial in experiments without block instructions was influenced by the outcome received on the first trial. The results also show that the effect of feedback is detectable, regardless of the nature of the outcome (zero or positive) received on the first trial (two-tailed WSRTs F vs. nF: $\text{Out}_{t=1} = 0$: $z = 2.28$, $P = 0.022$, $r = 0.32$, 95% CI [0.05, 0.57], $\text{Out}_{t=1} > 0$: $z = 4.54$, $P < 0.001$, $r = 0.43$, 95% CI [0.26, 0.58], $n = 166$; Figure 5B) and notably the effect did not differ across conditions (Bayesian paired t-test $\text{BF}_{01} = 9.8$, Bayesian sign test $\text{BF}_{01} = 8$; Figure 5D). This result further supports the idea of an outcome-independent attitudinal change in risk preferences. To sum up, not only do we demonstrate that the effect of feedback on risk preference precedes the reception of any feedback (and is, therefore, better understood as a change in attitude), but we also disprove any residual role for feedback-induced learning processes in the trial-by-trial dynamics by evidencing biased reactions to probabilistic and stochastic events akin to the gambler's fallacy⁴⁴⁻⁴⁷.

Verifying our findings in a previous dataset

We started our investigation by noting some discrepancies in the literature concerning the directionality of the effect of feedback in decision-making under risk, which was generally understood as stemming from a learning process (**Table S2 & S3**). Over 7 experiments, we found that the presence of feedback increases the propensity of taking risks, with no detectable consequence on the optimal choice rate. By manipulating block-wise instructions (present vs absence), we also found that the effects were mediated by a change of attitude of a different nature in the partial (consistent with curiosity) and complete (consistent with regret) feedback condition;

trial-by-trial dynamics analysis further ruled out that outcome-based learning plays a role in these processes.

We re-analyze a previously published dataset²⁶ that stands out as containing the largest sample size among the relevant studies (N=446) and proposing an influential cognitive model of decision-making (see next paragraph)^{26,28,31}, despite featuring some important limitations (it involved only complete feedback, the feedback and no-feedback conditions featured different number of trials and always appeared in the same order, there was no clear manipulation of the instructions). To replicate our analyses as comprehensively as possible, we restricted this re-analysis to the decision problems that feature identical or similar properties to ours, namely decisions opposing a sure to a risky option (to define risky choice rate), decisions involving options with different expected values (to define optimal choice rate) and decisions featuring non-extreme probabilities (excluding 1% or 99%) for the risky option. We also excluded trivial decision problems in which one option dominates the other (see Methods for more details about the study and the decision problem selection).

This re-analysis of Erev et al. (2017)²⁶ data was consistent with our own results on the absence of a positive effect of feedback on the optimal choice rate (without feedback 0.66 ± 0.21 , with feedback: 0.65 ± 0.17 ; Bayesian paired t-test: $BF_{01} = 17.1$, Bayesian sign test: $BF_{01} = 14.1$, $n = 446$; Figure 6A), as well as on the increase in risk choice rate (without feedback 0.38 ± 0.22 , with feedback 0.42 ± 0.18 ; two-tailed WSRT: $z = -6.34$, $P < 0.001$, $r = -0.36$, 95% CI [-0.46, -0.25; Figure 6B – see Table 1 for logistic regressions). Having in mind the fact that Erev et al. (2017) featured complete feedback, we looked at the effect of feedback specifically for different probability levels (low: $\text{prob} \leq 25\%$, medium: $25\% < \text{prob} < 75\%$, high: $\text{prob} \geq 75\%$) of the risky high-value outcome. As expected by the regret hypothesis, and consistent with our own findings, the effect of feedback monotonically scaled with the risky best-outcome probability (linear regression on the differences of participant-level average R-rates between F and nF conditions: $t(1095) = 8.31$, $P < 0.001$, $\beta = 0.093$, 95% CI [0.071, 0.114]; Figure 6C). Moreover, we looked at the trial-by-trial dynamics and found a negative recency pattern, consistent with a gambler's fallacy bias (logistic GLMM: $z = -8.86$, $P < 0.001$, OR = 0.49, 95% CI [0.42, 0.58], see Table 1; Figure 6D). Finally, while the manipulation of instructions was not present as such in Erev et al. (2017), the different orders of presentation of decision problems allowed us to perform an analogous analysis, which further

supports a first-trial/attitudinal effect (for the details, see Supplementary Note 3 and Figure S2). Overall, all the behavioral analyses that could be replicated in Erev et al. (2017) lead to similar results and conclusions as the ones performed on our own new data.

BEAST model and key behavioral results

Having verified that all the results that we could test in Erev et al. (2017) are replicated, we turned to test whether or not the model originally stemming from this dataset and later picked up by many other influential studies of human decision-making^{28,30,31,48} is able to capture the main behavioral patterns we observed.

The main idea of the Best Estimate and Sampling Tools (BEAST) model is that the attractiveness of one option is the sum of its Expected Value (assuming, for simplicity that no ambiguity is involved, which is the case here) and an average Sampling Value, which is generated by the use of four sampling tools (which correspond to four behavioral tendencies)²⁶. Three of the sampling tools are “biased” in the sense that they comprise mental draws from distributions that differ (in a specific, biased manner each) from the objective distributions of the prospects; they are not related to our discussion, so we omit the details here. The unbiased tool consists of mental draws either from the objective described distribution of the prospects, in case of no feedback, or from the realized one, in case of feedback. Crucially, the reliance on the observed history of outcomes when feedback is present makes the BEAST model a learning model. The unbiased tool captures the tendency to minimize immediate regret -or, to put it another way, it captures the tendency to prefer the option that gives the best outcome most of the time. The reliance on the unbiased tool increases as the participant receives more feedback, thereby making the impact of regret stronger.

We fitted the participants’ data from our experiments 1-6 to the BEAST model, and then we ran simulations with the fitted parameters to compute predicted choice probabilities for each participant. The model predictions replicate the main effect of feedback in increasing R-rates, but fail to capture the effect induced by the manipulation of block instructions -the effect of feedback is absent in the first trial, independently of the presence or absence of the block instructions (Figure 7A & 7B). This is an expected result since, in the BEAST, attitudinal changes occurring before any feedback is experienced are not contemplated. When it comes to the interaction of feedback with probability, model predictions capture very accurately the complete feedback case, but fail to capture the partial feedback case (Figure 7C, 7D & 7E). This is also an expected result, given that

the regret component of the BEAST model does not depend on the feedback regimen and is equally deployed in partial and complete feedback conditions. Finally, in contrast to our robust behavioral finding, BEAST predicts a positive recency pattern resulting from its learning nature (Figure 7F & 7G).

Consequently, the BEAST model is effective in capturing some of the behavioral patterns emerging in the setup of Erev et al. (2017) (a role for regret), while missing the negative recency effect also present in this dataset. Furthermore, we provide conclusive evidence that the model's predictions miss key behavioral findings, as soon as changes of the experimental setup are implemented, such as the presence of explicit block-wise instructions (attitudinal change) and partial feedback (role for curiosity).

ARTICLE IN PRESS

Discussion

In the present study, we aimed to clarify the effect of feedback on description-based decision-making. This represents a pressing question in behavioral decision-making research as many (if not most) real-life decisions feature both description (in various forms) and feedback (we enjoy – or suffer – the consequences of our choices). Our investigations aimed to address two (related) research questions. First, the directionality of the effect on two main outcome measures, namely risk aversion and value maximization. Second, the cognitive mechanisms underpinning the effect. To address these questions, we devised a series of original experiments designed to overcome some frequent (if not systematic) limitations identified in previous studies.

Regarding the directionality of the effect of feedback on decision-making under risk, we found that the presence of feedback increased the propensity to choose the risky option. Because it falls in line with the majority of studies surveyed in our literature review, we argue that this result can be considered credible, robust, and replicable^{26,35,49–51}. In light of the strength of our empirical evidence (including several experiments, different levels of outcomes, probabilities, gains and losses, and an overall extremely large sample size), we believe that the occasional reports of negative effects might be attributed to peculiarities of the design of the related studies -but further investigation would be needed to verify this.

Our experimental design also allowed us to confidently establish that the presence of feedback had no beneficial effect on the optimal (i.e., EV-maximizing) choice rate: risky choice rate was increased by feedback regardless of whether or not the risky option was more advantageous. This question was somehow overlooked in the literature, as most of the surveyed studies actually did not allow for testing this effect, either because they featured risky options that were consistently associated with the highest expected value^{52,53}, or because they used decision problems where the risky and the sure/safe option had the same expected values^{15,35,49,51,54}. The lack of clarity and investigation of the effect of feedback on choice optimality is surprising since feedback is often suggested as a possible way to correct decision-making biases and improve decision-making^{23,24}. Of note, once restricted to the subset of decision problems that are relevant to our investigation, a re-analysis confirmed our main results in a large dataset originally published by Erev et al. (2017)²⁶.

In addition to disambiguating the directionality (and amplitude) of feedback-induced changes in risk preferences, our study also provides further insights into the cognitive mechanisms underlying these effects. In the literature, the (more or less explicit) standard assumption is that experiencing decision outcomes affects the subjective beliefs characterizing the probability of realization of those outcomes^{15,51,52,55}. In other terms, the effects of feedback are traditionally conceived as the result of a learning process due to the sampling experience. Although predominant, the learning hypothesis has rarely been empirically challenged and compared to plausible alternatives, one of which being that the presence of feedback changes the attitude of a participant to make a risky decision. This attitudinal change can naturally be induced by the anticipation of the (informational and emotional) state that results from receiving the feedback. In order to evaluate the merit of this alternative hypothesis, we designed a new experimental manipulation that consisted of disclosing (or not) whether the upcoming block of trials would feature explicit feedback. This manipulation allowed us to reveal a simple but unambiguous behavioral signature of an attitudinal change: when participants were informed about the presence of feedback, its effect was present in the first trial, in decisions that preceded the disclosure of the first feedback. Furthermore, when looking at the experiments without block instructions, we observed that the effect of feedback on risk preference was present from the second trial regardless of which choice had been made or which outcome had been received in the first trial, further confirming that the change in risk preference was due to a sudden change in attitude, rather than being based on feedback integration. These results falsify the learning hypothesis as a sole determinant of the effect of feedback on risky decision-making, given that the effect of feedback is present before any actual learning could occur.

Several cognitive processes or psychological motives can actually underpin this attitudinal effect toward feedback. Our results, when restricted to the experiments featuring partial feedback (i.e., only the outcome of the chosen option was presented), are consistent with a curiosity-driven attitudinal change⁵⁶: since the outcome of the sure option is known with certainty before making a choice, choosing the risky option is the only way to resolve the uncertainty characterizing the outcome of the whole decision situation. These results are in line with the epistemic curiosity literature that shows that participants attribute a positive utility and actively seek uncertainty resolution³²⁻³⁶. However, while providing a satisfactory interpretation of the effects on risky decisions in the partial feedback condition, curiosity-driven motives cannot account for the fact that these effects persist under completed feedback conditions (i.e., when the outcomes of both the

chosen and the unchosen options were presented). To shed light on this incongruity, we analyzed the effect of feedback as a function of the probability level of the risky option. This analysis revealed that, while the effect of feedback on risk propensity seemed indistinguishable in the partial and complete feedback experiments at the aggregate level, clear differences emerged when splitting across probability levels. More specifically, in the complete feedback experiments, the feedback-induced increase in risk-taking was a monotonic function of the probability of the risky option, a pattern consistent with the idea that increased risk-taking is induced by the willingness to reduce the chance of experiencing regret⁵⁷. The rationale underlying this interpretation can be broken down into two main steps. First, regret is typically thought to be elicited – in the context of value-based decision-making – by the comparison between the obtained and the forgone outcome, and is, therefore, more salient in complete feedback environments (where this comparison can actually be directly observed). Second, inspired by the regret-minimizing tool of the BEAST model²⁶, we propose that the option that yields a superior outcome most of the time receives a regret premium. Our design implies a monotonically increasing relationship between the probability that the risky option gives its best outcome and its regret premium, and hence a positive correlation between the risky option's probability and the effect of feedback. This is exactly what we observed in all complete-feedback versions of our experiments, and the same held for the Erev et al. 2017²⁶ dataset (which featured complete feedback as well).

We should mention that the simple version of the regret-minimizing mechanism we propose here only captures the aforementioned monotonic pattern, but it is not enough to explain other aspects of our data. Notably, unequivocally relating the effect of feedback to a regret premium attributed to the option that gives the best outcome most of the time implies that in the low probability condition, we should observe a negative effect of feedback on R-rates, as the regret premium goes to the safe option in this case. While this is true in the Erev et al (2017)²⁶ dataset, in our complete-feedback experiments, the effect is negligible rather than negative ($BF_{01} \approx 8-12$; Figure 4D & 4E). This discrepancy is naturally accounted for assuming that the regret premium does not only depend on the probabilities (as we assumed so far) but also on outcome magnitudes. Specifically, we propose that this premium is further modulated by the ratio between the more frequent best outcome (by comparison to the alternative) and the range of the possible outcomes (i.e., the difference between the best and worst possible outcomes) in the decision problem. Note that according to this definition, in our design, the options with the greatest regret premium are those

associated with a high (90%) probability of delivering the risky outcome because: 1) they deliver the best possible outcome compared to the alternative 9 times out of 10, and 2) since the risky outcome is the largest in a decision pair, its relative value is maximum ($r=1$). Conversely, in decision problems where the probability of delivering the risky option is low (10%), the option that delivers the best outcome most of the time is the safe one. However, its regret premium is drastically reduced because the value of the safe outcome is much smaller than the range of the available outcomes ($r \ll 1$; because the sure option's magnitude is close to the Expected Value of the risky option, see Figure 1B).

Although explaining our current results requires assuming two different psychological processes operating in the two informational regimens, we note that regret and curiosity have been shown to interact in other experimental (and real-life) situations^{58–60}. Further research will be needed to better characterize the relation (cooperation or competition) between these two motives. Our results add up to the behavioral literature, spanning from reinforcement learning to behavioral economics, showing that partial and complete feedback situations may elicit (radically) different cognitive processes^{9,61–63}.

Evidence for an attitudinal influence (as manifested by the first-trial effect) does not rule out the possibility that learning processes operate in parallel and influence choices. To test this hypothesis, we looked at feedback-induced trial-by-trial choice adjustments. We reasoned that virtually all instantiations of learning, be they rooted in Bayesian update or Reinforcement learning⁶⁴, generate positive recency effects, i.e., predict that the probability of choosing a risky option should increase following a risky choice yielding the best possible outcome. Critically, this is also true for decision-by-sampling models, such as the BEAST, which supposes mental samples of the outcomes are drawn from the empirical (i.e., experienced) distribution. Accordingly, having experienced best possible outcomes in the past should increase the likelihood of repeating the same choice. In contrast with the prediction of the learning hypothesis, our analysis revealed a negative recency pattern: a positive feedback in the preceding trial reduced the chance of repeating a risky option. Our re-analysis of Erev et al. 2017 provided further support in favor of a negative recency (gambler's fallacy) effect. This prima facie puzzling behavioral effect can be understood as a manifestation of what is called the gambler's fallacy, i.e., the fact that human participants tend to misrepresent the independence of probabilistic outcomes^{44,46,65}. This interpretation is further

confirmed by the fact that this effect interacted with the probability of the risky outcome, such that the negative recency effect was stronger when the probability of the risky option giving its best outcome was lower: a situation where the gambler's fallacy intuition suggests that two consecutive lucky strikes appear almost impossible. These results not only rule out the possibility that residual feedback-induced learning effects are at play but also suggest that an explicit description of the probabilistic process may create (biased) prior expectations that prevent learning processes.

Taken together, our observations significantly challenge any cognitive model relying on a learning process as a mechanism of the effect of feedback on decision-making under risk, be it rooted in reinforcement learning⁶⁶, Bayesian⁶⁷, or decision-by-sampling theories²⁷. Concretely, we found the influential BEAST model, which supposes that once feedback is available, values are constructed by (mentally) sampling outcomes from the empirical distribution, falls short of capturing critical features of behavioral results, such as the first-trial effect and the negative recency effect. Also of note, while the BEAST model does capture the role of regret, it is not equipped to deal with the shape of effects highlighted in the partial feedback condition. Our results, therefore, suggest that the BEAST model—if it is to continue being used as a valid model of human decision-making—should undergo major restructuring with respect to several key structural features^{28,31,48}. Our findings also tell a cautionary tale against validating computational cognitive models using choice prediction, data-driven competitions, which, focusing on quantitative metrics and aggregate data, may miss critical features of human behavior that can be highly diagnostic of the underlying cognitive mechanisms^{26,68}.

To conclude, our findings shed new light on the behavioral effects of feedback in description-based decisions under risk, and on their underlying psychological mechanisms beyond learning. Because of the ubiquity of those situations, elucidating the effect of feedback in description-based scenarios can improve our understanding of apparent decision anomalies relevant to many real-life situations and allow us to improve our policies. In particular, our results suggest that, contrary to what common sense could dictate, providing feedback cannot be considered as a panacea to correct decision-making biases, notably because the effects of feedback are at least partially mediated by attitudinal changes rather than purely driven by learning processes^{23–25}. Our results add up to an increasing body of evidence highlighting complex interactions between description- and experience-based choices that are currently not well accounted for by standard models^{6,69}.

Methods

Participants

The INSERM Ethical Review Committee approved the study, and participants were provided written informed consent before their inclusion. The research was carried out following the principles and guidelines for experiments including human participants provided in the Declaration of Helsinki (1964, revised in 2013). For the six online experiments, we recruited a total of 620 participants (4x100 for Exp1-4, 104 for Exp5, and 116 for Exp6 | 300 females, 300 males, 20 gender not available | aged 29.29 ± 9.24 years) from an online platform (www.prolific.com).

For the laboratory experiment, 30 healthy participants completed the experiment (18 females | aged 28 ± 7.22). Given that Bellemare et al.⁷⁰ assert that 20 participants are needed in a within-subject analysis to achieve a power of 80%; a sample size of 30 is expected to have enough power to detect any effect. The choice of model to analyze the dataset, the generalized mixed effect model, also requires a smaller sample size than traditional ANOVA. Participants were contacted via the "Relais d'information en sciences de la cognition" (RISC), part of the French "Centre national de la recherche scientifique". Participants enroll on the platform and, as such, voluntarily accept to be contacted for scientific studies. Participants were recruited on the basis of their good understanding of French. They received an email in their mailing list containing a link to a questionnaire asking them for general information. After completion, they are contacted by the experimenter to agree on the day of the experiment. Participants, which, on average, lasted 45 minutes.

We note that gender information was determined based on self-reports. No statistical method was used to predetermine sample size, other than aligning with current practices in the literature for similar studies. The assignment of participants to a given experiment was not randomized, but the researchers played no role in it because it depended on Prolific's matching procedure. Nonetheless, we ensured that no participant took part in more than one experiment.

Exclusion criteria

To ensure the high quality of the data of the online experiments, we applied the following exclusion criteria (which were explicitly mentioned in the master thesis pre-planning document produced by the first author):

- participants with a missing trial or with a repetition of a trial in the test phase
- participants with two or more submissions of more than 100 out of 270 trials (the total number of trials, 270, includes training trials. So, we kept the complete submission of participants that had, for example, two submissions one with only a few trials and a complete one)
- participants with less than 9 (out of 10) correct answers in the catch block (see below for more information on the catch block)
- excessively long completion time [two standard deviations more than the average]
- right or left option >95% of the time [low quality data] - note that the position of the options was counterbalanced, so this behavior is aberrant

After applying the above exclusion criteria, we were left with a total of 510 participants for the main analyses (80, 95, 86, 85, 80, 84 for experiments 1-6, respectively). It is worth noting that all our main results replicate when we drastically lower our exclusion criteria (keeping participants with less than ten missing trials and with an accuracy in the catch block of at least 6/10), thereby including 609/620 participants in the analysis.

No participants were excluded from the analysis of the laboratory experiment.

Incentives

For the online experiments, participants received a fixed compensation of £3 for about half an hour of engagement (average completion time in minutes: 28.47 ± 8.68). Additionally, we incentivized participants to reveal their true preferences by offering a monetary bonus determined by the outcome of a randomly selected trial of the testing phase (average bonus won in British pounds: 2.68 ± 2.21).

For the laboratory experiment, participants received a show-up fee of 10€ for an average engagement of 45 minutes. To motivate the revelation of true preferences, an incentive system was settled on the basis of hypothetical gains or losses. Each decision gave a payoff in points — a draw from the selected option payoff distribution. They were told that their goal was to maximize their number of points in the gain domain and to minimize the number of losses. The total of points would determine their final payoff. The conversion rate between experimental units (EU) and euros was set according to the maximum and minimum possible payoffs at 0.02€/EU. The payoff

structure was determined such that no participant would incur a loss. On average, they received 15€. At the end of the experiment, participants received and signed the receipt for their payments.

Experimental design

The online experiments started with participants giving their consent to participate, followed by detailed instructions on the behavioral task, the structure of the experiment, and the compensation. Afterward, participants went through training, which was a mini version (four blocks of five trials each) of the actual experiment featuring similar yet not identical decision problems to those of the actual experiment.

Then, participants started the actual experiment, which consisted of 24 blocks (12 decision problems with and without feedback, as described in the main text) and a catch block (see below) in the middle of the actual experiment. The 24 blocks were comprised of 10 trials each and were randomized within and across participants. The actual experiment was divided into three sessions, between which the participants were allowed to take a self-paced break.

Each block started with a screen providing block instructions about the presence or absence of feedback in the upcoming block (Experiments 2, 4, 5, 6) or just prompting the participants to start the block by clicking on the “Start Block” button (Experiments 1 & 3). This step was self-paced. Then, the two options (their magnitude and probability) were presented side-by-side, with a clickable white square below each option. The position of the options (left or right) was randomized. Also, the relative vertical position of the magnitude and the probability (magnitude above and probability below or vice versa) was randomized across participants (but was constant within participants). Participants could make their choice at their own pace by clicking the white square below their preferred option. The outcome of the risky (or the safe) option was determined by an independent random draw –by definition, the outcome of the sure option was fixed. After the choice was made and the outcomes were determined, the outline of the selected square/option was highlighted, and inside the white square, the outcome of the chosen option was revealed (showing the obtained points) or hidden (showing a question mark) for 1500ms. In experiments where complete feedback was used, the forgone outcome was revealed as well (showing the points of the unchosen option) or hidden (showing a question mark) in a light gray font (in contrast to the standard black font for the obtained outcome). Note that partial feedback means that the outcome of only the chosen option is revealed; complete feedback means that the outcome of both the

chosen and the unchosen option is revealed. Then, the next trial started showing the same options (potentially in a different position). At the end of the block, a screen marking the end of the block was presented for 1500ms.

In the middle of the actual experiment, a catch block featuring the trivial choice between a sure option (probability = 100%) giving 5 points and a sure option giving 30 points was presented. The catch block was identical to all the rest blocks of the actual experiment. The related exclusion criterion (participants who scored below 9/10 correct answers in the catch block were excluded from the analysis) enabled us to ensure that participants understood and paid attention to the task.

At the end of the online experiments, participants were informed about the randomly selected trial and the associated monetary bonus, about their total compensation, and they were redirected to the recruiting platform (Prolific) to formally complete their participation.

The laboratory experiment was conducted at the laboratory of the Département des Etudes Cognitives at École Normale Supérieure, Paris. Participants were told their rights and gave their consent. They completed a mathematical questionnaire assessing their ability to perform multiplication. The experimenter explained the task using a visual support and read an example of a gamble that might be encountered by the decision maker to ensure that they understand the possible outcomes of the gamble. Participants were told to maximize their gains and minimize their losses.

On the computer, participants were presented with an instruction screen indicating the type of information that they would be facing in the block. Individuals read “réalisation révélée essai par essai - feedback provided after each trial” or “réalisation cachée - feedback not provided” and were informed that, at the end of the block, they would be told the number of points gained or lost at that block. After reading the instructions, participants were presented with two options side by side on the screen. Each was associated with a geometrical shape and two labels: one gave information on the probability of occurrence of the outcome, and the other, on the magnitude of the outcome. The shapes and their color varied randomly throughout the experiment and across participants. Colors and the brightness of the screen were adjusted to prevent eye fatigue and for easy reading. Labels describing the magnitude and the probability of each lottery figure were above and below the shapes. For half of the participants, the probabilities are always above and the magnitude below. For the other half of the sample, the labels are reversed. The position of the risky and the

sure lotteries was randomized within and across blocks. Participants could not anticipate their position and therefore had to stay attentive during the sequential choices. Choices were made using a mouse. After each one, an arrow indicated the localization of the choice, disappeared after 500ms to be replaced by a text at the place of the label of the magnitude of the chosen lottery. The text indicated either the number of points gained or lost (e.g, -32pts) or hid them (XX pts) –as in the online experiments, the outcome of the risky lottery was determined by an independent random draw. The disclosure of the points depended on the type of information provided at the beginning of the block. 1.5s after, another screen appeared with the second round of choice, providing the same gamble as before. At the end of the sequence of 10 choices, the accumulated points were shown on the screen.

Most steps of the task were self-paced, such that the participant could read and take as much time as needed for the instruction, the choice, and the bonus screens. The experiment was divided into three sessions to avoid fatigue. Participants were given feedback on their total points accumulated at the end of each session. At the end of the experiment, their payoff was computed on the basis of their total points.

Decision Problems

By experimental design, the magnitude and the probability of the risky option were determined, and by definition, the probability of the sure option was determined too. Given these, we computed the magnitude of the sure options for Experiments 1-4, so that the EV difference between the risky and the sure option was at 5%. Namely,

$$\text{mag}_{\text{sure}} := (1 + (1 - 2 * \text{riskyBetter})0.05)\text{mag}_{\text{risky}}\text{prob}_{\text{risky}}$$

In the end, we rounded the mag_{sure} towards the direction that agrees with the riskyBetter factor (if $\text{riskyBetter}=1$, we rounded mag_{sure} with the floor and if $\text{riskyBetter}=0$ with the ceiling function).

Safe options (Experiments 5 & 6) had also defined probabilities (50/50) and were set such that $\text{EV}_{\text{safe}} = \text{EV}_{\text{sure}}$. Additionally, one of the two magnitudes was set equal to the EV_{risky} , so that

both outcomes of the safe option were above (below) EV_{risky} when $\text{riskyBetter}=0$ ($=1$). So, the two outcomes for the safe option satisfied the following:

$$\text{mag}_{\text{safe1}} = EV_{\text{risky}} \ \& \ \text{mag}_{\text{safe2}} = 2EV_{\text{sure}} - \text{mag}_{\text{safe1}}$$

In the laboratory experiment (Experiment 7), the EV of the two options was set to be equal. Hence, the magnitude of the sure option was determined given the magnitude and the probability of the risky one. Rounding to the closest integer was used in this experiment.

Statistical analysis

The main analyses included two dependent variables: risky choice (1 if one chooses the risky option and 0 if one chooses the sure/safe one) and optimal choice (1 if one chooses the Expected Value-maximizing choice; 0 otherwise). We ran a regression analysis using a Generalized Linear Mixed-Effects model (GLMM). We used the canonical link function for response variables with the binomial distribution, namely the logit function. The method for estimating model parameters was maximum likelihood, and the analysis was run with R. The predictors were feedback (categorical: absent/present), riskyBetter (categorical: no/yes), the probability of the risky option (categorical: low, medium, high for 10%, 50%, 90% respectively), the magnitude of the risky option (categorical: low and high for 40 and 60 respectively) and trial (continuous: from 1 to 10). In the basic model, we included the main effects. Specific pairwise interactions were also added when they were relevant. To account for individual differences, we incorporated a random effects structure with random intercepts and random slopes at the participant level. The main effects of all predictors were included as random slopes (the only exception being the predictor trial in the optimal choice case, which had a very small variance and led to singular fitting; therefore, it was excluded). The magnitude of the risky option was not relevant for the Erev dataset, as the lotteries were not necessarily of the form $(m,p;0)$; therefore, this predictor was excluded from the analysis of this dataset. Additionally, in the Erev dataset, the trial was centered and scaled within each feedback condition to remove the dependency between feedback and trial (since trials 1–5 occurred without feedback and trials 6–25 with feedback) and to improve model interpretability and convergence.

We also ran similar regressions for the dependent variable `repeatRisky` (1 if one chooses the risky option and 0 if one chooses the sure/safe one at time t , conditional on having chosen the risky option at time $t-1$). The predictors were the probability of the risky option (specified as above) and whether the previous (at time $t-1$) risky outcome was the minimum or the maximum among the possible outcomes of the risky option (categorical: 0/1). As fixed effects, we included a basic model with the main effects only and a full model with the interaction too. The main effects of the predictors were used as random slopes. Note that the analysis was run only on feedback blocks. Because participants' choice frequencies and experienced outcomes were not experimentally controlled, not all participants contributed data to every combination of previous-outcome and probability conditions. Consequently, the dataset is unbalanced across participants and conditions. This does not pose a problem for the statistical analysis, as generalized linear mixed-effects models (GLMMs) are robust to unbalanced designs and appropriately accommodate unequal numbers of observations per participant. Accordingly, the related figures display variable sample sizes across conditions, with exact values reported in the legend of Figure 5.

To further examine whether the effect of feedback on risky choice rates depended on the probability of the risky option, we conducted a focused analysis testing the interaction between the probability of the risky option and the type of feedback (partial vs. complete). For each participant, we computed the difference in risky-choice rates between feedback and no-feedback conditions ($R(F) - R(nF)$), yielding a dependent variable ("Value") representing the feedback effect on risk-taking. This variable was analyzed using a linear model with the continuous probability of the risky option (`pRisky`) and feedback type (`TypeOfFeedback`, coded as 0 = partial, 1 = complete) as predictors, including their interaction term ($pRisky \times TypeOfFeedback$). The model was estimated using ordinary least squares in MATLAB's `fitlm` function. To decompose the interaction, we conducted simple-slope analyses estimating the effect of `pRisky` separately for each feedback type. When simple slopes were nonsignificant, we quantified the evidence for the null hypothesis using BIC-based Bayes factors comparing full and intercept-only models (BF_{01}), with values greater than 1 indicating evidence for the absence of an effect, and larger values reflecting stronger evidence.

To assess the difference between the feedback and no-feedback conditions at the participant level, we conducted a comprehensive paired-condition analysis. We first computed Bayesian paired t-tests and Bayesian sign tests (following JZS and Beta priors, respectively) to quantify the relative evidence for or against a difference between conditions. Sensitivity analyses were performed across a range of prior scales for the t-test ($r = 0.35, 0.71, 1.0$) and Beta priors for the sign test (Beta(0.5, 0.5), Beta(1, 1), Beta(2, 2)); all yielded confirmatory evidence consistent with the main analysis which is reported in the Results. Additional robustness checks, including a Wilcoxon signed-rank test, nonparametric bootstrap confidence intervals for the mean and trimmed mean (20%), and a permutation test on the mean difference, were also conducted and provided convergent confirmatory results.

Erev et al. (2017) re-analysis

This is one of the most extensive studies that includes description-only and description+experience choices, consisting of 150 decision problems and 446 participants. However, the main goal of the Erev et al. 2017 paper was not to compare these two paradigms. In an attempt to facilitate this comparison, we processed the raw data of the aforementioned paper and examined whether we could replicate with their data our main results. In doing so, we focused on a subset of their decision problems that enables us to answer the most important questions of our investigation.

First of all, we excluded from our analysis problems that involve ambiguity, more than two outcomes for one option, a correlation between the outcomes of the two lotteries, and decisions without lotteries (eg, coin-toss). Additionally, to be able to assess the effect of feedback on EV-maximization and on risk propensity, we further excluded decision problems in which the two options had equal EVs, problems that did not have the form of a sure vs a risky lottery, trivial problems (problems in which one option dominates the other) and problems with extreme probabilities for the risky option (1% or 99%); this left us with 31 decision problems, which constitute the pool for the related analysis.

Cognitive Modeling

The Best Estimation And Sampling Tools (BEAST) model contains six free parameters, with one capturing attitudes towards ambiguity, so it is irrelevant to us. The key idea behind the BEAST

model is that it assumes decisions to be determined by a combination of the best estimation of the expected value of an option and its subjective value determined by different sampling tools²⁷. Four parameters in the model are dedicated to quantifying the weight of different sampling tools (i.e., rules for drawing and comparing hypothetical outcomes) as well as parameters governing the probability of using one of these tools. In the original publication, the authors estimated the best-fit parameters on a population level, allowing parameters of individual agents to be drawn from $U(0, \text{param}_{\text{best-fit}})$ or $\{1, \dots, \text{param}_{\text{best-fit}}\}$ (depending on whether the parameter is continuous or discrete). Mean squared deviation and some additional constraints to capture various fundamental behavioral phenomena were used for estimating these population parameters. For all the details of the model and its optimization, we refer to the original publication²⁶.

We used a different approach. We fixed a search space on the population level which was quite wider than the one resulting from the original optimization (the best fits for the parameters $(\sigma, \beta, \gamma, \theta, \kappa)$ was (7, 2.6, .5, 1, 3) in Erev, et al. (2017) while we used (10, 10, 1, 5, 6) as the search space). We computed the best-fit parameters on the participant level using maximum likelihood estimation (note that we fixed the individual parameters on the participant level and simulated each decision problem one hundred times to obtain the predicted R-rates for each participant). Then we ran simulations using the fitted parameters of each participant to obtain the predictions of the BEAST model and compare them with the behavioral data (but note that we also produced simulations using the originally proposed best-fitting parameters, which led to the very same conclusions).

We also submitted our data to a modeling analysis, which involved a standard descriptive behavioral model—a variant of cumulative prospect theory⁷¹. Details about the model, the fitting procedures following standard guidelines^{72,73}, and the resulting statistical analyses can be found in the Supplementary Note 4.

Data availability

All data supporting the findings of this study are available on GitHub at <https://github.com/hrl-team/riskyDecisionMakingFeedback>, and a versioned, citable archive of the complete dataset is deposited on Zenodo⁷⁴ at <https://doi.org/10.5281/zenodo.17807046>. The file `Nasioulas2024_data.csv` contains the raw data from Experiments 1–7, and

Erev2017_data.csv contains processed data reproduced from Erev et al. (2017).

The original datasets from Erev et al. (2017) are publicly available at

<https://zenodo.org/records/321652>.

Code availability

All analysis code is openly available on GitHub at

<https://github.com/hrl-team/riskyDecisionMakingFeedback>, and a versioned, citable archive of the

code and data is provided on Zenodo⁷⁴ at <https://doi.org/10.5281/zenodo.17807046>.

The MATLAB script `generate_figures_main.m` reproduces all main-text and supplementary figures, and the R script `generate_glm_main.R` performs the logistic regression analyses reported in the manuscript.

References

1. Allais, M. Le Comportement de l'Homme Rationnel devant le Risque: Critique des Postulats et Axiomes de l'Ecole Americaine. *Econometrica* **21**, 503–546 (1953).
2. Kahneman, D. & Tversky, A. Prospect Theory: An Analysis of Decision under Risk. *Econometrica* **47**, 263–291 (1979).
3. Bell, D. E. Regret in Decision Making under Uncertainty. *Operations Research* **30**, 961–981 (1982).
4. Loomes, G. & Sugden, R. Regret Theory: An Alternative Theory of Rational Choice Under Uncertainty. *The Economic Journal* **92**, 805–824 (1982).
5. Wakker, P. & Tversky, A. An Axiomatization of Cumulative Prospect Theory. *Journal of Risk and Uncertainty* **7**, 147–175 (1993).
6. Erev, I., Yakobi, O., Ashby, N. J. S. & Chater, N. The impact of experience on decisions based on pre-choice samples and the face-or-cue hypothesis. *Theory Decis* **92**, 583–598 (2022).
7. Fantino, E. & Navarro, A. Description–experience Gaps: Assessments in Other Choice Paradigms. *Journal of Behavioral Decision Making* **25**, 303–314 (2012).
8. Plonsky, O. & Teodorescu, K. The influence of biased exposure to forgone outcomes. *Journal of Behavioral Decision Making* **33**, 393–407 (2020).
9. Weiss-Cohen, L., Konstantinidis, E. & Harvey, N. Timing of descriptions shapes experience-based risky choice. *Journal of Behavioral Decision Making* **34**, 66–84 (2021).
10. Garcia, B., Cerrotti, F. & Palminteri, S. The description–experience gap: a challenge for the neuroeconomics of decision-making under uncertainty. *Philosophical Transactions of the Royal Society B: Biological Sciences* **376**, 20190665 (2021).

11. Hertwig, R. & Erev, I. The description–experience gap in risky choice. *Trends in Cognitive Sciences* **13**, 517–523 (2009).
12. Hertwig, R. & Wulff, D. U. A Description–Experience Framework of the Psychology of Risk. *Perspect Psychol Sci* **17**, 631–651 (2022).
13. Lejarraga, T. & Hertwig, R. How experimental methods shaped views on human competence and rationality. *Psychological Bulletin* **147**, 535–564 (2021).
14. Prelec, D. The Probability Weighting Function. *Econometrica* **66**, 497–527 (1998).
15. Marchiori, D., Di Guida, S. & Erev, I. Noisy retrieval models of over- and undersensitivity to rare events. *Decision* **2**, 82–106 (2015).
16. Newell, B. R. & Rakow, T. The role of experience in decisions from description. *Psychonomic Bulletin & Review* **14**, 1133–1139 (2007).
17. Palminteri, S. & Lebreton, M. Context-dependent outcome encoding in human reinforcement learning. *Current Opinion in Behavioral Sciences* **41**, 144–151 (2021).
18. Palminteri, S. & Lebreton, M. The computational roots of positivity and confirmation biases in reinforcement learning. *Trends in Cognitive Sciences* **26**, 607–621 (2022).
19. Esponda, I., Vespa, E. & Yuksel, S. Mental Models and Learning: The Case of Base-Rate Neglect. *American Economic Review* **114**, 752–782 (2024).
20. Abel, M., Cole, S. A. & Zia, B. Debiasing on a Roll: Changing Gambling Behavior Through Experiential Learning. SSRN Scholarly Paper at <https://papers.ssrn.com/abstract=2579890> (2015).
21. Gigerenzer, G. The Bias Bias in Behavioral Economics. *RBE* **5**, 303–336 (2018).

22. Morewedge, C. K. *et al.* Debiasing Decisions: Improved Decision Making With a Single Training Intervention. *Policy Insights from the Behavioral and Brain Sciences* **2**, 129–140 (2015).
23. Fischhoff, B. Debiasing. in *Judgment under Uncertainty: Heuristics and Biases* (eds. Tversky, A., Kahneman, D. & Slovic, P.) 422–444 (Cambridge University Press, Cambridge, 1982). doi:10.1017/CBO9780511809477.032.
24. Grüne-Yanoff, T. & Hertwig, R. Nudge Versus Boost: How Coherent are Policy and Theory? *Minds & Machines* **26**, 149–183 (2016).
25. Vlaev, I. & Dolan, P. Action Change Theory: A Reinforcement Learning Perspective on Behavior Change. *Review of General Psychology* **19**, 69–95 (2015).
26. Erev, I., Ert, E., Plonsky, O., Cohen, D. & Cohen, O. From anomalies to forecasts: Toward a descriptive model of decisions under risk, under ambiguity, and from experience. *Psychological Review* **124**, 369–409 (2017).
27. Stewart, N., Chater, N. & Brown, G. D. A. Decision by sampling. *Cognitive Psychology* **53**, 1–26 (2006).
28. Peterson, J. C., Bourgin, D. D., Agrawal, M., Reichman, D. & Griffiths, T. L. Using large-scale experiments and machine learning to discover theories of human decision-making. *Science* **372**, 1209–1214 (2021).
29. Plonsky, O., Erev, I., Hazan, T. & Tennenholtz, M. Psychological Forest: Predicting Human Behavior. *Proceedings of the AAAI Conference on Artificial Intelligence* **31**, (2017).
30. Plonsky, O. *et al.* Predicting human decisions with behavioral theories and machine learning. in (2024).

31. Thomas, T. *et al.* Modelling dataset bias in machine-learned theories of economic decision-making. *Nat Hum Behav* **8**, 679–691 (2024).
32. Buyalskaya, A. & Camerer, C. F. The neuroeconomics of epistemic curiosity. *Current Opinion in Behavioral Sciences* **35**, 141–149 (2020).
33. Charpentier, C. J., Bromberg-Martin, E. S. & Sharot, T. Valuation of knowledge and ignorance in mesolimbic reward circuitry. *Proceedings of the National Academy of Sciences* **115**, E7255–E7264 (2018).
34. Cogliati Dezza, I., Maher, C. & Sharot, T. People adaptively use information to improve their internal states and external outcomes. *Cognition* **228**, 105224 (2022).
35. Rigoli, F., Martinelli, C. & Shergill, S. S. The role of expecting feedback during decision-making under risk. *Neuroimage* **202**, 116079 (2019).
36. Ruggeri, A., Stanciu, O., Pelz, M., Gopnik, A. & Schulz, E. Preschoolers search longer when there is more information to be gained. *Developmental Science* **n/a**, e13411 (2023).
37. Gottlieb, J. & Oudeyer, P.-Y. Towards a neuroscience of active sampling and curiosity. *Nat Rev Neurosci* **19**, 758–770 (2018).
38. Sharot, T. & Sunstein, C. R. How people decide what they want to know. *Nat Hum Behav* **4**, 14–19 (2020).
39. Zeelenberg, M., Beattie, J., van der Plight, J. & de Vries, N. K. Consequences of regret aversion: Effects of expected feedback on risky decision making. *Organizational Behavior and Human Decision Processes* **65**, 148–158 (1996).
40. Coricelli, G. *et al.* Regret and its avoidance: a neuroimaging study of choice behavior. *Nat Neurosci* **8**, 1255–1262 (2005).

41. Cohen, D., Plonsky, O. & Erev, I. On the impact of experience on probability weighting in decisions under risk. *Decision* **7**, 153–162 (2020).
42. Couto, J., Maanen, L. van & Lebreton, M. Investigating the origin and consequences of endogenous default options in repeated economic choices. *PLOS ONE* **15**, e0232385 (2020).
43. Tversky, A. & Kahneman, D. Rational Choice and the Framing of Decisions. *The Journal of Business* **59**, S251–S278 (1986).
44. Ayton, P. & Fischer, I. The hot hand fallacy and the gambler's fallacy: Two faces of subjective randomness? *Memory & Cognition* **32**, 1369–1378 (2004).
45. Barron, G. & Leider, S. The role of experience in the Gambler's Fallacy. *Journal of Behavioral Decision Making* **23**, 117–129 (2010).
46. Teoderescu, K., Amir, M. & Erev, I. The experience-description gap and the role of the inter decision interval. *Prog Brain Res* **202**, 99–115 (2013).
47. Clotfelter, C. T. & Cook, P. J. Notes: The "Gambler's Fallacy" in Lottery Play. *Management Science* <https://doi.org/10.1287/mnsc.39.12.1521> (1993) doi:10.1287/mnsc.39.12.1521.
48. Bourgin, D. D., Peterson, J. C., Reichman, D., Russell, S. J. & Griffiths, T. L. Cognitive model priors for predicting human decisions. in *Proceedings of the 36th International Conference on Machine Learning* 5133–5141 (PMLR, 2019).
49. Goyal, S. & Miyapuram, K. P. Feedback Influences Discriminability and Attractiveness Components of Probability Weighting in Descriptive Choice Under Risk. *Front. Psychol.* **10**, (2019).
50. Weiss-Cohen, L., Konstantinidis, E., Speekenbrink, M. & Harvey, N. Incorporating conflicting descriptions into decisions from experience. *Organizational Behavior and Human Decision Processes* **135**, 55–69 (2016).

51. Yechiam, E. & Barron, G. The Role of Personal Experience in Contributing to Different Patterns of Response to Rare Terrorist Attacks. *Journal of Conflict Resolution - J CONFLICT RESOLUT* **49**, (2005).
52. Jessup, R. K., Bishara, A. J. & Busemeyer, J. R. Feedback produces divergence from prospect theory in descriptive choice. *Psychol Sci* **19**, 1015–1022 (2008).
53. Lejarraga, T. & Gonzalez, C. Effects of feedback and complexity on repeated decisions from description. *Organizational Behavior and Human Decision Processes* **116**, 286–295 (2011).
54. Josephs, R. A., Larrick, R. P., Steele, C. M. & Nisbett, R. E. Protecting the self from the negative consequences of risky decisions. *J Pers Soc Psychol* **62**, 26–37 (1992).
55. Aydogan, I. & Gao, Y. Experience and rationality under risk: re-examining the impact of sampling experience. *Exp Econ* **23**, 1100–1128 (2020).
56. Loewenstein, G. The psychology of curiosity: A review and reinterpretation. *Psychological Bulletin* **116**, 75–98 (1994).
57. Gilovich, T. & Medvec, V. H. The experience of regret: What, when, and why. *Psychological Review* **102**, 379–395 (1995).
58. Caldwell, D. F. & Burger, J. M. Learning about unchosen alternatives: When does curiosity overcome regret avoidance? *Cognition and Emotion* **23**, 1630–1639 (2009).
59. Shani, Y. & Zeelenberg, M. When and why do we want to know? How experienced regret promotes post-decision information search. *Journal of Behavioral Decision Making* **20**, 207–222 (2007).
60. van Dijk, E. & Zeelenberg, M. When curiosity killed regret: Avoiding or seeking the unknown in decision-making under uncertainty. *Journal of Experimental Social Psychology* **43**, 656–662 (2007).

61. Bavard, S., Rustichini, A. & Palminteri, S. Two sides of the same coin: Beneficial and detrimental consequences of range adaptation in human reinforcement learning. *Science Advances* **7**, eabe0340 (2021).
62. Klein, T. A., Ullsperger, M. & Jocham, G. Learning relative values in the striatum induces violations of normative decision making. *Nat Commun* **8**, 16033 (2017).
63. Li, J. & Daw, N. D. Signals in Human Striatum Are Appropriate for Policy Update Rather than Value Prediction. *J. Neurosci.* **31**, 5504–5511 (2011).
64. Erev, I. & Haruvy, E. Learning and the Economics of Small Decisions (Chapter 10). in *The Handbook of Experimental Economics, Volume 2* vol. 2 638–716 (Princeton University Press, 2015).
65. Plonsky, O. & Teodorescu, K. The influence of biased exposure to forgone outcomes. *Journal of Behavioral Decision Making* **33**, 393–407 (2020).
66. Sutton, R. & Barto, A. *Reinforcement Learning: An Introduction*. (The MIT Press, 2018).
67. Ma, W. J., Kording, K. P. & Goldreich, D. *Bayesian Models of Perception and Action*. (The MIT Press, 2023).
68. Erev, I. *et al.* A choice prediction competition: Choices from experience and from description. *Journal of Behavioral Decision Making* **23**, 15–47 (2010).
69. Garcia, B., Lebreton, M., Bourgeois-Gironde, S. & Palminteri, S. Experiential values are underweighted in decisions involving symbolic options. *Nature Human Behaviour* **7**, (2023).
70. Bellemare, C., Bissonnette, L. & Kröger, S. *Statistical Power of Within and Between-Subjects Designs in Economic Experiments*. <https://econpapers.repec.org/paper/izaizadps/dp8583.htm> (2014).

71. Tversky, A. & Kahneman, D. Advances in prospect theory: Cumulative representation of uncertainty. *J Risk Uncertainty* **5**, 297–323 (1992).
72. Palminteri, S., Wyart, V. & Koechlin, E. The Importance of Falsification in Computational Cognitive Modeling. *Trends in Cognitive Sciences* **21**, 425–433 (2017).
73. Wilson, R. C. & Collins, A. G. Ten simple rules for the computational modeling of behavioral data. *Elife* **8**, e49547 (2019).
74. Nasioulas, A., Potier, E., Cerrotti, F., Lebreton, M. & Palminteri, S. Feedback-induced attitudinal changes in risk preferences. Zenodo <https://doi.org/10.5281/zenodo.17807047> (2025).

Acknowledgements

For stimulating and helpful discussion: Enrico Diecidue, Florent Meyniel, Marion Rouault, Ido Erev, and the members of the Human Reinforcement Learning team of ENS. For providing their data: Ryan Jessup. SP is supported by the European Research Council under the European Union's Horizon 2020 research and innovation program (ERC) (RaReMem: 101043804), and the Agence Nationale de la Recherche (CogFinAgent: ANR-21-CE23-0002-02; RELATIVE: ANR-21-CE37-0008-01; RANGE: ANR-21-CE28-0024-01). The Département d'Etudes Cognitives is funded by the Agence Nationale pour la Recherche (ANR-17-EURE-0017, ANR-10- IDEX-0001-02). This work has received support under the Major Research Program of PSL Research University "PSL-Neuro" launched by PSL Research University and implemented by ANR (ANR-10-IDEX-0001). ML and AN are supported by the European Research Council under the European Union's Horizon 2020 research and innovation program (ERC) (ERC StG INFORL 958671 awarded to ML). AN is additionally supported by the Foundation for Education and European Culture (IPEP).

Author contributions

SP provided the initial idea. AN, EP, FC, and SP developed the method. AN, EP, and FC validated the method and conducted the experiments. AN, ML, and SP proposed theoretical development and interpretation. AN processed and analyzed the data. AN realized the systematic literature review. AN and ML prepared figures and visualization. AN and SP wrote the original manuscript. ML reviewed and edited the manuscript. ML and SP acquired funding for the study.

Competing interests statement

The authors declare no competing interests.

Tables

Table 1: main results of the logistic regression analyses across Experiments 1–6 (analyzed both individually and in aggregate) and the Erev replication. Only the predictors discussed in the main text are included here. For the detailed tables, containing all the predictors and the exact P values, see the Supplementary Materials Tables S7-S9. For more details on the regressions, see the Statistical Analysis section of the Methods. All models were fit as binomial generalized linear mixed-effects models (GLMMs) with a logit link. Fixed effects were evaluated using two-tailed Wald z-tests, and P values were not adjusted for multiple comparisons.

ARTICLE IN PRESS

Figure Legends/Captions

Figure 1: Experimental design (A, B) and basic behavioral results (C, D). (A) Example task screens showing the start of a block and a typical trial. For each variant (with/without block instructions; partial/complete feedback information), we specify in which experiment it has been implemented. (B) Decision problems in Experiments 1-6. Top panel: specification of the risky options. The Y and the X axes present in bold the decision variables (non-zero magnitude [pts] and probability [prob], respectively) of the risky lottery (L_{risky}), which always has the form (pts, prob; 0, 1-prob), while the numbers contained in the cell of the matrix (in gray italics) represent the Expected Value of the risky option (rEV). Bottom panel: specification of the sure/safe lotteries ($L_{\text{sure}}/L_{\text{safe}}$). In Exp1-4, we have a sure lottery, and in Exp5-6, we have a 50/50 safe lottery. Magnitudes depend on three independent factors of the design (pts and prob of the risky option, and whether the risky or the sure/safe option has a higher EV). (C) R-Rate (risky-choice rate) and O-rate (optimal-choice rate, namely Expected Value maximizing rates) as a function of the presence or absence of feedback (F and nF, respectively) across the seven experiments. Note: In Exp 7, the two options had equal Expected Values (EV), so O-rates are not defined. (D) The effect of feedback on R-rates (namely, the difference of R-rates between Feedback and no-Feedback conditions) as a function of which option (risky or sure/safe) has higher EV across experiments 1-6 (as explained above, Exp 7 is not eligible for this analysis either). In (C) and (D), each data point represents the mean of participant-level averages (i.e., one observation per participant per condition). Sample sizes (biological replicates = independent participants): Exp. 1 = 80, Exp. 2 = 95, Exp. 3 = 86, Exp. 4 = 85, Exp. 5 = 80, Exp. 6 = 84, Exp. 7 = 30. Error bars denote \pm SEM (standard error of the mean). Data are presented as mean \pm SEM. Statistical significance: two-tailed Wilcoxon signed-rank tests (***P < 0.001, ns = not significant). All non-significant results are also supported by Bayes Factors in favor of the null; Exp. 4 O-rates comparison P = 0.032, Exp. 4 panel D comparison P = 0.031.

Figure 2: behavioral results of Experiments 1-2. (A) Risky choice rate (R-rate) as a function of the feedback (F) vs. no-feedback (nF) condition in Exp.1. (B) Optimal choice rate (O-rate) as a function of feedback condition and whether or not the optimal response was the sure ('sureBetter') or the risky ('riskyBetter') option in Exp.1. (C) The colored lines represent the risky choice rate

as a function of the feedback condition and the trial number within a block (Exp.1). **(D)** and **(E)** display the same variables as **(A)** and **(B)**, but for Exp.2. **(F)** The panel displays the same variables as panel **(C)**, but for Exp.2. In addition, the inset displays the effect of feedback on R-rates (namely, the difference of R-rates between Feedback and no-Feedback conditions) for the first trial for both Exp.1 (no effect) and Exp.2 (significant effect); significant difference between Exp.1 & Exp.2 (two-tailed Mann–Whitney U test $P = 0.0056$ (**)). In **(A)**, **(B)**, **(D)**, and **(E)**, points indicate participant-level averages, violin plots indicate probability density functions, line segments connect the values of the participants in different conditions, orange lines have the direction of the means, black the opposite, boxes indicate 95% confidence intervals, and error bars indicate s.e.m.. Sample sizes (biological replicates = independent participants): Exp. 1 = 80, Exp. 2 = 95. In **(C)** and **(F)**, the central bold line is the mean of the participant-level averages in each condition, and the shaded area above and below the mean is plus and minus, respectively, the s.e.m. of the individual averages. The solid grey line is drawn in trials displaying a significant difference ($p < 0.01$) between the two conditions.

Figure 3: behavioral results of Experiments 3-4 and comparison to Experiments 1-2. **(A)** The colored lines represent the risky choice rate (R-rate) as a function of the feedback condition and the trial number within a block (Exp.3: no block-wise instructions). **(B)** The panel displays the same variables as panel **(A)** but for Exp.4 (block-wise instructions present). In addition, the inset displays the difference between the feedback and no feedback conditions for the first trial for both Exp.3 (no effect) and Exp.4 (significant effect). **(C)** Risky choice rate as a function of the feedback condition and the probability level of the risky option in Experiments 1/2 (partial feedback). **(D)** Same as **(C)** but for Experiments 3/4 (complete feedback). **(E)** The panel displays the effect of feedback on R-rates (namely, the difference of R-rates between Feedback and no-Feedback conditions) as a function of probability levels in the partial (Exps 1/2) and the complete (Exps 3/4) experiments. In the inset of **(B)**, two-tailed MWU test $P = 0.0019$ (**). In **(C)** and **(D)**, points indicate participant-level averages, line segments connect the values of the participants in different conditions, orange lines have the direction of the means, black the opposite; mean and error bars (s.e.m.) correspond to individual experiments (with the number of the experiment displayed on top of them); violin plots indicate probability density functions and their bold areas indicate 95%

confidence interval over the grouped data. In (B, inset) and (F), means, error bars (s.e.m.), and boxes (95% confidence interval) are computed over the averages of participants of the related (individual for inset (B) and grouped for (F)) experiments. Sample sizes (biological replicates = independent participants): Exp. 1 = 80, Exp. 2 = 95, Exp. 3 = 86, Exp. 4 = 85, Exp. 1,2 = 175, Exp. 3,4 = 171. In (A) and (B), the central bold line is the mean of the participant-level averages in each condition, and the shaded area above and below the mean is plus and minus, respectively, the s.e.m. of the individual averages. The solid grey line is drawn in trials displaying a significant difference between the two conditions ($P < 0.01$, two-tailed Wilcoxon signed-rank tests, Benjamini–Hochberg adjusted).

Figure 4: behavioral results across experiments 1-6. (A) The colored lines represent the risky choice rate (R-rate) as a function of the feedback condition and the trial number within a block (Experiments 1/3: no block-wise instructions). (B) The panel displays the same variables as panel (A) but for Experiment 2/4/5/6 (block-wise instructions present). In addition, the inset displays the effect of feedback on R-rates (namely, the difference of R-rates between Feedback and no-Feedback conditions) for the first trial, for experiments with and without instructions (significant difference between the two conditions: two-tailed MWU test $P < 0.001$ (***)). (C) Risky choice rate as a function of the feedback condition and the probability level of the risky option in Experiments 1/2/5 (partial feedback). (D) Same as (C) but for Experiments 3/4/6 (complete feedback). (E) The panel displays the effect of feedback on R-rates (namely, the difference of R-rates between Feedback and no-Feedback conditions) as a function of the probability levels in the partial (Exps 1/2/5) and the complete experiments (Exps 3/4/6). In (C) and (D), points indicate participant-level averages, line segments connect the values of the participants in different conditions, orange lines have the direction of the means, black the opposite; mean and error bars (s.e.m.) correspond to individual experiments (with the number of the experiment displayed on top of them); violin plots indicate probability density functions and their bold areas indicate 95% confidence interval over the grouped data. In (B, inset) and (F), means, error bars (s.e.m.), and boxes (95% confidence interval) are computed over the averages of participants of the related experiments. Sample sizes (biological replicates = independent participants): Exp. 1 = 80, Exp. 2

= 95, Exp. 3 = 86, Exp. 4 = 85, Exp. 5 = 80, Exp. 6 = 84, Exp. 1,2,5 = 255, Exp. 3,4,6 = 255, Exp. no-Instructions = 166, Exp. Instructions = 344. In (A) and (B), the central bold line is the mean of the participant-level averages in each condition, and the shaded area above and below the mean is plus and minus, respectively, the s.e.m. of the individual averages. The solid grey line is drawn in trials displaying a significant difference between the two conditions ($P < 0.01$, two-tailed WSRT, BH-adjusted).

Figure 5: feedback-induced trial-by-trial dynamics. (A) Risky choice rate (R-rate) in the second trial as a function of the feedback condition and the choice made in the first trial (Choice_{t-1} : sure/safe versus risky) for the no-instructions experiments. (B) R-rate in the second trial as a function of the feedback condition and the outcome obtained in the first trial (Out_{t-1} : zero versus positive) for the no-instructions experiments. (C) and (D) result from (A) and (B) respectively, by collapsing the feedback dimension (n.s. indicates Bayesian Factor in favor of the null, i.e., $\text{BF}_{01} > 1$). (E) Theoretical predictions of repeating a risky choice at trial t as a function of the outcome of the risky option in trial $t-1$ (Risky_{t-1}). In the bottom graph, two accounts are presented (the learning and the gambler's fallacy). In the top part, the color coding of the two conditions (risky choice having being rewarded with its maximum, with yellow, or minimum, with orange, outcome) is explained. (F) Behavioral results of the rate of repeating a risky choice (repR-rate) using the data of experiments 1-6 and feedback blocks only, because the analysis cannot be performed in blocks without feedback. Repeat risky rate is plotted as a function of the outcome of the risky choice in the previous trial and of the probability level of the risky option. (G) results from (F) by collapsing the dimension of the previous risky outcome (R_{t-1} : max or min). In all cases (except for the theoretical predictions in panel E), means, error bars (\pm s.e.m.), and boxes (95% confidence intervals) are computed over participant-level averages for the corresponding experiments. The number of participants included in each panel was as follows {panel: [conditions from left to right]}: A [164, 164, 156, 158]; B [137, 144, 166, 166]; C [163, 153]; D [127, 166]; F [482, 343, 473, 474, 385, 473]; G [343, 466, 382].

Figure 6: re-analysis of Erev et al. (2017). (A) Optimal choice rate as a function of the feedback condition. (B) Risky choice rate as a function of feedback condition and the probability level of the best outcome of the risky option (low: $\text{prob} \leq 25\%$, medium: $25\% < \text{prob} < 75\%$, high: $\text{prob} \geq 75\%$).

(C) The difference in risky choice rates between the feedback(F) and no feedback(nF) conditions across probability levels of the best outcome of the risky option. (D) Repeat risky choice rate as a function of the outcome of the risky choice in the previous trial and of the probability level of the best outcome of the risky option. In all cases (except for the theoretical predictions in panel E), means, error bars (\pm s.e.m.), and boxes (95% confidence intervals) are computed over participant-level averages for the corresponding experiments. The number of participants included in each panel was as follows {panel: [conditions from left to right]}: A [446, 446]; B [285, 285, 446, 446, 366, 366]; C [285, 446, 366]; D [233, 210, 424, 413, 287, 310].

Figure 7: BEAST model predictions using fitted parameters of our behavioral data. (A) The colored lines represent the risky choice rate as a function of the feedback condition and the trial number within a block (Experiments without block instructions). (B) The panel displays the same variables as panel (A), but for Experiments with block instructions. In addition, the inset displays the effect of feedback on R-rates (namely, the difference of R-rates between Feedback and no-Feedback conditions) for the first trial for both Experiments without block instructions (no effect) and Experiments with block instructions (no effect); non-significant difference between the two conditions of the model predictions ($P = 0.27$ two-sample two-tailed t-test). (C) Risky choice rate as a function of the feedback condition and the probability level of the risky option for Experiments with partial feedback (1, 2, 5). (D) Same as (C) but for Experiments with complete feedback (3, 4, 6). (E) The panel displays the effect of feedback on R-rates (namely, the difference of R-rates between Feedback and no-Feedback conditions) as a function of probability levels for the partial and the complete feedback experiments. (F) The rate of repeating a risky choice is calculated using the data of feedback blocks only, because the analysis cannot be performed in blocks without feedback. Repeat risky rate is plotted as a function of the outcome of the risky choice in the previous trial and of the probability level of the risky option. (G) The panel displays the difference in the rates of repeating a risky choice between the case where the outcome of the risky choice in the previous trial was maximum (positive) and the case where it was minimum (zero), across probability levels. In the insets of (B), (E), and (G), the black dots show the behavioral data. In all cases, means, error bars (\pm s.e.m.), and boxes (95% confidence intervals) are computed over the simulated values whose sample size matched that of the corresponding behavioral dataset. In (A) and (B), the solid grey line is drawn in trials displaying a significant difference between the two conditions ($P < 0.05$, two-tailed Wilcoxon signed-rank tests, Benjamini–Hochberg adjusted).

ARTICLE IN PRESS

Editor Summary

Normative theory predicts that feedback should not affect decisions under risk, but past findings disagree. Here, the authors show that feedback shifts risk-taking by changing attitudes rather than through learning.

Peer Review Information: *Nature Communications* thanks anonymous reviewers for their contribution to the peer review of this work. A peer review file is available.

ARTICLE IN PRESS

Dependent Variable: Risky choice rate

EXP1 EXP2 EXP3 EXP4
EXP5 EXP6 EXP All EXP
Erev

	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)	
FEEDBACK1	0.620*** (0.121)	0.492*** (0.135)	0.720*** (0.091)	0.704*** (0.104)	0.352*** (0.091)	0.018 (0.106)	0.677*** (0.107)	0.122 (0.120)	0.660*** (0.086)	0.571*** (0.103)	0.462*** (0.077)	0.017 (0.091)	0.582*** (0.038)	- 0.354*** (0.058)	-0.304*** (0.079)
FEEDBACK1:P_RISKY1	-	0.143 (0.097)	-	-0.100 (0.084)	-	0.339*** (0.089)	-	0.713*** (0.092)	-	0.122 (0.094)	-	0.508*** (0.085)	-	-	0.634*** (0.065)
FEEDBACK1:P_RISKY2	-	0.253* (0.102)	-	0.164 (0.089)	-	0.607*** (0.092)	-	0.849*** (0.094)	-	0.134 (0.094)	-	0.824*** (0.087)	-	-	1.295*** (0.083)

Dependent Variable: Optimal choice rate

EXP1 EXP2 EXP3 EXP4 EXP5
EXP6 EXP All EXP Erev

	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)		
FEEDBACK1	0.027 (0.095)	-0.696*** (0.108)	-0.008 (0.045)	-0.614*** (0.056)	-0.087 (0.058)	-0.436*** (0.069)	-0.126 (0.080)	-0.708*** (0.090)	0.063 (0.069)	-0.545*** (0.079)	-0.112 (0.066)	-0.531*** (0.076)	-0.046 (0.027)	-0.582*** (0.032)	0.027 (0.056)	-0.199*** (0.059)
FEEDBACK1:R_BETTER1	-	1.210*** (0.076)	-	1.125*** (0.063)	-	0.660*** (0.065)	-	1.038*** (0.071)	-	1.102*** (0.070)	-	0.794*** (0.064)	-	0.973*** (0.028)	-	0.611*** (0.053)

Dependent Variable: Probability of repeating a risky choice

EXP1 EXP2 EXP3 EXP4 EXP5 EXP6 EXP All EXP
Erev

(1) (2) (1) (2) (1) (2) (1) (2) (1) (2) (1) (2) (1) (2) (1) (2)

(RISKY _{t-1} =MAX)1	-0.883*** (0.241)	-1.595*** (0.321)	-1.128*** (0.166)	-1.665*** (0.254)	-0.775*** (0.190)	-1.133*** (0.294)	-1.460*** (0.218)	-3.237*** (0.380)	-0.513* (0.204)	-1.629*** (0.308)	-0.939*** (0.205)	-1.545*** (0.305)	-0.973*** (0.083)	-1.746*** (0.123)	-0.710*** (0.080)	-1.390*** (0.139)
(RISKY _{t-1} =MAX)1:P_RISKY1	-	0.624* (0.317)	-	0.465 (0.264)	-	0.278 (0.301)	-	2.022*** (0.373)	-	1.162*** (0.316)	-	0.595* (0.302)	-	0.784*** (0.123)	-	0.724*** (0.132)
(RISKY _{t-1} =MAX)1:P_RISKY2	-	1.884*** (0.380)	-	1.494*** (0.325)	-	0.987** (0.363)	-	2.662*** (0.443)	-	2.042*** (0.378)	-	1.388*** (0.372)	-	1.670*** (0.150)	-	1.072*** (0.163)













