

Frequency Effects in Decision-Making Involving Loss Minimization

Darrell A. Worthy & Mianzhi Hu

*Texas A&M University*Forthcoming at *Cognition***Abstract**

Recent work provides evidence for frequency effects during decision-making, where less-rewarding options that are presented more frequently are selected more often than more-rewarding options presented less frequently. This is predicted by the Decay but not the Delta reinforcement-learning (RL) model. The Decay model assumes that higher-frequency options are preferred because their past outcomes are more available in memory than those of lower-frequency options. However, most of this research has involved decision-making with gains, rather than losses. In loss-minimization scenarios, the Decay model predicts a *reversed* frequency effect because it assumes greater memory for losses, for the more frequently encountered alternatives. We tested this prediction in three experiments and found that the Decay model provides a very poor fit to data in loss-minimization scenarios. In Experiment 2, where participants tried to minimize their expenditures in a hypothetical shopping scenario, we observed a modest frequency effect. In Experiments 1 and 3, where participants were asked to minimize losses as points, without the hypothetical shopping scenario context, frequency effects were attenuated, but not reversed. These effects were best-accounted for by two novel models, the Prospect-Valence Prediction-Error Decay model (PVPE-Decay), which assumes *relative* rather than absolute processing of rewards, and the Delta-Uncertainty model which assumes aversiveness to less frequent options that are higher in uncertainty. These results dovetail with recent work showing that people process reward outcomes in a context-dependent manner, and they suggest smaller losses can be perceived as relative gains if framed in familiar scenarios involving cost-minimization.

Keywords: Reinforcement learning; Prospect theory; Decision-making; Mathematical modeling; Choice Behavior

Word count: 12,791

Data and Analysis Code Available at:

https://osf.io/n43y5/?view_only=6634bab4ba3741e890f9b7304aae9917

Corresponding Author:

Darrell A. Worthy

Associate Professor

Texas A&M University

4235 TAMU

College Station, TX 77845-4235

worthyda@tamu.edu

worthylab.org

1. Introduction

There has been a long history in psychological research attempting to elucidate how people's decision-making strategies differ when the possible outcomes involve gains versus losses (Kahneman & Tversky, 1979; Gonzalez, Dana, Koshino, & Just, 2005; Pang, Blanco, Maddox, & Worthy, 2017; Yechiam & Hochman, 2013; Zeif & Yechiam, 2022; Yechiam, 2022). Research on framing and reflection effects has shown that people often behave markedly different when the same decision-making problems are framed in terms of gain-maximization versus loss minimization (Fischer, Jonas, Frey, & Kastenmuller, 2008; Kwak & Huettell, 2018; Gallagher & Updegraff, 2012; Fagley & Miller, 1997). For example, the classic behavioral economics literature suggests that people tend to be risk-averse in the context of gains, but risk-seeking in the context of losses (Kahneman & Tversky, 1979; March, 1996), although other researchers have questioned this generality (Schneider & Lopes, 1986).

More recently, decision-making studies have uncovered intriguing frequency effects, where people preferred an option that is slightly lower in average reward value, if it has been presented more frequently than the higher-average reward value alternative (Don et al., 2019; Don & Worthy, 2022; Hu et al., 2025)¹. However, this frequency-based preference has yet to be tested in loss contexts. Interestingly, a popular reinforcement learning (RL) model, the Decay model (Erev & Roth, 1998), which accurately predicts frequency effects in gain contexts, predicts a 'reversed frequency effect' under losses. That is, while frequently presented items tend to be favored in gain contexts, they may be avoided in loss contexts. In the current work, we examine whether this novel framing effect exists in decision-making when options are presented at unequal frequencies.

Previous research has shown that the Decay model correctly predicts an effect of reward frequency in both binary- and continuous-outcome tasks where all rewards were gains (Don et al., 2019; Hu et al., 2025). This model assumes that reward values accumulate, leading to higher expected value estimates for more frequently presented alternatives (Erev & Roth, 1998). Following Estes (1976), who showed extensive effects of reward frequency, Don and colleagues (2019) conducted an experiment which clearly demonstrated the impact of unequal reinforcement frequencies. In this study, participants selected between options AB or CD on separate trials during training. Options A and C were the best in each pair, providing a reward on 65% and 75% of trials respectively. While option C had a higher *average* reward rate, option A was associated with more *cumulative* reward because there were twice as many AB trials as CD trials in the task. During a later test phase, participants selected from options A or C, and there was a bias toward option A, the more frequently presented alternative, even though it had a lower objective reward value. This effect has been replicated using continuous rewards (Hu et al., 2025). Therefore, it was theorized that the Decay model effectively assumes that, when making a decision, people think of the previous rewards associated with each option (Don et al., 2019). More frequently presented items should be more available in memory, and those items will have a higher expected value because the memories of those past outcomes will accumulate, and more gains will be associated with them.

If cumulative rewarding experiences can make an option seem more valuable than less frequently rewarded alternatives, do repeated losses devalue frequently punished options more

¹ We use the term 'frequency effect' rather than 'mere exposure effect' because the tasks demonstrating frequency effects have involved reward-based outcomes, rather than mere exposure. As stated by Zajonc (1974): "When stimulus presentation is accompanied by an opportunity of forming particular associative bonds, we no longer have conditions satisfying the 'mere' exposure hypothesis."

than those encountered less often? For example, does paying per use for a service (e.g., a gym or music app) feel more costly than a higher-priced monthly membership, even if the latter could be objectively more expensive for infrequent users? The Decay model makes an interesting prediction in these scenarios that involve losses. As will be shown below, it predicts a *reversed* frequency effect, where the more frequently encountered item will be chosen less often because more losses are associated with that option. For the same reason that the Decay model predicts enhanced memory for previous rewards in a gains context, it also predicts enhanced memory for losses within a loss-minimization context. Knowing whether people show the same frequency effect under gains and losses is important because it helps us understand how people are remembering, or processing, past outcomes. People could process all the losses received as losses, or negative outcomes, which is assumed by the Decay model. Alternatively, they might process losses within their context, and view small losses as relative gains and large losses as relative losses (Rakow, Cheung, & Restelli, 2020; Brochard & Daunizeau, 2024). The experiments reported below will allow us to examine which of these two possibilities is supported by the data.

In addition to examining the predictions of the Decay RL model, we will also examine the predictions of six additional models. First, the Delta model (Sutton & Barto, 1998; Steingroever, Wetzels, & Wagenmakers, 2014), assumes that expected values are recency-weighted averages of the past outcomes associated with each alternative. Because the Delta model tracks average reward, it does not assume that options that are more frequently presented will be valued any more than less frequently presented alternatives. This model has been shown to provide a poorer account of frequency effects than the Decay model in tasks where the outcomes are gains (Don et al., 2019; Don & Worthy, 2022). We will also fit two variants of the Decay model that make different assumptions regarding how past outcomes are used to compute expected values for each alternative: The Decay-Win model (Hu & Worthy, 2025) assumes that participants' behavior is guided by relative 'wins,' or better than average outcomes, while the Decay-Loss model assumes that participants attend to relative 'losses,' or worse than average outcomes.

The Delta, Decay, Decay-Win, and Decay-Loss models each contain two free parameters. We will also fit four additional models that are more complex versions of the four models listed above; each of these additional models contains four free parameters. The Prospect-Valence Delta (PVL-Delta) and Prospect-Valence Decay (PVL-Decay) models are extensions of the Delta and Decay models, respectively, that have two additional parameters that are motivated by Prospect Theory (Ahn et al., 2008; Steingroever, Wetzels, & Wagenmakers, 2013). These models include a *shape* parameter, which allows for discounting of the magnitude of rewards, and a loss-aversion parameter, which allows the model to give more weight to either gains or losses. The Prospect Valence Prediction-Error Decay (PVPE-Decay) model is a novel model we developed to include the Decay-Win and Decay-Loss models nested as special cases. A similar model was used in another recent paper from our lab, and it fit gambling task data much better than the Delta model (Don et al., 2022). The PVPE-Decay model includes a shape parameter, just like the PVL-Delta and PVL-Decay models, and it also has a parameter that weights the effects of positive versus negative prediction errors. The final model is the Delta-Uncertainty model. This model tracks the uncertainty associated with each option, which is operationalized as a combination of the variance in rewards provided by each option, and how often each option has been selected in the past (i.e. familiarity). The Delta-Uncertainty model is designed to be better equipped to account for frequency effects than the basic Delta model because it penalizes options that are higher in uncertainty. A key difference between the models is that the Delta,

Decay, PVL-Delta, PVL-Decay and Delta-Uncertainty models all assume *absolute*, or context-free processing of reward outcomes, while the Decay-Win, Decay-Loss, and PVPE-Decay models all assume that rewards are processed in a *relative* manner, by being compared to the average reward provided across all outcomes. Comparing the fits of these two classes of models will allow us to determine whether losses are processed in an absolute or in a relative manner.

The models, along with their assumptions and predictions, will be detailed in the Model Formalisms section below.

1.1 Model Formalisms

All eight models compute expected values for each alternative presented in the task. These expected values are entered into the softmax rule shown in in Equation 1 to determine each model's probability of selecting each j alternative on trial t :

$$P|C_{j,t}| = \frac{e^{\beta \cdot EV_{j,t}}}{\sum_1^{N(J)} e^{\beta \cdot EV_{j,t}}} \quad (1)$$

Consistent with Yechiam & Ert (2007), $\beta = 3^c - 1$; $c \in (0,5)$, where c is an inverse temperature parameter that modulates how often the option with the higher expected value is chosen. As c approaches 0, choices are more random, inversely, choices are weighted more heavily towards the choice with the highest expected value as c increases.

1.1.1 Basic Learning Models

We divide our set of models into Basic models which have two free parameters, and Extended models which have four. The first basic model is the Delta model, which assumes that the expected value (EV) is updated for each j option on each t trial according to Equation 2:

$$EV_{j,t+1} = EV_{j,t} + \alpha \cdot (r_t - EV_{j,t}) \cdot I_j \quad (2)$$

Where I_j is an indicator variable that is set to 1 if option j is chosen on trial t , and 0 otherwise. This formulation ensures that only the expected value for the chosen option is updated, and the other options, whether seen or not, are not updated. Alpha (α) is the learning rate, or recency parameter. Higher α values indicate greater weight to more recent outcomes. To reduce multicollinearity between the learning rate and inverse temperature parameters we limited the range of alpha to $\alpha \in (0.01, 0.99)$ in all of our simulations and model fits.

The next basic learning model is the Decay model. This model tracks changes in expected value, but instead of updating the expected value by the prediction error, the raw reward value is used (r_t). On each trial, each j option will be modified by a decay parameter (A ; $A \in (0.01, 0.99)$) regardless of whether the option was seen or chosen. Critically, this means that the expected value for each option will decay over time and only increase when a reward for that option is received. Thus, the more frequent the reward, the greater the expected value. The Decay rule is updated according to Equation 3:

$$EV_{j,t+1} = EV_{j,t} \cdot (1 - A) + r_t \cdot I_j \quad (3)$$

Similar to Equation 2, I_j is an indicator variable that equals 1 if j is chosen, and zero otherwise. This means that all options decay toward zero on each trial, but the expected value of the chosen option is incremented by the reward given. As our simulations will confirm below, the Decay model predicts a bias toward more frequently presented options when gains are given, but a bias against more frequently presented options when losses are given.

The first new model we present is the Decay-Win model (Hu & Worthy, 2025). This model is also a basic, two-parameter model, and it assumes that rewards are processed *relative* to other rewards given in the same context. To provide an estimate of the average reward provided

across all options, this model tracks the recency-weighted average reward received on each trial according to Equation 4:

$$AV_{t+1} = AV_t + A \cdot (r_t - AV_t) \quad (4)$$

The Decay-Win model also assumes that only the valence of the outcome is used to guide choices, and specifically only the presence of positive outcomes. If the current reward is greater than the average reward ($r_t - AV_t$), then expected values are updated according to Equation 5, with I_j equal to 1:

$$EV_{j,t+1} = EV_{j,t} \cdot (1 - A) + 1 \cdot I_j \quad (5)$$

If the current reward is not greater than the average reward then the indicator variable, I_j , is set to 0; all expected values decay, but no expected value is incremented unless the reward is greater than average. The Decay-Win model thus assumes that a trial is considered a ‘win’ if the reward surpasses a threshold of being larger than average. The model does not track any information about the specific magnitudes of the rewards provided, it simply tracks the number of positive outcomes associated with each option. The Decay-Win model makes the same predictions for tasks involving gains or losses. Unlike the standard Decay model, the Decay-Win model does not predict a reversed frequency effect with losses but instead predicts a bias toward more frequently presented options in *both* gain and loss conditions.

The final Basic learning model, the Decay-Loss model assumes the opposite strategy of the Decay-Win model. Whereas the Decay-Win model tracks how often each option has provided a ‘win,’ or better than average reward, the Decay-Loss model tracks how often each option has provided a ‘loss,’ or worse than average reward. If the current reward is less than the average reward provided by Equation 4, then expected values are updated according to Equation 6, with I_j equal to 1:

$$EV_{j,t+1} = EV_{j,t} \cdot (1 - A) - 1 \cdot I_j \quad (6)$$

If the current reward is not less than the average reward then the indicator variable, I_j , is set to 0; all expected values decay upwards towards zero, but no expected value is decremented unless the reward is less than average. Thus, the Decay-Loss model assumes a loss-avoidant strategy, whereas the Decay-Win model assumes a gain-seeking strategy.

1.1.2 Extended Learning Models

As stated above, we also fit four models that were extensions of the basic models. The PVL-Delta model is an extension of the Delta model that includes two additional parameters that allow the model to account for discounting of large magnitude rewards, and for greater attention to losses versus gains. Rather than use the actual reward received on each trial (r_t) in Equation 2, the PVL-Delta model transforms the outcome received on trial t into a representation of subjective utility (u_t):

$$u_t = \begin{cases} r_t^\gamma & \text{if } r_t \geq 0 \\ -\lambda|r_t|^\gamma & \text{if } r_t < 0 \end{cases} \quad (7)$$

Here, the shape parameter γ ($0 \leq \gamma \leq 1$) determines the shape of the utility function. When $\gamma = 1$, all rewards are processed veridically, but as the shape parameter approaches 0 reward magnitudes are discounted. When $\gamma = 0$, all rewards are processed as the same amount (1), and the magnitude is completely disregarded. The loss aversion parameter λ ($0 \leq \lambda \leq 5$) allows for greater learning from losses or gains. When this parameter is set to 1, losses and gains receive equal weight, with values less than 1 indicating greater attention to gains than losses, and values greater than 1 indicating greater attention to losses than gains.

The utility is then entered into a Delta learning rule to update the expected value for the chosen option:

$$EV_{j,t+1} = EV_{j,t} + \alpha \cdot (u_t - EV_{j,t}) \cdot I_j \quad (8)$$

As in Equation 2, I_j , is an indicator variable that equals 1 if option j was chosen on trial t , and 0 otherwise.

The PVL-Decay model also uses Equation 7 to compute the utility of each outcome. The utility is then entered into a decay rule according to:

$$EV_{j,t+1} = EV_{j,t} \cdot (1 - A) + u_t \cdot I_j \quad (9)$$

Equation 9 is identical to Equation 3 except that actual reward outcome is replaced with subjective utility (u_t).

The third extended model, the PVPE-Decay model, is an extension of both the Decay-Win and Decay-Loss models, and each of these models are nested within the PVPE-Decay model as special cases. Like these simpler models, the PVPE-Decay model also assumes that rewards are processed relative to the overall average reward provided across all options, and the average value (AV_t) is computed using Equation 4. The AV is then used to compute the subjective utility of the outcome according to:

$$u_t = \begin{cases} (1 - w_L) \cdot (r_t - AV_t)^\gamma & \text{if } (r_t - AV_t) \geq 0 \\ w_L |(r_t - AV_t)|^\gamma & \text{if } (r_t - AV_t) < 0 \end{cases} \quad (10)$$

This utility function is similar to the utility function for the PVL-Delta and PVL-Decay models, except that it uses relative reward ($r_t - AV_t$), rather than the actual reward (r_t). Another difference is that this model uses a weight parameter for losses versus gains w_L ($0 \leq w_L \leq 1$). This allows the PVPE-Decay model to include the Decay-Win model nested as a special case when $w_L = 0$ and $\gamma = 0$, and the Decay-Loss nested as a special case when $w_L = 1$ and $\gamma = 0$.

Finally, the fourth extended model we used is the Delta-Uncertainty model. We fit this model because it is possible that frequency effects are due to lower uncertainty associated with the more frequent alternatives, compared to items encountered less often (Hu et al., 2025). This model learns expected values in the same way as the basic Delta model; however, the prediction error on each trial is used to track the variance, or uncertainty in rewards for the chosen option. One additional free parameter, Unc_0 , represents the initial uncertainty in reward for each option. This parameter is used to initialize uncertainty values for each j option according to:

$$UV_{j,0} = Unc_0^2 \quad (11)$$

The uncertainty value for the chosen option is then updated on each trial according to:

$$UV_{j,t+1} = UV_{j,t} + \alpha \cdot [(r_t - EV_{j,t})^2 - UV_{j,t}] \cdot I_j \quad (12)$$

In Equation 12, the squared prediction error from the basic Delta model is used to updated the uncertainty associated with the chosen option. We allowed the initial uncertainty parameter to vary from 0.5 to 5, which is greater than the standard deviation of rewards for each option. Therefore, the UVs associated with each option will generally decrease as they are selected more frequently. The uncertainty values were then converted into uncertainty estimates by taking their square root, and dividing by the number of times the chosen option had been selected:

$$Unc_{j,t} = \frac{\sqrt{UV_{j,t}}}{\sqrt{n_{j,t}}} \quad (13)$$

Dividing by the number of times each option has been chosen, allows the model to further reduce uncertainty associated with more frequent options.

The uncertainty values were then subtracted from the expected values, with a free parameter, w_{Unc} , weighting the degree to which the participant avoided options with high uncertainty:

$$QV_{j,t} = EV_{j,t} - Unc_{j,t} \cdot w_{Unc} \quad (14)$$

These Q-values were then used in Equation 1, in the place of expected values. Given that these uncertainty values were usually smaller than one, we allowed this free parameter (w_{Unc}) to vary between 0 and 1000 to allow the model to give greater weight to uncertainty than to expected values when this parameter was large. When w_{Unc} equals zero, the Delta-Uncertainty model is identical to the basic Delta model.

To summarize the Delta, Decay, PVL-Delta, PVL-Decay, and Delta-Uncertainty models all assume absolute, context-free processing of the gains or losses given by each option, while the Decay-Win, Decay-Loss, and PVPE-Decay models assume relative, or context-dependent processing where the outcomes provided by each option are processed relative to the overall average reward provided across all options. In the next section, we will show that the relative-reward processing models make similar predictions across gain-maximization and loss-minimization tasks; however, two of the absolute-reward processing models, the Decay and PVL-Decay models, predict a reversed effect of frequency under losses compared to gains. We will then present three experiments with human participants and evaluate which model provides the best fit and post-hoc recovery of participants' behavior.

1.2 A Priori Simulations

We simulated each of the above models in a task that was modified from that of Don and colleagues (2019), under both gains and losses conditions. There were a total of four options that the simulated agent chose from on different trials, labeled options A-D. The rewards given by each option were continuous and drawn from normal distributions, with the mean reward values in the gains task for options A-D equal to [0.65, 0.35, 0.75, 0.25]. The mean values for options A-D in the losses task were [-0.35, -0.65, -0.25, -0.75]. The values for the losses condition were simply the values from the gains condition subtracted by one. The standard deviation around the mean reward value was 0.43 for all options. This value was calculated based on the standard deviation from the binomial distribution for option C, the highest valued option: $(0.7 * 0.3)^{.5} = 0.43$. Using this value for the standard deviation around each mean value made the reward structure roughly equivalent to a continuous-rewards version of the binary-outcome task from Don et al. (2019), and was also implemented in Hu et al. (2025) and Hu & Worthy (2025).

For all models, expected values for all j options were initialized at the first reward or utility value given on trial 1. This restricted the initial expected values to be on the same scale as the rewards or utilities used to update the model's expected values on each trial. During training, the models selected from options AB or CD on different trials. There were 100 AB trials, and 50 CD trials. The trial types were interspersed randomly for each simulation. During the test phase, the models selected from novel option pairs, CA, CB, AD, and BD, each for 25 trials. The CA test trials are of most interest because the models must choose between the two high-value options within each training pair. Although option A has a lower mean value of 0.65, it is selected more than option C, which has a higher average value of 0.75, in a binary outcome task (Don et al., 2019; Don & Worthy, 2022).

We simulated each task 1,000 times with each model, with parameter values randomly drawn from a uniform distribution across the ranges presented above. Figure 1, shows the average proportion of C choices on the critical CA test phase trials, averaged across all simulations for each model.




Figure 1. Predicted proportion of C choices on CA test phase trials for a priori simulations for each model. Black horizontal line indicates equal choices for options A and C.

The Delta and PVL-Delta models clearly predict a preference for the more rewarding option C, across both conditions. The Decay and PVL-Decay models predict fewer C, or more A choices in the gains condition; however, as expected, these models predict more C choices in the losses condition. As described above, the Decay model predicts a frequency effect under gains, but a reversed frequency effect under losses, and the PVL-Decay model makes similar predictions. The Decay-Win model predicts a frequency effect across both the gains and losses conditions. In contrast, the Decay-Loss model predicts a reversed frequency effect in both gains and losses conditions, where the simulated agents consistently preferred option C over option A in both conditions, similar to the Delta and PVL-Delta models. Finally, the PVPE-Decay and Delta-Uncertainty models predict roughly equal choices of A and C, because the PVPE-Decay model is flexible enough to mimic both the Decay-Win and Decay-Loss models, and the Delta-Uncertainty model can mimic the Delta model if the weight to uncertainty is low, but it can also predict frequency effects if the weight to uncertainty is high.

2. Experiment 1

In Experiment 1, we ran participants in the two conditions simulated above to examine whether they displayed a reversed frequency effect in the losses condition, as predicted by the Decay model, or whether participants showed a similar frequency effect under both gains and losses, where they preferred the more frequently rewarded alternative during the test phase.

2.1 Method

2.1.1 Participants

We conducted an a priori power analysis using G*Power software (Kang, 2021) to estimate the appropriate sample size to conduct *t*-tests on the proportion of C choices made on the critical test trials between the gains and losses conditions. Assuming an effect size of $d = 0.5$, an alpha threshold of 0.01,² we would need 96 participants in each condition for 80% power. Based on past studies from our lab, we reasoned that some participants may show little evidence of learning during the training phase, particularly in the loss-minimization conditions which are sometimes confusing to participants. Therefore, we planned to run approximately 120 participants in each condition to account for “noise” in the data from participants who did not sufficiently learn which choices were optimal during training.

Our final sample size was 252 participants. Participants were randomly assigned to one of the two between-subjects conditions, gains or losses: There were 120 participants in the Gains condition (83 females, 35 males, 2 other), and 132 in the Losses condition (85 females, 46 males, 1 other).

2.1.2 Materials and Procedure

Participants performed the task on PC computers in a laboratory setting. They first completed a series of questionnaires which were added as part of a pilot study, and will not be analyzed here. These scales are listed and briefly described in the Supplemental Materials. Participants then performed the main task. During the first 150 trials, referred to as the training phase participants selected from options A-B or from options C-D. In the gains condition, option A provided an average reward of 0.65 points, while option B had an average reward of only 0.35 points. Within the CD pair, option C was dominant over option D, with an average reward of 0.75 compared to 0.25 for D. Rewards were drawn on each trial from continuous normal distribution with a standard deviation of 0.43. This value was based on the standard deviation from a binomial distribution for option C $(0.75 * 0.25)^{.5}$, and our goal was to create a continuous version of the binomial task used in Don et al. (2019). Thus, the average reward values and the variance are roughly equivalent to that in Don et al.’s 2019 study, although we used continuously distributed, rather than binomial rewards in the present study. This reward structure has also been shown to produce regular frequency effects in gains contexts (Hu et al., 2025).

The losses conditions were created as analogues of the gains task, with average rewards for decks A-D equal to: -0.35, -0.65, -0.25, -0.75. The variance in rewards was the same as in the gains condition. Participants were told that they would lose points on most trials, and that their goal was to lose as few points as possible. It is important to note that gains and losses were sometimes given in both conditions, due to the high variance in rewards. Figure 2 shows example screenshots from the experiment. Participants were allowed to make choices at their own pace, and they were not given information about how many trials were left in the experiment. They were only shown the outcome for the option selected on each trial; foregone outcomes were not presented.

² We planned to conduct Bayesian *t*-tests, which are generally more conservative in rejecting the null hypothesis than frequentist *t*-tests with alpha = .05 (Wetzel et al., 2011), therefore we used an alpha level of .01. We acknowledge that conducting a fully Bayesian power analysis would have been more appropriate, but we are unaware of any software similar to G*Power for Bayesian power analysis, and we felt conducting a frequentist power analysis in this manner suited our purpose.




Figure 2. Example screenshots for the training and test phases of the experiment.

Participants in the both conditions performed twice as many AB training trials as CD trials (100 compared to 50). AB and CD trials were randomly interspersed across training. In the transfer phase participants completed 25 trials of each of the remaining combinations of options: CA, CB, AD, and BD. These trial types were randomly interspersed across the 100-trial transfer phase. Feedback was given on each trial of the training phase, but for the test phase, participants were not shown how many points they earned. In the gains condition, participants were told that their job was to figure out which option within each pair was most rewarding. In the losses condition, participants were told that their job was to figure out which option gave the smallest losses. For the transfer phase, participants were told that they would not be shown how many points they received after each trial, but to try to pick the best option based on what they had learned so far. Participants were not given a specific goal, and they were not given monetary rewards in the task, but were simply asked to do their best to pick the most rewarding option.

2.1.3 Data Analysis

We ran Bayesian general linear mixed-effects regression models to analyze the data using the R package *brms* (Bürkner, 2017). Brms provides parameter estimates for both fixed and random effects. We examined the fixed-effect coefficient values from models where condition variables (e.g. reward condition) were used as predictors for the proportion of optimal choices. We considered an effect to exist or be ‘significant’ if the 95% highest credible interval (HCl) for the predictor did not include zero (Nalborczyk et al., 2019; Byrne et al., 2020).

We also conducted Bayesian *t*-tests on the key dependent variables such as the proportion of C choices on the critical CA transfer trials. We used JASP software for our Bayesian analysis (jasp-stats.org; version 0.17.2.1) using the default Cauchy prior (.707). We report Bayes Factors in terms of evidence supporting the alternate (BF_{10}) hypothesis. A Bayes Factor (BF_{10}) of 3 or more is considered to indicate moderate support for the alternate hypothesis, and a Bayes Factor (BF_{10}) less than 1/3 is considered moderate support for the null (Wagenmakers et al., 2018; Jeffreys, 1961), although Bayes Factors can be interpreted continuously on an odds scale. For

example, a Bayes Factor (BF_{10}) of 3 suggests that the alternate hypothesis is three times more likely than the null hypothesis, given the data.³ Bayes Factors greater than 10 are considered strong support for the alternate hypothesis, and Bayes Factors greater than 100 indicate extreme support (Wagenmakers et al., 2018; Jeffreys, 1961).

2.2 Results

We first computed the proportion of optimal choices during training, or the proportion of A and C choices in the AB and CD choice pairs, respectively. These are shown in Figure 3a. We ran a Bayesian multilevel logistic regression model that predicted whether participants made the optimal choice on each trial (coded as 1 for optimal, 0 for not-optimal) based on condition, with random intercepts for each participant. There was a significant effect of reward type, $b = -0.42$, $SE = 0.12$, $95\% HCl = [-0.65, -0.19]$, $OR = 0.65$, which suggests that learning was better in the gains compared to the losses condition. The odds ratio for selecting the optimal choice decreased by a factor of 0.66 for the losses compared to the gains condition. A similar model that included the interaction term between trial type (AB coded as 0; CD coded as 1) and condition indicated that there was a significant interaction, $b = 0.20$, $SE = 0.05$, $95\% HCl = [0.11, 0.30]$, $OR = 1.21$. This suggests that for participants in the losses condition, the odds of selecting the optimal choice increased by a factor of 1.21 on CD, compared to AB trials. This could have been due to the larger difference in average loss between options C and D than A and B, but it is notable that learning was equivalent for AB and CD pairs in the gains condition.




Figure 3 **a.)** Proportion of optimal choices during training for each trial type. **B.)** Proportion of C choices on CA transfer trials. Error bars represent standard errors of the mean.

We next examined the proportion of C choices on the critical CA test trials, which are shown in Figure 3b. Visual inspection of the graph suggests that there was a frequency effect in the gains condition, where participants selected option C less often than chance. A one-sample Bayesian t-test with 0.5 set as the test value indicated a significant difference, $BF_{10} = 20.27$, $d = .31$. However, in the losses condition participants selected from options A and C roughly equally, on average. A one-sample Bayesian t -test indicated support for the null hypothesis that the proportion of C choices did not differ from 0.5, $BF_{10} = 0.11$, $d = .05$. We ran a similar Bayesian multilevel model as in the training data above, with C choices (coded as 1, 0 for A choices) regressed on the effect of condition, with random intercepts for participants. This model suggested no difference based on condition, $b = 0.48$, $SE = 0.28$, $95\% HCl = [-0.06, 1.03]$,

³ Bayes Factors for the alternate and null hypotheses are inverse of each other ($BF_{10} = 1 / BF_{01}$).

$OR = 1.61$. A Bayesian independent samples t-test on the average proportion of C choices between conditions did not indicate at least moderate support for either the null or alternate hypothesis, $BF_{10} = 0.891$, $d = .25$. Thus, although the proportion of C choices were significantly below chance in the gains condition, the difference between the gains and losses condition was not significant, and participants selected option A slightly more than C in the losses condition.

Figure 4 shows the distribution of C choices on the critical CA test trials. Values to the left of each plot indicate a preference for A, the more frequently rewarded option, and values on the right indicate a preference for C. Interestingly the modal value in each condition is close to 0% C choices, indicating that many participants showed a strong effect of frequency in both conditions. However, in the losses condition there are considerably more participants who selected C almost exclusively, than there are in the gains condition, as indicated by the cluster of participants to the extreme right, within the losses plot. There are also more losses participants who chose A and C roughly equally often.




Figure 4 Distributions of the average proportion of C choices on the critical CA test trials for participants in each condition in Experiment 1.

2.2.1 Model Fits

For each model, we fit each participant's data individually by estimating maximum likelihood for the eight models presented above. All choices except for the very first trial were fit. To compare the models, we computed the BIC value for each model (Schwarz, 1978), which penalizes models based on their number of free parameters. Table 1 shows the average best-fitting parameter values, along with the average BIC values. Lower BIC values indicate a better fit, and BIC values can be transformed into Bayes Factors, favoring one model over the other by exponentiating the difference between the poorer fitting model and the best-fitting model and dividing by two (Wagenmakers, 2007). This means that a BIC difference of 3 indicates moderate

support for the better fitting model ($BF_{10} = 4.48$); a BIC difference of 5 indicates strong support ($BF_{10} = 12.18$). Thus, BIC differences less than 3 indicate that neither model is substantially more supported by the data than the other, and these values are presented in bold in Table 1 to indicate that the model did not fit significantly worse than the best-fitting model. We also conducted group-level random-effects Variational Bayesian model selection (VBMS; Stephan et al., 2009), which treats models as random variables that may vary across individuals. In this framework, model frequencies are estimated by fitting a Dirichlet distribution, which is then used to define a multinomial distribution representing the probability that any given model generated the data for a randomly selected subject. Specifically, the posterior Dirichlet parameters, α , represent the estimated population frequency with which each model generated individual data. The posterior multinomial parameter, r_k , describes the probability that data from a randomly chosen subject is generated by a specific model k . Finally, the exceedance probability, φ_k , quantifies the likelihood that a particular model k is more likely than any other model in the comparison set to generate group-level data. In both conditions, the Decay-Win, PVPE-Decay, and Delta-Uncertainty models provided a substantially better fit to the data, on average, than any of the other models, with the exception of the Decay model in the gains condition. The PVPE-Decay and Delta-Uncertainty models fit substantially better than the other extended models, the PVL-Delta and PVL-Decay models. In both conditions, the Decay-Win model provided the best fit of the Basic, two-parameter models; however, the Delta and Decay models fit almost as well in the Gains condition. In the Losses condition, the PVPE-Decay model fit substantially better than both the Delta ($BF_{10} = 25.79$) and Decay models ($BF_{10} > 10K$). Thus, two of the relative-reward processing models, Decay-Win and PVPE-Decay, were the best-fitting basic and extended models. Other than the Delta-Uncertainty model, the absolute-reward processing models received less support, particularly in the Losses condition.

Table 1. Average best-fitting BIC values and BIC-weights in Experiment 1

	Mean BIC	BF _{BestModel,M}	BIC-weight	VB α	VB r_k	VB φ_k
Gains						
Delta	285.75	136.32	.18	25.51	0.20	0.39
Decay	278.12	3.00	.17	20.79	0.16	0.10
Decay-Win	275.92	-	.18	26.56	0.21	0.50
Decay-Loss	352.09	>10K	.08	8.52	0.07	< .001
PVL-Delta	290.47	1442.75	.10	11.32	0.09	< .001
PVL-Decay	285.37	112.73	.12	15.25	0.12	0.01
PVPE-Decay	277.69	2.42	.06	7.73	0.06	< .001
Delta-Uncertainty	278.82	4.26	.10	12.32	0.10	< .001
Losses						
Delta	310.41	25.79	.28	46.61	0.33	0.41
Decay	348.40	>10K	.08	6.98	0.05	< .001
Decay-Win	304.61	1.41	.31	48.84	0.35	0.59
Decay-Loss	348.79	>10K	.11	12.48	0.09	< .001
PVL-Delta	315.99	419.89	.02	1.09	0.01	< .001
PVL-Decay	358.70	>10K	.03	1.24	0.01	< .001
PVPE-Decay	303.91	-	.10	12.44	0.09	< .001
Delta-Uncertainty	308.50	4.26	.10	10.32	0.07	< .001

Note: Bayes Factors in bold indicate a model fit close to the best-fitting model.

The third column in Table 1 lists the average BIC weights for each model (Wagenmakers & Farrell, 2004). These values are similar to the proportion of data sets best fit by each model, however, some weight is given to models that fit nearly as well as the best-fitting model.

Interestingly, the Decay-Win model has the highest BIC-weight in each condition, while the PVPE-Decay model receives less weight. This pattern is also replicated in VBMS results, which suggests that a substantial proportion of participants' data are best fit by the Decay-Win model alone, and the inclusion of the shape and weight to relative losses parameters does not provide a substantially better fit. Following the Decay-Win model, the Delta model had the next highest BIC-weights and VBMS indices in both conditions, and the Decay model received more weight in the gains than in the losses condition.

Table 2. Average best-fitting parameter and BIC values in Experiment 1

Parameter:	a or A	c	γ or Unc_0	λ or w_L or w_{Unc}
Gains				
Delta	.33 (.35)	1.49 (1.06)		
Decay	.23 (.31)	.45 (.39)		
Decay-Win	.17 (.23)	.42 (.35)		
Decay-Loss	.74 (.35)	.14 (.22)		
PVL-Delta	.33 (.36)	1.36 (.89)	.50 (.43)	2.73 (1.41)
PVL-Decay	.25 (.31)	.43 (.34)	.57 (.44)	2.47 (1.36)
PVPE-Decay	.16 (.24)	.61 (.41)	.26 (.34)	.29 (.33)
Delta-Uncertainty	.31 (.32)	1.05 (.78)	1.53 (1.84)	106.31 (273.26)
Losses				
Delta	.31 (.36)	1.43 (1.33)		
Decay	.68 (.38)	.22 (.29)		
Decay-Win	.22 (.32)	.29 (.27)		
Decay-Loss	.69 (.37)	.22 (.30)		
PVL-Delta	.31 (.37)	1.26 (.81)	.46 (.44)	2.63 (1.35)
PVL-Decay	.70 (.37)	.22 (.27)	.74 (.41)	2.33 (1.31)
PVPE-Decay	.24 (.33)	.54 (.35)	.27 (.37)	.42 (.37)
Delta-Uncertainty	.24 (.32)	.92 (.85)	1.35 (1.69)	180.52 (354.64)

Note: Values in parentheses are standard deviations.

Table 2 lists the average parameter values for each model. Because the PVPE-Decay model provided the best fit, and includes the Decay-Win and Decay-Loss models as special cases, we focus on its best-fitting parameter estimates. Of interest, the weight to relative losses parameter (w_L) is below 0.5 in each condition, indicating a stronger reliance on relative gains than relative losses. Figure 5 shows the distribution of best-fitting parameter values for each of the four parameters from the PVPE-Decay model. The distribution of best-fitting inverse-temperature, or sensitivity parameters (c) is fairly normally distributed (Figure 5b). However, there is strong bimodality in the distributions of the other three parameters, where there are large subclusters of participants whose data are best fit by extreme values at the parameter bounds. For the decay parameter (Figure 5a) this distribution suggests that many participants had little or no decay of past outcomes, while another group of participants showed almost complete decay of past outcomes, where their decisions were based mainly on the outcome from the last trial. For the shape parameter (Figure 4c) a value of 1 indicates full processing of the magnitude of relative

rewards, while values of 0 indicate that all reward magnitude are processed as either +1 or -1, depending on whether the relative reward is positive or negative, as in the Decay-Win and Decay-Loss models. Finally, the distribution of weight-to-loss parameters suggests that a substantial proportion of participants attended solely to relative gains or solely to relative losses, similar to the strategy assumed by the Decay-Win and Decay-Loss models. Overall, the pattern suggests that the PVPE-Decay model is flexible enough to account for a variety of strategies, and that a substantial portion of participants were using specific strategies assumed by the basic Decay-Win or Decay-Loss models.




Figure 5 Distributions of the average proportion of C choices on the critical CA test trials for participants in each condition.

2.2.2 Post-hoc Simulations

We next conducted post-hoc simulations using the best-fitting parameters for each model (Ahn et al., 2008; Busemeyer & Stout, 2002). For each participant's best fitting parameters, we conducted 200 simulations and averaged the model's predicted choices across those simulations. Figure 6 shows the average proportion of C choices on CA test trials, along with the same data for participants. Qualitatively, the Decay, Decay-Win, PVL-Decay, and PVPE-Decay models are the only ones that can reproduce the frequency effect within the gains condition, where participants select option C less than option A. For the losses condition, the PVPE-Decay model appears to come closest to reproducing the observed proportion of C choices, with the Decay-Win model underpredicting C choices, and the other models overpredicting C choices. Despite fitting the data well, the Delta-Uncertainty model did not reproduce the frequency effect.




Figure 6 Average proportion of C choices on critical CA test trials from post-hoc simulations for Experiment 1.

To quantify each model's performance, we computed the root mean squared deviation between each model's predicted choice proportions to those observed from human participants, across all trials. For each trial type we computed each model's predicted proportion of optimal choices in the order that they were presented. For example, across all model simulations we computed the average C choices made on the first CA trial, then choices for the second CA trial and so on, for all trial types. These RMSD values are shown in Table 3. For CA trials, the PVPE-Decay model had the lowest RMSD values in the gains condition, and it had slightly higher RMSD than the PVL-Decay model in the losses condition. In the gains condition, where a significant frequency effect was observed, the Delta and PVL-Delta models had the poorest RMSD.

Table 3: RMSD Values From Post-hoc Simulations

	Gains	Losses
CA Trials		
Delta	.144	.057
Decay	.088	.061
Decay-Win	.080	.124
Decay-Loss	.124	.061
PVL-Delta	.133	.065
PVL-Decay	.081	.028
PVPE-Decay	.037	.032
Delta Uncertainty	.116	.044
All Trials		

Delta	.073	.051
Decay	.058	.101
Decay-Win	.056	.079
Decay-Loss	.171	.104
PVL-Delta	.072	.051
PVL-Decay	.056	.056
PVPE-Decay	.058	.069
Delta Uncertainty	.082	.048

Table 3 also lists the average RMSD values across all six trial types. The Decay-Win and PVL-Decay models had the lowest overall RMSD in the gains condition, followed closely by the PVPE-Decay, and basic Decay models. Interestingly, in the losses condition, the Delta-Uncertainty model had the lowest overall RMSD, followed by the basic Delta and PVL-Delta models. Thus overall, the PVPE-Decay model provided the best post-hoc recovery of the frequency effect observed in the gains condition, and it provided good recovery across all trials. The Decay-Loss model provided the worst post-hoc recovery across all trials, and most of the other models provided a good account for trials other than the critical CA trials in the gains condition. The RMSD values for each individual trial type are shown in the Supplemental Materials (Table S1), along with plots of the observed and simulated data.

2.3 Discussion

The results of Experiment 1 clearly do not support the prediction of a reverse frequency effect, made by the Decay and PVL-Decay models in the losses condition. However, while we did not observe a reverse frequency effect under losses, we also did not observe a strong frequency effect. One possibility for the ambiguous results in the losses condition is that we observed significantly poorer learning compared to the gains condition. In an effort to improve learning, and enhance participants' understanding of the purpose of the task, we designed a second experiment, which included only the losses conditions, where we created a more engaging cover story where participants were told to imagine that they were workers in a dog shelter, and on each trial they picked from one of two stores from which they could purchase dog food for the shelter. Their goal, on each trial, was to try to pick the store that would provide the cheapest food, so as to minimize the total money spent on dog food.

We predicted that this would be a more engaging cover story, or scenario for the task, and that participants would be more likely to understand that large numerical values were worse than small numerical values, than in Experiment 1. We also modified the reward structure in Experiment 2, so that all outcomes were losses, whereas in Experiment 1 there were some rare gains in the losses condition, and rare losses in the gains condition. We believed this modification might further strengthen participants' understanding of the task.

3. Experiment 2

Experiment 2 included two conditions, both with a losses reward structure: a control condition where there were equal AB and CD trials during training, and a frequency condition similar to Experiment 1. In the control condition, we predicted that people would show a preference for option C on CA test trials, as found in prior work (Don & Worthy, 2022).

3.1 Method

3.1.1 Participants

Based on the power analysis reported for Experiment 1 above, we planned to run approximately 120 participants in each condition. Our final sample size was 244 participants,

122 in each condition. Within the control condition there were 79 females, 42 males, and one other; within the frequency condition there were 73 females, 48 males, and one other.

3.1.2 Materials and Procedure

Participants performed the experiment on the same computers and used the same software as in Experiment 1. Figure 7 shows example screen shots from the training phase of the experiment. Participants were told that they would be playing the role of a person who works at a dog shelter who is tasked with buying food for the dogs each day. On each trial they picked one of two options that represented the stores they could buy from. Each time they made a choice they were shown how much the dog food cost that day. Their job was to figure out which option in each pair had the lowest price. The total amount spent was shown at the top, and participants were told to try their best to minimize that amount.



Figure 7. Example screenshots for Experiment 2.

The trial structure of the task was identical to that from Experiment 1. In the control condition, participants performed 75 AB trials and 75 CD trials during training, while participants in the frequency condition performed 100 AB trials and 50 CD trials. Trial types were randomly interspersed and randomized separately for each individual. During the test phase all participants performed 25 trials of each of the novel trial types: CA, CB, AD, and BD, in a randomly interspersed order.

The reward structure was a linear transformation of the losses reward structure used in Experiment 1. Average reward values for options A-D were multiplied by 10, and then 10 points

was subtracted from each value. These yielded average losses for options A-D of (-\$13.50, -\$16.50, -\$12.50, and -\$17.50). The standard deviation in rewards from Experiment 1 was .43, and this value was multiplied by 10 for the current experiment, yielding a standard deviation around the mean reward values of 4.3. Options A and C had the lowest average cost for dog food within each training pair, and option C had the lowest cost overall (-\$12.50 versus -\$13.50 for A).

3.1.3 Data Analysis

We used the same data analysis methods as reported in Experiment 1.

3.2 Results

Figure 8a shows the proportion of optimal choices during the training phase. A Bayesian mixed effects logistic regression model predicting optimal choices during training from condition, with random intercepts for participants, indicated no effect of condition, $b = -0.07$, $SE = 0.14$, $95\% HCl = [-0.34, 0.22]$, $OR = 0.93$. We also ran a similar model predicting optimal choices from the interaction between condition and trial type, with random intercepts for participants. There was no interaction effect, $b = -0.00$, $SE = 0.05$, $95\% HCl = [-0.09, 0.10]$, $OR = 0.99$. Thus, there was no difference between conditions, and participants in each condition showed similar levels of learning for each trial type.




Figure 8 a.) Proportion of optimal choices during training for each trial type in Experiment 2. b.) Proportion of C choices on CA transfer trials. Error bars represent standard errors of the mean.

We next examined the proportion of C choices on the critical CA transfer trials, which are shown in Figure 8b. A Bayesian mixed effects logistic regression model with optimal choices predicted from condition, with random intercepts for participants, indicated a significant effect of condition, $b = 1.04$, $SE = 0.35$, $95\% HCl = [0.37, 1.77]$, $OR = 2.86$. The odds ratio indicates that on any CA test trial, the odds of selecting option C were 2.86 times higher for participants in the control condition. A Bayesian t-test with average C choices as the dependent variable indicates a moderate effect of condition, $BF_{10} = 6.33$, $d = .37$. Participants in the control condition selected option C on 57% of trials, while participants in the frequency condition selected A on only 44% of trials.

We next examined whether the proportion of C choices within the frequency condition was significantly different than 0.5. A one-sample Bayesian t-test indicated no significant difference from 0.5, $BF_{10} = 0.54$, $d = .169$, although there was also no conclusive support for the null hypothesis. The 95% credible interval included 0.5 within its upper bound, $95\% HCl = [0.383, 0.503]$. Thus, although the difference between conditions was significant, there was not a

significant frequency effect observed in the frequency condition when the proportion of C choices (0.44) is compared to chance (0.50).

We also compared the proportion of C choices on CA trials between the frequency condition in Experiment 2 and the gains condition from Experiment 1, where a significant frequency effect was observed, with the proportion of C choices equaling 0.40. An independent sample Bayesian t-test indicated support for the null hypothesis, $BF_{10} = 0.22$, $d = .126$, which suggests that there is no difference in C choices between the two conditions across experiments. Overall, these results suggest a moderate frequency effect in Experiment 2 that was slightly attenuated compared to the gains condition in Experiment 1.

Distribution of C Choices on CA Test Trials




Figure 9 Distributions of the average proportion of C choices made on CA test trials, within each condition in Experiment 2. Left panel: Control condition, Right panel: Frequency condition. Values to the left within each panel indicate more A choices, and values to the right indicate more C choices.

Figure 9 shows the distribution of C choices on CA test trials for each group. On the left panel, for Control group participants, the modal value is on the extreme right edge of the distribution, indicating a large group of participants who selected option C on nearly every trial. On the left panel, participants in the frequency condition showed a slight bias overall for option A over C, with the modal group clustered toward the left side of the plot, indicating almost zero C choices for this group of participants. However, there is also a cluster of participants at the far right of the plot for the frequency condition who selected option C on nearly every trial. Overall, the pattern of C choices on CA trials reveals a modest frequency effect.

3.2.1 Model Fits

Table 4 shows the average BIC values, Bayes Factors, BIC-weights, and VBMS statistics for each model. The Delta-Uncertainty, Decay-Win and PVPE-Decay models provided the best fit in the Frequency condition, all within less than 1 BIC unit in average fit. The BIC weights suggests that about 38% of participants' data sets are best fit by the Decay-Win model, with the Delta model fitting 25% of participants' data, and the Delta-Uncertainty fitting 11%. The other

models all had BIC weights less than 0.10, indicating that less than 10% of data sets were best fit by these models. This pattern is replicated in the VBMS results and suggests that most participants used a strategy represented by the Decay-Win or the Delta model. In the control condition, the Decay-Win provided the best fit to the data, followed by the PVPE-Decay and the Delta-Uncertainty models. The Decay-Win and Delta model again had the highest BIC-weights, followed by the Delta-Uncertainty model. Intriguingly, despite providing a poor fit on average, the Decay-Loss model had a BIC weight of .11, which suggests that around 11% of participants were using a relative loss minimization strategy in the control condition.

Table 4. Average best-fitting BIC values and BIC-weights in Experiment 2

	Mean BIC	BF _{BestModel,M}	BIC-weight	VB α	VB r_k	VB φ_k
Frequency						
Delta	294.18	79.84	.25	37.24	0.29	0.04
Decay	353.78	>10K	.06	8.37	0.06	<.001
Decay-Win	285.21	1.08	.38	53.81	0.41	0.96
Decay-Loss	352.35	>10K	.03	2.50	0.02	<.001
PVL-Delta	298.53	716.95	.03	1.01	0.01	<.001
PVL-Decay	358.91	>10K	.01	1.02	0.01	<.001
PVPE-Decay	285.90	1.31	.07	6.92	0.05	<.001
Delta-Uncertainty	285.04	-	.16	19.14	0.15	<.001
Control						
Delta	297.45	229.29	.23	33.99	0.26	0.03
Decay	351.31	>10K	.04	7.09	0.06	<.001
Decay-Win	286.45	-	.33	51.14	0.39	0.97
Decay-Loss	347.97	>10K	.11	12.78	0.10	<.001
PVL-Delta	301.20	1,540.71	.04	1.09	0.01	<.001
PVL-Decay	351.68	>10K	.05	2.58	0.02	<.001
PVPE-Decay	289.73	7.17	.08	7.97	0.06	<.001
Delta-Uncertainty	293.64	36.41	.13	13.36	0.10	<.001

Note: Values in parentheses are standard deviations.

Table 5 shows the best-fitting parameter values for each model. For the best-fitting Decay-Win model, the average decay parameter value is relatively low (~0.20), which suggests that participants weighed outcomes from many recent trials when developing expectations about the outcomes provided by each option. As in Experiment 1, for the PVPE-Decay model the weight-to-relative losses parameter was less than 0.5, indicating an average bias toward relative gain outcomes. The average shape parameter values were also low, with many participants discounting relative reward magnitudes, and focusing only on whether rewards were relative wins or losses. One additional point to note is that the Decay and PVL-Decay models had very low best-fitting parameter values for the inverse temperature parameter in both conditions. This suggests that the model was often best fit by assuming near random responding in the task. As discussed above, as well as further below, the Decay model makes unrealistic predictions when only losses are provided.

Table 5. Average best-fitting parameter and BIC values in Experiment 2

Parameter:	a or A	c	γ or Unc_0	λ or w_L or w_{Unc}
Frequency				
Delta	.31 (.38)	.90 (1.44)		
Decay	.87 (.25)	.01 (.01)		
Decay-Win	.19 (.29)	.37 (.41)		
Decay-Loss	.77 (.33)	.15 (.24)		
PVL-Delta	.33 (.37)	2.09 (2.17)	.47 (.47)	
PVL-Decay	.89 (.24)	.07 (.17)	.28 (.45)	
PVPE-Decay	.17 (.27)	.37 (.37)	.28 (.41)	.25 (.31)
Delta-Uncertainty	.22 (.32)	.28 (.34)	2.15 (1.94)	143.81 (274.38)
Control				
Delta	.23 (.31)	.88 (1.35)		
Decay	.79 (.31)	.01 (.02)		
Decay-Win	.20 (.31)	.41 (.53)		
Decay-Loss	.65 (.39)	.21 (.28)		
PVL-Delta	.25 (.32)	1.69 (2.10)	.50 (.47)	
PVL-Decay	.80 (.30)	.07 (.17)	.28 (.45)	
PVPE-Decay	.21 (.28)	.46 (.54)	.30 (.40)	.36 (.40)
Delta-Uncertainty	.24 (.32)	.92 (.85)	1.35 (1.69)	180.52 (354.64)

Note: Values in parentheses are standard deviations.

3.2.2 Post-Hoc Simulations

We next conducted post-hoc simulations, identical to the procedures outlined in Experiment 1 above. Figure 10 shows the predicted and observed choices made on the critical CA test trials. Participants' data are plotted in the top left of the plot, then the average simulated proportion of C choices on each CA test trial is shown for each model. The Decay-Win and PVPE-Decay models most clearly predict the pattern of the data where participants in the control condition selected option C much more than participants in the frequency condition, who preferred option A. However, the Decay-Win model predicts a larger frequency effect than was observed. The Delta and Delta-Uncertainty models predict smaller frequency effects than were observed, and the PVL-Delta model reproduced a reversed frequency effect that was opposite to the pattern of the data. The Decay, Decay-Loss, and the PVL-Decay models all predict chance behavior because the models learn poorly under all-losses conditions.




Figure 10 Post-hoc simulated C choices on each CA test trial in Experiment 2, along with participants' data in the top left.

The top section of Table 6 shows the RMSD between the simulated and observed C choice on CA trials for each model. In the frequency condition the PVPE-Decay model had the lowest RMSD. Interestingly, the Decay-Win model had the second-highest RMSD because it predicted much fewer C choices than observed. The PVPE-Decay model was likely better able to reproduce the overall pattern of the data than the Decay-Win model because it is more flexible, and can account for participants who did not exclusively attend to relative gains, as assumed by the Decay-Win model. In the control condition, the Delta-Uncertainty model had the lowest RMSD, followed by the PVL-Delta, Delta, and PVPE-Decay models.

Table 6: RMSD Values From Post-hoc Simulations, Exp. 2

	Frequency	Control
CA Trials		
Delta	.094	.047
Decay	.071	.069
Decay-Win	.114	.048
Decay-Loss	.077	.063
PVL-Delta	.263	.042
PVL-Decay	.072	.070
PVPE-Decay	.048	.061
Delta-Uncertainty	.056	.041
All Trials		
Delta	.055	.061
Decay	.137	.124

Decay-Win	.065	.072
Decay-Loss	.132	.111
PVL-Delta	.159	.070
PVL-Decay	.137	.125
PVPE-Decay	.069	.077
Delta-Uncertainty	.046	.048

We next examined post-hoc simulations across all trials. These are plotted in Figures S6 and S7, and Table 4 lists the RMSD values at the bottom. In both conditions the Delta-Uncertainty model had the lowest overall RMSD, followed by the Delta, Decay-Win and PVPE-Decay models. The PVL-Delta model had a comparatively low RMSD in the control, but not the frequency condition, and the Decay, PVL-Decay and Decay-Loss models had the highest RMSD across all trials. Table S2 indicates that the Delta-Uncertainty model had a much lower RMSD than the other models for both training trial types, but it did not have the lowest RMSD for any of the remaining test trials, with the Decay-Win and PVPE-Decay models usually providing the best post-hoc recovery of the test-trial data.

3.3 Discussion

Experiment 2 differed from Experiment 1 in that a cover story was introduced which provided a context for the loss-minimization scenario (buying dog food), and the reward structure was shifted to where all outcomes were losses. The average losses for each option were between -12.50 and -17.50, compared to -.25 and -.75 for Experiment 1. These changes resulted in a significant frequency effect in Experiment 2. Model-based analyses suggest that the PVPE-Decay model can best account for the frequency effect by assuming that participants interpreted smaller losses as relative wins, or positive outcomes, and that more frequent options were associated with more of these relative wins.

We believe that the context manipulation added in Experiment 2, where participants viewed each trial as trying to minimize a shopping expenditure was the key aspect of the manipulation that led to the frequency effect in Experiment 2, because it made it more likely that participants viewed smaller losses as relative gains. However, we also changed the reward structure in Experiment 2 to where all outcomes were losses. To examine whether the change in reward structure alone can produce a frequency effect, we ran a modified version of Experiment 2, where the reward structure was the same, but the cover story about shopping for dog food for a local shelter was removed. Instead, participants were given the same instructions as Experiment 1 where they were simply asked to try to minimize the number of points lost. We predicted that the lack of a contextual cover story would lead to attenuated frequency effects because participants would be less able to interpret smaller losses as relative gains.

4. Experiment 3

4.1 Method

4.1.1 Participants

Given time constraints for completing Experiment 3, we planned to run approximately 100 participants in each condition. The computer program randomly assigned each participant to a condition. Our final sample size was 203 participants, 97 participants in the control condition and 106 participants in the frequency condition. Within the control condition there were 70 females, 26 males, and one other; within the frequency condition there were 78 females, 27

males, and one other. Participants completed the experiment online for partial fulfilment of a course credit.

4.1.2 Materials and Procedure

The materials and procedures were identical to Experiment 2, except that the cover story about buying dog food for a local shelter was removed, and participants were simply asked to pick which option they thought would lead to the smallest losses. This task framing was identical to the losses condition in Experiment 1.

4.1.3 Data Analysis

We used the same data analysis methods as reported in Experiment 1.

4.2 Results

Figure 11a shows the proportion of optimal choices during the training phase. A Bayesian mixed effects logistic regression model predicting optimal choices during training from condition, with random intercepts for participants, indicated no effect of condition, $b = -0.26$, $SE = 0.25$, $95\% HCl = [-0.77, 0.23]$, $OR = 0.77$. We also ran a similar model predicting optimal choices from the interaction between condition and trial type, with random intercepts for participants. There was no interaction effect, $b = -0.02$, $SE = 0.07$, $95\% HCl = [-0.18, 0.12]$, $OR = 0.98$. Training accuracy was also much greater than chance overall at 78% for AB trials and 80% for CD trials.




Figure 11 a.) Proportion of optimal choices during training for each trial type in Experiment 3. **b.)** Proportion of C choices on CA transfer trials. Error bars represent standard errors of the mean.

We next examined the proportion of C choices on the critical CA transfer trials, which are shown in Figure 11b. A Bayesian mixed effects logistic regression model with optimal choices predicted from condition indicated no effect of condition, $b = -.54$, $SE = 0.56$, $95\% HCl = [-1.63, .56]$, $OR = 0.58$. A Bayesian t-test with average C choices as the dependent variable indicated a null effect of condition, $BF_{10} = .25$, $d = .15$. The Bayes Factor in support of the null hypothesis (BF_{01}) was 3.97, indicating moderate support. Thus, there was no effect of reward frequency.

We also conducted a one-sample t-test within the frequency condition with 0.5 as the test statistic. This suggested moderate support for the null hypothesis that the proportion of C choices did not differ from 0.5, $BF_{10} = .20$, $d = .11$. Figure 12 shows the distribution of C choices on CA test trials for each group. Slightly more participants preferred option A in the frequency than the control condition, but the difference is very small.




Figure 12 Distributions of the average proportion of C choices made on CA test trials, within each condition of Experiment 3. Left panel: Control condition, Right panel: Frequency condition. Values to the left within each panel indicate more A choices, and values to the right indicate more C choices.

We also compared the data to Experiment 2, which differed only in the framing of the task. A 2 (experiment) X 2 (frequency condition) Bayesian ANOVA with the proportion of C choices on CA test trials as the dependent variable indicated a null effect of the experiment X frequency condition interaction, $BF_M = .26$, $\eta^2_p = .002$, which suggests that the pattern of the difference between the control and frequency conditions is consistent across Experiments 2 and 3. The best supported model included only the effect of frequency condition, $BF_M = .2.94$, $\eta^2_p = .015$. A t-test comparing the Frequency condition data across Experiments 2 and 3 indicates only anecdotal support for the hypothesis that the two groups differed in their proportion of C choices, $BF_{10} = 1.11$, $d = .28$, despite an 11% difference (44% C choices in Experiment 2 versus 55% in Experiment 3). However, a similar comparison between the frequency condition from Experiment 3 and the gains condition from Experiment 1, indicated significantly fewer C choices in Experiment 1 ($M=.40$) than in Experiment 3 ($M=.55$), $BF_{10} = 9.04$, $d = .40$. As reported above, there was no difference in C choices between the frequency condition from Experiment 2 and the gains condition from Experiment 1. While the lack of an interaction between Experiments 2 and 3 suggests that the pattern of the effect of frequency did not differ between the two experiments, there was no support for a frequency effect in Experiment 3 when tested against chance (0.5), and when compared against the frequency effect found in the gains condition of Experiment 1.

4.2.1 Model Fits

Table 7 shows the average BIC values, Bayes Factors, BIC-weights, and VBMS results for each model. The Delta-Uncertainty, and Decay-Win models provided the best fit in the Frequency condition, within less than 1 BIC unit in average fit. The BIC weights suggests that about 28% of participants' data sets are best fit by the Delta-Uncertainty model, with the Delta model fitting 24% of participants' data, the Decay-Win model fitting 23% of participants' data,

and the PVPE-Decay 16%. The other models all had BIC weights less than 0.10, indicating that less than 10% of data sets were best fit by these models. In the control condition, the Delta-Uncertainty model provided the best fit to the data, followed by the Delta and Decay-Win model. The Delta, Delta-Uncertainty, and Decay-Win models also received the most support according to the VBMS statistics, such as the exceedance probability.

Table 7. Average best-fitting BIC values and BIC-weights in Experiment 3

	Mean BIC	BF _{BestModel,M}	BIC-weight	VB α	VB r_k	VB φ_k
Frequency						
Delta	209.03	221.41	.24	29.64	0.26	0.35
Decay	354.79	>10K	.03	3.98	0.04	< .001
Decay-Win	199.02	1.48	.23	26.57	0.23	0.16
Decay-Loss	324.56	>10K	.03	3.44	0.03	< .001
PVL-Delta	214.51	3,428.92	.01	1.04	0.01	< .001
PVL-Decay	360.18	>10K	.00	1.04	0.01	< .001
PVPE-Decay	201.45	5.00	.16	17.08	0.15	< .001
Delta-Uncertainty	198.23	-	.28	31.20	0.27	0.49
Control						
Delta	206.77	3.24	.31	34.86	0.33	0.55
Decay	356.18	>10K	.01	1.02	0.01	< .001
Decay-Win	211.29	31.03	.31	33.74	0.32	0.44
Decay-Loss	315.28	>10K	.07	6.30	0.06	< .001
PVL-Delta	211.54	35.16	.02	1.05	0.01	< .001
PVL-Decay	361.59	>10K	.02	1.00	0.01	< .001
PVPE-Decay	216.74	473.43	.16	4.62	0.04	< .001
Delta-Uncertainty	204.42	-	.28	22.41	0.21	0.02

Note: Values in parentheses are standard deviations.

Table 8 shows the best-fitting parameter values for each model. For the best-fitting Decay-Win model, the average decay parameter value is relatively low (~0.15), which suggests that participants weighed outcomes from many recent trials when developing expectations about the outcomes provided by each option. As in the first two experiments, for the PVPE-Decay model the weight-to-relative losses parameter was less than 0.5, indicating an average bias toward relative gain outcomes. The average shape parameter values were also low (~0.20), with many participants discounting relative reward magnitudes, and focusing only on whether rewards were relative wins or losses. As in the first experiments, the Decay and PVL-Decay models had very low best-fitting parameter values for the inverse temperature parameter suggesting that these models can, at best, predict random performance.

Table 8. Average best-fitting parameter and BIC values in Experiment 3

Parameter:	a or A	c	γ or Unc_0	λ or w_L or w_{Unc}
Frequency				
Delta	.37 (.39)	.99 (1.06)		
Decay	.95 (.19)	.00 (.01)		
Decay-Win	.15 (.23)	.53 (.42)		
Decay-Loss	.63 (.46))	.30 (.64)		
PVL-Delta	.40 (.39)	1.90 (1.90)	.59 (.48)	
PVL-Decay	.91 (.26)	.07 (.50)	.10 (.29)	
PVPE-Decay	.17 (.28)	.66 (.49)	.20 (.36)	.26 (.41)
Delta-Uncertainty	.24 (.33)	.56 (.59)	1.56 (1.79)	174.39 (306.87)
Control				
Delta	.43 (.38)	.75 (.89)		
Decay	.95 (.17)	.00 (.00)		
Decay-Win	.16 (.24)	.50 (.56)		
Decay-Loss	.50 (.47)	.31 (.44)		
PVL-Delta	.50 (.39)	2.13 (2.03)	.52 (.49)	
PVL-Decay	.97 (.12)	.00 (.03)	.05 (.22)	
PVPE-Decay	.16 (.24)	.66 (.65)	.23 (.39)	.25 (.39)
Delta-Uncertainty	.23 (.30)	.50 (.64)	2.08 (1.97)	203.20 (309.41)

Note: Values in parentheses are standard deviations.

4.2.2 Post-Hoc Simulations

Figure 13 shows the predicted and observed choices made on the critical CA test trials. The PVPE-Decay, Delta and Delta-Uncertainty models most clearly predict the pattern of the data where participants in the control condition selected option C slightly more than participants in the frequency condition. However, the Decay-Win model predicts a larger frequency effect than was observed. The PVPE-Decay model predicted a slightly larger frequency effect than was observed. As in Experiment 2, the PVL-Delta model again reproduced a reversed frequency effect that was opposite to the pattern of the data. The Decay, Decay-Loss, and the PVL-Decay models all predict chance behavior because the models learn poorly under all-losses conditions.




Figure 13 Post-hoc simulated C choices on each CA test trial in Experiment 3, along with participants' data in the top left.

The top section of Table 9 shows the RMSD between the simulated and observed C choice on CA trials for each model. In the frequency condition the Delta-Uncertainty model had the lowest RMSD. As in Experiment 2, the Decay-Win model had the second-highest RMSD because it predicted much fewer C choices than observed. In the control condition, the PVL-Delta model had the lowest RMSD, but all models except the Decay and PVL-Decay models had similarly low RMSD values. Overall, the pattern of the Delta-Uncertainty model appears to best account for the data in the frequency condition.

We next examined post-hoc simulations across all trials. The RMSD values are shown at the bottom of Table 9, and the simulation results are plotted in the Supplemental Materials. The basic Delta had the lowest RMSD in both conditions for both training trial types (AB and CD). The Decay-Win, PVPE-Decay, and the basic Delta model had the lowest RMSD values for the remaining test trial types (CB, AD, and BD).

Table 9: RMSD Values From Post-hoc Simulations, Exp. 3

CA Trials	Frequency	Control
Delta	.044	.034
Decay	.049	.108
Decay-Win	.225	.031
Decay-Loss	.083	.052
PVL-Delta	.164	.030
PVL-Decay	.047	.108
PVPE-Decay	.079	.040
Delta-Uncertainty	.030	.035
All Trials		

Delta	.065	.059
Decay	.250	.290
Decay-Win	.103	.119
Decay-Loss	.243	.246
PVL-Delta	.168	.136
PVL-Decay	.249	.291
PVPE-Decay	.101	.119
Delta-Uncertainty	.113	.153

5. General Discussion

The results of our experiments do not support the prediction of the Decay model, of a reversed frequency effect under losses, and a standard frequency effect under gains. We did not find a reversed frequency effect in Experiments 1 or 3, and in Experiment 2 we found a moderate frequency effect under losses, in the same direction as the gains condition from Experiment 1. Although the Decay model correctly predicted a frequency effect in the gains condition from Experiment 1, it provided very poor fits and simulations for the losses conditions across all three experiments. In a recent paper from our lab (Don et al., 2019), we found that the Decay model provided the best account of the observed frequency effect in a binary outcome task involving gains, but it appears that the Decay model's predictions are qualitatively incorrect under loss-minimization scenarios. Under losses, the Decay model predicts that as an option is chosen more often, it will become worse in value because the chosen option will become more negative, while the non-chosen options will decay toward zero. This prediction runs counter to many experiments that indicate that people have a strong tendency to persevere, or repeatedly pick the same option (Worthy, Pang & Byrne, 2013; Senftleben, Schoeman, Rudolf, & Scherbaum, 2021; Gershman, 2020)

Our theoretical interpretation of the Decay model is that it assumes that when participants make a decision on each trial, they think of all the past outcomes associated with each option, in a recency-weighted manner. Under gains, the model assumes participants recall the past gains for each option, biasing choices towards more frequently presented options, but under losses the model assumes that the losses for each option will be recalled, and there will be more losses associated with the more frequently presented alternative. These additionally recalled losses will bias choices away from the more frequently encountered items, creating the reversed frequency effect. In contrast, the Decay-Win and PVPE-Decay models, which provided the best fits and post-hoc recovery of the data, particularly for the critical CA test trials. These models assume that reward outcomes are processed in a relative rather than an absolute manner (Rakow et al., 2020). In tasks involving roughly equal numbers of gains and losses, such as the Iowa Gambling task (Bechara et al., 1994), absolute and relative losses and gains would tend to be the same.

Our data suggest that losses are not necessarily processed as aversive, particularly when additional context is provided which helps participants view smaller losses as positive outcomes, as in Experiment 2. The context-dependent PVPE-Decay and Decay-Win models best reproduced the frequency effect observed under losses, which was similar to the gains condition in Experiment 1. In Experiments 1 and 3, where the only context given to participants was to minimize the number of points lost, our data suggest that participants used strategies represented by the either delta models such as the Delta and Delta-Uncertainty models, or contextual decay models such as the Decay-Win and PVPE-Decay models. The delta and contextual decay

models may represent separate value-based versus frequency-based decision-making strategies, or modes that are modulated by the amount uncertainty associated with each strategy (Hu et al., 2025). One speculative interpretation of our findings is that in the abstract loss-minimization scenario it is more difficult for participants to perceive which losses are better than others. This may create uncertainty for the frequency-based strategy because it is unclear what outcomes are relative wins, and this could shift some participants towards a value-based strategy represented by delta models. Alternatively, when given a decision-making context that helps to frame smaller losses as wins or “bargains”, such as the shopping manipulation in Experiment 2, this may enhance confidence in the frequency-based system and lead to stronger frequency effects.

The Decay-Win model, which predicted the overall pattern of the data, and fit the majority of participants the best, makes three important assumptions about how people make decisions. First, it assumes similar behavior under gains and losses because it tracks the average reward provided across all options, and updates its expected values relative to that. This is similar to recent models that have demonstrated the importance of taking the reward context into account (Molinaro & Collins, 2023; Hayes & Wedell, 2023; Brochard & Daunizeau, 2024). Second, it assumes that people only attend to the relative valence of the reward (i.e. whether it was better than expected), and disregard the magnitude. Third, the model assumes that people only focus on the relative gains provided by each option, and they disregard the relative losses.

These last two assumptions are much stronger than the first one, and they should be further tested further in future work. However, these assumptions are consistent with the results of studies using tasks like the Iowa and Soochow gambling tasks (Bechara, Damasio, Damasio, & Anderson, 1994; Chiu et al., 2008). A common finding in these studies is that people initially display a strong preference for the options that give consistent, small gains, and they disregard the rare losses provided by the frequently-rewarding options (Byrne & Worthy, 2016, Don et al., 2022; Aite et al., 2012). In a recent paper from our lab, we found that a model that was a combination of the Decay-Win and Decay-Loss models, the Prediction-Error Decay model, provided the best fit to both younger and older adult data (Don et al., 2022). This model included the first two assumptions listed above, but it incremented +1 for a ‘win’ and -1 for a ‘loss.’ Recent work has challenged the idea that ‘losses loom larger than gains’ (Yechiam, 2019). For example, Hao and colleagues found that people learn better from wins than from losses (Hao, Cabrera-Haro, Lin, Reuter-Lorenz, & Lewis, 2023). Our results suggest that relative gains, or positive outcomes, have greater weight in determining future choices than relative losses.

We also found that the Delta-Uncertainty model provided a good fit to the data, and good post-hoc recovery of the training data. However, it did not perform well in reproducing the critical frequency effect observed in Experiment 2. This model was formalized as an alternative way in which frequency effects might manifest, compared to the memory-based mechanisms assumed by the Decay model. Frequency effects could be caused by aversion to uncertainty. One of the key differences between the Delta-Uncertainty and the Decay-Win models is that the Delta-Uncertainty model assumes a preference for more frequently encountered options simply because they have been *selected* more often, whereas the Decay-Win model values more frequent options because they have been *rewarded* more often. Although, this paper was aimed at testing the predictions of the Decay model under losses, future studies are likely needed with tasks that are better suited to address whether frequency effects are due to memory for past rewards or for familiarity due to more frequent selection. Future studies could also examine whether *mere*

exposure to some options, without any associated outcomes, also produces frequency effects (Zajonc, 1974).

It is important to note that we used a model comparison approach where eight different models that made non-overlapping predictions about behavior in the task were compared. Four of these models were considered basic models, each containing only two free-parameters, thus they were not overly flexible (Roberts & Pashler, 2000). Although the Decay-Win model provided the best fit to a large proportion of participants' data, a smaller group of participants' data were best fit by the Delta model, which assumes optimal responding on the critical CA transfer trials. An extended version of the Decay-Win model, the PVPE-Decay model, is flexible enough to also account for optimal behavior on the critical CA trials, similar to the Delta model. Depending on a given researcher's goals the PVPE-Decay model may be more useful than the Decay-Win model. If the goal is to describe participants' behavior, then the PVPE-Decay model is likely more useful than the simpler Decay-Win model, because its parameters can provide information about the degree of discounting of the magnitude of rewards and the attention to relative gains versus losses. Alternatively, if the researcher is attempting to test competing theories, simpler, more falsifiable models like the Decay-Win, Delta, and other models will be more useful, as the extended models we used are much more flexible.

It is also worth noting that we obtained a frequency effect under losses only in Experiment 2, where all outcomes were losses, and using a scenario that participants were probably familiar with (purchasing items and trying to minimize the cost). The frequency effect in Experiment 2 was also weaker than in the gains condition for Experiment 1, as the proportion of C choices on critical transfer trials (0.44) was not significantly different from 0.5; however, there was a significant difference compared to the control condition (0.57). We believe the comparison with the control condition is most appropriate, since option C was objectively more valuable than option A; however, the frequency effect observed in Experiment 2 did not pass the stronger test of departing significantly from chance behavior. We also reported comparisons across Experiments that suggest that the pattern of the data was consistent across Experiments 2 and 3, and a comparison of C choices in the frequency conditions from Experiments 2 and 3 did not reach significance. However, there was no difference in C choices between the gains condition from Experiment 1 and the frequency condition for Experiment 2, but there was a significant difference between Experiments 1 and 3, with the gains condition from Experiment 1 yielding the strongest frequency effect, and the frequency condition from Experiment 3 yielding the weakest effect.

The familiarity with the situation of purchasing dog food in Experiment 2 may have enhanced participants' ability to view smaller losses as relative gains or 'wins.' Losses may be viewed as more uniformly negative outcomes in situations where the context of the reward decision-making scenario is less clear. One limitation of the current set of experiments is that we did not run Experiment 1's losses condition with the dog food purchasing framing. Based on our interpretation of the data, we would predict a similar frequency effect as that in Experiment 2, if the losses task was viewed as a shopping expenditure minimization task. Future work can more broadly test how framing loss-minimization scenarios in more interpretable ways enhances frequency effects and other reward processing mechanisms. For example, does familiarly framing a loss-minimization scenario improve memory for past outcomes?

It is also unclear exactly what information about past outcomes is stored in memory on each trial. Do participants simply remember outcomes as positive or negative experiences, or do they remember numerical information about the rewards received and compare it to a

representation of the average reward at the time they are making choices? The superior fits of the Decay-Win over the Decay-Loss model also suggests that people focus more on positive outcomes than on negative outcomes. An open question is whether this is due to enhanced attention to positive outcomes, or better memory for positive outcomes. The current experiments cannot address these issues, but they could inspire clever experiments designed to address them in future work.

6. Conclusion

We tested the Decay model's prediction of a reversed frequency effect under losses, compared to one previously reported for gains (Don et al., 2019; Don & Worthy, 2022; Hu et al., 2025). The Decay model's predictions were supported only under a gains reward structure; the model performed very poorly under losses. We found a frequency effect under gains in Experiment 1, and under losses in Experiment 2. A Decay-Win model, that tracked the number of relative gains provided the best qualitative account of the observed data, and an extended model, the PVPE-Decay model accounted for alternative strategies used in the task. A basic Delta model, as well as the Delta-Uncertainty model provided good fits and post-hoc recovery for most trials; however, these models could not recover the frequency effects as well as the contextual decay models. Theoretically, this suggests that frequency effects can occur under gain-maximization and loss-minimization scenarios, that framing loss-minimization scenarios in familiar ways can cause people to attend more to the *relative* valence of each outcome, and attend less to reward magnitude, and that behavior is driven more by relative gains than by losses.

References

Ahn, W. Y., Busemeyer, J. R., Wagenmakers, E. J., & Stout, J. C. (2008). Comparison of decision learning models using the generalization criterion method. *Cognitive science*, 32(8), 1376-1402.

Aïte, A., Cassotti, M., Rossi, S., Poirel, N., Lubin, A., Houdé, O., & Moutier, S. (2012). Is human decision making under ambiguity guided by loss frequency regardless of the costs? A developmental study using the Soochow Gambling Task. *Journal of Experimental Child Psychology*, 113(2), 286-294.

Bechara, A., Damasio, A. R., Damasio, H., & Anderson, S. W. (1994). Insensitivity to future consequences following damage to human prefrontal cortex. *Cognition*, 50(1-3), 7-15.

Brochard, J., & Daunizeau, J. (2024). Efficient value synthesis in the orbitofrontal cortex explains how loss aversion adapts to the ranges of gain and loss prospects. *eLife*, 13, e80979.

Bürkner, P. C. (2017). brms: An R package for Bayesian multilevel models using Stan. *Journal of statistical software*, 80, 1-28.

Busemeyer, J. R., & Stout, J. C. (2002). A contribution of cognitive decision models to clinical assessment: decomposing performance on the Bechara gambling task. *Psychological assessment*, 14(3), 253.

Byrne, K. A., Peters, C., Willis, H. C., Phan, D., Cornwall, A., & Worthy, D. A. (2020). Acute stress enhances tolerance of uncertainty during decision-making. *Cognition*, 205, 104448.

Byrne, K. A., & Worthy, D. A. (2016). Toward a mechanistic account of gender differences in reward-based decision-making. *Journal of Neuroscience, Psychology, and Economics*, 9(3-4), 157.

Chiu, Y. C., Lin, C. H., Huang, J. T., Lin, S., Lee, P. L., & Hsieh, J. C. (2008). Immediate gain is long-term loss: Are there foresighted decision makers in the Iowa Gambling Task?. *Behavioral and Brain Functions*, 4(1), 1-10.

Don, H. J., Davis, T., Ray, K. L., McMahon, M. C., Cornwall, A. C., Schnyer, D. M., & Worthy, D. A. (2022). Neural regions associated with gain-loss frequency and average reward in older and younger adults. *Neurobiology of Aging*, 109, 247-258.

Don, H. J., Otto, A. R., Cornwall, A. C., Davis, T., & Worthy, D. A. (2019). Learning reward frequency over reward probability: A tale of two learning rules. *Cognition*, 193, 104042.

Don, H. J., & Worthy, D. A. (2022). Frequency effects in action versus value learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 48(9), 1311.

Erev, I., & Roth, A. E. (1998). Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American Economic Review*, 848-881.

Estes, W. K. (1976). The cognitive side of probability learning. *Psychological Review*, 83(1), 37.

Fagley, N. S., & Miller, P. M. (1997). Framing effects and arenas of choice: Your money or your life?. *Organizational behavior and human decision processes*, 71(3), 355-373.

Gallagher, K. M., & Updegraff, J. A. (2012). Health message framing effects on attitudes, intentions, and behavior: a meta-analytic review. *Annals of behavioral medicine*, 43(1), 101-116.

Gershman, S. J. (2020). Origin of perseveration in the trade-off between reward and complexity. *Cognition*, 204, 104394.

Gonzalez, C., Dana, J., Koshino, H., & Just, M. (2005). The framing effect and risky decisions: Examining cognitive functions with fMRI. *Journal of economic psychology*, 26(1), 1-20.

Hayes, W. M., & Wedell, D. H. (2023). Testing models of context-dependent outcome encoding in reinforcement learning. *Cognition*, 230, 105280.

Hu, M., Don, H. J., & Worthy, D. A. (2025). Distributional dual-process model predicts strategic shifts in decision-making under uncertainty. *Communications Psychology*, 3(1), 61.

Hu, M., & Worthy, D.A. (2025). The Rational Irrational: Better Learners Show Stronger Reward Frequency Biases. *Journal of Experimental Psychology: Learning, Memory, and Cognition*.

Jeffreys, H. (1961). *Theory of Probability*. Oxford University Press, UK.

Kahneman, D., & Tversky, A. (1979). Prospect Theory: An analysis of decision making under risk. *Econometrica*, 47, 263-291.

Kühberger, A. (1995). The framing of decisions: A new look at old problems. *Organizational Behavior and Human Decision Processes*, 62(2), 230-240.

Kang, H. (2021). Sample size determination and power analysis using the G* Power software. *Journal of educational evaluation for health professions*, 18.

Kwak, Y., & Huettel, S. (2018). The order of information processing alters economic gain-loss framing effects. *Acta Psychologica*, 182, 46-54.

March, J. G. (1996). Learning to be risk averse. *Psychological review*, 103(2), 309.

Molinaro, G., & Collins, A. G. (2023). Intrinsic rewards explain context-sensitive valuation in reinforcement learning. *PLoS Biology*, 21(7), e3002201.

Nalborczyk, L., Batailler, C., Loevenbruck, H., Vilain, A., & Bürkner, P. C. (2019). An introduction to Bayesian multilevel models using brms: A case study of gender effects on vowel variability in standard Indonesian. *Journal of Speech, Language, and Hearing Research*, 62(5), 1225-1242.

Pang, B., Blanco, N. J., Maddox, W. T., & Worthy, D. A. (2017). To not settle for small losses: evidence for an ecological aspiration level of zero in dynamic decision-making. *Psychonomic bulletin & review*, 24, 536-546.

Rakow, T., Cheung, N. Y., & Restelli, C. (2020). Losing my loss aversion: The effects of current and past environment on the relative sensitivity to losses and gains. *Psychonomic bulletin & review*, 27, 1333-1340.

Roberts, S., & Pashler, H. (2000). How persuasive is a good fit? A comment on theory testing. *Psychological review*, 107(2), 358.

Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT press.

Schneider, S. L., & Lopes, L. L. (1986). Reflection in preferences under risk: Who and when may suggest why. *Journal of Experimental Psychology: Human perception and performance*, 12(4), 535.

Schwarz, G. (1978). Estimating the dimensions of a model. *Annals of Statistics*, 6, 461-464.

Senftleben, U., Schoemann, M., Rudolf, M., & Scherbaum, S. (2021). To stay or not to stay: The stability of choice perseveration in value-based decision making. *Quarterly Journal of Experimental Psychology*, 74(1), 199-217.

Steingroever, H., Wetzels, R., & Wagenmakers, E. J. (2013). Validating the PVL-Delta model for the Iowa gambling task. *Frontiers in psychology*, 4, 898.

Steingroever, H., Wetzels, R., & Wagenmakers, E. J. (2014). Absolute performance of reinforcement-learning models for the Iowa Gambling Task. *Decision*, 1(3), 161.

Stephan, K. E., Penny, W. D., Daunizeau, J., Moran, R. J., & Friston, K. J. (2009). Bayesian model selection for group studies. *Neuroimage*, 46(4), 1004-1017.

Wagenmakers, E. J. (2007). A practical solution to the pervasive problems of p values. *Psychonomic Bulletin & Review*, 14(5), 779-804.

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

Wagenmakers, E. J., Love, J., Marsman, M., Jamil, T., Ly, A., Verhagen, J., ... & Morey, R. D. (2018). Bayesian inference for psychology. Part II: Example applications with JASP. *Psychonomic bulletin & review*, 25, 58-76.

Wetzels, R., Matzke, D., Lee, M. D., Rouder, J. N., Iverson, G. J., & Wagenmakers, E. J. (2011). Statistical evidence in experimental psychology: An empirical comparison using 855 t tests. *Perspectives on Psychological Science*, 6(3), 291-298.

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 & review*, 20, 364-371.

Worthy, D. A., Pang, B., & Byrne, K. A. (2013). Decomposing the roles of perseveration and expected value representation in models of the Iowa gambling task. *Frontiers in psychology*, 4, 640.

Yechiam, E. (2019). Acceptable losses: The debatable origins of loss aversion. *Psychological research*, 83(7), 1327-1339.

Yechiam, E., & Ert, E. (2007). Evaluating the reliance on past choices in adaptive learning models. *Journal of Mathematical Psychology*, 51(2), 75-84.

Yechiam, E., & Hochman, G. (2013). Losses as modulators of attention: review and analysis of the unique effects of losses over gains. *Psychological bulletin*, 139(2), 497.

Zeif, D., & Yechiam, E. (2022). Loss aversion (simply) does not materialize for smaller losses. *Judgment and Decision Making*, 17(5), 1015-1042.