©American Psychological Association, 2021. This paper is not the copy of record and may not exactly replicate the authoritative document published in the APA journal. The final article is available, upon publication, at: <https://doi.org/10.1037/dec0000156>

Regret in Experience-Based Decisions:

The Effects of Expected Value Differences and Mixed Gains and Losses

William M. Hayes and Douglas H. Wedell

University of South Carolina

Author Note

Department of Psychology, University of South Carolina, 1512 Pendleton Street, Columbia SC 29208, USA

Correspondence concerning this article should be addressed to William Hayes (Email: wmh1@email.sc.edu).

This research was supported by a SPARC grant from the Office of the Vice President for Research at the University of South Carolina.

We are grateful to Stefano Palminteri, who suggested testing a model with asymmetric learning rates. We would also like to acknowledge the helpful feedback from reviewers and the editor, David Budescu, who suggested the possibility of a rejoice-stay, regret-shift heuristic.

Abstract

Previous research on experience-based decisions with complete feedback supports the idea that people generally prefer options that produce better outcomes most of the time. The current study explored whether this preference is modulated by differences in expected value (EV) and the presence or absence of occasional losses. Participants ($n = 52$) recruited through a crowdsourcing platform completed an online experiment that involved repeated choices between a safer and a riskier option while receiving complete feedback. The riskier option yielded a better outcome on 80% of draws so that choosing it minimized the probability of regret. Preference for the riskier option was high when it had the same EV as the safer option and all outcomes were gains, but it decreased when the safer option had a higher EV and when both options included occasional losses. These findings replicated the results of a preliminary experiment with undergraduate participants ($n = 105$). Outcome ratings obtained on 50% of trials showed large effects of regret and rejoicing, confirming that participants were sensitive to relative comparisons between obtained and forgone outcomes. Reinforcement-learning modeling indicated that the effects of unequal EVs and mixed outcomes could be accounted for by assuming combined encoding of absolute and relative outcomes and unequal weighting of gains and losses. Overall, our results suggest that minimizing the probability of regret is an important motivational factor in experience-based decisions, but structural features of the choice environment can modulate the extent to which decision makers follow this strategy.

Keywords: reinforcement-learning models; complete feedback; regret theory; relative valuation; loss aversion

Regret in Experience-Based Decisions: The Effects of Expected Value Differences and Mixed Gains and Losses

In many decision scenarios, people learn the outcomes for both their course of action as well as what they would have received if they had chosen differently. For example, investors can monitor the value of stocks they choose to buy as well as ones they forgo. Such information naturally leads to relative or counterfactual comparisons between obtained and forgone outcomes. Feelings of *regret* may occur when obtained outcomes are worse than forgone outcomes. Conversely, feelings of *rejoicing* may occur when the opposite is true (Mellers et al., 1997). A desire to avoid feelings of regret can explain why people bet on gambles that have a very small probability of winning, such as horse races and lotteries, and at the same time purchase home or automobile insurance to protect against improbable large losses (Bell, 1982). Anticipated feelings of regret and rejoicing may play a key role in guiding choice under risk and uncertainty (Bell, 1982; Coricelli et al., 2005; 2007; Loomes & Sugden, 1982; Mellers, 2000), and can lead to either risk aversion, as with the insurance example, or risk seeking, as with the lottery example (Zeelenberg & Beattie, 1997; Zeelenberg et al., 1996).

Most of the literature on the role of anticipated feelings in decision-making has concerned decisions from description, in which possible outcomes and their probabilities are known in advance (Hertwig et al., 2004). Two types of counterfactual comparisons are possible with this kind of prior knowledge (Mellers et al., 1997; 1999). Comparing obtained outcomes against other potential outcomes that the *same* choice could have produced under a different set of circumstances can lead to feelings of *elation* or *disappointment* when the obtained outcome is better or worse, respectively (Bell, 1985; Loomes & Sugden, 1986). Note that these feelings are based on explicit knowledge of the chosen alternative's potential outcomes. On the other hand, feelings of rejoicing or regret can occur whenever individuals have access to what would have happened if they had made a *different* choice (Bell, 1982; Loomes & Sugden, 1982). For example, choosing an option that wins \$10 would normally be a pleasurable

experience, but learning that choosing a different option would have resulted in \$100 would likely make the overall experience negative. In the current study, we focus exclusively on these feelings of regret and rejoicing that result from relative valuation to the unchosen options.

Consider an individual facing a single choice between two gambles and who has explicit knowledge of the possible outcomes and probabilities. Further, the individual knows that after they make a choice, they will receive complete feedback on the outcomes of both gambles. According to the classic formulation of regret theory by Loomes and Sugden (1982, p. 810), the individual will have a weak preference for gamble i over gamble j if

$$\sum_{k=1}^n p_k \cdot Q(u(x_{ik}) - u(x_{jk})) \geq 0,$$

where x_{ik} is the outcome from gamble i under the k th combination of outcomes and x_{jk} is the outcome from gamble j under the k th combination of outcomes. There are n possible combinations of outcomes from gambles i and j with corresponding probabilities p_k . The two main components of regret theory are the continuous strictly increasing utility function $u(\cdot)$, which captures the individual's sensitivity to reward, and the continuous strictly increasing skew-symmetric function $Q(\cdot)$, which depends on the difference between the utilities of the obtained and forgone outcomes.¹ Rejoicing occurs when obtained outcomes are better than forgone ones and regret occurs when obtained outcomes are worse than forgone ones. Regret theory assumes that individuals anticipate these feelings when deciding between risky prospects and seek to minimize regret (or, conversely, maximize rejoicing) in addition to maximizing expected utility (Loomes & Sugden, 1982).

The predictions of regret theory primarily depend on the form of the function $Q(\cdot)$. If $Q(\cdot)$ is linear, regret theory predicts the same pattern of preferences as expected utility theory, and if $Q(\cdot)$ is concave, it predicts violations of expected utility theory but in the opposite direction to

¹ Skew-symmetry means that for any positive outcome ξ , $Q(\xi) = -Q(-\xi)$.

those typically observed in decisions from description. It is only when $Q(\cdot)$ is convex that regret theory predicts regret averse behavior, such as the simultaneous desire to gamble for improbable large gains and purchase insurance to protect against improbable large losses (Loomes & Sugden, 1982). Thus, $Q(\cdot)$ is typically assumed to be convex, implying high sensitivity to the magnitude of the difference between obtained and forgone outcomes. In decisions from description, regret theory predicts risk seeking for low probabilities of gains and risk aversion for high probabilities of gains if $Q(\cdot)$ is convex (Somasundaram & Diecidue, 2017). For an example of the latter, consider a choice between a safer gamble $S = (.80, \$90; .20, \$65)$ and a riskier gamble $R = (.80, \$95; .20, \$45)$, where the outcomes of both gambles are assumed to be independent. In this choice scenario, both gambles have equal expected value (EV) but R yields a better outcome than S with $p = .80$. A regret-averse individual should prefer S because they wish to avoid the possibility of choosing R , receiving the lower outcome (\$45), and finding out that they would have received the higher outcome from S (\$90). On the other hand, the difference between the lower outcome from S (\$65) and the higher outcome from R (\$95) is smaller and should therefore elicit less regret. More generally, given a choice between a risky gamble and a safer option, where the former is likely to produce a higher outcome but has a small probability of producing a substantially lower outcome, a regret-averse individual is expected to prefer the safer option. This specific prediction from regret theory has been empirically supported in studies using description-based choice tasks (Somasundaram & Diecidue, 2017).²

However, a very different pattern of behavior emerges when individuals make these types of decisions from experience; that is, when instead of having explicit knowledge of the possible outcomes and their probabilities up front, they learn this information gradually by

² The other prediction of risk-seeking for small probabilities of gains has received less empirical support (see Somasundaram & Diecidue, 2017 for details). In the current study, we focus on cases in which the riskier option has a high probability of yielding a better outcome than the safer option, where the results from studies of description-based decisions are mostly consistent with regret theory.

sampling from the options and receiving feedback (Hertwig et al., 2004). In these situations, individuals often behave as if their goal is to minimize the *probability* of regret by selecting options that produce better outcomes most of the time (Barron & Erev, 2003; Erev et al., 2017; Hertwig et al., 2004; Yechiam & Busemeyer, 2006). In a study by Ahn et al. (2012), participants made repeated choices between a sure reward and a risky option that yielded a higher or lower reward with probabilities that had to be learned through experience. Importantly, the two options always had equal EVs but the probability that the risky option yielded a higher reward increased from .20 to .40 to .60 to .80 across four separate games. Selections of the risky option increased as its probability of yielding a higher reward increased, and computational modeling suggested that participants sought to avoid making choices that had a high expected probability of resulting in regret (Ahn et al., 2012; see also Haines et al., 2021). In the choice between the hypothetical gambles R and S discussed above, minimizing the probability of regret would lead to a preference for R since it produces a better outcome than S on 80% of draws. Although this kind of risk-seeking preference conflicts with the standard predictions of regret theory, it could arise under the assumption of a concave $Q(\cdot)$ function (see Supplemental Figure S1).

When EVs are equal, choosing the riskier option that yields better outcomes most of the time seems like a reasonable strategy. However, a reliance on frequencies can sometimes lead to suboptimal selections of risky options that have lower EVs simply because they often produce the highest outcomes (“big eyes effect”; Grosskopf, Erev, & Yechiam, 2006; Yechiam & Busemeyer, 2006; Yechiam et al., 2005). It is unclear to what extent individuals would still prefer risky options that minimize the probability of regret when they have a lower EV than a safer option. A second question concerns the presence of occasional negative outcomes. Losses have been shown to have a variety of behavioral and physiological effects in decision-making contexts. For example, losses lead to greater increases in autonomic arousal compared to gains (Löw et al., 2008; Studer & Clark, 2011; Wright & Rakow, 2017). Although the heightened arousal may be associated with behavioral loss aversion, or avoidance of choice options that

yield the largest losses (Hayes & Wedell, 2020a; Sokol-Hessner et al., 2009), this is not always the case (Hochman & Yechiam, 2011). Yechiam and colleagues' loss attention model posits that minor losses act as attention focusers and may promote EV-maximization when long-term payoffs differ (Yechiam & Hochman, 2013a; 2013b; Yechiam et al., 2015). One proposed mechanism accounting for this is an increase in the choice consistency or inverse temperature parameter in loss contexts, which heightens sensitivity to expectancies for each option and decreases random exploration (Yechiam & Hochman, 2013a; 2013b). This would imply that losses could either enhance EV-maximization or regret-minimization strategies, depending on how expectancies are formed.

Current Study

The current study focuses on experience-based choice problems involving a riskier option (R) and a safer option (S) with complete outcome feedback. The key feature of our paradigm is that R produces better outcomes than S on 80% of draws so that it minimizes the probability of regret. However, due to the infrequent low outcomes from R, regret theory would predict risk averse preferences under the standard assumption that $Q(\cdot)$ is convex. The main experiment ($n = 52$) included four conditions that were manipulated within-subjects. In two of the conditions, both options had equal EVs, and outcomes were either gains only (Gains-Equal-EVs) or mixed gains and losses (Mixed-Equal-EVs). In the other two conditions, the EV of the riskier option dropped below that of the safer option, and the outcomes were either gains only (Gains-Unequal-EVs) or mixed gains and losses (Mixed-Unequal-EVs). Note that the unequal EV conditions pitted long-term payoffs against the probability of experiencing immediate regret. We also present results from a preliminary experiment ($n = 105$) which included three out of the four conditions manipulated within-subjects (additional details about the methodology for the preliminary experiment can be found in Supplemental Note 1). The overall results from both experiments were highly similar.

We tested a set of reinforcement-learning models (RL; Sutton & Barto, 1998) on their

ability to account for the choice patterns observed in the main experiment. RL models decompose experience-based decisions into at least three stages: outcome encoding, updating, and action selection (Rangel et al., 2008). Each stage can be represented by a set of equations with their own parameters. We selected nine candidate models that were formed by crossing three different hypotheses about the way outcomes are encoded with three hypotheses about the role of losses. General equations for the various model components are shown in Table 1, along with the fixed and free parameters for each of the nine models.

Table 1*Reinforcement-Learning Models*

Model Component	Equation					
Outcome Encoding	$u_i(t) = \begin{cases} x_i(t)^A & \text{if } x_i(t) \geq 0 \\ -\lambda \cdot x_i(t) ^A & \text{otherwise} \end{cases} \quad (1)$					
	$V_i(t) = (1 - \omega) \cdot u_i(t) + \omega \cdot r_{ij}(t) \quad (2)$					
Updating	$E_i(t + 1) = E_i(t) + \alpha \cdot [\delta_i(t) + (1 - \delta_i(t)) \cdot \gamma] \cdot (V_i(t) - E_i(t)) \quad (3)$					
Action Selection	$\text{Pr}_i(t + 1) = \frac{\exp \{\theta \cdot E_i(t + 1)\}}{\exp \{\theta \cdot E_i(t + 1)\} + \exp \{\theta \cdot E_j(t + 1)\}} \quad (4)$					
	Absolute Encoding		Diminishing Sensitivity		Partially Relative Encoding	
	Fixed	Free	Fixed	Free	Fixed	Free
Loss Neutrality	$A = 1$ $\lambda = 1$ $\omega = 0$	α γ θ	$\lambda = 1$ $\omega = 0$	A α γ θ	$A = 1$ $\lambda = 1$	ω α γ θ
Loss Aversion	$A = 1$ $\omega = 0$	λ α γ θ	$\omega = 0$	A λ α γ θ	$A = 1$	λ ω α γ θ
Loss Attention	$A = 1$ $\lambda = 1$ $\omega = 0$	α γ θ_{Gains} θ_{Mixed}	$\lambda = 1$ $\omega = 0$	A α γ θ_{Gains} θ_{Mixed}	$A = 1$ $\lambda = 1$	ω α γ θ_{Gains} θ_{Mixed}

Note. A = utility function shape; λ = loss weighting; ω = relative encoding; α = learning rate; γ = counterfactual learning; θ = inverse temperature (may differ for gains and mixed contexts). In Equation 2, $r_{ij}(t)$ is an indicator variable that equals 1 if option i 's outcome was better than option j 's on trial t and 0 otherwise. In Equation 3, $\delta_i(t)$ is an indicator that equals 1 if option i was chosen on trial t and 0 otherwise. The bottom half of the table shows the fixed parameters and free parameters in each model. The parameters ω , α , and γ were constrained to values between 0 and 1, A was constrained to values between 0 and 1.5, λ was constrained to values between 0 and 5, and $\theta = 3^c - 1$ ($0 \leq c \leq 5$) was constrained to values between 0 and 242.

For the outcome encoding component, the first and simplest possibility we tested is that the subjective utility of the outcome produced by option i on trial t , $u_i(t)$, is equal to the objective value of the outcome, $x_i(t)$ (absolute encoding). A second possibility, derived from prospect theory (Kahneman & Tversky, 1979), is that subjective utilities are nonlinearly related to objective values, and that individuals are less sensitive to differences in large magnitudes than to differences in small magnitudes (diminishing sensitivity). When the utility parameter A is below 1, individuals become risk averse in the domain of gains and risk seeking in the domain of losses, as predicted by prospect theory. The opposite pattern arises when A is above 1. A third possibility is that when individuals receive complete feedback from options i and j on trial t , they encode a weighted combination of each outcome's utility and whether each outcome is better or worse than the other (partially relative encoding; see Bavard et al., 2018 for a similar model). In Table 1, $r_{ij}(t)$ is an indicator variable that equals 1 if option i 's outcome is better than option j 's and 0 otherwise (see Equation 2). As the relative encoding parameter ω approaches 1, individuals attend more to these ordinal comparisons and less to the actual magnitudes of the outcomes. The partially relative encoding hypothesis is consistent with an alternate version of regret theory in which the only information encoded is the sign of the difference between obtained and forgone outcomes (i.e., an extremely concave or step function for $Q(\cdot)$). In essence, this model incorporates regret into experience-based decisions by assuming individuals consider the *probability* that one option will produce a better outcome than the other when deliberating their next choice.

The three outcome encoding hypotheses above were crossed with three hypotheses regarding the role of losses. The first is simply that losses have no unique effects on behavior relative to gains (loss neutrality). However, according to prospect theory, losses and gains of the same magnitude are weighted unequally, with losses typically having greater impact (loss aversion; Kahneman & Tversky, 1979). When the loss weighting parameter λ exceeds 1, losses

have greater subjective impact than gains; on the other hand, λ below 1 produces loss insensitivity. The third possibility that we tested is that the heightened arousal elicited by losses serves to focus attention on the task at hand, which increases choice consistency (or decreases randomness) without producing loss aversion (loss attention; Yechiam & Hochman, 2013a; 2013b; Yechiam et al., 2015). To accomplish this, separate inverse temperature parameters θ_{Gains} and θ_{Mixed} were estimated in the gains and mixed outcomes conditions. For the models that assumed loss neutrality or loss aversion, only a single θ was estimated. Higher values of θ lead to more deterministic choice behavior while lower values lead to more random behavior in the softmax choice rule (see Table 1, Equation 4).

All models assumed that feedback-driven updating of expectancies for each option occurs via a delta learning rule (Rescorla & Wagner, 1972; Yechiam & Busemeyer, 2005) (see Table 1, Equation 3). The learning rate parameter α controls the degree to which expectancies E_i are updated by prediction errors, or the differences between realized outcomes and prior expectancies. Higher values of α lead to faster learning and greater recency effects, while lower values lead to slower, more gradual learning. We included a counterfactual learning parameter γ that modulates the learning rate for nonchosen options relative to the learning rate for chosen options (Yechiam & Busemeyer, 2005). Expectancies were initialized to zero on the first trial of each condition.

The nine RL models formed by crossing absolute encoding, diminishing sensitivity, and partially relative encoding with loss neutrality, loss aversion, and loss attention were each simulated 1,000 times in our choice task using parameters that were specifically chosen to emphasize the differences between models. The results of these simulations are shown in Figure 1. One observation that can be made from this figure is that the partially relative encoding models were the only ones that produced a choice share for R above 0.50 when EVs are equal under the chosen parameter settings (see figure caption). This is because both

options have equal long-run payoffs and produce high outcomes at the same rate, but R produces better *relative* outcomes more frequently than S. A second observation is that all nine models predict a decrease in preference for R when EVs are unequal. Note that if ω is set close to 1, implying nearly exclusive attention to relative comparisons, then the partially relative encoding hypothesis could account for a consistent high preference for R even when it is suboptimal (not shown in Figure 1). A third observation is that the loss aversion and loss attention models make noticeably different predictions about the effect of mixed outcomes. Whereas loss aversion produces a pattern resembling a main effect of mixed outcomes, loss attention predicts an interaction, with greater sensitivity to unequal EVs when outcomes are mixed.

Lastly, we collected affective ratings of outcomes throughout the task in order to verify that regret and rejoicing effects occur even without explicit knowledge of outcomes and probabilities. We expected regret and rejoicing effects would be present in all conditions such that the same outcome would be evaluated more positively when the counterfactual outcome was lower in value and more negatively when the counterfactual outcome was higher (Mellers et al., 1997; 1999).

Figure 1*Reinforcement-Learning Model Simulations*

Note. Each panel depicts the average simulated proportion of times (across 1,000 simulations) that the riskier, regret-minimizing option (R) was chosen in each condition using one of the nine models under consideration. The learning rate and counterfactual learning parameters were the same for all models: $\alpha = .2$ and $\gamma = .5$. The relative encoding parameter (ω) was set to .7 for the models that assume partially relative encoding and 0 for all others. The shape parameter (A) was set to 0.5 for the models that assume diminishing sensitivity and 1 for all others. Loss weighting (λ) was set to 2.5 for the models that assume loss aversion and 1 for all others. For the models that assume loss attention, the choice sensitivity parameters for gains (θ_G) and mixed conditions (θ_M) were set to 0.5 and 2.5, respectively. For all other models, the single consistency parameter (θ) was set to 1.5.

Method

Participants

A total of 58 participants were recruited through the crowdsourcing platform Prolific (<http://www.prolific.ac>) to take part in an online experiment. Each provided informed consent and all aspects of this study were approved by the Institutional Review Board at the University of South Carolina. Participants were paid just over \$6.00 each for their time. Six were excluded for taking too long to finish (> 3 SDs from the mean completion time) or using a smartphone (the task was optimized for a tablet or desktop computer), resulting in a final sample of 52 participants (18 men, 32 women, 2 non-binary; ages 18-63, $M = 32.69$, $SD = 10.69$; median completion time = 55 min). All participants selected the U.S. as their nationality and had an approval rate on Prolific of at least 75%. An *a priori* power analysis indicated that 46 participants would be required to detect a difference between conditions as large as that observed in the preliminary experiment ($d = 0.49$) with 90% power (Supplemental Note 1).

Design

There were four experimental conditions that were presented in blocks of 80 trials and manipulated within-subjects (Table 2). The order of blocks was randomized for each participant. In all four conditions, R and S yielded higher outcomes (points) with $p = .80$ and lower outcomes with $p = .20$, but the higher and lower outcomes from R were always more extreme than those from S. On any given draw, the probability that R would produce a better outcome than S was .80. Two of the conditions (Gains) included strictly positive outcomes while the other two (Mixed) included both gains and losses. In each domain, the long-term payoffs yielded by R and S were either the same on average (Equal EVs) or the EV of R was 10 points less than S (Unequal EVs) so that minimizing the probability of regret would result in lower long-run payoffs. Outcomes were generated prior to the experiment according to the distributions in Table 2 and rounded to the nearest integer. The sequence of outcomes within each condition was randomized for each participant.

Table 2*Outcomes, Probabilities, and Expected Values for the Choice Options in each Condition*

	p	Gains Equal EVs	Gains Unequal EVs	Mixed Equal EVs	Mixed Unequal EVs
Safer Option (S)	.80	$90 + \varepsilon$	$100 + \varepsilon$	$10 + \varepsilon$	$20 + \varepsilon$
	.20	$65 + \varepsilon$	$75 + \varepsilon$	$-15 + \varepsilon$	$-5 + \varepsilon$
	EV	85	95	5	15
Riskier Option (R)	.80	$95 + \varepsilon$	$105 + \varepsilon$	$15 + \varepsilon$	$25 + \varepsilon$
	.20	$45 + \varepsilon$	$5 + \varepsilon$	$-35 + \varepsilon$	$-75 + \varepsilon$
	EV	85	85	5	5

Note. The riskier option minimized the probability of regret in all four conditions. A small amount of Gaussian noise ε was added to each outcome, $\varepsilon \sim N(0,1)$. All outcomes were presented as points rounded to the nearest integer. p = probability, EV = expected value, computed as the sum of outcomes weighted by their associated probabilities.

Procedure

The experiment was administered online in the form of a Qualtrics survey. The instructions stated that participants would be choosing repeatedly between two options represented by symbols in four separate games. Choices would result in either winning or losing points, and the points earned in each game would be added to their total. Their goal was to finish each game with as many points as possible. Participants were told that they would get to see the points won or lost after each choice as well as the points they *would have* won or lost had they chosen the other option. Only the points resulting from the chosen option would affect their total (an example was provided). Next, participants were informed that they would sometimes be asked to rate how they felt about an outcome using a two-dimensional grid that ranged from negative (left) to positive (right) on the horizontal axis and low arousal (bottom) to high arousal (top) on the vertical axis (see Hayes & Wedell, 2020b, for a similar application). Several examples of how to use the grid were shown. Finally, to increase motivation, participants were told that one of the four games would be randomly selected at the end of the

experiment and the total points earned in that game would be compared to an unknown bonus threshold. If their total earnings exceeded the threshold, \$0.50 would be added to their participation payment. In reality, all participants were told that they had received the bonus at the end.

At the beginning of each block, participants saw the two symbols that they would be choosing between for the next several trials and their point total was reset to zero. Participants were not informed of the total number of trials in a block. The symbols and screen positions for R and S in each block were randomly assigned to each participant, but both were held constant across trials. Trials began with a 0.5 s fixation, followed by the choice display with a cue to make a selection (Figure 2A). The current point total was always displayed at the top of the screen. After making a choice, participants saw the points they won or lost below the chosen symbol (emphasized with a black border), along with the points they would have won or lost below the nonchosen symbol (Figure 2B). Trials were self-paced.

Outcome ratings occurred on 50% of trials. To accomplish this, a random number between 0 and 1 was generated on each trial and the program decided whether to obtain a rating based on whether the number exceeded 0.5. On rating trials, participants rated the valence and arousal of the outcome they just viewed by clicking a location within the 480 × 480 pixel affect rating grid (Figure 2C). At the top of the screen were the words, “How do you feel about the outcome you received?” Their response was echoed by a dot on the grid until they advanced to the next trial, and they could change their response as many times as they wanted before advancing. The software recorded the x-y coordinates of their response (in pixels). These coordinates were transformed prior to analysis so that valence ratings ranged between -240 and 240.

Figure 2*Choice and Affect Rating Procedure*

Note. (A) Each trial began with a cue to choose between two symbols, with the participant's current point total displayed at the top of the screen. (B) Choices were immediately followed by feedback on both obtained and forgone outcomes, and the running total was updated. (C) On 50% of trials, participants rated how they felt about the outcome using the affect grid. The horizontal axis ranged from "Negative" (left) to "Positive" (right) and the vertical axis ranged from "Low Arousal" (bottom) to "High Arousal" (top).

Model Fitting and Comparison

The nine RL models discussed in the Introduction were fit to the sequence of 320 choices and outcomes (80 trials per condition) for each participant using individual Bayesian estimation. Choice outcomes were scaled down by a factor of 10 to avoid numerical issues arising in the softmax choice rule. All parameters were first rescaled to the [0,1] interval and given Uniform(0,1) priors. We used Stan software (version 2.21; Carpenter et al., 2017) to obtain 5,000 samples from the joint posterior distribution of the model parameters for each participant. Convergence to the posterior was monitored by ensuring that all $\hat{R}$ values, a relative index of between- to within-chain variability, were below 1.05 (Gelman et al., 2013). An

importance sampling technique was used to estimate the marginal likelihood of each participant's data, i.e. the likelihood of the data averaged over the parameter space, where the averaging weights are determined by the prior (Steingroever et al., 2016).³ This process was repeated for each model. Then, the estimated marginal likelihoods from each model were used to compute Bayes factors (BFs) favoring one model over another. Let M_1 and M_2 denote two competing models. The Bayes factor in favor of M_1 over M_2 is simply a ratio of their marginal likelihoods: $BF_{12} = \Pr(\text{Data}|M_1)/\Pr(\text{Data}|M_2)$. A BF of 10, for example, would indicate that the observed data are 10 times more likely under M_1 than under M_2 . So that more than two models could be compared at once, we also transformed marginal likelihoods into posterior model probabilities using the formula $\Pr(M_i|\text{Data}) = \Pr(\text{Data}|M_i)/\sum_k \Pr(\text{Data}|M_k)$, where k ($k = 1, \dots, K$) indexes the K models under consideration (Wasserman, 2000). Note that this formula assumes equal prior model probabilities ($\Pr(M_1) = \Pr(M_2) = \dots = \Pr(M_K)$). The nine RL models were compared to a baseline model that assumes a constant probability of choosing R in each of the four conditions to verify that they could account for trial-to-trial learning effects.⁴

After performing model comparison, the best model was fit hierarchically to the data for all 52 participants to draw inferences from parameters. Hierarchical estimation has several advantages over independent estimation, including borrowing information from the group to constrain individual estimates (Scheibehenne & Pachur, 2015; Wetzels et al., 2010). In our implementation, the individual-level parameters were first rescaled to the $[0,1]$ interval and then probit-transformed, e.g., $\lambda' = \text{Probit}(\lambda/5)$ (the probit function is the inverse-CDF for the normal distribution). The probit-scale parameters were given normal group-level priors with group mean

³ We drew 50,000 samples from the beta mixture importance densities described in Steingroever et al. (2016) to compute approximate marginal likelihoods, which resulted in estimates with extremely low variance.

⁴ This is simply a Bernoulli model that includes four parameters corresponding to the probability of choosing R in each condition. If y_k is the observed number of R choices in the k^{th} condition and n_k is the total number of trials in the k^{th} condition, then under the assumption of uniform priors on the parameters, the marginal likelihood of the baseline model is given by $B(y_k + 1, n_k - y_k + 1)$, where $B(\cdot)$ denotes the beta function.

$\mu_{(c)}$ and standard deviation $\sigma_{(c)}$. Group means were given Normal(0, 1) hyperpriors and group standard deviations were given Half-Normal(0, 0.2) hyperpriors, which results in roughly uniform priors for the individual-level parameters over their allowable ranges (e.g., Haines et al., 2018). Using Stan, we obtained 2,500 posterior samples from four separate chains after discarding the first 2,500 samples as warmup, resulting in 10,000 usable samples. Convergence to the posterior was confirmed through visual inspection and by checking that all $\hat{R} < 1.01$. Code for reproducing the modeling analyses (R and Stan) is available at <https://osf.io/bp3cx/>.

Results

Preliminary Experiment

The preliminary experiment was conducted online and included 105 undergraduate participants who received partial course credit for participating. The primary difference from the main experiment is that there was no Mixed-Equal-EVs condition. All participants completed the Gains-Equal-EVs condition first, followed by the remaining two conditions in one of two randomly assigned orders. Each condition had 100 choice trials instead of 80, and outcome ratings occurred on all trials. Full details about the method can be found in Supplemental Note 1 and Supplemental Table S1 includes a summary of the three conditions.

Valence Ratings

To determine whether counterfactual outcomes elicited feelings of regret and rejoicing in this paradigm, we focused on the valence dimension of the outcome ratings. Regret and rejoicing effects were especially prominent when S was chosen: The same outcomes were rated more positively when the forgone outcome from R was low compared to when it was high (Supplemental Figure S2). Random-intercepts regression confirmed that the effect of forgone outcomes on valence ratings was larger when S was chosen, but the difference was smaller in the Mixed-Unequal-EVs condition compared to other two conditions with only gains. One possible interpretation of this result is that the presence of losses decreased the intensity of

counterfactual emotions overall because of a greater attention to EV differences (Yechiam & Hochman, 2013a; 2013b; Yechiam et al., 2015). Additional details, including an analysis of arousal ratings, are available in the Supplemental Material (see Figures S2 and S3 for plots of mean ratings and Tables S2 and S3 for regression statistics).

Choice Patterns

The analysis of choice behavior focused on the proportion of times R was chosen across five blocks of 20 trials in each condition (Supplemental Figure S4). An overall preference for R developed quickly in the Gains-Equal-EVs condition and remained constant after the second block of trials. That is, the average participant was risk seeking when EVs were equal and all outcomes were gains. In the Gains-Unequal-EVs condition, R was slightly preferred overall but the average learning curve remained flat across all five blocks. It was only in the Mixed-Unequal-EVs condition that the proportion of R selections decreased across blocks, suggesting a gradually increasing preference for the option with a higher EV. An analysis of variance (ANOVA) indicated a significant main effect of Condition and a Condition $\times$ Block interaction on the proportion of R choices. There were no significant differences between conditions in the first block, but in the subsequent blocks, R was selected more frequently in the Gains-Equal-EVs condition compared to the other two conditions. Further, in the fourth and fifth blocks, R was chosen more frequently in the Gains-Unequal-EVs condition than in the Mixed-Unequal-EVs condition (see Supplemental Note 1 for ANOVA statistics).

Main Experiment

Because the preliminary experiment did not include a condition with equal EVs and mixed outcomes, it is difficult to draw strong conclusions from the above results alone. Thus, we now turn to the main experiment which used a complete 2 (EV: equal or unequal) $\times$ 2 (Domain: gains or mixed) within-subjects design. Overall, the results were very similar between the two experiments, suggesting that our findings are robust. We also include a computational modeling analysis to provide an algorithmic explanation of choice behavior in the main experiment.

Valence Ratings

We computed average valence ratings as a function of EV (equal or unequal), Domain (gains or mixed), Choice (S or R), Obtained Outcome (high or low), and Forgone Outcome (high or low) for each participant. The mean ratings for each participant were then submitted to a random-intercept regression using the factors above as predictors.⁵ For conciseness, only slopes and their associated p values for key effects are presented here (the Supplemental Material contains full results and a parallel analysis of arousal; see Tables S4 and S5 and Figure S5).

As shown in Figure 3, regret and rejoicing effects were especially apparent when S was chosen: the same outcomes from S were rated consistently lower when the forgone outcome from R was high compared to when it was low. The regression indicated that with all else constant, valence ratings were higher in the gains domain ($b = 21.32$, $p < .001$), when S was chosen ($b = 16.88$, $p < .001$), when the obtained outcome was high ($b = 50.81$, $p < .001$), and when the forgone outcome was low ($b = -15.39$, $p < .001$). The effect of EV was nonsignificant ($p = .97$). There were two significant two-way interactions of particular interest, one for Domain and Obtained Outcome ($b = -17.50$, $p < .001$) and one for Choice and Forgone Outcome ($b = -14.33$, $p < .001$). The former was due to the effect of Obtained Outcome being stronger in the mixed domain ($b = 68.32$, $p < .001$; Figure 3, right panels) than in the gains domain ($b = 33.31$, $p < .001$; Figure 3, left panels), which might suggest that occasional losses increased participants' sensitivity to value differences. The latter was due to the effect of Forgone Outcome being significant when S was chosen ($b = -29.72$, $p < .001$; Figure 3, top panels), but nonsignificant when R was chosen ($b = -1.06$, $p = .73$; Figure 3, bottom panels). This means that when participants won 105 points from R in the Gains-Unequal-EVs condition, for example,

⁵ Because ratings were obtained on 50% of trials, randomly selected for each participant, not all participants had data for each combination of factors. We used a random intercepts regression approach instead of a repeated measures ANOVA to analyze the data since the latter cannot handle empty cells.

their feelings were hardly affected by whether S produced 100 or 75 points. On the other hand, the way they felt after winning 100 points from S was highly dependent on whether R produced 105 or 5 points. This result is consistent with regret and rejoicing being more sensitive to the sign of the counterfactual difference than the magnitude.

Figure 3

Outcome Valence Ratings

Note. Ratings were centered around their midpoint so that they ranged from -240 to 240. Error bars represent ± 1 SE. S Chosen = safer option was chosen, R Chosen = riskier option was chosen.

Choice Patterns

The average choice proportion for R in each condition is shown in Figure 4A, and the evolution of risk preference across trials is shown in Figure 4B. We analyzed choice proportions for R using a 2 (EV: equal or unequal) $\times$ 2 (Domain: gains or mixed) $\times$ 4 (Block: trials 1-20, 21-40, 41-60, or 61-80) repeated measures analysis of variance (ANOVA).⁶ Greenhouse-Geisser ϵ

⁶ Note that when arcsine-transformed choice proportions were analyzed instead of raw proportions, the overall conclusions were the same.

values and corrected p values are presented in cases where sphericity was violated.

Figure 4

Choice Patterns and Model Fit in each Condition

Note. (A) Overall proportion of riskier, regret-minimizing (R) choices in each condition, averaged across participants and trials. Error bars represent ± 1 SE. The circles represent the posterior predictions from the hierarchical RL model. (B) Proportion of R choices across trials in each condition, averaged across participants. The black line represents the observed data and the gray ribbon represents ± 1 SE. The thick blue line represents model posterior predictions and the thin blue lines represent the 95% Bayesian credible interval.

There were significant main effects of EV, $F(1, 51) = 9.68$, $MSE = 0.18$, $p = .003$, $\eta_p^2 = .16$, and Domain, $F(1, 51) = 5.23$, $MSE = 0.28$, $p = .026$, $\eta_p^2 = .09$, that were qualified by significant EV $\times$ Block, $F(3, 153) = 8.29$, $MSE = 0.03$, $\varepsilon = .62$, $p < .001$, $\eta_p^2 = .14$, and Domain $\times$ Block interactions, $F(3, 153) = 5.32$, $MSE = 0.03$, $\varepsilon = .72$, $p = .002$, $\eta_p^2 = .09$. All other effects were nonsignificant [Block: $F(3, 153) = 0.20$, $MSE = 0.02$, $\varepsilon = .86$, $p = .90$; EV $\times$ Domain: $F(1, 51) = 0.02$, $MSE = 0.27$, $p = .89$; EV $\times$ Domain $\times$ Block: $F(3, 153) = 1.02$, $MSE = 0.02$, $\varepsilon = .79$, $p = .39$]. Follow-up t -tests with Bonferroni correction (four comparisons) indicated that the

differences between the EV conditions emerged in the third [$M_{\text{equal}} = .65$, $M_{\text{unequal}} = .51$; $t(51) = 4.13$, $p < .001$, $d = 0.58$] and fourth blocks of trials [$M_{\text{equal}} = .66$, $M_{\text{unequal}} = .50$; $t(51) = 3.65$, $p = .002$, $d = 0.51$], with more frequent risky selections when EVs were equal. The differences in the first [$M_{\text{equal}} = .59$, $M_{\text{unequal}} = .58$; $t(51) = 0.28$, $p > .99$] and second blocks [$M_{\text{equal}} = .60$, $M_{\text{unequal}} = .54$; $t(51) = 1.94$, $p = .23$] were nonsignificant. For the Domain $\times$ Block interaction, there was a marginally significant difference between gains and mixed domains in the second block [$M_{\text{gains}} = .63$, $M_{\text{mixed}} = .51$; $t(51) = 2.48$, $p = .066$, $d = 0.35$] and significant difference in the fourth block of trials [$M_{\text{gains}} = .65$, $M_{\text{mixed}} = .51$; $t(51) = 3.19$, $p = .01$, $d = 0.45$], with more frequent risky selections in the gains domain. The differences in the first [$M_{\text{gains}} = .59$, $M_{\text{mixed}} = .58$; $t(51) = 0.54$, $p > .99$] and third blocks [$M_{\text{gains}} = .61$, $M_{\text{mixed}} = .55$; $t(51) = 1.52$, $p = .54$] were nonsignificant.

Reinforcement-Learning Modeling

Table 3 summarizes the results of the model comparison that was carried out at the individual level. Three indices are shown for each of the nine RL models: (1) the median Bayes factor (BF) for that specific model versus the baseline (i.e., no learning) model, (2) the number of participants fit best by that model, and (3) the median posterior model probability across participants (see Supplemental Table S6 for a comparison using WAIC, a measure of posterior predictive deviance; Vehtari et al., 2017). Across all three criteria, the model that assumed partially relative outcome encoding and loss aversion was superior to the others. The median BF for this model over the baseline model (BF = 123.34) indicates that it was extremely favored over baseline for at least half of the participants in our sample (see Lee & Wagenmakers, 2013 for BF classifications). Posterior model probabilities can be summed across models that share a common component to give the posterior probability of that component at the individual level (Wagenmakers & Farrell, 2004). Across participants, the median posterior probabilities of absolute encoding, diminished sensitivity, and partially relative encoding were .003, .24, and .52, respectively. Similarly, the median posterior probabilities of loss neutrality, loss aversion, and loss attention were .30, .35, and .14, respectively.

Table 3*Model Comparison Results*

		Absolute Encoding	Diminishing Sensitivity	Partially Relative Encoding	Row Summary
Loss Neutrality	Median BF vs. Baseline	0.01	0.45	21.19	
	Individuals Fit Best (%)	1 (1.9 %)	8 (15.4 %)	9 (17.3 %)	18 (34.6 %)
	Median Posterior Prob.	< .01	.01	.08	.30
Loss Aversion	Median BF vs. Baseline	0.06	0.36	123.34 *	
	Individuals Fit Best (%)	2 (3.8 %)	5 (9.6 %)	15 (28.8 %)	22 (42.3 %)
	Median Posterior Prob.	< .01	.02	.14	.35
Loss Attention	Median BF vs. Baseline	< .01	1.06	23.63	
	Individuals Fit Best (%)	2 (3.8 %)	2 (3.8 %)	8 (15.4 %)	12 (23.1 %)
	Median Posterior Prob.	< .01	.03	.05	.14
Column Summary		5 (9.6 %) < .01	15 (28.8 %) .24	32 (61.5 %) .52	

Note. The best-fitting model for each participant was defined as the one with the highest marginal likelihood. The baseline model was included for assessing the relative performance of each model via Bayes factors, but was excluded when determining which model fit best and when computing posterior model probabilities. The column and row summaries show the number of participants fit best and the median posterior probability of each of the outcome encoding and loss weighting components, summing across the specific models that include those components. For interpretation purposes, we did not aggregate Bayes factors across models. BF = Bayes factor.

* Best model overall.

Having selected the partially relative encoding model with loss aversion as the best out of the set of candidates, we fit the data for all participants simultaneously using hierarchical Bayesian methods. The fit of the model to the group-averaged data is shown in Figure 4 (see Supplemental Figure S6 for fits to individual choice data). We were particularly interested in drawing inferences from the group-level mean parameters, $\mu_{(\cdot)}$ (Figure 5). Based on the 95% highest density interval (HDI) for μ_{ω} (.61 to .79), we can infer that the average participant

encoded a combination of absolute outcome values and ordinal comparisons between obtained and forgone outcomes during feedback presentation. As individuals integrate this information across repeated experiences, their choices become driven to some degree by the expected probability of experiencing regret (i.e., receiving an outcome that is lower than the forgone outcome from the nonchosen option). This then leads to a pronounced preference for R when EVs are equal, and a reduction in that preference when R becomes suboptimal due to its infrequent, low outcomes. The 95% HDI for μ_λ (1.03 to 2.21) suggests that losses were more impactful than gains, which could account for the lower selections of R in both conditions with mixed outcomes (recall that R always produced larger losses than S). The delta learning rule parameters suggest that learning occurred gradually (μ_α : .04 to .10) and that expectancies for nonchosen options were updated more slowly than for chosen options (μ_γ : .03 to .23).

Figure 5*Posterior Distributions for Group-Level Mean Parameters*

Note. The axis limits were chosen to facilitate visual inspection and do not necessarily correspond to the full allowable range for each parameter. HDI = highest density interval.

Alternative Models

In addition to the nine models above, we tested a model that uses asymmetric learning rates for positive (α^+) and negative (α^-) prediction errors: one version with loss neutrality, another with loss aversion, and a third with loss attention. The reason for testing the asymmetric learning rate hypothesis in a task that involves choices between a risky and safe option is that it predicts risk seeking when $\alpha^+ > \alpha^-$ and risk aversion when $\alpha^+ < \alpha^-$ (Niv et al., 2012).

Consistent with the overall preference for the riskier option (Figure 4), the mean estimate of α^+ was greater than the mean estimate of α^- across participants for each of the three asymmetric learning rate models. We also tested a simple two-parameter “rejoice-stay, regret-shift” heuristic which assumes a constant probability of staying with an option after a positive counterfactual

comparison [$p(\text{stay}|\text{rejoice})$] and a constant probability of switching to the other option after a negative counterfactual comparison [$p(\text{shift}|\text{regret})$]. Note that the rejoice-stay, regret-shift heuristic depends on the previous trial outcomes alone and does not assume any form of long-term learning. Such simple strategies have been shown to account for the choice behavior of a considerable subset of individuals in similar RL tasks (Worthy et al., 2013). Additional details about these alternative models can be found in Supplemental Note 2.

Table 4 shows the percentage of participants whose choices were better characterized by one model over the other. Thirteen total models are listed, including the asymmetric learning rate models (M_{10} , M_{11} , and M_{12}) and the rejoice-stay, regret-shift heuristic (M_{13}). The model with asymmetric learning rates and loss aversion (M_{11}) compared favorably to the absolute encoding models (M_1 , M_2 , and M_3) and was on par with the diminishing sensitivity models (M_4 , M_5 , and M_6), but it was outperformed by the partially relative encoding models (M_7 , M_8 , and M_9). The rejoice-stay, regret-shift heuristic was the best-fitting model for a sizeable number of participants (38%); however, without any learning mechanisms or unequal weighting of gains and losses, it has no way of accounting for the effects of EV differences or mixed outcomes that were observed in this study. Instead, it can only account for sequential dependencies resulting from the outcomes on the previous trial. Table 4 also shows that the partially relative encoding models with loss neutrality (M_7) and loss aversion (M_8) were the only ones that were favored for a majority of participants when compared to all other candidate models.

Table 4

Percentage of Participants whose Choices were Better Fit by the Column Model Compared to the Row Model

	M ₁	M ₂	M ₃	M ₄	M ₅	M ₆	M ₇	M ₈	M ₉	M ₁₀	M ₁₁	M ₁₂	M ₁₃
M ₁	--												
M ₂	42	--											
M ₃	48	65	--										
M ₄	10	27	14	--									
M ₅	19	19	17	50	--								
M ₆	27	33	12	58	62	--							
M ₇	29	21	27	39	37	40	--						
M ₈	<u>23</u>	<u>25</u>	<u>17</u>	<u>37</u>	<u>35</u>	<u>25</u>	<u>50</u>	--					
M ₉	25	27	21	40	44	42	52	<u>65</u>	--				
M ₁₀	50	48	42	64	64	58	83	<u>79</u>	71	--			
M ₁₁	39	40	35	50	54	48	64	<u>77</u>	64	44	--		
M ₁₂	52	60	46	65	69	67	79	<u>87</u>	87	65	77	--	
M ₁₃	37	42	39	46	48	44	54	<u>54</u>	52	42	50	50	--

Note. Each cell shows the percentage of participants ($n = 52$) whose choices were better characterized by the column model compared to the row model, determined by comparing marginal likelihoods. The comparisons involving the best model are underlined.

M₁ = absolute encoding + loss neutrality, M₂ = absolute encoding + loss aversion, M₃ = absolute encoding + loss attention, M₄ = diminishing sensitivity + loss neutrality, M₅ = diminishing sensitivity + loss aversion, M₆ = diminishing sensitivity + loss attention, M₇ = partially relative encoding + loss neutrality, M₈ = partially relative encoding + loss aversion, M₉ = partially relative encoding + loss attention, M₁₀ = asymmetric learning rates + loss neutrality, M₁₁ = asymmetric learning rates + loss aversion, M₁₂ = asymmetric learning rates + loss attention, M₁₃ = rejoice-stay, regret-shift heuristic.

Discussion

Regret theory posits that individuals try to minimize the regret they expect to feel after choosing between options and comparing the outcomes obtained to the outcomes forgone (Bell, 1982; Loomes & Sugden, 1982). In decisions from description, where the potential outcomes and their probabilities are known in advance, and assuming a convex $Q(\cdot)$ function, regret

theory predicts an aversion toward risky options that have a high probability of producing a gain but a small probability of producing a low outcome that would elicit strong feelings of regret (Somasundaram & Diecidue, 2017). However, a different picture has emerged from research on experience-based decision making. When individuals repeatedly sample from the same set of options and have access to forgone outcomes, they tend to prefer those options that yield better outcomes most of the time, consistent with a desire to minimize the *probability* of regret (Barron & Erev, 2003; Erev et al., 2017; Hertwig et al., 2004; Yechiam & Busemeyer, 2006). This can lead to risk seeking for gambles that produce bad outcomes on an infrequent basis, in conflict with standard predictions of regret theory based on gambles with described (rather than experienced) outcomes and probabilities.

In the current study, we found that when making decisions from experience, individuals are indeed risk seeking overall if the riskier and safer options have equal EVs and all outcomes are gains, conceptually replicating prior work (Ahn et al., 2012; Haines et al., 2021). Because the riskier option occasionally produced low outcomes that should have elicited the strongest feelings of regret in our task, this finding conflicts with the assumption of a convex $Q(\cdot)$ function (Loomes & Sugden, 1982). However, our findings are consistent with minimizing the probability of regret given that the riskier option always had a higher chance of producing a better outcome. The current study further demonstrated that the predominate risk seeking preference decreased under two conditions: (1) when the riskier option that frequently yielded better outcomes had a lower EV, and (2) when occasional losses were present. Moreover, the effects of unequal EVs and mixed outcomes were additive rather than multiplicative (i.e., main effects with no interaction).

To directly measure rejoicing and regret, we obtained subjective ratings of choice outcomes on 50% of trials for each participant (in the preliminary experiment, ratings were obtained on all trials). The ratings were highly consistent with regret theory (Camille et al., 2004; Loomes & Sugden, 1982; Mellers et al., 1997; 1999): Participants felt better about the same

gains or losses when the forgone outcome was lower compared to when the forgone outcome was higher, and especially when rating the moderate outcomes produced by the safer option. Further, in the main experiment the effect of forgone outcomes did not vary across conditions, suggesting that regret and rejoicing were experienced in all contexts and did not depend on EV differences or the domain of the outcomes.

If one's only goal were minimizing the probability of regret, then the riskier option should have been highly preferred in all four conditions. However, we observed that risk seeking decreased when the riskier option's EV was lower than that of the safer option, as well as when occasional losses were experienced from both options. RL modeling suggested that two cognitive mechanisms were important for explaining these effects. First, partially relative encoding leads to a preference for options that minimize the probability of regret if long-run payoffs are equal. If those same options have a lower EV, this model predicts a reduced preference, with the size of the reduction dependent on the magnitude of the EV difference and the attention allocated to relative outcomes (ω). With larger EV differences and lower values of ω , the model predicts a greater shift away from minimizing regret probability and toward maximizing EV. The second mechanism, loss aversion, made the regret-minimizing option less attractive in both mixed outcome conditions due to its infrequent large losses. Taken together, the two mechanisms predict that the preference for actions that minimize the probability of regret will be reduced whenever these actions have lower EVs or the potential to produce the largest losses. When EVs are equal and all outcomes are gains, options that minimize the probability of regret will generally be favored (Ahn et al., 2012).

For most participants, loss aversion was found to be a better explanation for the effect of mixed outcomes than the loss attention hypothesis (Yechiam & Hochman, 2013a; 2013b). In one study, Yechiam et al. (2015) found that minor losses had no effect on preference for a riskier option when it had the same EV as a safer option. This conflicts with our finding that preference for R was reduced in both conditions with mixed outcomes, even when EVs were

equal. However, it may be that the losses in the Mixed-Equal-EVs condition (-35 points, on average) were large enough to induce loss aversion in our participants, whereas the very small losses in Yechiam et al.'s study (-1 Israeli new shekel) were minor enough to focus attention without leading to aversion. In our Mixed-Equal-EVs condition, the losses were the most extreme outcomes and were thus not readily comparable to other, more extreme outcomes that may have lessened loss aversion (Mukherjee et al., 2017).

One possibility, suggested by one of our reviewers, is that individuals recall small samples of similar past experiences when deciding between choice alternatives and choose whichever one tended to produce the better outcomes on average, as in Hochman and Erev's (2013) contingent-sampling (CS) model. The CS model is an exemplar-based model with a single integer-valued parameter (m) representing the number of previous trials that the decision maker tracks. We simulated the model in the current paradigm, allowing m to range from 1 to 8, and tested whether it could reproduce the average choice proportions that were observed in each condition. Without any additional assumptions, the CS model considerably underestimated choices of R in the gains-only conditions and could not produce the observed effect of mixed outcomes. Only when partially relative encoding (ω) and loss aversion (λ) were added into the model could it produce a pattern that closely matched the observed data (see Supplemental Figure S7). Thus, the CS model does not appear to be capable of accounting for the choice patterns in the current study without auxiliary assumptions that make it more similar to the best-fitting RL model.⁷

In summary, although there is a general tendency to choose options that minimize the probability of regret in experience-based decisions, structural features of the choice scenario, such as EV differences and the sign of outcomes, can diminish this tendency. More generally,

⁷ Hochman and Erev (2013) noted that additional assumptions could be added to the CS model (choice inertia, generalization, noisy action selection, and bounded memory), but were not necessary for reproducing the pattern of results in their study. It is possible that these additional assumptions could improve the fit of the CS model in our paradigm.

our results add to a growing body of research suggesting that relative comparisons between obtained and forgone outcomes play a key role in guiding decision-making (Ahn et al., 2012; Coricelli et al., 2005; 2007; Haines et al., 2021; Mellers et al., 1997; 1999). However, our findings diverge from standard regret theory applied to decisions from description in which $Q(\cdot)$ is assumed to be convex (Loomes & Sugden, 1982). A concave $Q(\cdot)$ function that is more sensitive to the probability of regret than the magnitude appears to be more consistent with our data and the findings from previous work on experience-based decisions (Erev et al., 2017). In our case, we demonstrated that a step-like function for comparing obtained and forgone outcomes (i.e., the partially relative encoding model) was sufficient to account for the overall choice patterns.

References

- Ahn, W.-Y., Rass, O., Shin, Y.-W., Busemeyer, J. R., Brown, J. W., & O'Donnell, B. F. (2012). Emotion-based reinforcement learning. In *Proceedings of the 34th Annual Meeting of the Cognitive Science Society* (Vol. 34, pp. 78–83).
- Barron, G., & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making*, 16, 215-233.
- Bavard, S., Lebreton, M., Khamassi, M., Coricelli, G., & Palminteri, S. (2018). Reference-point centering and range-adaptation enhance human reinforcement learning at the cost of irrational preferences. *Nature Communications*, 9. <https://doi.org/10.1038/s41467-018-06781-2>
- Bell, D. E. (1982). Regret in decision making under uncertainty. *Operations Research*, 30, 961–981.
- Bell, D. E. (1985). Disappointment in decision making under uncertainty. *Operations Research*, 33, 1-27.
- Camille, N., Coricelli, G., Sallet, J., Pradat-Diehl, P., Duhamel, J.-R., & Sirigu, A. (2004). The involvement of the orbitofrontal cortex in the experience of regret. *Science*, 304, 1167-1170.
- Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M., Guo, J., Li, P., & Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of Statistical Software*, 76.
- Coricelli, G., Critchley, H. D., Joffily, M., O'Doherty, J. P., Sirigu, A., & Dolan, R. J. (2005). Regret and its avoidance: A neuroimaging study of choice behavior. *Nature Neuroscience*, 8, 1255–1262.
- Coricelli, G., Dolan, R. J., & Sirigu, A. (2007). Brain, emotion and decision making: the paradigmatic example of regret. *Trends in Cognitive Sciences*, 11, 258–265.
- Erev, I., Ert, E., Plonsky, O., Cohen, D., & Cohen, O. (2017). From anomalies to forecasts:

- Toward a descriptive model of decisions under risk, under ambiguity, and from experience. *Psychological Review*, 124, 369–409.
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2013). *Bayesian data analysis* (3rd ed.). Boca Raton: CRC Press.
- Grosskopf, B., Erev, I., & Yechiam, E. (2006). Foregone with the wind: Indirect payoff information and its implications for choice. *International Journal of Game Theory*, 34, 285–302.
- Haines, N., Rass, O., Shin, Y.-W., Brown, J. W., & Ahn, W.-Y. (2021). *Negative affect induces rapid learning of counterfactual representations: A model-based facial expression analysis approach*. bioRxiv. <https://doi.org/10.1101/560011>
- Haines, N., Vassileva, J., & Ahn, W.-Y. (2018). The Outcome-Representation Learning model: A novel reinforcement learning model of the Iowa Gambling Task. *Cognitive Science*, 42, 2534–2561.
- Hayes, W. M., & Wedell, D. H. (2020a). Autonomic responses to choice outcomes: Links to task performance and reinforcement-learning parameters. *Biological Psychology*, 156. <https://doi.org/10.1016/j.biopsycho.2020.107968>
- Hayes, W. M., & Wedell, D. H. (2020b). Modeling the role of feelings in the Iowa Gambling Task. *Decision*, 7, 67–89.
- Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, 15, 534–539.
- Hochman, G., & Erev, I. (2013). The partial-reinforcement extinction effect and the contingent-sampling hypothesis. *Psychonomic Bulletin & Review*, 20, 1336–1342.
- Hochman, G., & Yechiam, E. (2011). Loss aversion in the eye and in the heart: The autonomic nervous system's responses to losses. *Journal of Behavioral Decision Making*, 24, 140–156.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk.

Econometrica, 47, 263-292.

Lee, M. D., & Wagenmakers, E.-J. (2013). *Bayesian cognitive modeling: A practical course*. Cambridge University Press.

Loomes, G., & Sugden, R. (1982). Regret theory: An alternative theory of rational choice under uncertainty. *The Economic Journal*, 92, 805-824.

Loomes, G., & Sugden, R. (1986). Disappointment and dynamic consistency in choice under uncertainty. *Review of Economic Studies*, 53, 271-282.

Löw, A., Lang, P. J., Smith, J. C., & Bradley, M. M. (2008). Both predator and prey: Emotional arousal in threat and reward. *Psychological Science*, 19, 865–873.

Mellers, B. A. (2000). Choice and the relative pleasure of consequences. *Psychological Bulletin*, 126, 910–924.

Mellers, B. A., Schwartz, A., Ho, K., & Ritov, I. (1997). Decision affect theory: Emotional reactions to the outcomes of risky options. *Psychological Science*, 8, 423–429.

Mellers, B., Schwartz, A., & Ritov, I. (1999). Emotion-based choice. *Journal of Experimental Psychology: General*, 128, 332–345.

Mukherjee, S., Sahay, A., Pammi, V. S. C., & Srinivasan, N. (2017). Is loss-aversion magnitude-dependent? Measuring prospective affective judgments regarding gains and losses. *Judgment and Decision Making*, 12, 81–89.

Niv, Y., Edlund, J. A., Dayan, P., & O'Doherty, J. P. (2012). Neural prediction errors reveal a risk-sensitive reinforcement-learning process in the human brain. *Journal of Neuroscience*, 32, 551-562.

Rangel, A., Camerer, C., & Montague, P. R. (2008). A framework for studying the neurobiology of value-based decision making. *Nature Reviews Neuroscience*, 9, 545–556.

Rescorla, R., & Wagner, A. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black, & W. F. Prokasy (Eds.), *Classical conditioning II: Current research and theory* (pp. 64–99). New York:

Appleton-Century-Crofts.

Scheibehenne, B., & Pachur, T. (2015). Using Bayesian hierarchical parameter estimation to assess the generalizability of cognitive models of choice. *Psychonomic Bulletin & Review*, 22, 391–407.

Sokol-Hessner, P., Hsu, M., Curley, N. G., Delgado, M. R., Camerer, C. F., & Phelps, E. A. (2009). Thinking like a trader selectively reduces individuals' loss aversion. *Proceedings of the National Academy of Sciences of the United States of America*, 106, 5035–5040.

Somasundaram, J., & Diecidue, E. (2017). Regret theory and risk attitudes. *Journal of Risk and Uncertainty*, 55, 147–175.

Steingroever, H., Wetzels, R., & Wagenmakers, E. J. (2016). Bayes factors for reinforcement-learning models of the Iowa Gambling Task. *Decision*, 3, 115–131.

Studer, B., & Clark, L. (2011). Place your bets: Psychophysiological correlates of decision-making under risk. *Cognitive, Affective and Behavioral Neuroscience*, 11, 144–158.

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.

Vehtari, A., Gelman, A., & Gabry, J. (2017). Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Statistics and Computing*, 27, 1413–1432.

Wagenmakers, E.-J., & Farrell, S. (2004). AIC model selection using Akaike weights. *Psychonomic Bulletin & Review*, 11, 192–196.

Wasserman, L. (2000). Bayesian model selection and model averaging. *Journal of Mathematical Psychology*, 44, 92–107.

Wetzels, R., Vandekerckhove, J., Tuerlinckx, F., & Wagenmakers, E. J. (2010). Bayesian parameter estimation in the Expectancy Valence model of the Iowa gambling task. *Journal of Mathematical Psychology*, 54, 14–27.

Worthy, D. A., Hawthorne, M. J., & Otto, A. R. (2013). Heterogeneity of strategy use in the Iowa gambling task: A comparison of win-stay/lose-shift and reinforcement learning models.

Psychonomic Bulletin and Review, 20, 364-371.

Wright, R. J., & Rakow, T. (2017). Don't sweat it: Re-examining the somatic marker hypothesis using variants of the Balloon Analogue Risk Task. *Decision*, 4, 52–65.

Yechiam, E., & Busemeyer, J. R. (2005). Comparison of basic assumptions embedded in learning models for experience-based decision making. *Psychonomic Bulletin & Review*, 12, 387–402.

Yechiam, E., & Busemeyer, J. R. (2006). The effect of foregone payoffs on underweighting small probability events. *Journal of Behavioral Decision Making*, 19, 1–16.

Yechiam, E., & Hochman, G. (2013a). Loss-aversion or loss-attention: The impact of losses on cognitive performance. *Cognitive Psychology*, 66, 212–231.

Yechiam, E., & Hochman, G. (2013b). Losses as modulators of attention: Review and analysis of the unique effects of losses over gains. *Psychological Bulletin*, 139, 497–518.

Yechiam, E., Retzer, M., Telpaz, A., & Hochman, G. (2015). Losses as ecological guides: Minor losses lead to maximization and not to avoidance. *Cognition*, 139, 10–17.

Yechiam, E., Stout, J. C., Busemeyer, J. R., Rock, S. L., & Finn, P. R. (2005). Individual differences in the response to forgone payoffs: An examination of high functioning drug abusers. *Journal of Behavioral Decision Making*, 18, 97-110.

Zeelenberg, M., & Beattie, J. (1997). Consequences of regret aversion 2: Additional evidence for effects of feedback on decision making. *Organizational Behavior and Human Decision Processes*, 72, 63–78.

Zeelenberg, M., Beattie, J., Van Der Pligt, J., & De Vries, N. K. (1996). Consequences of regret aversion: Effects of expected feedback on risky decision making. *Organizational Behavior and Human Decision Processes*, 65, 148–158.

Supplementary Material

“Regret in Experience-Based Decisions: The Effects of Expected Value Differences and Mixed Gains and Losses”

William M. Hayes and Douglas H. Wedell

Supplementary Note 1: Preliminary Experiment**1.1 Participants**

Participants were 117 undergraduates who signed up for the study via an online participant pool in exchange for partial course credit. The experiment was administered online in Qualtrics. Data from 12 participants were excluded because they took longer than two hours to finish ($n = 11$) or because of excessive missing data ($n = 1$), leaving 105 participants in the final sample (91 women, 14 men; ages 18-56, $M = 20.40$, $SD = 4.19$; mean completion time = 51 min., $SD = 16.28$). We stopped collecting data at the end of the semester. All participants provided informed consent and all aspects of this study were approved by the Institutional Review Board at the University of South Carolina.

1.2 Design

Participants completed the three conditions described in Table S1, each consisting of 100 choice trials, in one of two randomly-assigned orders. The Gains-Equal-EVs condition was completed first, with the order of the remaining two conditions balanced across participants. There were no statistical differences in choice patterns between the two condition orders, and thus we excluded this variable from subsequent analyses. Note that there was no condition with equal EVs and mixed outcomes in the preliminary experiment. Further, whereas the options always had positive EVs in the main experiment, R had a negative EV (-5) in the Mixed-Unequal-EVs condition of the preliminary experiment.

In all three conditions, R and S yielded high outcomes (points) on 80% of draws and low outcomes on 20% of draws, but the high and low outcomes from R were always more extreme than those from S. On any given draw, the probability that R would produce a better outcome than S was 80%. Outcomes were randomly generated at the start of each trial and rounded to the nearest integer.

1.3 Procedure

The procedure was identical to that described in the text for the main experiment with the following exceptions. First, during feedback presentation, the words, “You [win/lose] X points!” were presented at the top of the display, and the words, “You chose” were displayed underneath the chosen option’s symbol. Second, high outcomes were printed in boldface and

low outcomes in italics. Third, participants rated how they felt about the outcome on all trials instead of on 50% of trials. Finally, at the conclusion of the block of trials for each condition, participants were given their total number of points earned as well as the “average” points earned by other players, which was computed simply as the average EV of the two options multiplied by the number of trials (included to increase task motivation). All other aspects of the procedure were identical to those of the main experiment.

1.4 Valence rating results

Prior to analysis, we computed each participant’s average valence rating for each combination of Condition (Gains-Equal-EVs, Gains-Unequal-EVs, or Mixed-Unequal-EVs), Choice (safer = 1, riskier = -1), Obtained Outcome (high = 1, low = -1), and Forgone Outcome (high = 1, low = -1). The mean valence ratings were analyzed using a random-intercepts regression model.

Figure S1 shows that regret and rejoicing effects were especially prominent when S was chosen (top row): The same outcomes from S were rated more positively when the forgone outcome from R was low compared to when it was high. Table S2 contains the regression slopes as well as the estimated intercept and residual variance from the random-intercepts regression. Main effects and interactions were tested using Type III Wald chi-square tests. There were significant main effects of Condition [$\chi^2(2) = 161.66, p < .001$], Choice [$\chi^2(1) = 161.16, p < .001$], Obtained Outcome [$\chi^2(1) = 2021.18, p < .001$], and Forgone Outcome [$\chi^2(1) = 412.93, p < .001$]. Valence ratings were higher in the conditions with gains only, when S was chosen, when obtained outcomes were high, and when forgone outcomes were low. These effects were qualified by significant two- and three-way interactions. In particular, the Choice $\times$ Obtained Outcome $\times$ Forgone Outcome interaction [$\chi^2(1) = 8.65, p = .003$] indicated that while the negative effect of forgone outcomes was slightly reduced when the obtained outcome was high, this effect was larger when S was chosen. The Condition $\times$ Choice $\times$ Forgone Outcome interaction [$\chi^2(2) = 6.49, p = .039$] indicated that while the negative effect of forgone outcomes was larger when S was chosen, the difference was smaller in the Mixed-Unequal-EVs condition than in the other two conditions. We speculate that the presence of losses might have decreased the intensity of counterfactual emotions overall because of greater allocation of attention to expected value (Yechiam & Hochman, 2013a; 2013b; Yechiam et al., 2015). Finally, the Condition $\times$ Choice $\times$ Obtained Outcome interaction [$\chi^2(2) = 18.30, p < .001$] indicated that the greater positive effect of obtained outcomes for R choices compared to S choices was most pronounced in the Gains-Unequal-EVs condition.

1.5 Arousal rating results

A parallel regression analysis was conducted for outcome arousal ratings (Table S3). Figure S2 shows the mean arousal ratings as a function of Condition, Choice, Obtained Outcome, and Forgone Outcome. There was a significant Condition $\times$ Choice $\times$ Obtained Outcome interaction [$\chi^2(2) = 18.23, p < .001$]. Although arousal ratings were generally higher following high obtained outcomes, low obtained outcomes from R received higher arousal ratings than high obtained outcomes in the Mixed-Unequal-EVs condition. Recall that the low outcomes from the riskier option in this condition were relatively large losses (-85 points on average). We speculate that the higher self-reported arousal might reflect increased attention due to the presence of losses (Yechiam & Hochman, 2013a; 2013b; Yechiam et al., 2015). There was also a significant Obtained $\times$ Forgone interaction [$\chi^2(1) = 16.70, p < .001$]: Arousal ratings were generally higher when forgone outcomes were low, but especially when the obtained outcome was high. This likely reflects the excitement of receiving a positive outcome and avoiding a much worse outcome. Finally, there was a significant Choice $\times$ Forgone interaction [$\chi^2(1) = 42.75, p < .001$], due to the much larger negative effect of forgone outcomes when the safer option was chosen. Larger forgone outcome effects for safer choices were found for both valence and arousal and are likely due to the safer option's outcomes being compared to the more extreme forgone outcomes from the riskier option.

1.6 Choice results

The average choice proportion for R across blocks of 20 trials in each condition is shown in Figure S3. We analyzed choice proportions for R using a 3 (Condition: Gains-Equal-EVs, Gains-Unequal-EVs, or Mixed-Unequal-EVs) $\times$ 5 (Block: trials 1-20, 21-40, 41-60, 61-80, or 81-100) repeated measures ANOVA.⁸ Greenhouse-Geisser ε values and corrected p values are presented in cases where sphericity was violated.

There was a significant main effect of Condition, $F(2, 208) = 11.24, MSE = 0.36, p < .001, \eta_p^2 = .10$, and a significant Condition $\times$ Block interaction, $F(8, 832) = 7.86, MSE = 0.02, \varepsilon = .24, p < .001, \eta_p^2 = .07$. The main effect of Block was nonsignificant, $F(4, 416) = 0.70, MSE = 0.02, \varepsilon = .45, p = .55$. To probe the interaction, we analyzed the simple effect of Condition within each Block separately. There were no differences between conditions in the first block ($p > .99$), but there were significant differences between conditions in Block 2 [$F(2, 208) = 8.69, MSE = 0.09, p = .001, \eta_p^2 = .08$], Block 3 [$F(2, 208) = 9.34, MSE = 0.10, p < .001, \eta_p^2 = .08$], Block 4 [$F(2, 208) = 13.24, MSE = 0.10, p < .001, \eta_p^2 = .11$], and Block 5 [$F(2, 208) = 16.86, MSE = 0.10$,

⁸ The results were unchanged when arcsine-transformed choice proportions were analyzed instead of raw proportions.

$p < .001$, $\eta_p^2 = .14$] (Bonferroni-corrected p values). Follow-up paired t -tests with Tukey correction indicated that after the first block, R was selected more frequently in the Gains-Equal-EVs condition compared to the other two conditions (all $ps < .022$, $ds > 0.26$). There were no significant differences between the two unequal EV conditions until Block 4 ($p = .035$, $d = 0.25$) and Block 5 ($p = .015$, $d = 0.28$), with R selected more frequently in the Gains-Unequal-EVs condition compared to the Mixed-Unequal-EVs condition.

1.7 Power analysis

To determine the required sample size for the main experiment, we took the largest standardized mean difference (d) between conditions in Block 4 as an estimate of effect size (Gains-Equal-EVs vs. Mixed-Unequal-EVs: $d = 0.49$). We chose Block 4 because each condition in the main experiment contained 80 trials instead of 100, and we expected the largest differences to emerge in the final block of 20 trials. Based on this effect size, 46 subjects would be required to achieve a power of .90 ($\alpha = .05$, paired t -test, two-tailed).

1.8 Conclusions

Preference for the riskier option that frequently yielded better outcomes was highest when its EV was equal to that of the safer option and when all outcomes were gains, replicating previous studies (Ahn et al., 2012; Haines et al., 2020). Participants chose this option significantly more often than in the other two conditions with unequal EVs and the differences emerged early in the task, after the first block of 20 trials. Additionally, in the last 40 trials, participants exhibited some sensitivity to loss outcomes: the riskier option was chosen significantly less often in the Mixed-Unequal-EVs condition than in the Gains-Unequal-EVs condition, even though the EV differences were the same. Outcome valence ratings suggested that participants experienced regret and rejoicing in all three conditions, especially after choosing the safer option and evaluating its outcome relative to the forgone outcome from the riskier option.

Supplemental Note 2: Alternative Models

2.1 Asymmetric learning rates

For the models with asymmetric learning rates, the updating of each option's expectancy depends on whether the prediction error for that option (i.e., the difference between the received and expected outcome on trial t) is positive or negative:

$$E_i(t+1) = \begin{cases} E_i(t) + \alpha^+ \cdot [\delta_i(t) + (1 - \delta_i(t)) \cdot \gamma] \cdot (V_i(t) - E_i(t)), & \text{if } V_i(t) - E_i(t) \geq 0 \\ E_i(t) + \alpha^- \cdot [\delta_i(t) + (1 - \delta_i(t)) \cdot \gamma] \cdot (V_i(t) - E_i(t)), & \text{if } V_i(t) - E_i(t) < 0 \end{cases}$$

In the equation above, α^+ ($0 \leq \alpha^+ \leq 1$) is the learning rate for positive prediction errors and α^- ($0 \leq \alpha^- \leq 1$) is the learning rate for negative prediction errors. The indicator variable $\delta_i(t)$ equals 1 if option i was selected on trial t and 0 otherwise, and the counterfactual learning parameter γ ($0 \leq \gamma \leq 1$) modulates the rate of learning for nonchosen options. When $\alpha^+ > \alpha^-$, updating is faster for positive prediction errors and the model predicts risk-seeking behavior. When $\alpha^+ < \alpha^-$, updating is faster for negative prediction errors and the model predicts risk-averse behavior. The remaining components are the same as those in the absolute encoding model. Thus, the free parameters are α^+ , α^- , γ , and θ for the version with loss neutrality; λ , α^+ , α^- , γ , and θ for the version with loss aversion; and α^+ , α^- , γ , θ_{Gains} , and θ_{Mixed} for the version with loss attention.

Note that for these models, the expectancies E_i were initialized to the average of the outcomes in each condition, instead of zero, on the first trial of each condition. Initializing expectancies near the middle of the outcome distribution can help avoid over- or underestimating the learning rate asymmetries (Stefano Palminteri, personal communication, May 19, 2021).

2.2 Rejoice-stay, regret-shift heuristic

Unlike the RL models that we considered, the rejoice-stay, regret-shift heuristic does not track separate expectancies for each option by integrating information across several trials. Instead, the choice on trial t depends only on the relative outcomes on trial $t-1$. Suppose the riskier option R was chosen on trial $t-1$ and the outcome from R was greater than the outcome from the safer option S (i.e., rejoicing). Then the probability of choosing R again on trial t would be $p(stay|rejoice)$ and the probability of switching to S on would be $1 - p(stay|rejoice)$. If R was chosen on trial $t-1$ and the outcome from R was less than the outcome from S (i.e., regret), then the probability of switching to S on trial t would be $p(shift|regret)$ and the probability of staying with option R would be $1 - p(shift|regret)$. The same steps apply when S

was chosen on trial $t - 1$. Thus, the two parameters of the rejoice-stay, regret-shift heuristic are the conditional probabilities $p(\text{stay}|\text{rejoice})$ and $p(\text{shift}|\text{regret})$.

Supplemental References

- Ahn, W.-Y., Rass, O., Shin, Y.-W., Busemeyer, J. R., Brown, J. W., & O'Donnell, B. F. (2012). Emotion-based reinforcement learning. In *Proceedings of the 34th Annual Meeting of the Cognitive Science Society* (Vol. 34, pp. 78–83).
- Haines, N., Rass, O., Shin, Y.-W., Brown, J. W., & Ahn, W.-Y. (2020). Negative affect induces rapid learning of counterfactual representations: A model-based facial expression analysis approach. *bioRxiv*, <https://doi.org/10.1101/560011>
- Hochman, G., & Erev, I. (2013). The partial-reinforcement extinction effect and the contingent-sampling hypothesis. *Psychonomic Bulletin and Review*, 20, 1336–1342.
- Yechiam, E., & Hochman, G. (2013a). Loss-aversion or loss-attention: The impact of losses on cognitive performance. *Cognitive Psychology*, 66, 212–231.
- Yechiam, E., & Hochman, G. (2013b). Losses as modulators of attention: Review and analysis of the unique effects of losses over gains. *Psychological Bulletin*, 139, 497–518.
- Yechiam, E., Retzer, M., Telpaz, A., & Hochman, G. (2015). Losses as ecological guides: Minor losses lead to maximization and not to avoidance. *Cognition*, 139, 10–17.

Table S1

Outcomes, Probabilities, and Expected Values for the Choice Options in the Preliminary Experiment

	p	Gains Equal EVs	Gains Unequal EVs	Mixed Unequal EVs
Safer Option (S)	.80	$90 + \varepsilon$	$100 + \varepsilon$	$10 + \varepsilon$
	.20	$65 + \varepsilon$	$75 + \varepsilon$	$-15 + \varepsilon$
EV		85	95	5
Riskier Option (R)	.80	$95 + \varepsilon$	$105 + \varepsilon$	$15 + \varepsilon$
	.20	$45 + \varepsilon$	$5 + \varepsilon$	$-85 + \varepsilon$
EV		85	85	-5

Note. The riskier option minimized the probability of regret in all three conditions. A small amount of Gaussian noise ε was added to each outcome, $\varepsilon \sim N(0,1)$. All outcomes were presented as points rounded to the nearest integer. p = probability, EV = expected value, computed as the sum of outcomes weighted by their associated probabilities.

Table S2*Random-Intercepts Regression for Valence Ratings in Preliminary Experiment*

	<i>b</i>	<i>SE</i>	<i>z</i>
Fixed effects			
Intercept	13.87	3.94	3.52***
Cond1	14.65	2.23	6.58***
Cond2	-27.92	2.20	-12.71***
Choice	-19.95	1.57	-12.70***
Obtained	70.28	1.56	44.96***
Forgone	-31.77	1.56	-20.32***
Cond1 × Choice	-9.01	2.23	-4.04***
Cond2 × Choice	2.48	2.20	1.13
Cond1 × Obtained	-8.07	2.22	-3.63***
Cond2 × Obtained	16.45	2.19	7.50***
Choice × Obtained	36.05	1.56	23.05***
Cond1 × Forgone	0.80	2.22	0.36
Cond2 × Forgone	8.58	2.19	3.92***
Choice × Forgone	24.66	1.57	15.76***
Obtained × Forgone	4.69	1.56	3.00**
Cond1 × Choice × Obtained	9.49	2.23	4.27***
Cond2 × Choice × Obtained	-4.07	2.19	-1.86
Cond1 × Choice × Forgone	2.83	2.23	1.27
Cond2 × Choice × Forgone	-5.59	2.19	-2.55*
Cond1 × Obtained × Forgone	-2.71	2.22	-1.22
Cond2 × Obtained × Forgone	0.31	2.19	0.14
Choice × Obtained × Forgone	-4.60	1.56	-2.94**
Cond1 × Choice × Obtained × Forgone	-1.39	2.22	-0.63
Cond2 × Choice × Obtained × Forgone	-2.11	2.19	-0.96
Random effects			
Intercept variance	1375.00		
Residual variance	5391.00		

Note. Valence ratings ranged between -240 and 240. Cond1 was coded as Gains-Unequal-EVs = 1, Mixed-Unequal-EVs = 0, and Gains-Equal-EVs = -1. Cond2 was coded as Mixed-Unequal-EVs = 1, Gains-Unequal-EVs = 0, and Gains-Equal-EVs = -1. * $p < .05$ ** $p < .01$ *** $p < .001$.

Table S3*Random-Intercepts Regression for Arousal Ratings in Preliminary Experiment*

	<i>b</i>	<i>SE</i>	<i>z</i>
Fixed effects			
Intercept	239.02	7.77	30.75***
Cond1	-8.91	2.35	-3.80***
Cond2	11.62	2.31	5.02***
Choice	0.88	1.66	0.53
Obtained	3.70	1.65	2.25*
Forgone	-11.17	1.65	-6.79***
Cond1 × Choice	-6.24	2.35	-2.65**
Cond2 × Choice	15.52	2.32	6.69***
Cond1 × Obtained	0.80	2.34	0.34
Cond2 × Obtained	-11.26	2.31	-4.88***
Choice × Obtained	-0.84	1.65	-0.51
Cond1 × Forgone	-3.34	2.34	-1.43
Cond2 × Forgone	7.44	2.31	3.22**
Choice × Forgone	10.77	1.65	6.54***
Obtained × Forgone	-6.72	1.64	-4.09***
Cond1 × Choice × Obtained	2.11	2.34	0.90
Cond2 × Choice × Obtained	-9.39	2.31	-4.07***
Cond1 × Choice × Forgone	3.77	2.34	1.61
Cond2 × Choice × Forgone	-5.26	2.31	-2.28*
Cond1 × Obtained × Forgone	1.60	2.34	0.68
Cond2 × Obtained × Forgone	-2.18	2.31	-0.94
Choice × Obtained × Forgone	-0.10	1.64	-0.06
Cond1 × Choice × Obtained × Forgone	-2.28	2.34	-0.98
Cond2 × Choice × Obtained × Forgone	1.32	2.31	0.57
Random effects			
Intercept variance	6056.00		
Residual variance	5971.00		

Note. Arousal ratings ranged between 0 and 480. Cond1 was coded as Gains-Unequal-EVs = 1, Mixed-Unequal-EVs = 0, and Gains-Equal-EVs = -1. Cond2 was coded as Mixed-Unequal-EVs = 1, Gains-Unequal-EVs = 0, and Gains-Equal-EVs = -1. * $p < .05$ ** $p < .01$ *** $p < .001$.

Table S4*Random-Intercepts Regression for Valence Ratings in Main Experiment*

	<i>b</i>	<i>SE</i>	<i>z</i>
Fixed effects			
Intercept	34.50	8.09	4.26***
EV	0.09	2.28	0.04
Domain	21.32	2.29	9.32***
Choice	16.89	2.32	7.29***
Obtained	50.81	2.28	22.30***
Forgone	-15.39	2.28	-6.75***
EV × Domain	1.94	2.28	0.85
EV × Choice	-8.47	2.30	-3.68***
Domain × Choice	-0.25	2.30	-0.11
EV × Obtained	0.52	2.28	0.23
Domain × Obtained	-17.50	2.28	-7.67***
Choice × Obtained	-23.78	2.28	-10.42***
EV × Forgone	-0.22	2.28	-0.10
Domain × Forgone	-2.35	2.29	-1.03
Choice × Forgone	-14.33	2.29	-6.25***
Obtained × Forgone	-2.82	2.28	-1.24
EV × Domain × Choice	0.26	2.30	0.11
EV × Domain × Obtained	-1.10	2.28	-0.48
EV × Choice × Obtained	8.35	2.29	3.65***
Domain × Choice × Obtained	0.81	2.29	0.35
EV × Domain × Forgone	1.06	2.28	0.47
EV × Choice × Forgone	1.57	2.28	0.69
Domain × Choice × Forgone	-1.08	2.29	-0.47
EV × Obtained × Forgone	0.41	2.28	0.18
Domain × Obtained × Forgone	3.80	2.28	1.67
Choice × Obtained × Forgone	2.50	2.28	1.10
EV × Domain × Choice × Obtained	-0.05	2.28	-0.02
EV × Domain × Choice × Forgone	-2.34	2.29	-1.03
EV × Domain × Obtained × Forgone	0.33	2.28	0.15
EV × Choice × Obtained × Forgone	-1.54	2.28	-0.67
Domain × Choice × Obtained × Forgone	0.72	2.28	0.32
EV × Domain × Choice × Obtained × Forgone	3.92	2.28	1.72

Random effects

Intercept variance	3129.00
Residual variance	6197.00

Note. Valence ratings were obtained on 50% of trials and ranged between -240 and 240. EV was coded equal = 1, unequal = -1. Domain was coded gains = 1, mixed = -1. Choice was coded safer = 1, riskier = -1. Obtained and Forgone were coded high = 1, low = -1. * $p < .05$ ** $p < .01$ *** $p < .001$.

Table S5*Random-Intercepts Regression for Arousal Ratings in Main Experiment*

	<i>b</i>	<i>SE</i>	<i>z</i>
Fixed effects			
Intercept	211.09	11.16	18.92***
EV	11.13	2.45	4.54***
Domain	-7.54	2.46	-3.07**
Choice	2.09	2.49	0.84
Obtained	-6.21	2.45	-2.54*
Forgone	11.42	2.45	4.66***
EV × Domain	-2.49	2.45	-1.02
EV × Choice	-1.60	2.47	-0.65
Domain × Choice	-8.49	2.47	-3.43***
EV × Obtained	-3.46	2.45	-1.41
Domain × Obtained	-0.86	2.45	-0.35
Choice × Obtained	-1.62	2.45	-0.66
EV × Forgone	-3.82	2.45	-1.56
Domain × Forgone	2.35	2.45	0.96
Choice × Forgone	4.87	2.46	1.98*
Obtained × Forgone	3.13	2.45	1.28
EV × Domain × Choice	2.03	2.47	0.83
EV × Domain × Obtained	2.63	2.45	1.07
EV × Choice × Obtained	0.97	2.45	0.39
Domain × Choice × Obtained	3.03	2.45	1.24
EV × Domain × Forgone	2.28	2.45	0.93
EV × Choice × Forgone	-1.26	2.45	-0.51
Domain × Choice × Forgone	2.78	2.46	1.13
EV × Obtained × Forgone	3.85	2.45	1.57
Domain × Obtained × Forgone	-0.65	2.45	-0.27
Choice × Obtained × Forgone	3.12	2.45	1.27
EV × Domain × Choice × Obtained	-0.73	2.45	-0.30
EV × Domain × Choice × Forgone	-1.04	2.45	-0.43
EV × Domain × Obtained × Forgone	-3.46	2.45	-1.41
EV × Choice × Obtained × Forgone	-0.04	2.45	-0.02
Domain × Choice × Obtained × Forgone	-1.32	2.45	-0.54
EV × Domain × Choice × Obtained × Forgone	-1.13	2.45	-0.46

Random effects

Intercept variance	6154
Residual variance	7142

Note. Arousal ratings ranged between 0 and 480. EV was coded equal = 1, unequal = -1.

Domain was coded gains = 1, mixed = -1. Choice was coded safer = 1, riskier = -1. Obtained and Forgone were coded high = 1, low = -1. * $p < .05$ ** $p < .01$ *** $p < .001$.

Table S6*WAIC Values for each RL Model*

	Absolute Encoding	Diminishing Sensitivity	Partially Relative Encoding
Loss Neutrality	365.26	347.52	348.07
Loss Aversion	358.92	343.31	332.28
Loss Attention	353.48	337.88	336.90

Note. WAIC = widely applicable information criterion. Lower numbers indicate lower deviance, or higher estimated predictive accuracy.

Figure S1*Regret Theory Predictions*

Utility function: $u(x) = x^a$

Q-function: $Q(\xi) = \xi + \text{sign}(\xi) \cdot |\xi|^b - \text{sign}(-\xi) \cdot |-\xi|^b$

Expectation of Q
over n states: $\sum_{k=1}^n p_k \cdot Q(u(x_{ik}) - u(x_{jk}))$

Note. To generate predictions from regret theory (Loomes & Sugden, 1982) for the current study, we graphed the expectation of the Q-function across a range of values for two key parameters. The curvature of the utility function (a) was allowed to range between .05 and 1 in increments of .05 and the curvature of the regret function (b) was allowed to range between 0 and 2 in increments of .05. Predictions were generated using the equations above for the Gains-Equal-EVs (A and C) and Gains-Unequal-EVs (B and D) conditions, assuming individuals have full knowledge of the possible outcomes and probabilities for each option (decisions from

description). Outcomes were scaled down by a factor of 10 for visualization purposes. When the Q-function is convex ($b > 1$), regret theory predicts a preference for the safer option. On the other hand, when the Q-function is sufficiently concave ($b \ll 1$), the theory predicts a preference for the riskier option. The concavity would need to be even more extreme to predict risk seeking when the EV of the riskier option is lower than the EV of the safer option.

Figure S2*Outcome Valence Ratings in the Preliminary Experiment*

Note. Mean ratings are shown as a function of condition, chosen option, obtained outcome, and forgone outcome. Ratings were centered around their midpoint so that they ranged from -240 to 240. Error bars represent ± 1 SE.

Figure S3*Outcome Arousal Ratings in the Preliminary Experiment*

Note. Mean ratings are shown as a function of condition, chosen option, obtained outcome, and forgone outcome. Ratings ranged from 0 to 480. Error bars represent ± 1 SE.

Figure S4*Choice of the Riskier Option in the Preliminary Experiment*

Note. Mean choice proportions are shown across five blocks of 20 trials in each condition. Error bars represent ± 1 SE.

Figure S5*Outcome Arousal Ratings in the Main Experiment*

Note. Mean ratings are shown as a function of condition, chosen option, obtained outcome, and forgone outcome. Note that arousal was generally higher when the forgone outcome was high, which conflicts with the findings for arousal in the preliminary experiment. This suggests that arousal ratings were idiosyncratic to each experiment. Ratings ranged from 0 to 480. Error bars represent ± 1 SE.

Figure S6*Hierarchical RL Model Fits to Individual Choice Data in the Main Experiment*

Note. Each point represents the data for a single individual. The x-axis shows the observed R choice proportion and the y-axis shows the model-predicted R choice proportion in each condition. Point estimates for the model predictions were obtained using posterior means and uncertainty was quantified using 95% Bayesian credible intervals.

Figure S7

Simulations of the Contingent-Sampling (CS) Model with the Additional Assumptions of Partially Relative Encoding and Loss Aversion

Note. According to the CS model (Hochman & Erev, 2013), decision makers recall the relative advantages of one of the two options over the last m trials (results with $m = 6$ are shown here). Suppose, for example, that R produced a higher outcome than S on trials $t - 4$, $t - 3$, and $t - 2$, but a lower outcome on trial $t - 1$. The decision maker would then retrieve all previous experiences of that same sequence from memory and the outcomes that immediately followed them. Whichever option had a higher average outcome on those trials would be chosen. In a simulation study, we tested whether the CS model could account for the observed choice patterns in the main experiment and whether the auxiliary assumptions of partially relative outcome encoding (ω) and loss aversion (λ) could improve the fit. We tested all values of m from 1 to 8, 10 values of ω spanning the interval $[0, 1]$, and 10 values of λ spanning the interval

[0.25, 2.50]. The model was simulated 1,000 times for each parameter combination and the proportion of R choices was averaged across simulations. Choices of R increased as attention to relative outcomes increased and as loss aversion decreased (in the conditions with mixed outcomes). The parameters that resulted in the smallest discrepancy between actual and simulated choice proportions (root mean squared error) were $m = 6$, $\omega = .44$, and $\lambda = 2.50$ (RMSE = 0.01). In the figure above, the solid horizontal lines represent the observed mean choice proportions in each condition. The dotted horizontal lines represent the simulated choice proportions when $m = 6$, $\omega = 0$, and $\lambda = 1$ (RMSE = 0.05; $m = 6$ resulted in the closest fit under the original CS model assumptions, i.e. $\omega = 0$ and $\lambda = 1$). Without partially relative encoding and loss aversion, the CS model produces an effect of unequal EVs, but fails to produce an effect of mixed outcomes and underestimates R selections in the gains-only conditions.