Attention in Choices with Aversive or Ambivalent Attributes

Brenden Eum* Stephanie M. Smith†

November 17, 2026

Abstract

Attention is commonly thought to attract choice toward whatever is looked at, but nearly all evidence for this assertion comes from choices with appetitive outcomes. In two preregistered mouse-tracking experiments, participants (N = 1,506) chose between options with appetitive, aversive, and ambivalent attributes, with the range of values manipulated between participants. We estimated the attentional effects of looking at an attribute value (for an attribute type and range condition) and found that attention can either attract or repel choice. For aversive attributes, attention to the lowest (highest) values in a range was associated with a lower (higher) probability of choice. Critically, the range of the attribute mattered. Attention to the same value can predict opposite effects on choice: repelling choice in the Narrow range condition, but attracting choice in the Wide range condition. For appetitive and ambivalent attributes, attention predicted greater increases in choice probability as value rose. These patterns can be explained by range-dependent reference points. Attention attracts choice for attribute values above this reference point, and repels below it. To illustrate, we extend a Bayesian sequential sampling model with a prior calibrated to the range, but not the attribute type. This suggests that attention acts on beliefs about value, not value itself, and so attentional effects depend on expectations. This has implications for attentional nudges in environments where choice architecture is crucial, such as e-commerce, public health, and computational psychiatry.

Keywords: neuroeconomics, attention, choice process, attributes, aversive

*Rotman School of Management, University of Toronto. Corresponding author: eum.brenden@gmail.com.
†Booth School of Business, University of Chicago.

Statement of Relevance

Many interventions work by drawing attention to an attribute (e.g., a calorie count on a menu, a side effect on a drug label, fuel economy on a window sticker). The usual finding is that attention favors whatever it lands on. We show this is not always true. Specifically, it does not hold for attributes people know to be bad. In two experiments, attention to an aversive attribute repelled choice away from an option when its value was about as bad as it could be, and attracted choice toward the option when it was about as good as it could be (even though the value was still negative). Where the switch occurred depended on the range of values participants had learned to expect. From this, we suggest that attention sharpens what people believe about an attribute value, and whether that attracts choices to or repels choices from that option depends on what they expected. Designers of labels, disclosures, and nudges therefore need to know their audience's expectations before directing its attention.

1 Introduction

> "We find beauty not in the thing itself but in the patterns of shadows, the light and the darkness, that one thing against another creates."
>
> Jun'ichirō Tanizaki, *In Praise of Shadows* (1933)

Every day, we make decisions that require us to integrate different sources of information in order to compare options. Choosing a laptop means trading processing power against weight, size, and price. Choosing a medication means trading efficacy against side effects. In these multi-attribute choices, visual attention has repeatedly been shown to bias choice toward the option whose attribute is currently being attended (Busemeyer et al., 2019; Cohen et al., 2017; Fisher, 2017, 2021a; Gluth et al., 2024; March & Gluth, 2025; Yang & Krajbich, 2023). The same regularity holds in simple choices between whole options, where the item that is looked at longer, and especially the item that is looked at last, is more likely to be chosen (Eum et al., 2023; Krajbich & Rangel, 2011; Krajbich et al., 2010; Shimojo et al., 2003; S. M. Smith & Krajbich, 2018). The relationship between attention and choice is not only correlational, but also causal (Fisher, 2021b; Gwinn et al., 2019; Pärnamets et al., 2015; S. M. Smith & Spiller, 2026; for a review, see Bhatnagar & Orquin, 2022).

Almost all of this evidence comes from stimuli with positive values.[1] Participants choose between snack foods they like to varying degrees, between free lottery tickets, or between consumer goods they would be happy to receive. A lingering gap in the literature is what happens when we attend to something we know can (or will) be aversive. Many consequential attributes in the world are inherently aversive (fees, emissions, adverse side effects, wait times,

[1]

calories), in that they can not take on a positive value. Many attributes are ambivalent in the sense that they can take on positive or negative values, depending on the option (you may prefer a laptop screen that is not too small but not too large). What role does attention to ambivalent and aversive attributes play in the choice process?

Canonical process models make a sharp prediction here. In the Attentional Drift Diffusion Model and its relatives, attention discounts the value of unattended alternatives (Busemeyer & Townsend, 1993; Krajbich & Rangel, 2011; Krajbich et al., 2010). If an option has negative value, attending to it should repel choice because discounting the unattended option makes the attended option look comparatively worse (Armel et al., 2008). Empirically, this is not what happens in choices with free response time. Eum et al. (2026) find that when all options are aversive, attention still biases choice toward the attended option, and reconcile this with a model in which values are range normalized before entering the integration process. Under that account, the sign of the evidence that attention amplifies is not the sign of the objective value but the sign of the reference-dependent value.

To investigate how attention to known aversive or ambivalent attributes affects the choice process, we ran two mouse-tracking experiments. Participants made binary choices between two options described by two attributes, laid out as a 2×2 grid of boxes (Figure 1A). Attribute values were hidden until the participant hovered over a box with the mouse, which gives us a continuous record of what was being examined and for how long (S. Smith & Dolifka, 2026). Critically, each attribute row was labeled in advance by the sign of the values it could take: P (appetitive, or always positive), N (aversive, or always negative), or B (ambivalent, or both positive and negative). This mirrors the structure of many real attributes, where decision makers know (or don't know, in the ambivalent case) the valence before they know the realized value. We also manipulated the range of values that attributes could take (Narrow versus Wide), which lets us ask whether the mapping from attribute value to attentional influence is anchored in absolute value or in the local range (Louie et al., 2013; Padoa-Schioppa, 2009; Rangel & Clithero, 2012; Webb et al., 2021).

To preview the results, attention does not predict choice in a fixed direction. Attention to some of the lowest possible values of aversive attributes was associated with a decrease in choice probability, but this relationship reversed for attention to the highest possible values of aversive attributes. Moreover, these attentional effects were dependent on range condition, such that attention to the same realized value predicted opposite effects: a decrease in choice in the Narrow range condition, but an increase in choice in the Wide range condition. For appetitive and ambivalent attributes, attention to low values did not predict an increase or decrease in choice probability, but attention to higher values was increasingly predictive of choice as the value rose. Between range conditions, attentional effects were less sensitive to value in the Wide condition. To explain the mechanism behind these behavioral signatures, we adapt a two-alternative Bayesian sequential sampling model (Frömer et al., 2025) to account for multi-attribute decision making. Our model departs from the original in two crucial ways: (1) it extends the model to a case where the value of an alternative is the sum of independent attributes, and (2) it calibrates prior beliefs using choices, response times, and attention metrics, rather than assuming the prior. This free prior is separately calibrated by range condition (but not by attribute type) and sets a reference point by which attributes

are evaluated. Consequently, attention attracts choice when the attended attribute value is better than the reference point and repels choice when it is worse. Under this account, attention acts on belief about value rather than on value itself, and it's direction depends on the decision-maker's expectations.

The paper proceeds as follows. Section 2 describes the methods. Section 3 presents the behavioral, attentional, and modeling results. Section 4 concludes and discusses implications.

2 Methods

2.1 Task

To investigate how attention affects choice when attributes can be aversive or ambivalent, we ran two mouse-tracking studies involving multi-attribute choices. Participants made binary choices between two options, each described by two attributes, presented as a 2×2 grid of boxes (Figure 1A). Options occupied the left and right columns; attributes occupied the top and bottom rows. Each attribute row was labeled to the left of the grid with its type, which was visible for the entire trial: P for an appetitive attribute, whose values on that trial were always positive; N for an aversive attribute, whose values were always negative; and B for an ambivalent attribute, whose values could be positive and/or negative. Labeling the attribute type is intended to mimic real attributes, whose valence is typically known in advance of the magnitude. For example, when searching for a laptop, more processing power is good, more weight is bad, and more screen size could be either good or bad. The location of attribute types and options were randomized across trials.

At the start of a trial, the boxes are empty. An attribute value is revealed only while the participant's mouse cursor is inside its box (see Figure 1B). The box reverts to empty as soon as the cursor leaves the box, but participants are allowed to hover over a box for as long as they'd like, and they can re-visit boxes as many times as they wish. Only one value is therefore visible at any moment, which mutes some peripheral processing of unattended values and gives us a less ambiguous record of what was being examined (Eum et al., 2023). Values are displayed as colors rather than numbers, on a continuous scale running from dark green through white to dark pink, with color saturation indicating magnitude (see Figure 1C) . The direction of the scale was randomized at the participant level; for half of participants, green indicated negative values and pink indicated positive values, while for the other half, green indicated positive values and pink indicated negative values. Participants can click anywhere within an option's column to select it. There is no time limit and no trial-by-trial feedback.

The value of an option is the sum of its two attribute values. The task is incentive compatible. Points were accumulated across all trials and converted to money at the end of the experiment. In Study 1, participants received a $2.50 show-up fee and points were converted at 400 points per US dollar, so that a participant choosing at chance earned approximately $2.50 and a participant choosing perfectly earned approximately $4.50. In Study 2, the show-up fee was raised to $5.00, but other details were the same.

The task was implemented in Qualtrics using QMT (Qualtrics Mouse-Tracking) (S. Smith & Dolifka, 2026).

Range Manipulation (Between-Subjects). At the start of the experiment, each participant was randomly assigned to a Narrow or a Wide range condition, which fixed the set of attribute values they would see for the entire experiment. In the Narrow condition, appetitive values were drawn from $\{5, 10, 15, 20, 25, 30\}$, aversive values from $\{-30, -25, -20, -15, -10, -5\}$, and ambivalent values from $\{-15, -10, -5, 5, 10, 15\}$. In the Wide condition, the corresponding sets were $\{5, 10, 20, 30, 40, 50\}$, $\{-50, -40, -30, -20, -10, -5\}$, and $\{-20, -10, -5, 5, 10, 20\}$. The Narrow and Wide conditions therefore share four values per attribute type (see Figure 1D for a visualization). Pre-task instructions displayed the color scale and its mapping onto values. A pre-task quiz required participants to view a color and reproduce the corresponding value on a slider, which both verified comprehension and familiarized them with the range they had been assigned. Participants were given three attempts to pass the quiz. If they did not pass on the third try, they were allowed to proceed through the study, but we excluded their data (in line with our preregistrations). To see the order of operations in the experiment, see Figure 1E.

Cells. Throughout the paper, we refer to a cell as an attribute value for a given attribute type and range. For example, the tuple (10, P, Narrow) is a cell and denotes when the value of 10 is seen in the appetitive attribute row during the Narrow range condition.

Stimulus Construction. Attribute values were randomly sampled, subject to three constraints: (1) no option first order dominated the other; (2) the two attribute values within an option could not be identical, which would render the option a single solid color; and (3) attribute values could not repeat across options, which would reduce the comparison to a single attribute.

Differences Between Experiments. In Study 1, participants completed 30 choice trials in 3 blocks of 10 (one per trial type). We use these data to illustrate key behavioral patterns in the data. In Study 2, participants completed 60 choice trials organized into three blocks of 20 (one per trial type). Value sets, constraints, instructions, quizzes, and the number of practice trials were identical across experiments. We increased the number of trials in Study 2 to provide sufficient precision in our estimation of attentional weights. We use data from both studies for our modeling procedure.

2.2 Participants

Participants were recruited through CloudResearch Connect. The study was approved by the Institutional Review Board at the University of Chicago, and all participants provided informed consent.

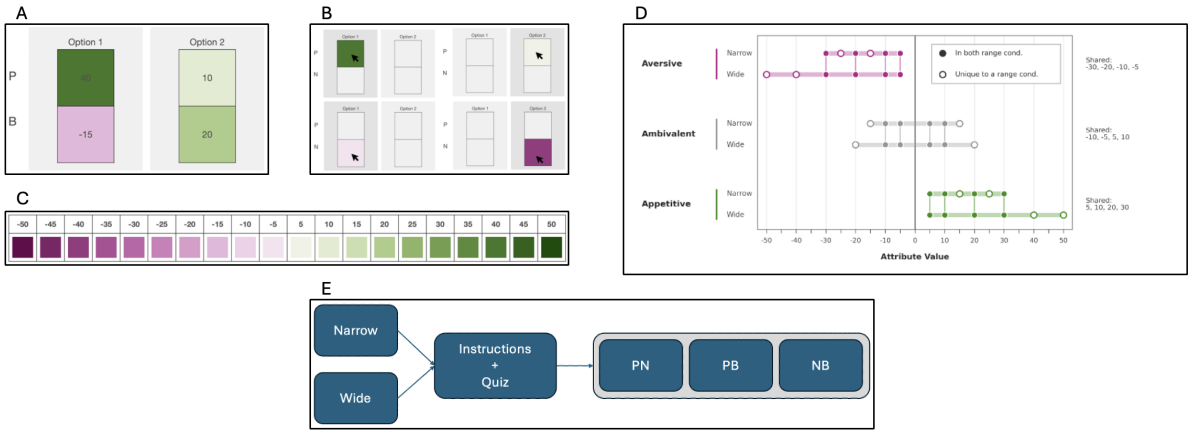


Figure 1: **Task.** (A) Participants chose between two options, each described by two attributes. Options occupy columns and attributes occupy rows. Row labels on the left give the attribute type: P (appetitive), N (aversive), or B (ambivalent). Numbers and all boxes are shown here for illustration only. (B) A value is visible only while the mouse pointer is inside its box, and the box returns to empty when the pointer leaves. (C) An example of the color-value map. The range of values depends on the Narrow or Wide condition. (D) Attribute values in the Narrow and Wide condition, separately by attribute type. 4 of 6 values per attribute type and range condition are shared across range conditions. (E) Participants were randomly assigned to the Narrow or Wide condition, completed instructions and a color-to-value quiz, then made choices in randomized blocks of PN, PB, and NB trials.

We hoped to retain 1000 participants after exclusions in Study 1 and 500 participants after exclusions in Study 2. For each study, based on our preregistered data collection process, we first collected 50 participants and calculated the exclusion rate for those 50 participants. We then collected enough additional participants, based on the exclusion rate, to meet our desired sample sizes of 1000 and 500. Therefore, we recruited a total of 1515 participants for Study 1 and 961 participants for Study 2.

As a pre-commitment to high data quality, participants were excluded if they failed any of the following pre-registered criteria: mouse-tracking data was recorded on fewer than 90% of trials; they did not fixate every attribute box at least once on at least 90% of trials; their mean response time exceeded 30,000 ms; they failed the instructions quiz more than three times; or their choice accuracy was indistinguishable from chance. This resulted in the exclusion of 534 participants in Study 1 (35% of the sample) and 436 participants in Study 2 (44% of the sample). Trial-level exclusions were applied during data processing: (1) trials in which the participant never revealed any box; and (2) trials in which total attention time was more than two standard deviations above the participant's mean log response time, suggesting they were not paying attention. This resulted in the exclusion of 784 trials in Study 1 (2.7% of the trials) and 977 trials in Study 2 (3.1% of the trials). After exclusions, the final samples comprised 981 participants in Study 1 and 525 in Study 2, for a total of N = 1,506.[2]

[2]In Study 2, six additional participants were removed before analysis due to a coding error. A subset of their trials carried attribute type labels that were inconsistent with the signs of the displayed values.

The analyzed data comprise 28,646 trials in Study 1 (mean of 29.2 out of 30 per participant) and 30,523 trials in Study 2 (mean of 58.1 out of 60 per participant). Trials are distributed across trial types as follows in Study 1: (PN=9,570, PB=9,581, NB=9,495). In Study 2: (PN=10,204, PB=10,201, NB=10,118).

2.3 Mouse Tracking

We treat mouse hovering as our measure of attention and use the vocabulary of eye-tracking throughout: a *dwell* is a continuous period during which the cursor remained inside one of the four attribute boxes, and *dwell time* is the summed duration of all dwells to a given box within a trial. Periods in which the cursor was outside all four boxes are recorded as transition periods and are not assigned to any box. Because a value is visible only while its box is hovered-over, and the task cannot be performed without hovering, the cursor position indexes which information the participant is currently sampling (Lohse & Johnson, 1996; S. Smith & Dolifka, 2026; Stillman et al., 2020).

2.4 Mixed-Effects and Hierarchical Regressions

All regressions include subject-level random effects. Non-Bayesian mixed models were fit with lme4 and glmmTMB using default settings. We report 95% Wald confidence intervals.

Bayesian models were fit with brms (Bürkner, 2018) using the cmdstanr backend, with four chains and weakly informative priors centered at the null result: $N(0, 2.5)$ on regression coefficients and $N(0, 2)$ on intercepts. Models were run for 4,000 iterations per chain with adapt delta of 0.95 and max treedepth of 12. Following brms conventions, we report the posterior median estimate and the 95% highest density interval for each parameter. All Bayesian models converged with no divergent transitions. All group-level parameters and subject-level random effects had $\hat{R} \leq 1.01$ and effective sample size > 1000.

2.5 Computational Model

We adapt the Bayesian sequential-sampling model from Frömer et al. (2025) to attributes. Let $i \in \{L, R\}$ denote options and $j \in \{1, 2\}$ denote attributes. Each box (option $\times$ attribute) is a value the agent estimates by sampling while the box is attended to. Before any sampling, an agent starts with a prior belief about the value of a box $\mathcal{N}(m_0, \sigma_p^2)$. Their prior is calibrated to a reference distribution D of values, which is all eighteen values that the attributes can take on in a given range condition. Later in the text and supplementary materials, we discuss variants of this model and their results.

The mean of the prior m_0 can be interpreted as the raw value of the "reference point" with which attribute values are evaluated against. Departing from Frömer et al. (2025), we allow this value to be fit to the data, rather than assume it lies at the mean value in a context. To compare this reference point across different contexts, we define a parameter q that determines the position of the reference point m_0 within the range of D. For example, $q = 0.5$ is the

middle of the range of D, and if $D = \{1, 4, 5\}$, then $m_0 = 3$. In the spirit of the best-fitting model (unbiased prior) from Frömer et al. (2025), we fix the variance of the prior (σ_p^2) to the variance of values in D.

Formally, we say that an agent attends to box ij with value a_{ij} for t_{ij} seconds. Doing so yields a sample mean, $\bar{x}_{ij} \sim \mathcal{N}(a_{ij}, \sigma_x^2/t_{ij})$. Attention therefore enters into this model by increasing the precision of samples. The sampling noise σ_x differs across participants, and we treat it as a random effect. Within a range condition, $\log \sigma_x$ is normally distributed across participants with mean μ and standard deviation τ:

$$\log \sigma_x \sim \mathcal{N}(\mu, \tau^2) \tag{1}$$

The agent uses their prior and newly sampled information to form a posterior belief about the value of the box. Their estimate of the value is the mean of their posterior belief, m_{ij}, which can be expressed as a convex combination of their prior mean and sample mean:

$$m_{ij} = m_0 + \pi(t_{ij})\,(\bar{x}_{ij} - m_0) \tag{2}$$

where $\pi(t_{ij})$ is the relative precision of the sampled information (t_{ij}/σ_x^2) compared to the combined precision of the prior ($1/\sigma_p^2$) and sampled information:

$$\pi(t_{ij}) = \frac{t_{ij}/\sigma_x^2}{1/\sigma_p^2 + t_{ij}/\sigma_x^2} \tag{3}$$

The relative precision of the sampled information is increasing in time spent attending to that box (t_{ij}), meaning that the more an agent attends to a box, the more they will rely on the information they have sampled, relative to what they previously believed before looking.

Because we are investigating losses in this paradigm, we consider the role of loss aversion in this model. The value estimate is valued relative to the prior mean, with a free loss aversion parameter (λ):

$$u(m) = \begin{cases} m_0 + (m - m_0), & \text{if } m \geq m_0, \\ m_0 + \lambda(m - m_0), & \text{if } m < m_0. \end{cases} \tag{4}$$

To make a choice, the agent adds up the utilities of each option's two boxes, $U_i = u(m_{i1}) + u(m_{i2})$, and picks the option with the larger total. Because option utilities are themselves random variables, the choice is stochastic, and therefore each choice has a probability (a likelihood) which we can use to fit the model. When $\lambda = 1$, we can calculate this likelihood using the Probit $\Phi(\delta/\sqrt{v})$, where δ and v are the mean and variance of the difference between the two options' totals. But when $\lambda \neq 1$, that difference is no longer normally distributed, so we numerically approximate the probability with a grid of utility values instead.[3] See

[3]This procedure is called discrete numerical convolution (Panjer, 1981), and is similar to the grid-based likelihood computations for sequential sampling models with time-dependent drift (Liu et al., 2025; Tavares et al., 2017).

Supplementary A for details on computing trial likelihoods.

The key prediction of this model is that attention attracts choice toward an attribute whose estimate lies above the prior mean and repels it from an attribute whose estimate lies below. The strengths of these attentional effects depend on the difference between the true attribute value and the prior mean, the precision of the signal, and the precision of the prior.

We chose this model because we find evidence in our data of range sensitivity, attentional attraction effects, and attentional repulsion effects, and this model is capable of simultaneously predicting all three. It attributes these effects to two mechanisms: the prior beliefs that participants hold about the distribution of values in a range condition, and the precision of sampled information under these different range conditions.

2.6 Model Fitting

We fit the model to the data from both studies via maximum likelihood estimation, then evaluate it via cross-validation. Each participant's trials are cut into chunks of five in the order they were completed, and each of the six (Study 1) or twelve (Study 2) folds holds out one chunk per participant while the model is refit to the rest.

There are 3 kinds of group-level parameters in the model. The first, q, describes where in the range of values the reference point lies. The second, λ, describes the degree of loss aversion. The third, (μ, τ), describes the population distribution of sampling noise σ_x. The noise parameter σ_x is not estimated per participant. Instead, we estimate the distribution it follows across participants, separately for each range condition (i.e., $(\mu_{\text{Narrow}}, \tau_{\text{Narrow}})$ and $(\mu_{\text{Wide}}, \tau_{\text{Wide}})$). To get the marginal likelihood for a participant, we compute the probability of their full set of choices at each value σ_x could take, then take the weighted average of these probabilities using the population distribution for that range condition.[4] The reference point q is estimated at the group-level, conditional on range condition and pooled across both experiments (i.e., q_{Narrow} and q_{Wide}). Loss aversion λ is estimated at the group-level across all participants, regardless of range condition or experiment (i.e., just λ).

Since the two studies used the same value sets, constraints, and instructions, we pooled these parameters in this way for the purpose of identification. One of the benefits of pooling is that we can better identify q and λ: when q is relatively low (or high) compared to the range of attribute values in a condition, then there are only a few trials per participant where utilities are negative, and the model may not have enough data to pin down q or λ. Parameter recovery exercises on simulated data with our stimuli and dwell durations show that q, λ, μ, and τ are recovered accurately, and that σ_x (a participant's posterior mean given their choices) is adequately recovered at the participant level.

[4]This is what mixed-effects regressions do. Treating σ_x as a random effect prevents (a) highly accurate participants from being assigned near-perfect sampling precision (i.e., $\sigma_x \approx 0$), and (b) near random participants from being assigned near-∞ sampling precision. Otherwise, a single held-out prediction error can dominate a participant's out-of-sample likelihood in these cases. Averaging over the population distribution puts little weight on such values, and how far a participant's own trials can pull their σ_x away from the population is set by τ. Another approach would be to include a lapse rate, but this absorbs informative errors on easy trials caused by attention, which in turn biases q and λ.

Fitting proceeds in three stages. In Stage 1, we fix $\lambda = 1$ and iterate through a coarse grid of q. At each grid value, we estimate each range condition's population distribution of sampling noise (μ and τ) by maximizing the summed marginal likelihood of that condition's participants. In Stage 2, starting from the best q value in the coarse grid for each range condition, we refine q by running a tightly bounded search, keeping λ fixed at 1 and re-estimating the population distributions at each new candidate q in the tighter range. In Stage 3, we use Nelder-Mead (starting at the refined q for each range condition and $\lambda = 1$) to jointly refine q in both range conditions and λ, again re-estimating the population distributions at each step. Because λ is shared over participants, ranges, and studies, this maximization runs over the whole dataset. In the first two stages $\lambda = 1$, so the choice probabilities come from the Probit. In the third stage, λ is fit, so choice probabilities must be numerically approximated.

Trials in which an attribute was never hovered are excluded from the fitting process (1,047 across the whole dataset). See Supplementary C for more fitting details.

2.7 Out-of-Sample Simulations

We compare variants of the model by cross-validation. Fold k holds out the k-th chunk of five trials for every participant, while fitting the model on the remaining trials. Summing over all folds, we compute the out-of-sample log-likelihood (OOS-LL) of a variant.

Since we are particularly interested in whether the model can reproduce the observed attentional attraction and repulsion effects from our data, we also compute an additional measure of out-of-sample model fit more tailored for this objective. We compute the out-of-sample mean absolute error (OOS-MAE) between "attentional weights" predicted by the model versus observed in the data. To do this, we simulate ten datasets using the out-of-sample choice probabilities per fold, and report the averaged metric over these ten datasets. Section 3.4 explains these attentional weights.

3 Results

3.1 Basic Psychometrics

Figure 2 and Table S2 show basic psychometric properties of the choice data.

The top row of Figure 2 shows the probability of choosing the left option as a function of the value difference between the options, separately by study and range condition. Choice accuracy did not meaningfully differ in the Narrow condition across studies, though there was a slight but significant difference in the Wide condition across studies ($\beta^{\text{Value Diff., Narrow}}_{\text{Study 2 - Study 1}} = -0.027\ [-0.135, 0.082]$; $\beta^{\text{Value Diff., Wide}}_{\text{Study 2 - Study 1}} = -0.146\ [-0.239, -0.054]$). Choices were sensitive to value differences, though more so in the Narrow condition than the Wide condition by quite a significant margin ($\beta^{\text{Value Diff., Wide - Narrow}}_{\text{Study 1}} = -0.433\ [-0.527, -0.340]$; $\beta^{\text{Value Diff., Wide - Narrow}}_{\text{Study 2}} = -0.553\ [-0.661, -0.445]$). Because a given value difference is a smaller fraction of the available range in the Wide condition, this is a key indication of range sensitivity in these data. Across

various models of value normalization (e.g., range normalization, divisive normalization, rank normalization), simulated predictions show that value sensitivity is higher in the Narrow condition compared to the Wide condition (see Supplementary E). Supplementary F verifies that this range sensitivity is present regardless of the attribute type in our data.

The bottom row of Figure 2 shows response times as a function of the absolute value difference between options (i.e., inverse choice difficulty), separately by study and range condition. Response times decreased with the absolute value difference between options. On average, response times in Study 2 were faster than in Study 1 ($\Delta_{\text{Study 2 - Study 1}} = -2.588$ $[-3.676, -1.500]$ s), though sensitivities to absolute value differences did not statistically differ ($\beta^{\text{Difficulty}}_{\text{Study 2 - Study 1}} = 0.229$ $[-0.005, 0.464]$). With more trials per subject in Study 2, part (but not all) of the difference in response times can be attributed to increased familiarity with the task in later trials ($\Delta^{\text{Trials 31–60 vs. 1–30}}_{\text{Study 2 - Study 1}} = -2.71$ s, $d = -0.39$, $t(1502.2) = -8.5$, $p = 4e{-}17$; $\Delta^{\text{Trials 1–30 vs. 1–30}}_{\text{Study 2 - Study 1}} = -1.43$ s, $d = -0.20$, $t(1385.3) = -4.0$, $p = 5e{-}05$). Decisions in the Wide range condition were longer on average than decisions in the Narrow range condition, but these level differences were not significant ($\Delta^{\text{Wide - Narrow}}_{\text{Study 1}} = 0.326$ $[-0.590, 1.243]$; $\Delta^{\text{Wide - Narrow}}_{\text{Study 2}} = 0.823$ $[-0.413, 2.058]$).

Overall, there was a negligible difference in the accuracy of choices across studies, however, range sensitivity is clearly evident. Response times differed across studies but not range conditions.

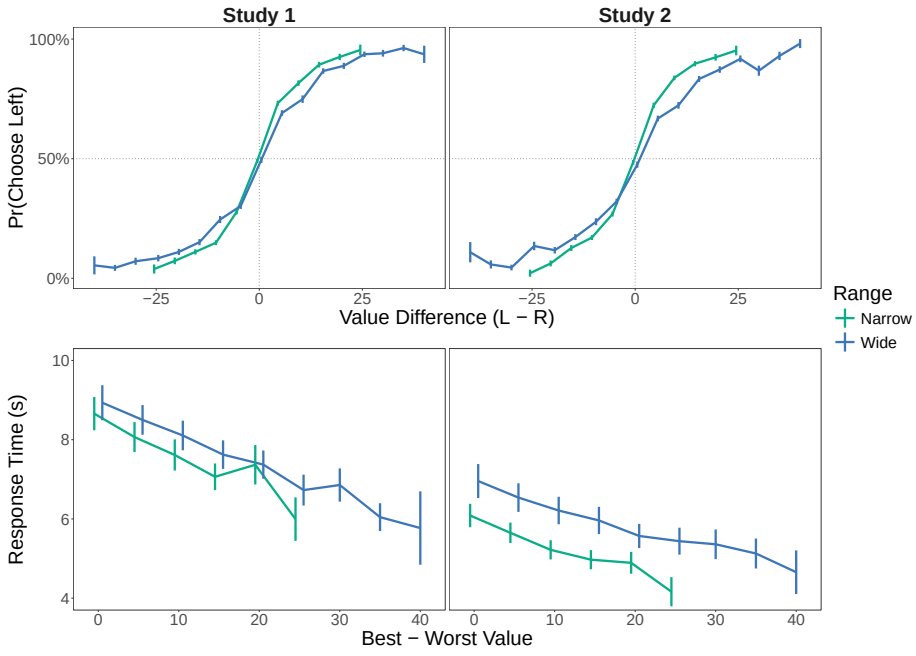


Figure 2: **Basic Psychometrics.** (Top Row) Probability of choosing the left option, as a function of the left minus right value, separately by range condition. (Bottom Row) Response time, as a function of best minus worst value, separately by range condition. Columns indicate study. Error bars depict standard error of the mean across participant means.

3.2 Dwell Properties

Figure 3 and Table S4 explore the dwell process in more detail.

The first row of Figure 3 shows the probability of first looking at the best option as a function of the value difference between the best and worst option, separately by study and range condition. The probability is at chance across studies and conditions. This is largely a sanity check, since participants had no way of knowing a priori which option would have higher value.

The second row of Figure 3 shows mean dwell duration across attribute values, separately by study, range condition, and attribute type. In Studies 1 and 2, mean dwell durations were typically between 600 to 800 ms and 500 to 700 ms, respectively. They did not differ by attribute type ($\beta^{\text{N - P}}_{\text{Study 1}} = -0.014$ $[-0.043, 0.015]$, $\beta^{\text{B - P}}_{\text{Study 1}} = 0.013$ $[-0.009, 0.036]$, $\beta^{\text{N - P}}_{\text{Study 2}} = -0.008$ $[-0.038, 0.023]$, $\beta^{\text{B - P}}_{\text{Study 2}} = -0.005$ $[-0.028, 0.018]$) or attribute value ($\beta^{\text{value, N}}_{\text{Study 1}} = -0.002$ $[-0.013, 0.008]$, $\beta^{\text{value, B}}_{\text{Study 1}} = -0.007$ $[-0.015, 0.001]$, $\beta^{\text{value, P}}_{\text{Study 1}} = -0.008$ $[-0.018, 0.002]$, $\beta^{\text{value, N}}_{\text{Study 2}} = 0.008$ $[-0.003, 0.019]$, $\beta^{\text{value, B}}_{\text{Study 2}} = 0.007$ $[-0.001, 0.015]$, $\beta^{\text{value, P}}_{\text{Study 2}} = -0.012$ $[-0.023, -0.001]$), however dwell durations were shorter in the Narrow condition ($\Delta^{\text{Wide - Narrow}}_{\text{Study 1}} = 0.052$ $[-0.014, 0.118]$, $\Delta^{\text{Wide - Narrow}}_{\text{Study 2}} = 0.063$ $[-0.023, 0.150]$).

The third row of Figure 3 shows the mean share of dwells across attribute values, separately by study, range condition, and attribute type. Subjects exhibited a tendency to more frequently revisit higher-value attributes ($\beta^{\text{value, N}}_{\text{Study 1}} = 0.0026$ $[0.0013, 0.0038]$, $\beta^{\text{value, B}}_{\text{Study 1}} = 0.0041$ $[0.0031, 0.0051]$, $\beta^{\text{value, P}}_{\text{Study 1}} = 0.0010$ $[-0.0002, 0.0023]$, $\beta^{\text{value, N}}_{\text{Study 2}} = 0.0030$ $[0.0018, 0.0042]$, $\beta^{\text{value, B}}_{\text{Study 2}} = 0.0029$ $[0.0019, 0.0039]$, $\beta^{\text{value, P}}_{\text{Study 2}} = 0.0024$ $[0.0012, 0.0036]$). This was also true across attribute types, with appetitive attributes receiving more revisits than ambivalent attributes, which in turn had more revisits than aversive attributes ($\Delta^{\text{P - B}}_{\text{Study 1}} = 0.53$ pp, $d = 0.29$, $t(980) = 9.1$, $p = 5e-19$; $\Delta^{\text{B - N}}_{\text{Study 1}} = 0.30$ pp, $d = 0.16$, $t(980) = 5.1$, $p = 4e-07$; $\Delta^{\text{P - B}}_{\text{Study 2}} = 0.52$ pp, $d = 0.35$, $t(524) = 8.0$, $p = 7e-15$; $\Delta^{\text{B - N}}_{\text{Study 2}} = 0.35$ pp, $d = 0.25$, $t(524) = 5.8$, $p = 9e-09$).

Taken together, these results suggest that where value shapes attention, it does so by bringing the participant back to an attribute more often rather than by holding them there longer. This is consistent with accounts in which the dwell process is a search policy over uncertain values rather than a passive dwell process (Callaway et al., 2021; Gluth et al., 2024; Jang et al., 2021). We note that the effects of value-driven attention are small, and that for the remainder of this paper, we treat attention as exogenous.

3.3 Option-Level Attention

Figure 4 and Table S5 explore attentional choice biases at the option-level. First, we show a positive relationship between net attention to an option and the probability of choosing that option. Then, we document a bias toward the last-seen option.

The top row of Figure 4 shows the corrected probability of choosing the left option as a function of the net dwell time to the left option, separately by study and range condition. The

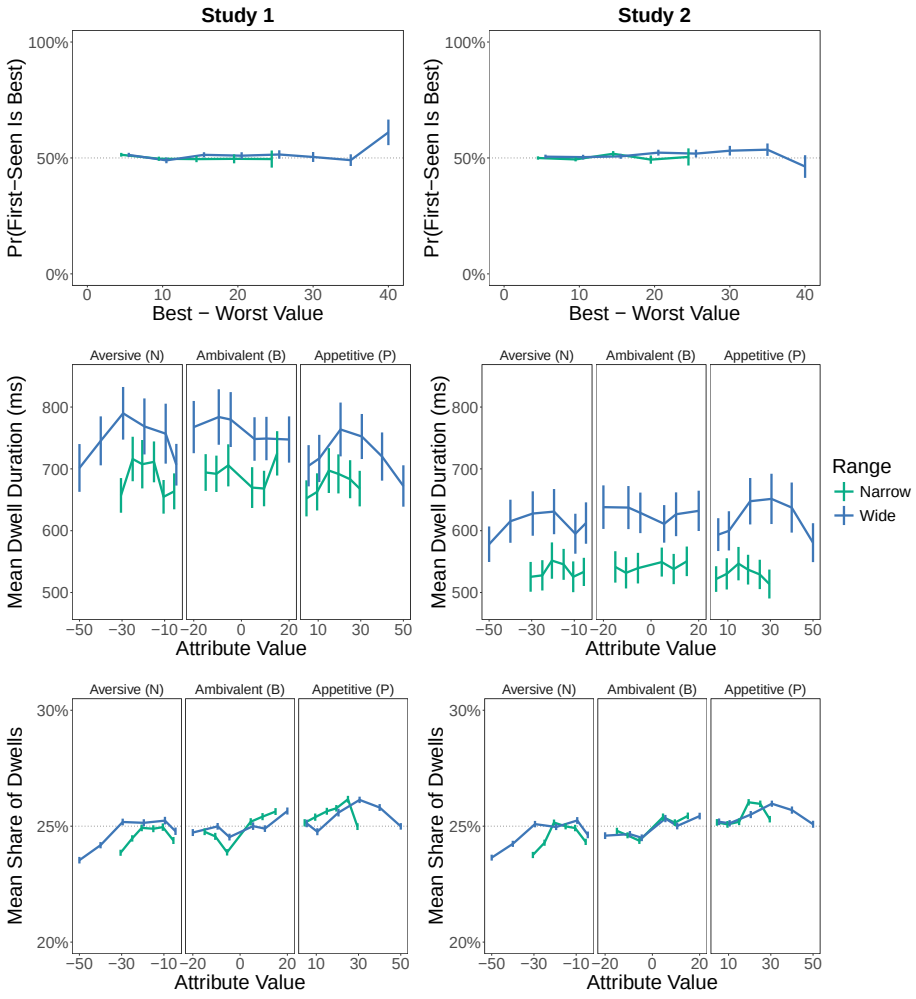


Figure 3: **Dwell Properties.** (Row 1) Probability of choosing the left option, as a function of the left minus right value, separately by range condition. (Row 2) Mean dwell time, as a function of best minus worst value, separately by range condition. (Row 3) Mean share of dwells, as a function of best minus worst value, separately by range condition. Columns indicate study. Error bars depict standard error of the mean across participant means.

probability of choosing the left option was corrected by running a logistic regression of choice on value difference, then subtracting the predicted probability from the binary choice variable. In the absence of an attentional bias, this corrected probability should be 0 for all values of net dwell time to left. Instead, we find a positive and significant relationship in both studies ($\beta^{\text{Dwell}}_{\text{Study 1}} = 0.013\ [0.003, 0.023]$, $\beta^{\text{Dwell}}_{\text{Study 2}} = 0.024\ [0.007, 0.040]$), which does not differ by range condition ($\beta^{\text{Dwell}\times\text{Wide}}_{\text{Study 1}} = -0.009\ [-0.023, 0.006]$, $\beta^{\text{Dwell}\times\text{Wide}}_{\text{Study 2}} = 0.001\ [-0.020, 0.022]$).

The bottom row of Figure 4 shows the probability of choosing the last-seen option as a function of the value difference between the last-seen option and the other option, separately by study, range condition, and the last-seen attribute type. In the absence of an attentional choice bias, the lines should cross 50% when value difference is 0. Instead, we find evidence of final-dwell bias in both studies, regardless of the range condition or which attribute was seen last. ($\beta^{\text{Last, B, Wide}}_{\text{Study 1}} = 0.918\ [0.824, 1.011]$, $\beta^{\text{Last, B, Narrow}}_{\text{Study 1}} = 1.087\ [0.993, 1.181]$, $\beta^{\text{Last, P, Narrow}}_{\text{Study 2}} = 0.878\ [0.779, 0.978]$, $\beta^{\text{Last, P, Wide}}_{\text{Study 2}} = 1.032\ [0.931, 1.133]$).

These results are consistent with previous studies (Eum et al., 2026; Krajbich & Rangel, 2011; Krajbich et al., 2010; S. M. Smith & Krajbich, 2018; Yang & Krajbich, 2023) and demonstrate that option-level attention plays a role in the choice-process.

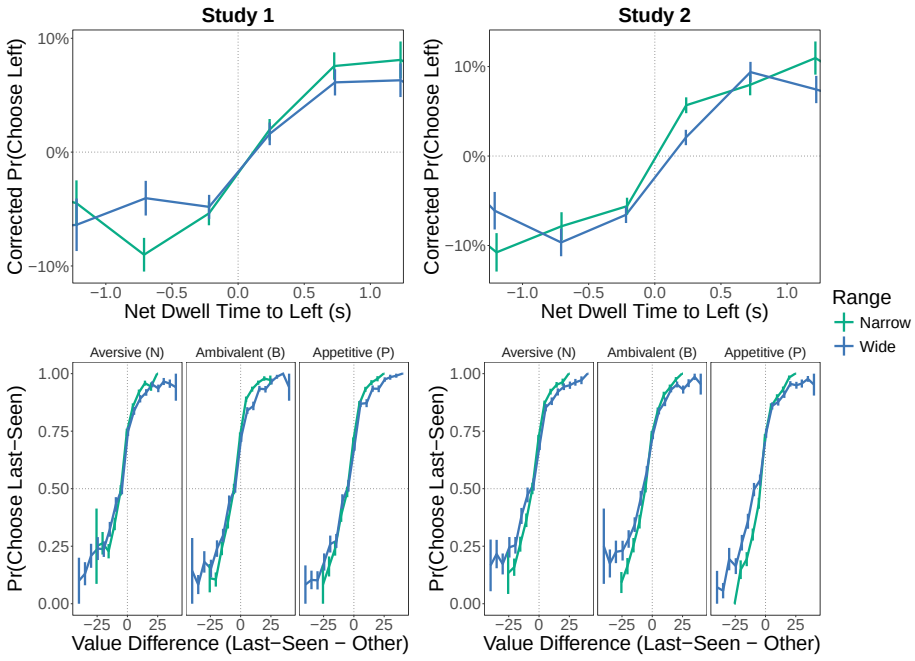


Figure 4: **Option-Level Attention.** (Top Row) Corrected probability of choosing the left option, as a function of the net dwell time to the left option, separately by study and range condition. The probability of choosing the left option was corrected by running a logistic regression of choice on value difference, then subtracting the predicted probability from the binary choice variable. (Bottom Row) Probability of choosing the last-seen option, as a function of the value difference between the last-seen and other option, separately by study, range condition, and attribute type. Columns indicate study. Error bars depict standard error of the mean across participant means.

3.4 Attribute-Level Attention

Figure 5 and Table S6 explore attentional choice biases at the attribute-level.

In order to explore the effects of attribute-level attention on choice without a structural model of utilities and the role of attention in the decision process, we estimate attention weights with a Bayesian hierarchical logistic regression of choice on the four boxes of the trial. Each box belongs to one of 36 cells ($c \in C$; attribute value $\times$ attribute type $\times$ range condition) and enters the regression through two regressors. The first is signed presence, I_c, which is $+1$ if the cell is in the left option, -1 if it is in the right option, and 0 for cells that are not in the current trial. The second is signed dwell time, which is the dwell time to the box in seconds (G_c) multiplied by I_c (so that it is positive if the cell is in the left option, negative if it is in the right option, and 0 otherwise). For participant s in trial n:

$$\text{logit}[\Pr(\text{choose left}_{sn})] = \underbrace{\alpha_1 + \alpha_2\, 1_{\text{Study 2}} + \sum_{c \in C} (u_c + u'_{c,s})\, I_{c,sn}}_{\text{value encoding effects}} + \underbrace{\sum_{c \in C} (\omega_c + \omega'_{c,s})\, G_{c,sn} I_{c,sn}}_{\text{attentional effects}} \quad (5)$$

Where α_1 and α_2 are study fixed effects, and $u'_{c,s}$ and $\omega'_{c,s}$ are random effects. The coefficient u_c is the group-level mean utility of that cell, or the re-encoding effects. It quantifies the change in the log-odds of choosing an option when it carries that attribute value in that attribute type and range condition. Because each attribute type has one box in each option, only differences between the utilities of a type are identified, so we set the lowest value of each type to be the reference with utility 0. The coefficient ω_c is the group-level mean attention weight of the cell, or the attentional effects. It is the change in the log-odds of choosing an option per second of dwell on that box, holding the four attribute values and the other three dwell times constant. The ω_c's, our attentional weights, are our primary estimates of interest. Though we do include separate study fixed effects to capture level differences in choices between the two studies, we chose to pool the two studies together for the utility and attentional weights. We do this because (a) the studies used identical value sets, constraints, and instructions, and (b) fits with separate weights per experiment are consistent, but noisier (see Supplementary H). Thus, in the spirit of meta-analysis, estimating a model jointly on the two studies improves the precision of attentional weight estimates.

The vertical line ranges in Figure 5 plot the 95% highest density interval (HDI) of the group-level attention weights (ω_c) as a function of attribute value, separately by attribute type and range condition. In the absence of an attribute-level attentional bias, attention weights would be 0. Instead, we find that attention weights depend on three things: the attribute type, the attribute value, and the range condition. In the following sub-sections, we discuss the pattern of attention weights for each attribute type, in turn.

Aversive Attributes. Attention weights exhibit three striking patterns. For the two lowest attribute values in each range condition, attention weights are significantly (or almost significantly) negative ($\omega_{\text{N}}^{\text{Narrow}}(-30) = -0.211$ $[-0.287, -0.134]$, $\omega_{\text{N}}^{\text{Narrow}}(-25) = -0.027$ $[-0.071, 0.016]$, $\omega_{\text{N}}^{\text{Wide}}(-50) = -0.178$ $[-0.266, -0.091]$, $\omega_{\text{N}}^{\text{Wide}}(-40) = -0.044$ $[-0.091, 0.000]$). This suggests that attention repels choices away from

the option whose attribute value is attended, which is consistent with predictions from the standard Attentional Drift Diffusion Model (aDDM) (Eum et al., 2026; Krajbich et al., 2010). For the two highest (but still negative) attribute values in each range condition, the attention weights are positive ($\omega_{\text{N}}^{\text{Narrow}}(-10) = 0.105$ $[0.040, 0.178]$, $\omega_{\text{N}}^{\text{Narrow}}(-5) = 0.222$ $[0.129, 0.315]$, $\omega_{\text{N}}^{\text{Wide}}(-10) = 0.122$ $[0.058, 0.188]$, $\omega_{\text{N}}^{\text{Wide}}(-5) = 0.115$ $[0.062, 0.172]$). This suggests that attention attracts choices towards the option with the attended attribute when its value is relatively high, which is inconsistent with predictions from the standard aDDM.

For the intermediate attribute values in each range condition, the attention weight is not significantly different from 0 ($\omega_{\text{N}}^{\text{Narrow}}(-20) = -0.006$ $[-0.046, 0.037]$, $\omega_{\text{N}}^{\text{Narrow}}(-15) = 0.026$ $[-0.024, 0.078]$, $\omega_{\text{N}}^{\text{Wide}}(-30) = 0.012$ $[-0.024, 0.048]$, $\omega_{\text{N}}^{\text{Wide}}(-20) = 0.031$ $[-0.006, 0.071]$). Crucially, intermediate values differ across range conditions, and so the same value might have different attentional effects on choice under different ranges. For example, -30 is the lowest possible value in the Narrow condition, whereas it is a middle value in the Wide condition. In the Narrow condition, it exhibits an attentional repelling effect on choice, while in the Wide condition, it exhibits little to no attentional effect on choice, and these varying effects are statistically different ($\omega_{\text{N}}^{\text{Narrow}-\text{Wide}}(-30) = -0.224$ $[-0.310, -0.141]$).

These attentional attraction and repelling effects are as-if values are evaluated with respect to a reference point somewhere in the middle of the range of negative values, prior to integration in an attentional process.

Ambivalent Attributes. Ambivalent values overlap with appetitive and aversive values in an interesting way. In our data, the same positive values in an ambivalent attribute sometimes have a stronger attentional attraction effect than their appetitive counterparts ($\omega_{\text{B-P}}^{\text{Narrow}}(5) = 0.094$ $[-0.004, 0.183]$, $\omega_{\text{B-P}}^{\text{Narrow}}(10) = 0.100$ $[0.013, 0.182]$, $\omega_{\text{B-P}}^{\text{Narrow}}(15) = 0.116$ $[0.043, 0.198]$, $\omega_{\text{B-P}}^{\text{Wide}}(5) = 0.076$ $[0.011, 0.144]$, $\omega_{\text{B-P}}^{\text{Wide}}(10) = 0.025$ $[-0.035, 0.087]$, $\omega_{\text{B-P}}^{\text{Wide}}(20) = 0.060$ $[0.004, 0.120]$). Because negative values close to 0 have an attentional attraction effect in aversive attributes, and these are the same negative values that ambivalent attributes take on, the same negative value in an ambivalent attribute sometimes has a weaker attentional attraction effect than its aversive counterpart ($\omega_{\text{B-N}}^{\text{Narrow}}(-15) = -0.044$ $[-0.119, 0.033]$, $\omega_{\text{B-N}}^{\text{Narrow}}(-10) = -0.056$ $[-0.146, 0.031]$, $\omega_{\text{B-N}}^{\text{Narrow}}(-5) = -0.169$ $[-0.274, -0.067]$, $\omega_{\text{B-N}}^{\text{Wide}}(-20) = -0.010$ $[-0.072, 0.051]$, $\omega_{\text{B-N}}^{\text{Wide}}(-10) = -0.115$ $[-0.194, -0.040]$, $\omega_{\text{B-N}}^{\text{Wide}}(-5) = -0.102$ $[-0.174, -0.031]$).

These results are consistent with the idea that the same value can be evaluated differently depending on the attribute type, and that this evaluation can affect whether attention positively or negatively predicts choice.

Appetitive Attributes. For the two highest attribute values in a range condition, attention weights were significantly positive ($\omega_{\text{P}}^{\text{Narrow}}(25) = 0.088$ $[0.032, 0.152]$, $\omega_{\text{P}}^{\text{Narrow}}(30) = 0.148$ $[0.072, 0.222]$, $\omega_{\text{P}}^{\text{Wide}}(40) = 0.064$ $[0.013, 0.121]$, $\omega_{\text{P}}^{\text{Wide}}(50) = 0.296$ $[0.201, 0.391]$). This is consistent with the predictions of the aDDM and suggests an attentional attraction effect. For the two lowest attribute values in a range condition, attention weights were indistinguishable from 0 ($\omega_{\text{P}}^{\text{Narrow}}(5) = 0.005$ $[-0.068, 0.076]$, $\omega_{\text{P}}^{\text{Narrow}}(10) = 0.009$ $[-0.042, 0.057]$,

$\omega_P^{Wide}(5) = -0.026$ $[-0.075, 0.023]$, $\omega_P^{Wide}(10) = 0.009$ $[-0.037, 0.052]$), suggesting that attention to the lowest values for a positive attribute is not a significant predictor of choice. Crucially, while the same attribute values sometimes have different attention weights under different range conditions ($\omega_P^{Narrow-Wide}(30) = 0.125$ $[0.045, 0.210]$), we do not find concrete evidence of a sign flip in attention weights for positive attributes.

These as-if positive values are seemingly evaluated with respect to a reference point at or below the minimum possible value in a context, prior to integration in an attentional process, which is consistent with previous findings (Eum et al., 2026).

In our pre-registration, we specified a different model for estimating attentional weights. That model does not account for value re-encoding effects, and so it potentially inflates attentional weights away from 0. The model provided here is therefore a more conservative test. We discuss this in more detail and relegate the results of the pre-registered model to Supplementary I.

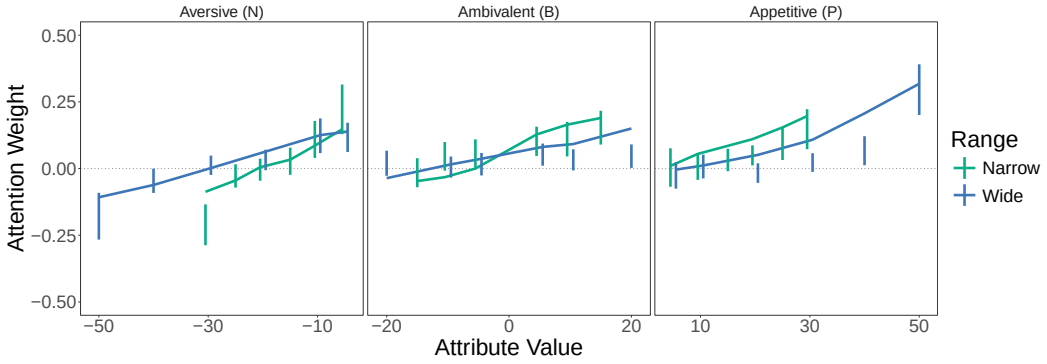


Figure 5: **Attention Weights by Attribute Value.** Vertical line ranges show the 95% HDI of the group-level attention weights (ω_c) as a function of attribute value, separately by attribute type and range condition. Horizontal lines depict the attention weights that the main Bayesian sequential sampling model produces out-of-sample. Error bars depict 95% highest density intervals.

3.5 Bayesian Sequential Sampling

The two-alternative Bayesian sequential sampling model from Frömer et al. (2025), which we adapt to multi-attribute decision making, is an ideal candidate for modeling our data because it can predict attentional attraction and repulsion effects on choice, as well as range sensitivity, which are all present in our data. Our model achieves these predictions through manipulations to the prior, which is analogous to our instructions and quizzes for the different range conditions in our paradigm. Rather than assuming prior beliefs like in Frömer et al. (2025), our model calibrates the prior to participants' behavior (i.e., the model freely estimates the reference point, instead of assuming it).

In this model, prior beliefs and precision from attention combine to produce attraction and repulsion effects. Before looking at an attribute, a decision maker expects its value to be

near a reference point (the mean of their prior beliefs). Attending to an attribute shifts beliefs away from that expectation towards the value presented on the screen, in proportion to the time spent attending. When an attribute value is worse than the reference point, attention makes the associated option less attractive, and when an attribute value is better than the reference point, attention makes the associated option more attractive. At the same time, attending to an attribute makes beliefs about the value of that attribute more precise, which should increase the accuracy of the choice. Because better options tend to have better attribute values within an attribute type, attention tends to exhibit stronger attraction effects as attribute value rises within an attribute type. Thus, attentional weights can still rise within an attribute type, even if the reference point is not in the range of that attribute type.

The lines in Figure 5 show the attention weights that the Bayesian sequential sampling model produces out-of-sample. There are five features of the attentional weights in the data that we want to emphasize: (1) the monotonic rise of the weights with value in every attribute type; (2) the crossover point from attentional repulsion to attraction inside the aversive value range; (3) weights that rise from null to positive in the ambivalent and appetitive attribute types; (4) a shallower rise of the weights with value in the Wide range condition; and (5) the degree of attentional repulsion at the lowest aversive values. The model reproduces features (1) through (4). It reproduces feature (5) in direction but not quite in magnitude.

The monotonic rise of attentional weights with value in every attribute type occurs because dwell increases the precision of the estimate of the attended box, which pushes choice toward the de facto better option, and a higher value is more likely to belong to the better option.

The crossover point from attentional repulsion to attraction within the aversive value range comes from the prior distribution. Attention moves the estimate of a box from the prior mean toward the box's value, so it lowers the estimate of a box worth less than the prior mean (a repulsion effect) and raises the estimate of a box worth more (the attraction effect). With $q \approx 0.2$ (i.e., a prior mean at the 20th percentile between the minimum and the mean), the prior mean sits at $m_0 = -17$ in the Narrow condition and at $m_0 = -31$ in the Wide condition. Only the lowest aversive values lie below it, so these are the only values at which attention repels, and the model's predicted crossover point in attentional weights falls inside the aversive range.

The weights that rise from null to positive in the ambivalent and appetitive attribute types occur because of two opposing channels. Every ambivalent and appetitive value lies above the prior mean, so the prior is imposing an attentional attraction effect. However, because the lowest values of each type are more likely to belong to the worse option, improved precision from attention repels choice away from options with these values. At the lowest values, the two channels roughly counteract each other and attentional weights become indistinguishable from zero. As value rises, precision from attention switches to attracting choice towards options, since these higher values are more likely to belong to the better option, and thus the weights become positive. The model's predicted attentional weights start near zero at the lowest ambivalent and appetitive values and end positive at the highest, as in the data.

The shallower rise in the weights in the Wide range condition comes from the sampling noise σ_x, which the model estimates separately for each range condition. The group-level

mean sampling noise per second ($\exp(\mu)$) is larger in the Wide condition than in the Narrow condition ($\Delta^{\text{Wide}-\text{Narrow}}_{\exp(\mu)} = 2.84$ [2.37, 3.29]). Noisier samples make neighboring values harder to tell apart, so a second of dwell on a value shifts choice less. Thus, the model's predicted weights rise more slowly with value in the Wide condition.

The degree of attentional repulsion at the lowest aversive values is only partially reproduced by the model. The lowest values lie below the prior mean and are the worst values of their type, and so the model predicts attentional repelling effects. But in the model, the repulsion grows in proportion to the distance below the prior mean, which is itself estimated to be somewhat close to the minimum value ($q \approx 0.2$), and therefore the model underestimates the degree of attentional repulsion at the lowest values. The free loss aversion parameter does not help in predicting this feature of the data.

In Supplementary G, we report results four variants of the model: the Per-Condition Prior Model of the main text with $\lambda = 1$; a Per-Type Prior Model, which calibrates a separate prior to each attribute type within a range condition; a Zero Prior Model, which fixes the reference point at zero; and a Flat Prior Model, which removes the reference point altogether. For each variant, we report group-level estimates, OOS-LL, the OOS-MAE of their attentional weights, and we plot their attentional weights. For OOS-LL, the main model fits best. For OOS-MAE, it ties with the Per-Condition Prior and $\lambda = 1$ variant.[5] We include a discussion of the different out-of-sample performances of the variants.

Taken together, these modeling results highlight the importance of prior beliefs on the role of attention in the choice process. The prior can adapt to the range of values in a context, and in turn, this can affect how attention influences how value is encoded during the choice process. This means that attention is not necessarily amplifying/discounting the raw values, but rather, it's moving value estimates away from the prior and towards the reference-dependent value through increased sampling precision. Under this account, the degree of attentional attraction or repulsion is a function of the moments of the prior, not just the difference, since these moments determine the crossover point in attentional weights, as well as the relative weighting of sampled information.

4 Discussion

While attention is typically thought to attract choice toward the attended option, it is unclear whether attention to known aversive or ambivalent attributes should exhibit the same attraction effects, or instead exhibit repulsion effects. In two mouse-tracking experiments in which each attribute's valence was known in advance, attention exhibited a mixture of attraction and repulsion effects. For aversive attributes, the sign of the attentional effect depended on the attended value. For the lowest aversive values, attention exhibited repulsion effects on choice. For intermediate aversive values, attentional effects were null, and for

[5] Despite this, we report the main model because, in settings with stronger loss aversion (Sokol-Hessner et al., 2009; Tom et al., 2007; Tversky & Kahneman, 1992), a Bayesian model without λ could only fit the asymmetry below the reference point by incorrectly shifting q or σ_x, attributing loss aversion to the prior or to sampling precision, when instead it emerges from preferences.

the highest aversive values, attention exhibited attraction effects on choice. Therefore, the attentional crossover point between repulsion and attraction fell roughly in the middle of the aversive value range, though the raw value at which this occurred depended on the range. For example, attention to the value of -30 repelled choice in the Narrow range condition (where it was the worst value) and had no detectable attentional effect in the Wide condition (where it was an intermediate value). For ambivalent and appetitive attributes, the lowest values in a range exhibited no detectable attentional effects, but that as value rose, stronger attentional attraction effects emerged. Overall, the magnitudes of attentional effects were less sensitive to value in the Wide condition compared to the Narrow condition.

A Bayesian sequential sampling model provides a mechanistic explanation of these attraction and repulsion effects, as well as the difference in attentional value sensitivity. Decision makers start from a prior, calibrated to the range of values in their context, and the prior mean acts as a reference point. Attention attracts choice toward an attribute whose value is better than the reference point and repels it away from one that is worse. The range modulates where the reference point lies and how noisy the sampling process is, and therefore sets both where attentional effects reverse and how sensitive choices are to value. We find that a Bayesian sequential sampling model with a prior distribution that varies by range condition, but not by attribute type, performs best out-of-sample. Additionally, the freely-fit reference point was estimated to be at the same relative position in a given range of values, which means that the raw value of this reference point differs across ranges.

These results suggest that attention does not directly attract or repel choices to or from the attended option, but instead, its effect on the precision of beliefs about value during the choice process can indirectly attract or repel choices. Attentional effects depend on whether the seen value exceeds or falls below the decision maker's expectations (which may not be the mean), and these expectations are affected by context effects, like the range.

These results provide nuance to the concept of attentional discounting in the choice process. In models like the Attentional Drift Diffusion Model, attention discounts the value of the option that is not currently looked at (Krajbich et al., 2010; S. M. Smith & Krajbich, 2018). However, when options contained aversive attributes, or were themselves aversive, different studies found varying directions of attentional biases, sometimes attracting and sometimes repelling choice from the attended option (Armel et al., 2008; Eum et al., 2026; Fisher, 2017; Krajbich et al., 2012; Sheng et al., 2020; Yang & Krajbich, 2023). Each of these findings can be individually reconciled with one of two channels: (1) either attentional discounting becomes attentional amplification in some contexts, or (2) values are re-encoded prior to the choice process. Our results favor the latter interpretation. When aversive values are reference-dependent, and the reference point is allowed to vary by context, all of these seemingly contradictory attentional results are explainable. In some studies, the reference point could be at or below the lowest possible value, and therefore attentional attraction effects are present. This explains why attentional attraction effects are typically found in most studies. In other studies, the reference point may fall somewhere within the range of possible values, and therefore attentional attraction and repulsion effects are both present. Either way, with this re-encoding of value, attentional discounting will push choices in the same direction as the Bayesian sequential sampling model, and thus operates as an approximation

of improving value-belief precision via attention.

The crucial difference between the two models is mechanistic: attentional discounting assumes attention operates directly on value, while in Bayesian sequential sampling, attention operates on sampling precision. Differences in fitted attentional parameters across domains (S. M. Smith & Krajbich, 2018), across individuals (Thomas et al., 2019), or across ranges (like in our study) could thus reflect differences in reference points (i.e., prior beliefs), rather than differences in attentional effects themselves. Disentangling these two accounts requires measuring and/or manipulating prior beliefs.

Additionally, under Bayesian sequential sampling, attention should exhibit diminishing returns on sampling precision, so early attention matters more than late attention. This would have the opposite effect of as collapsing boundaries in sequential sampling models (Glickman & Usher, 2019; Hawkins et al., 2015; Tajima et al., 2019). This is outside the scope of this paper, but grounds for future work.

Process models usually summarize attention with a single parameter, such as the discount in the Attentional Drift Diffusion Model, and compare it across domains (S. M. Smith & Krajbich, 2018), individuals (Thomas et al., 2019), and options/conditions (S. M. Smith et al., 2025). In our data, a second of looking has no single effect on choice. It is positive for values above the reference point, negative below the reference point, and small near the reference point. Critically, the position of the reference point depends on the range of possible values. A parameter estimated with the reference point fixed at zero, or at the minimum of the context, averages over these regions, weighted by how the values on offer are distributed around the reference the decision maker actually uses. The same attentional mechanism would therefore appear as a strong bias when most values lie far above the reference point, as a weak one when most values lie near it, and as a reversed one when most values lie below it.

Our fitted reference point sat at about the same relative position of the range in both conditions, which means that its raw value moved with the range. Because participants learned their range from the instructions and associated quiz, the range manipulation can be interpreted as a manipulation of the prior. We found that in turn, choices were less sensitive to values when ranges were wider. These behavioral signatures are consistent with range-adapting value codes in the orbitofrontal cortex (Padoa-Schioppa, 2009), value normalization models (Louie et al., 2013; O'Donoghue & Sprenger, 2018; Oprea & Woodford, 2026; Rangel & Clithero, 2012), signals of deviations from reference points in the ventral striatum (De Martino et al., 2009), and reward signals for avoiding aversive outcomes in the medial prefrontal cortex (Lim et al., 2011). We demonstrate that this context-dependent re-encoding of value is a determinant in the role of attention in the choice process.

The distribution of values is not the only context effect that can influence the attentional choice process. Recent evidence has suggested that goals also remap values, not choice policies (Frömer et al., 2019; Sepulveda et al., 2020; Shi et al., 2026). For example, when asked to minimize rather than maximize outcomes, participants do not simply reverse their choices. Instead, their behavior is better explained by a goal-dependent remapping of the underlying values. It remains untested whether goals remap aversive values the same way that they do appetitive values. With a reference point estimated within the range of aversive values, it may be that attention once again demonstrates attraction and repulsion effects under different

goals. This is a potential avenue for future work.

Our participants exhibited relatively little loss aversion compared to other studies (Sokol-Hessner et al., 2009; Tom et al., 2007; Tversky & Kahneman, 1992). This is potentially because our stakes were small and our value-to-color mapping enforced an additional layer of abstraction from raw values. Alternatively, it is possible that without accounting for preferential gaze towards losses, estimates of loss aversion might conflate attentional and valuation biases (Lejarraga et al., 2019; Sheng et al., 2020; Yechiam & Hochman, 2013). Regardless, we kept the loss aversion parameter in our model since it is possible that in future studies with other contexts and different participants, loss aversion might be more prevalent than it was in the present data. If so, then failing to account for loss aversion would artificially push attentional weights (as we have estimated them) away from 0, introducing spurious results.

The model underestimates the magnitude of the attentional weights at the lowest aversive values. These may potentially be accounted for by a prior that cannot place any mass below the minimum of the range (Bedi et al., 2025). Truncating the prior at its lower tail shifts the prior mass away from the minimum of the range. With a reference point near this minimum, sampling values below the reference point may produce a larger downward revision in its estimated value, thereby amplifying the attentional repulsion effect on choice. Truncating the prior at the maximum value would work the same way in principle, but in our case, this matters little because the fitted prior mean is near the bottom of the range of values. To maintain a clear set of results without incorporating too many changes to the Bayesian sequential sampling model, we chose to leave this to future work.

One potential future direction would be to directly test whether attention can identify participant-level reference points. We did not elicit reference points from participants in our studies, and instead infer the reference points based on our model of attention and choice. In our model, the estimated mean of prior beliefs depends on the amount of time an attribute has been sampled, the value of that attribute, and the resulting choice. Without accounting for these attentional patterns, tasks that estimate priors from choice or self-report alone may misattribute the effects of attention to other parameters, like loss aversion or decision error.

Computational accounts of psychiatric conditions sometimes look to the location of the prior and the weight placed on new evidence as metrics for diagnoses or treatment. For instance, depression has been linked to more negative prior beliefs about oneself (Gagne et al., 2022). It has also been linked to elevated and inflexible reference points (which then encourage values to be negatively evaluated) and that this tracks with symptom severity (Vittala et al., 2025). Unusual, delusion-like beliefs, which are central to psychotic disorders such as schizophrenia, have been linked to sequential base-rate neglect and to noisier representations of prior beliefs (Ashinoff et al., 2022). Our results suggest that the same attentional intervention could have different, even opposite, effects in people who differ in prior beliefs, such as different clinical populations. Raising the reference, as in depression, puts more stimuli below it, so that attending to them repels rather than attracts choice. Additionally, tasks aimed at estimating priors in clinical populations may want to consider measuring and controlling for attentional effects, as differences in attention-based choices could result from differences in attention rather than differences in the prior belief. Our sample was not clinical in nature and we

collected no measures of psychological state, so we leave this to future work.

Attentional interventions are not limited to clinical populations, and many interventions function by directing attention to an option or attribute, either through saliency manipulations or changes in display (Johnson et al., 2012). Directing attention to the health aspects of foods can improve dietary choice (Hare et al., 2011), and making taxes salient changes how consumers respond to prices (Chetty et al., 2009). Our results imply that such interventions have no fixed sign, and instead depend on beliefs in that context. Directing attention to an attribute shifts beliefs towards its value and away from what the decision maker expected a priori, so it favors an option when the attended value is better than expected and counts against it when the value is worse.

In public health, a label that draws attention to a high calorie count may attract an audience whose reference is set by menus with far higher counts and repel an audience whose reference is set by lighter counts. Consistent with the possibility that informational interventions have heterogeneous effects across populations, Wisdom et al. (2010) found that providing daily calorie recommendations reduced calories ordered by non-overweight individuals, but not by overweight individuals. The same holds for the side effects of medical treatments. A highlighted mild side effect may attract a patient whose reference is set by the severe side effects they have heard of or repel a patient whose reference is set by a standard of no side effects.[6] In consumer settings, the same product attribute may be perceived as favorable or unfavorable depending on the alternatives consumers are accustomed to evaluating. For example, Voichek and Novemsky (2025) found that the same product could be evaluated more or less favorably depending on whether it was compared with superior or inferior alternatives. Our results suggest a potential mechanism through which attention may contribute to such context-dependent evaluations: by shifting value estimates away from prior expectations, attention can attract choice toward an option when its attributes are better than expected and repel choice when they are worse. This is particularly relevant for online marketplaces, where designers have control over nearly every aspect of the choice architecture.

For practical purposes, we recommend considering the following four steps before implementing an attentional intervention. First, learn the audience's prior beliefs. Second, treat prior beliefs as a design variable, since it can be adjusted (in certain populations) through training, experience, or instructions. Third, test the attentional intervention at the low and the high ends of the values it will display, because its effect can reverse. Fourth, audiences also differ in their priors and in how they update them, so the same intervention can have different effects on different groups. Consumer research now records attention directly with eye-trackers, webcam eye-tracking, and mouse-tracking (Mormann et al., 2020; Orquin & Mueller Loose, 2013; S. Smith & Dolifka, 2026; Yang & Krajbich, 2021), which makes these predictions testable in the field (with relatively low cost).

Our study has some limitations. Our task used color-coded abstract values with small stakes and binary choices between options with two attributes, so whether the same pattern holds with larger stakes or numeric displays remains to be tested. The range manipulation was between participants and used two ranges. Our design does not separate between range,

[6] Perhaps this is what happens when medical treatments are presented with different frames (McNeil et al., 1982).

divisive, and rank normalization (using a skewed distribution of values could do this), but we use the different sensitivity to value across ranges evidence that values are encoded relative to context, and our focus is on attentional effects in aversive and ambivalent attributes. We estimate the reference point and loss aversion as group-level parameters, so as to focus on an average effect in a large sample rather than individual-level effects in a smaller sample with more trials. Heterogeneity between individuals remains to be studied. We treat attention as exogenous, but models of gaze allocation (Callaway et al., 2021; Gluth et al., 2024; Jang et al., 2021) can be combined with the account developed here. Because viewing was free, attention and choice may influence each other, so we use "attracts" and "repels" to describe correlational attentional effects on choice. Experiments that control exposure durations (Armel et al., 2008; Frömer et al., 2025; S. M. Smith & Spiller, 2026) could test the direction of causation, though these designs make it difficult to control for decisions made prior to the end of a trial.

To understand whether looking at an attribute will attract or repel choice, it is not enough to simply know the value of that attribute. We need to take into consideration the decision-maker's beliefs about that value. This means that attentional effects are set by the context, so that the same attribute, looked at for the same time, can attract choice in one context and repel in another. Beauty is not in the thing itself, but in how it stands against what surrounds it.

Transparency

Open Practices. Studies were preregistered at https://aspredicted.org/fs9cs5.pdf and https://aspredicted.org/ed5k6t.pdf. Code is available at https://github.com/BrendenEum/aversive-attributes. Data is available upon request.

Funding. The authors gratefully acknowledge funding from the Booth School of Business at the University of Chicago.

Declaration of Conflicting Interests. The authors declare no conflicts of interests.

Author Contributions. Brenden Eum: Conceptualization; Formal Analysis; Methodology; Validation; Visualization; Writing - original draft; Writing - review & editing. Stephanie Smith: Conceptualization; Funding Acquisition; Methodology; Data Curation; Investigation; Project Administration; Resources; Writing - review & editing.

ORCID. Brenden Eum https://orcid.org/0000-0002-5484-495X. Stephanie Smith. https://orcid.org/0000-0001-7064-9209.

Acknowledgments. We thank Romy Fromer for suggesting the Bayesian sequential sampling model in this paper.

References

Armel, C., Beaumel, A., & Rangel, A. (2008). Biasing simple choices by manipulating relative visual attention. *Judgment and Decision Making*, *3*(5), 396–403. https://doi.org/10.1017/S1930297500000413

Ashinoff, B. K., Buck, J., Woodford, M., & Horga, G. (2022). The effects of base rate neglect on sequential belief updating and real-world beliefs. *PLOS Computational Biology*, *18*(12), e1010796. https://doi.org/10.1371/journal.pcbi.1010796

Bedi, S., Hollander, G. de, & Ruff, C. C. (2025, September 11). *Probability weighting arises from boundary repulsions of cognitive noise*. https://doi.org/10.1101/2025.09.11.675565

Bhatnagar, R., & Orquin, J. L. (2022). A meta-analysis on the effect of visual attention on choice. *Journal of Experimental Psychology: General*, *151*(10), 2265–2283. https://doi.org/10.1037/xge0001204

Bürkner, P.-C. (2018). Advanced Bayesian Multilevel Modeling with the R Package brms. *The R Journal*, *10*(1), 395–411. https://doi.org/10.32614/RJ-2018-017

Busemeyer, J. R., Gluth, S., Rieskamp, J., & Turner, B. M. (2019). Cognitive and Neural Bases of Multi-Attribute, Multi-Alternative, Value-based Decisions. *Trends in Cognitive Sciences*, *23*(3), 251–263. https://doi.org/10.1016/j.tics.2018.12.003

Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, *100*(3), 432–459. https://doi.org/10.1037/0033-295X.100.3.432

Callaway, F., Rangel, A., & Griffiths, T. L. (2021). Fixation patterns in simple choice reflect optimal information sampling. *PLOS Computational Biology*, *17*(3), e1008863. https://doi.org/10.1371/journal.pcbi.1008863

Chetty, R., Looney, A., & Kroft, K. (2009). Salience and taxation: Theory and evidence. *American Economic Review*, *99*(4), 1145–1177. https://doi.org/10.1257/aer.99.4.1145

Cohen, A. L., Kang, N., & Leise, T. L. (2017). Multi-attribute, multi-alternative models of choice: Choice, reaction time, and process tracing. *Cognitive Psychology*, *98*, 45–72. https://doi.org/10.1016/j.cogpsych.2017.08.001

De Martino, B., Kumaran, D., Holt, B., & Dolan, R. J. (2009). The Neurobiology of Reference-Dependent Value Computation. *The Journal of Neuroscience*, *29*(12), 3833–3842. https://doi.org/10.1523/JNEUROSCI.4832-08.2009

Eum, B., Dolbier, S., & Rangel, A. (2023). Peripheral Visual Information Halves Attentional Choice Biases. *Psychological Science*, *34*(9), 984–998. https://doi.org/10.1177/09567976231184878

Eum, B., Gonzalez, S., & Rangel, A. (2026, August 11). *Attention and Range-Normalized Sequential Integration in Aversive Value-Based Choices*. 7280000. https://doi.org/10.2139/ssrn.7280000

Fisher, G. (2017). An attentional drift diffusion model over binary-attribute choice. *Cognition*, *168*, 34–45. https://doi.org/10.1016/j.cognition.2017.06.007

Fisher, G. (2021a). A multiattribute attentional drift diffusion model. *Organizational Behavior and Human Decision Processes*, *165*, 167–182. https://doi.org/10.1016/j.obhdp.2021.04.004

Fisher, G. (2021b). Intertemporal Choices Are Causally Influenced by Fluctuations in Visual Attention. *Management Science*, *67*(8), 4961–4981. https://doi.org/10.1287/mnsc.2020.3732

Frömer, R., Callaway, F., Griffiths, T. L., & Shenhav, A. (2025). Considering what we know and what we don't know: Expectations and confidence guide value integration in value-based decision-making. *Open Mind*, *9*, 791–813. https://doi.org/10.1162/opmi.a.3

Frömer, R., Dean Wolf, C. K., & Shenhav, A. (2019). Goal congruency dominates reward value in accounting for behavioral and neural correlates of value-based decision-making. *Nature Communications*, *10*(1), Article 1, 4926. https://doi.org/10.1038/s41467-019-12931-x

Gagne, C., Agai, S., Ramiro, C., Dayan, P., & Bishop, S. (2022). Biased belief priors versus biased belief updating: Differential correlates of depression and anxiety. *PLoS Computational Biology*, *18*(8), e1010176. https://doi.org/10.1371/journal.pcbi.1010176

Glickman, M., & Usher, M. (2019). Integration to boundary in decisions between numerical sequences. *Cognition*, *193*, 104022. https://doi.org/10.1016/j.cognition.2019.104022

Gluth, S., Deakin, J., & Rieskamp, J. (2024, November 27). *A theory of multi-attribute search and choice*. https://doi.org/10.31234/osf.io/3qzak

Gwinn, R., Leber, A. B., & Krajbich, I. (2019). The spillover effects of attentional learning on value-based choice. *Cognition*, *182*, 294–306. https://doi.org/10.1016/j.cognition.2018.10.012

Hare, T., Malmaud, J., & Rangel, A. (2011). Focusing Attention on the Health Aspects of Foods Changes Value Signals in vmPFC and Improves Dietary Choice. *Journal of Neuroscience*, *31*(30), 11077–11087. https://doi.org/10.1523/JNEUROSCI.6383-10.2011

Hawkins, G. E., Forstmann, B. U., Wagenmakers, E.-J., Ratcliff, R., & Brown, S. D. (2015). Revisiting the Evidence for Collapsing Boundaries and Urgency Signals in Perceptual Decision-Making. *The Journal of Neuroscience*, *35*(6), 2476–2484. https://doi.org/10.1523/JNEUROSCI.2410-14.2015

Jang, A. I., Sharma, R., & Drugowitsch, J. (2021). Optimal policy for attention-modulated decisions explains human fixation behavior. *eLife*, *10*, e63436. https://doi.org/10.7554/eLife.63436

Johnson, E. J., Shu, S. B., Dellaert, B. G. C., Fox, C., Goldstein, D. G., Häubl, G., Larrick, R. P., Payne, J. W., Peters, E., Schkade, D., Wansink, B., & Weber, E. U. (2012). Beyond nudges: Tools of a choice architecture. *Marketing Letters*, *23*(2), 487–504. https://doi.org/10.1007/s11002-012-9186-1

Krajbich, I., Armel, C., & Rangel, A. (2010). Visual fixations and the computation and comparison of value in simple choice. *Nature Neuroscience*, *13*(10), 1292–1298. https://doi.org/10.1038/nn.2635

Krajbich, I., Lu, D., Camerer, C., & Rangel, A. (2012). The Attentional Drift-Diffusion Model Extends to Simple Purchasing Decisions. *Frontiers in Psychology*, *3*. https://doi.org/10.3389/fpsyg.2012.00193

Krajbich, I., & Rangel, A. (2011). Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. *Proceedings of the National Academy of Sciences*, *108*(33), 13852–13857. https://doi.org/10.1073/pnas.1101328108

Lejarraga, T., Schulte-Mecklenbeck, M., Pachur, T., & Hertwig, R. (2019). The attention–aversion gap: How allocation of attention relates to loss aversion. *Evolution and Human Behavior*, *40*(5), 457–469. https://doi.org/10.1016/j.evolhumbehav.2019.05.008

Lim, S.-L., O'Doherty, J. P., & Rangel, A. (2011). The Decision Value Computations in the vmPFC and Striatum Use a Relative Value Code That is Guided by Visual Attention. *The Journal of Neuroscience*, *31*(37), 13214–13223. https://doi.org/10.1523/JNEUROSCI.1246-11.2011

Liu, S., Fengler, A., Frank, M. J., & Harrison, M. T. (2025, December 11). *Time-averaged drift approximations are inconsistent for inference in drift diffusion models*. arXiv: 2512.10250 [stat]. https://doi.org/10.48550/arXiv.2512.10250

Lohse, G. L., & Johnson, E. J. (1996). A Comparison of Two Process Tracing Methods for Choice Tasks. *Organizational Behavior and Human Decision Processes*, *68*(1), 28–43. https://doi.org/10.1006/obhd.1996.0087

Louie, K., Khaw, M. W., & Glimcher, P. W. (2013). Normalization is a general neural mechanism for context-dependent decision making. *Proceedings of the National Academy of Sciences*, *110*(15), 6139–6144. https://doi.org/10.1073/pnas.1217854110

March, J., & Gluth, S. (2025). The hungry lens: Hunger shifts attention and attribute weighting in dietary choice. *Elife*, *13*. https://doi.org/10.7554/eLife.103736.2

McNeil, B. J., Pauker, S. G., Sox, H. C., & Tversky, A. (1982). On the elicitation of preferences for alternative therapies. *New England Journal of Medicine*, *306*(21), 1259–1262. https://doi.org/10.1056/NEJM198205273062103

Mormann, M., Griffiths, T., Janiszewski, C., Russo, J. E., Aribarg, A., Ashby, N. J. S., Bagchi, R., Bhatia, S., Kovacheva, A., Meissner, M., & Mrkva, K. J. (2020). Time to pay attention to attention: Using attention-based process traces to better understand consumer decision-making. *Marketing Letters*, *31*(4), 381–392. https://doi.org/10.1007/s11002-020-09520-0

O'Donoghue, T., & Sprenger, C. (2018). Reference-Dependent Preferences. In *Handbook of Behavioral Economics: Applications and Foundations 1* (pp. 1–77, Vol. 1). Elsevier. https://doi.org/10.1016/bs.hesbe.2018.07.003

Oprea, R., & Woodford, M. (2026, September). *Behavioral signatures of cognitive noise*. 35714. https://doi.org/10.3386/w35714

Orquin, J. L., & Mueller Loose, S. (2013). Attention and choice: A review on eye movements in decision making. *Acta Psychologica*, *144*(1), 190–206. https://doi.org/10.1016/j.actpsy.2013.06.003

Padoa-Schioppa, C. (2009). Range-Adapting Representation of Economic Value in the Orbitofrontal Cortex. *Journal of Neuroscience*, *29*(44), 14004–14014. https://doi.org/10.1523/JNEUROSCI.3751-09.2009

Panjer, H. H. (1981). Recursive evaluation of a family of compound distributions. *ASTIN Bulletin*, *12*(1), 22–26. https://doi.org/10.1017/S0515036100006796

Pärnamets, P., Johansson, P., Hall, L., Balkenius, C., Spivey, M. J., & Richardson, D. C. (2015). Biasing moral decisions by exploiting the dynamics of eye gaze. *Proceedings of*

the National Academy of Sciences, *112*(13), 4170–4175. https://doi.org/10.1073/pnas.1415250112

Rangel, A., & Clithero, J. A. (2012). Value normalization in decision making: Theory and evidence. *Current Opinion in Neurobiology*, *22*(6), 970–981. https://doi.org/10.1016/j.conb.2012.07.011

Sepulveda, P., Usher, M., Davies, N., Benson, A. A., Ortoleva, P., & De Martino, B. (2020). Visual attention modulates the integration of goal-relevant evidence and not value. *eLife*, *9*, e60705. https://doi.org/10.7554/eLife.60705

Sheng, F., Ramakrishnan, A., Seok, D., Zhao, W. J., Thelaus, S., Cen, P., & Platt, M. L. (2020). Decomposing loss aversion from gaze allocation and pupil dilation. *Proceedings of the National Academy of Sciences*, *117*(21), 11356–11363. https://doi.org/10.1073/pnas.1919670117

Shi, K., Heffner, J., & Rutledge, R. (2026, June 30). *Goals remap values and not choices*. Retrieved October 26, 2026, from https://osf.io/preprints/psyarxiv/74mcw_v1/

Shimojo, S., Simion, C., Shimojo, E., & Scheier, C. (2003). Gaze bias both reflects and influences preference. *Nature Neuroscience*, *6*(12), 1317–1322. https://doi.org/10.1038/nn1150

Smith, S., & Dolifka, D. (2026, August 5). *Measuring attention in qualtrics: Method and validation*. 7656078. https://doi.org/10.2139/ssrn.7656078

Smith, S. M., & Krajbich, I. (2018). Attention and choice across domains. *Journal of Experimental Psychology: General*, *147*(12), 1810–1826. https://doi.org/10.1037/xge0000482

Smith, S. M., & Spiller, S. A. (2026). Consumers undervalue multi-option alternatives. *Journal of Consumer Research*, ucag024. https://doi.org/10.1093/jcr/ucag024

Smith, S. M., Spiller, S. A., & Krajbich, I. (2025). The role of visual attention in opportunity cost neglect and consideration. *Cognition*, *261*, 106145. https://doi.org/10.1016/j.cognition.2025.106145

Sokol-Hessner, P., Hsu, M., Curley, N. G., Delgado, M. R., Camerer, C. F., & Phelps, E. A. (2009). Thinking like a trader selectively reduces individuals' loss aversion. *Proceedings of the National Academy of Sciences*, *106*(13), 5035–5040. https://doi.org/10.1073/pnas.0806761106

Stillman, P. E., Krajbich, I., & Ferguson, M. J. (2020). Using dynamic monitoring of choices to predict and understand risk preferences. *Proceedings of the National Academy of Sciences*, *117*(50), 31738–31747. https://doi.org/10.1073/pnas.2010056117

Tajima, S., Drugowitsch, J., Patel, N., & Pouget, A. (2019). Optimal policy for multi-alternative decisions. *Nature Neuroscience*, *22*(9), Article 9, 1503–1511. https://doi.org/10.1038/s41593-019-0453-9

Tavares, G., Perona, P., & Rangel, A. (2017). The Attentional Drift Diffusion Model of Simple Perceptual Decision-Making. *Frontiers in Neuroscience*, *11*, 468. https://doi.org/10.3389/fnins.2017.00468

Thomas, A. W., Molter, F., Krajbich, I., Heekeren, H. R., & Mohr, P. N. C. (2019). Gaze bias differences capture individual choice behaviour. *Nature Human Behaviour*, *3*(6), 625–635. https://doi.org/10.1038/s41562-019-0584-8

Tom, S. M., Fox, C. R., Trepel, C., & Poldrack, R. A. (2007). The Neural Basis of Loss Aversion in Decision-Making Under Risk. *Science*, *315*(5811), 515–518. https://doi.org/10.1126/science.1134239

Tversky, A., & Kahneman, D. (1992). Advances in Prospect Theory: Cumulative Representation of Uncertainty. *Journal of Risk and Uncertainty*, *5*(4), 297–323. Retrieved August 2, 2023, from https://www.jstor.org/stable/41755005

Vittala, A., Wu, L., Yan, D., Liebers, D., Tell, E., Song, X., Dashti, D., Louie, K., Raio, C., Iosifescu, D. V., & Glimcher, P. (2025, April 23). *Decisional reference point pathology: A mechanism and marker for major depressive disorder in humans*. https://doi.org/10.1101/2025.04.18.649538

Voichek, G., & Novemsky, N. (2025). Positive contrast scope insensitivity. *Journal of Consumer Research*, *52*(1), 157–178. https://doi.org/10.1093/jcr/ucae052

Webb, R., Glimcher, P. W., & Louie, K. (2021). The Normalization of Consumer Valuations: Context-Dependent Preferences from Neurobiological Constraints. *Management Science*, *67*(1), 93–125. https://doi.org/10.1287/mnsc.2019.3536

Wisdom, J., Downs, J. S., & Loewenstein, G. (2010). Promoting healthy choices: Information versus convenience. *American Economic Journal: Applied Economics*, *2*(2), 164–178. https://doi.org/10.1257/app.2.2.164

Yang, X., & Krajbich, I. (2021). Webcam-based online eye-tracking for behavioral research. *Judgment and Decision Making*, *16*(6), 1485–1505. https://doi.org/10.1017/S1930297500008512

Yang, X., & Krajbich, I. (2023). A dynamic computational model of gaze and choice in multi-attribute decisions. *Psychological Review*, *130*(1), 52–70. https://doi.org/10.1037/rev0000350

Yechiam, E., & Hochman, G. (2013). Loss-aversion or loss-attention: The impact of losses on cognitive performance. *Cognitive Psychology*, *66*(2), 212–231. https://doi.org/10.1016/j.cogpsych.2012.12.001

Supplementary Material

A Computing Choice Probability

On each trial, there are two options, left and right ($i \in \{L, R\}$), each with two attributes ($j \in \{1, 2\}$), and the agent chooses left if the total utility of the left option's two attributes exceeds that of the right option's. The model assigns the choice a probability: the fraction of the samples the agent might have drawn on that trial for which the left total comes out higher. This probability enters the likelihood of every trial. Below we consider one attribute at a time. a_{ij} is the value of the attribute and t_{ij} the time it was attended to.

The posterior estimate of the attribute value is $m_{ij} = m_0 + \pi(t_{ij})(\bar{x}_{ij} - m_0)$. It is a weighted average of the reference point m_0 and the sample mean $\bar{x}_{ij}$, with the relative precision of the sampled information, $\pi(t_{ij})$, as the weight on the sample mean, where

$$\pi(t_{ij}) = \frac{t_{ij}/\sigma_x^2}{1/\sigma_p^2 + t_{ij}/\sigma_x^2}.$$

Here σ_x is the noise in the agent's sampling (the standard deviation of one second of sampling) and σ_p is the standard deviation of the agent's prior. The precision of a belief is one over its variance, so t_{ij}/σ_x^2 is the precision of the sample and $1/\sigma_p^2$ is the precision of the prior. The relative precision $\pi(t_{ij})$ is the share of the total precision that comes from the sample. It is between 0 and 1, and it grows with the time t_{ij} spent attending to the attribute: with little attention the estimate stays close to the reference point, and with more attention it moves toward the sample. The sample mean is normally distributed around the attribute's value, $\bar{x}_{ij} \sim \mathcal{N}(a_{ij}, \sigma_x^2/t_{ij})$, which is a bell curve centered on a_{ij} that gets narrower the longer the attribute is attended to. A weighted average of a constant and a normal variable is also normal, so the estimate is normally distributed. Let $X_{ij} = m_{ij} - m_0$ be the distance of the estimate from the reference point. Then

$$X_{ij} \sim \mathcal{N}(\eta_{ij}, \kappa_{ij}^2), \qquad \eta_{ij} = \pi(t_{ij})(a_{ij} - m_0), \qquad \kappa_{ij}^2 = \frac{\pi(t_{ij})^2 \sigma_x^2}{t_{ij}}.$$

Here η_{ij} is the average distance of the estimate from the reference point: on average, the estimate sits a fraction $\pi(t_{ij})$ of the way from the reference point to the attribute's value a_{ij}. If η_{ij} is positive, the estimate is on average above the reference point, and if η_{ij} is negative, it is on average below it. The quantity κ_{ij} is the standard deviation of the estimate, which is how much it varies from one sample to the next. The utility of the attribute relative to the reference point is X_{ij} when the estimate falls above the reference point and λX_{ij} when it falls below. Thus, the utility of one attribute has a known shape: a bell curve that is stretched by a factor λ below the reference point.

The utility of an option is the sum of its two attributes. When $\lambda = 1$, each attribute's utility is a bell curve, the sum of two bell curves is another bell curve, and the choice probability has a simple formula (the Probit). When $\lambda \neq 1$, each attribute's utility is a bell curve with a stretched lower half, and the sum of two such curves has no analytical solution. We therefore

compute its distribution numerically, in the spirit of Tavares et al. (2017). We cut the utility axis into narrow bins of equal width h (defined below) and, for each attribute, work out the probability that its utility falls in each bin. This probability is approximately the height of the attribute's utility curve at the middle of the bin times the width of the bin. The reference point sits at the edge between two bins, so the bend in the utility curve never falls inside a bin.

The two attributes of an option are sampled independently, so the probability that one attribute lands in one bin and the other lands in another bin is the product of their two bin probabilities. Assigning to each bin its middle value, the option's utility for that pair of bins is the sum of the two values. For example, suppose the bins have width 1, one attribute has a utility of 0, 1, or 2 with probabilities .2, .5 and .3, and the other has a utility of 0, 1 or 2 with probabilities .4, .4 and .2. The option's utility is 2 if the attributes land on $(0, 2)$, $(1, 1)$ or $(2, 0)$, so the probability that its utility is 2 is $.2 \times .2 + .5 \times .4 + .3 \times .4 = .36$. In general, many pairs of bins add up to the same value, and the probability that the option's utility falls in a given bin is the sum of the products over all such pairs (a "discrete convolution"). Doing this for every bin gives a histogram of the option's possible utilities. We build one histogram for the left option and one histogram for the right option on the same grid so that their bins line up. An attribute that was never attended to has an estimate equal to the reference point, so its utility is 0 and contributes nothing to the option's utility histogram.

To get the probability of choosing left, we compare the two histograms. For each bin the right option might be in, we multiply its probability by the probability that the left option is in a higher bin, and add these up over all bins. When both options are in the same bin, we count half of the bins, since their true utilities are equally likely to be on either side. We compute the probability of choosing right in the same way.

The accuracy of this computation depends on how narrow the bins are. We set the bin width to $h = \kappa_{\min}/2$, where $\kappa_{\min}$ is the smallest standard deviation κ_{ij} among the trial's attributes, so that the most precisely estimated attribute still spans several bins, and we extend each attribute's grid 7.5 standard deviations to either side of its mean, which leaves out a negligible amount of probability. We compute the probability with bins of width h, $h/2$ and $h/4$ and extrapolate from the three results to what infinitely narrow bins would give, a technique called Richardson extrapolation (Press, 2007). Writing p_h for the probability computed with bins of width h, the extrapolated choice probability for a trial is $(p_h - 20p_{h/2} + 64p_{h/4})/45$. These weights cancel out the two largest error terms.

The grid is used only when it is needed. When $\lambda = 1$, the difference between the options' totals is normally distributed, and the Probit $\Phi(\delta/\sqrt{v})$ described in the main text gives the exact choice probability. The same is true on trials where, for every attended attribute, the estimate is more than 8.3 standard deviations from the reference point. The estimate then almost surely stays on one side of the reference point (the chance of crossing is below 10^{-16}), meaning that each attribute's utility is very well approximated by a bell curve. On these trials, we also use the Probit to speed up computation.

B Parameter Recovery

We simulated 20 datasets of 260 participants each that use the same stimuli and dwell durations as the real experiments. For each dataset, we drew "true" values of:

- the population distribution of sampling noise for each range condition, $(\mu_{\text{Narrow}}, \tau_{\text{Narrow}})$ and $(\mu_{\text{Wide}}, \tau_{\text{Wide}})$
- the reference point for each range condition, q_{Narrow} and q_{Wide}
- the loss aversion parameter, λ

For each range condition, we drew μ uniformly between $\log 4$ and $\log 12$ (a median σ_x between 4 and 12 points per second) and τ uniformly between 0.3 and 0.8, then drew each participant's σ_x from that log-normal distribution. We drew q uniformly between -0.2 and 0.6 for each range condition, and λ log-uniformly between 0.8 and 3. We then simulated one choice for every trial from the model's choice probability and refit the model to the simulated choices with the same procedure we used for the real data.

The model itself does not estimate a σ_x for each participant, only the population distribution of σ_x in each range condition. We obtain a participant's σ_x after the fit by Bayes' rule, treating the fitted population distribution as the prior and the probability of the participant's choices given each possible value of σ_x as the likelihood, much like the participant-level random effects of a mixed-effects regression. A participant's recovered σ_x is the mean of their posterior given their simulated choices and the recovered population distribution. The correlations between true and recovered values are $r^q = 1.00$, $r^\mu = 1.00$, $r^\tau = 0.99$, and on the log scale $r^\lambda = 1.00$, $r^{\sigma_x}_{\text{Study 1}} = 0.84$, $r^{\sigma_x}_{\text{Study 2}} = 0.90$. Estimates of σ_x become less accurate as they get larger, which is a side-effect of noisier choices from noisier sampling.

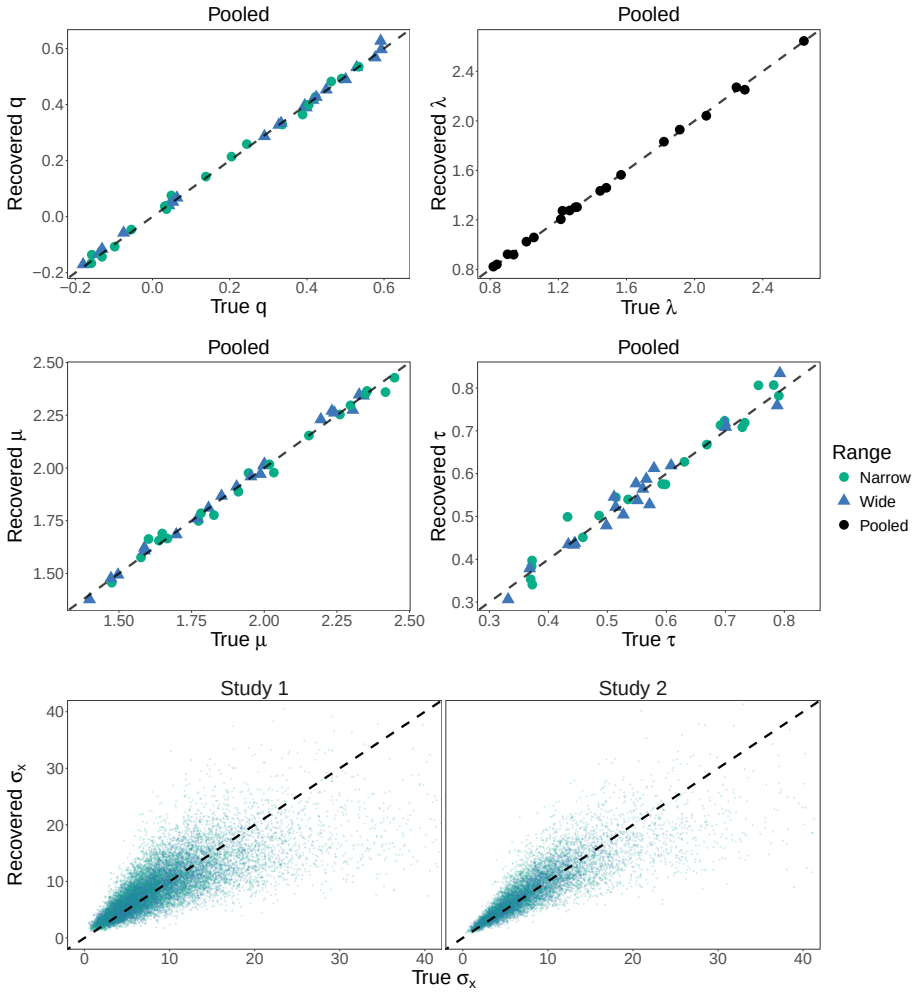


Figure S1: **Parameter Recovery.** Each point compares a parameter recovered by refitting the model to simulated choices (y-axis) with the true value used to simulate them (x-axis). Dashed lines mark perfect recovery. Color and shape show the range condition. (Top Left) Reference point q, with one point per range condition and dataset. (Top Right) Loss aversion λ, which is shared by both range conditions, with one point per dataset. (Middle Row) The population mean μ and standard deviation τ of $\log \sigma_x$, with one point per range condition and dataset. (Bottom Row) Sampling noise σ_x, with one point per participant and simulated dataset, separately for each study.

C Model Fitting Details

The group-level parameters are q_{Narrow}, q_{Wide}, λ, μ_{Narrow}, τ_{Narrow}, μ_{Wide}, and τ_{Wide}. Participants are independent given q and λ, so each condition's (μ, τ) can be estimated from that condition's participants alone at any candidate q and λ. The fitting process is carried out in three stages.

Stage 1: Coarse. The coarse grid for q is $\{-0.2, -0.1, 0, 0.1, \ldots, 0.7\}$. We fix $\lambda = 1$ for now so we can use the Probit for choice probability. Because q places m_0 as a fraction of the way from the smallest to the largest value in D, $q = 0$ and $q = 1$ are the two ends of the span, and $q = 0.5$ is its midpoint. For exploratory purposes, the grid reaches a fifth of a span below the smallest attribute value.[7] At each grid value, we estimate $(\mu_{\text{Narrow}}, \tau_{\text{Narrow}}, \mu_{\text{Wide}}, \tau_{\text{Wide}})$ and record the negative log-likelihood separately for the Narrow and Wide conditions.

Stage 2: Refine. For both range conditions, we refine q with a targeted search procedure over the interval from 0.3 below to 0.3 above that condition's best grid value. We search using `optimize` in R with the default tolerance (1.2×10^{-4}). We keep λ fixed at 1 for now so we can continue to use the Probit for choice probability. $(\mu_{\text{Narrow}}, \tau_{\text{Narrow}}, \mu_{\text{Wide}}, \tau_{\text{Wide}})$ are still re-estimated at every candidate q.

Stage 3: Optimize. Lastly, we run Nelder-Mead (`optim` in R) over all of the group-level parameters, starting from the two refined q's and $\lambda = 1$. We set relative tolerance to 10^{-7} and `maxit` to 200, which bounds the computation time.

In some places, the likelihood can be nearly flat for a group parameter (e.g., for λ when few posterior estimates fall below m_0). In a flat region, the simplex can wander dramatically toward values that do not reflect the fit to the data, which could potentially bias estimates. To avoid this, we impose bounds on the group-level parameters. However, Nelder-Mead has no built-in way to respect bounds, so we use a logistic transformation to keep the group parameters within reasonable intervals. Instead of running Nelder-Mead directly on q and λ, we run it on an unconstrained scale from negative infinity to infinity and then map this to the bounded region for each group-level parameter. The mapping is smooth and one-to-one, so the optimum on the unconstrained scale is the optimum on the original scale. This simply keeps the edges of the interval from being reached, so the simplex cannot drift away through a flat region. We bound q to $[-1, 2]$, from one full span below the smallest attribute value to one full span above the largest. We bound λ to $[0.25, 8]$, from strong gain-seeking (0.25) to eight-fold loss aversion (2). We apply a log transform to λ prior to search so that we can better search over values closer to 1.

Estimating the Distribution of Sampling Noise. Within a range condition, $\log \sigma_x$ is normally distributed across participants with mean μ and standard deviation τ. Participant

[7] Eum et al. (2026) found that for some individuals, option-level reference points fall at or below the minimum in a context.

s's contribution to the likelihood is then the probability of their choices averaged over the σ_x's they might have:

$$\text{Likelihood}_s = \int \Pr(\text{choices}_s \mid \sigma_x, q, \lambda)\, \mathcal{N}(\log \sigma_x;\, \mu, \tau^2)\, d \log \sigma_x$$

This is the marginal likelihood of a random-effects model. We numerically estimate this using a grid of σ_x values.

Grid Resolution. There are two grids used in Stage 3 of the fitting procedure. First, there is the grid of σ_x values over which each participant's likelihood is averaged, and second, there is the grid of utility values on which each trial's choice probability is computed. In both cases, we check that the grid is fine enough by refining it and recomputing the log-likelihood of the full data set at the fitted parameters of the main model. Doubling the σ_x grid from 60 to 120 values changes log-likelihoods by 2.0×10^{-5}. Halving the spacing of the utility grid changes log-likelihoods by 2.3×10^{-5}.

Priors. The prior standard deviation σ_p is not a free parameter. It is the standard deviation of the reference distribution D (i.e., the 18 attribute values of a range condition). In the per-type variant, D is instead the set of values that a given attribute type can take within a range condition, so m_0 and σ_p are specific to the attribute type and the range condition, while q still has one value per range condition. This is possible because m_0 is the true value of the reference point, while q is just the fraction that represents where in the range the reference point falls. The zero and flat priors fix q at 0.5, the midpoint of D. Because the attribute values in each range condition are symmetric around zero, the midpoint is exactly 0, so the zero and flat prior variants place the prior mean and the loss-aversion kink at zero. The flat prior additionally sets σ_p to 1,000 times the standard deviation of D, so the prior carries essentially no weight in the posterior. With q fixed in the zero and flat prior variants, the three-stage fitting process simplifies. Stage 1 is a single population step at $q = 0.5$, Stage 2 has nothing to refine, and Stage 3 is a bounded search over λ within the same bounds, again re-estimating $(\mu_{\text{Narrow}}, \tau_{\text{Narrow}}, \mu_{\text{Wide}}, \tau_{\text{Wide}})$ at every candidate.

Table S1 lists which group parameters are free and which are fixed in each variant.

Prior	D	q	σ_p	λ
Per condition (main text)	Range condition	Free per condition	SD of D	Free
Per condition	Range condition	Free per condition	SD of D	1 (fixed)
Per type	Type × condition	Free per condition	SD of D	Free
Zero	Range condition	0.5 (0 points)	SD of D	Free
Flat	Range condition	0.5 (0 points)	1,000× SD of D	Free

Table S1: Free and fixed group parameters in each model variant. D is the set of values the prior is calibrated to (the 18 values of a range condition, or the values of one attribute type within a range condition), q is the fraction of the way from the smallest to the largest value in D that m_0 falls, σ_p is the prior standard deviation, and λ is the degree of loss aversion.

Cross-Validation. Each participant's trials were cut into blocks of five. A participant whose trial count is not a multiple of five (due to trial exclusions) has a shorter last block. Fold k holds out the k-th block of every participant. Participants have at most 30 trials in Study 1 or 60 in Study 2. With 12 folds, folds 7 to 12 only hold out trials for Study 2 participants. In each fold, the model is fit from scratch to the training trials using the same procedure as above. Crucially, nothing from the full data fit influences the estimates in a fold. Although we clip choice probabilities to $[10^{-12}, 1 - 10^{-12}]$ before taking logs and summing so that no trial can produce an infinite log-likelihood, the clip never binds (0 of 58,122 trials fall below the floor or above the ceiling).

Convergence. In all cases (fit to the full data, all 12 folds of every variant, all 20 parameter-recovery datasets), the joint stage stopped on the tolerance criterion and not at the cap.

Software and Computation Time. Most of the fitting was done in R 4.5.0 with the base functions `optim` and `optimize`, with the choice probabilities computed in C++ through Rcpp, and with tidyverse (2.0.0) for data handling. Fits, folds, and recovery datasets were run on an Apple M2 laptop with 24 GB of RAM with 7 parallel processes. In total, the fits, folds, and recovery datasets took 13 hours. For the Bayesian regressions estimating attentional weights from the model variants' OOS simulations, we used the Nibi cluster of the Digital Research Alliance of Canada. Using R 4.4.0 with brms (2.23.0) and CmdStan (2.40.0), we submitted each OOS dataset from a variant as a separate job with 8 cores. Run in parallel, all jobs finished in a little over 5 hours.

D Regressions

We use lme4 for linear outcomes and glmmTMB for logistic outcomes for the analyses in Tables S2, S3, S4, and S5, because these models have few fixed effects and at most a handful of random slopes per participant, for which these methods are fast and stable. Each row in these tables is one predictor. Because every model crosses its predictors with range condition and study, the columns give the predictor's coefficient at the reference levels (Narrow, Study 1) and its interactions with the Wide condition, with Study 2, and with both, so the effect of a predictor in the Wide condition of Study 2 is the sum of its four columns. Values and value differences are in units of 10 points and dwell times are in seconds.

Table S2 reports frequentist mixed-effects regressions, pooled across studies, with study, range condition, and their interactions as fixed effects and subject-level random effects. Intervals are 95% Wald confidence intervals. Value Diff. is the value of the left option minus the value of the right option, and Abs. Value Diff. is the best option value minus the worst option value.

Table S3 reports a frequentist mixed-effects logistic regression with subject-level random slopes on the three value differences. Value Diff. (P), Value Diff. (N), and Value Diff. (B) are the left value minus right value for each attribute type.

Table S4 reports frequentist mixed-effects regressions. Dwell durations are modeled on the log scale. Type N and Type B compare aversive and ambivalent boxes with appetitive ones, and Value is the value of the looked-at box (Panel B) or of the box (Panel C).

Tables S5 reports frequentist mixed-effects logistic regressions. Value Diff. is the value of the left option minus the value of the right option in Panel A and the value of the last-seen option minus the value of the other option in Panel B. Net Dwell is the dwell time on the left minus the dwell time on the right option, and Last N and Last B compare a last dwell on an aversive or ambivalent box with a last dwell on an appetitive one.

Table S6 reports the Bayesian hierarchical logistic regressions fit with brms, one per range condition (the columns), each pooled over the two studies with a Study 2 term on the intercept. Panel A gives the utility of each attribute value relative to the lowest value of its type (a reference value for pinning down utility), and Panel B gives the attention weight of each cell per second of dwell. Every participant receives a random slope for each of the 15 utilities and each of the 18 cells. Weakly informative priors and partial pooling regularize the cell-level weights, so these are the only models we estimate in a Bayesian framework. Cells show the posterior median and 95% highest density interval.

Table S2: Basic Psychometrics.

	Narrow, Study 1	× Wide	× Study 2	× Wide × Study 2
Panel A: Choice				
Intercept	-0.020	0.023	0.004	-0.059
	[-0.061, 0.022]	[-0.037, 0.083]	[-0.054, 0.062]	[-0.143, 0.025]
Value Diff.	1.688	-0.433	-0.027	-0.120
	[1.615, 1.761]	[-0.527, -0.340]	[-0.135, 0.082]	[-0.262, 0.023]
Num.Obs.	59169			
R^2 Marg.	0.452			
R^2 Cond.	0.516			
Estimation	Frequentist (glmmTMB)			
Panel B: RT				
Intercept	8.584	0.326	-2.588	0.497
	[7.941, 9.227]	[-0.590, 1.243]	[-3.676, -1.500]	[-1.042, 2.035]
Abs. Value Diff.	-0.922	0.105	0.229	0.035
	[-1.085, -0.758]	[-0.095, 0.306]	[-0.005, 0.464]	[-0.255, 0.326]
Num.Obs.	59169			
R^2 Marg.	0.018			
R^2 Cond.	0.618			
Estimation	Frequentist (lme4)			

Note. Random-effect SDs across subjects. Panel A: Intercept 0.119; Value Diff. 0.492. Panel B: Intercept 7.120; Abs. Value Diff. 0.682; Residual 5.674.

Table S3: Range Sensitivity.

	Narrow, Study 1	× Wide	× Study 2	× Wide × Study 2
Intercept	-0.021	0.042	-0.002	-0.076
	[-0.065, 0.023]	[-0.023, 0.107]	[-0.064, 0.060]	[-0.166, 0.014]
Value Diff. (P)	2.122	-0.546	-0.001	-0.221
	[2.023, 2.221]	[-0.670, -0.421]	[-0.148, 0.147]	[-0.414, -0.029]
Value Diff. (N)	1.928	-0.460	0.021	-0.185
	[1.835, 2.022]	[-0.577, -0.342]	[-0.118, 0.160]	[-0.366, -0.004]
Value Diff. (B)	1.848	-0.286	-0.013	-0.168
	[1.764, 1.932]	[-0.398, -0.174]	[-0.138, 0.112]	[-0.339, 0.003]
Num.Obs.	59169			
R^2 Marg.	0.464			
R^2 Cond.	0.644			
Estimation	Frequentist (glmmTMB)			

Note. Random-effect SDs across subjects. Intercept 0.101; Value Diff. (P) 0.674; Value Diff. (N) 0.627; Value Diff. (B) 0.590.

Table S4: Dwell Properties.

	Narrow, Study 1	× Wide	× Study 2	× Wide × Study 2
Panel A: First-Seen Is Best				
Intercept	0.057	-0.049	-0.075	0.085
	[-0.024, 0.137]	[-0.154, 0.057]	[-0.188, 0.038]	[-0.062, 0.232]
Abs. Value Diff.	-0.056	0.066	0.081	-0.067
	[-0.128, 0.016]	[-0.016, 0.149]	[-0.019, 0.181]	[-0.182, 0.047]
Num.Obs.	47593			
R^2 Marg.	0.000			
R^2 Cond.	0.002			
Estimation	Frequentist (glmmTMB)			
Panel B: log Dwell Duration				
Intercept	5.901	0.052	-0.063	0.011
	[5.854, 5.948]	[-0.014, 0.118]	[-0.141, 0.015]	[-0.098, 0.120]
Type N	-0.014	-0.016	0.006	0.005
	[-0.043, 0.015]	[-0.054, 0.023]	[-0.036, 0.049]	[-0.050, 0.060]
Type B	0.013	-0.013	-0.019	0.021
	[-0.009, 0.036]	[-0.042, 0.017]	[-0.051, 0.014]	[-0.021, 0.064]
Value	-0.008	-0.003	-0.004	0.012
	[-0.018, 0.002]	[-0.015, 0.008]	[-0.019, 0.011]	[-0.005, 0.029]
Type N × Value	0.006	0.003	0.014	-0.019
	[-0.009, 0.021]	[-0.014, 0.020]	[-0.007, 0.035]	[-0.043, 0.005]
Type B × Value	0.001	0.009	0.018	-0.029
	[-0.012, 0.014]	[-0.006, 0.025]	[-0.001, 0.037]	[-0.051, -0.006]
Num.Obs.	627979			
R^2 Marg.	0.001			
R^2 Cond.	0.179			
Estimation	Frequentist (lme4)			
Panel C: Share of Dwells				
Intercept	0.2532	-0.0012	-0.0028	0.0041
	[0.2507, 0.2557]	[-0.0045, 0.0021]	[-0.0063, 0.0007]	[-0.0005, 0.0086]
Type N	-0.0025	0.0061	0.0032	-0.0067
	[-0.0061, 0.0010]	[0.0015, 0.0108]	[-0.0018, 0.0081]	[-0.0132, -0.0002]
Type B	-0.0043	0.0017	0.0035	-0.0050
	[-0.0070, -0.0016]	[-0.0019, 0.0053]	[-0.0003, 0.0073]	[-0.0100, 0.0000]
Value	0.0010	-0.0003	0.0014	-0.0017
	[-0.0002, 0.0023]	[-0.0017, 0.0011]	[-0.0004, 0.0031]	[-0.0037, 0.0003]
Type N × Value	0.0016	0.0011	-0.0010	0.0005
	[-0.0002, 0.0033]	[-0.0009, 0.0031]	[-0.0034, 0.0015]	[-0.0023, 0.0033]
Type B × Value	0.0031	-0.0018	-0.0026	0.0032
	[0.0015, 0.0047]	[-0.0037, 0.0001]	[-0.0048, -0.0004]	[0.0005, 0.0058]
Num.Obs.	232488			
R^2 Marg.	0.004			
R^2 Cond.				
Estimation	Frequentist (lme4)			

Note. Random-effect SDs across subjects. Panel A: Intercept 0.073; Abs. Value Diff. 0.000. Panel B: Intercept 0.480; Residual 1.031. Panel C: Intercept 0.000; Residual 0.075.

Table S5: Option-Level Attention.

	Narrow, Study 1	× Wide	× Study 2	× Wide × Study 2
Panel A: Choose Left				
Intercept	-0.014	0.030	0.003	-0.068
	[-0.056, 0.029]	[-0.031, 0.091]	[-0.058, 0.063]	[-0.154, 0.018]
Value Diff.	1.694	-0.442	-0.025	-0.100
	[1.620, 1.768]	[-0.537, -0.347]	[-0.135, 0.086]	[-0.245, 0.045]
Net Dwell	0.013	-0.009	0.011	0.010
	[0.003, 0.023]	[-0.023, 0.006]	[-0.009, 0.030]	[-0.016, 0.035]
Value Diff. × Net Dwell	0.008	0.006	-0.005	-0.018
	[-0.007, 0.024]	[-0.013, 0.025]	[-0.033, 0.023]	[-0.051, 0.014]
Num.Obs.	58122			
R^2 Marg.	0.456			
R^2 Cond.	0.520			
Estimation	Frequentist (glmmTMB)			
Panel B: Choose Last-Seen				
Intercept	1.012	-0.058	-0.134	0.212
	[0.920, 1.104]	[-0.191, 0.074]	[-0.269, 0.002]	[0.018, 0.406]
Value Diff.	1.746	-0.467	-0.015	-0.110
	[1.626, 1.866]	[-0.615, -0.319]	[-0.186, 0.155]	[-0.322, 0.101]
Last N	-0.006	0.015	0.142	-0.195
	[-0.119, 0.107]	[-0.148, 0.179]	[-0.015, 0.300]	[-0.422, 0.031]
Last B	0.075	-0.111	0.015	0.007
	[-0.038, 0.188]	[-0.273, 0.050]	[-0.141, 0.171]	[-0.215, 0.229]
Value Diff. × Last N	-0.119	0.003	-0.014	0.082
	[-0.275, 0.036]	[-0.184, 0.191]	[-0.230, 0.202]	[-0.176, 0.341]
Value Diff. × Last B	-0.033	0.000	-0.044	0.015
	[-0.188, 0.121]	[-0.187, 0.187]	[-0.256, 0.169]	[-0.241, 0.271]
Num.Obs.	58122			
R^2 Marg.	0.422			
R^2 Cond.	0.517			
Estimation	Frequentist (glmmTMB)			

Note. Random-effect SDs across subjects. Panel A: Intercept 0.133; Value Diff. 0.498. Panel B: Intercept 0.531; Value Diff. 0.479.

Table S6: Attribute-Level Attention.

	Narrow	Wide
Panel A: Utility of each value, relative to the lowest value of its type		
Intercept	-0.029 [-0.078, 0.020]	0.010 [-0.038, 0.060]
Study 2	0.012 [-0.053, 0.080]	-0.076 [-0.142, -0.009]
N, -50		0 (reference)
N, -40		1.086 [0.951, 1.236]
N, -30	0 (reference)	2.347 [2.198, 2.491]
N, -25	0.616 [0.476, 0.756]	
N, -20	1.933 [1.795, 2.078]	3.543 [3.380, 3.710]
N, -15	2.928 [2.758, 3.081]	
N, -10	3.724 [3.537, 3.913]	4.603 [4.402, 4.805]
N, -5	4.200 [3.990, 4.410]	5.020 [4.802, 5.226]
B, -20		0 (reference)
B, -15	0 (reference)	
B, -10	0.778 [0.634, 0.914]	1.309 [1.181, 1.431]
B, -5	1.584 [1.439, 1.736]	1.779 [1.645, 1.921]
B, 5	3.820 [3.642, 3.996]	3.317 [3.164, 3.485]
B, 10	4.428 [4.241, 4.624]	3.682 [3.515, 3.853]
B, 15	5.327 [5.116, 5.548]	
B, 20		5.215 [5.009, 5.425]
P, 5	0 (reference)	0 (reference)
P, 10	0.710 [0.567, 0.861]	0.283 [0.140, 0.427]
P, 15	1.600 [1.451, 1.741]	
P, 20	2.796 [2.648, 2.954]	1.877 [1.744, 2.024]
P, 25	4.223 [4.036, 4.422]	
P, 30	4.981 [4.758, 5.194]	3.089 [2.929, 3.232]
P, 40		4.391 [4.204, 4.579]
P, 50		5.156 [4.926, 5.392]
Panel B: Attention weight per second of dwell		
N, -50		-0.178 [-0.266, -0.091]
N, -40		-0.044 [-0.091, 0.000]
N, -30	-0.211 [-0.287, -0.134]	0.012 [-0.024, 0.048]
N, -25	-0.027 [-0.071, 0.016]	
N, -20	-0.006 [-0.046, 0.037]	0.031 [-0.006, 0.071]
N, -15	0.026 [-0.024, 0.078]	
N, -10	0.105 [0.040, 0.178]	0.122 [0.058, 0.188]
N, -5	0.222 [0.129, 0.315]	0.115 [0.062, 0.172]
B, -20		0.020 [-0.027, 0.067]
B, -15	-0.017 [-0.070, 0.039]	
B, -10	0.048 [-0.008, 0.100]	0.006 [-0.034, 0.045]
B, -5	0.052 [0.003, 0.110]	0.013 [-0.026, 0.058]
B, 5	0.099 [0.047, 0.157]	0.050 [0.011, 0.094]
B, 10	0.108 [0.045, 0.175]	0.034 [-0.007, 0.072]
B, 15	0.150 [0.090, 0.217]	
B, 20		0.045 [0.002, 0.090]
P, 5	0.005 [-0.068, 0.076]	-0.026 [-0.075, 0.023]
P, 10	0.009 [-0.042, 0.057]	0.009 [-0.037, 0.052]
P, 15	0.034 [-0.010, 0.074]	
P, 20	0.048 [0.012, 0.087]	-0.016 [-0.054, 0.020]
P, 25	0.088 [0.032, 0.152]	
P, 30	0.148 [0.072, 0.222]	0.022 [-0.012, 0.057]
P, 40		0.064 [0.013, 0.121]
P, 50		0.296 [0.201, 0.391]
Num.Obs.	29100	29022
R^2 Marg.	0.377	0.378
R^2 Cond.	0.461	0.452
Estimation	Bayesian (brms)	

Note. A blank cell is a value that does not occur in that range condition. Random-effect SDs across subjects (posterior medians). Narrow, utilities: N, -25 0.602; N, -20 0.315; N, -15 0.456; N, -10 0.552; N, -5 0.841; B, -10 0.401; B, -5 0.596; B, 5 0.555; B, 10 0.468; B, 15 0.759; P, 10 0.362; P, 15 0.132; P, 20 0.138; P, 25 0.683; P, 30 1.016. Narrow, weights: N, -30 0.275; N, -25 0.068; N, -20 0.072; N, -15 0.101; N, -10 0.172; N, -5 0.252; B, -15 0.164; B, -10 0.161; B, -5 0.092; B, 5 0.095; B, 10 0.200; B, 15 0.173; P, 5 0.220; P, 10 0.050; P, 15 0.059; P, 20 0.033; P, 25 0.123; P, 30 0.159. Wide, utilities: N, -40 0.400; N, -30 0.349; N, -20 0.128; N, -10 0.625; N, -5 0.798; B, -10 0.227; B, -5 0.416; B, 5 0.530; B, 10 0.527; B, 20 0.796; P, 10 0.550; P, 20 0.237; P, 30 0.334; P, 40 0.715; P, 50 1.079. Wide, weights: N, -50 0.427; N, -40 0.111; N, -30 0.064; N, -20 0.042; N, -10 0.150; N, -5 0.060; B, -20 0.111; B, -10 0.057; B, -5 0.058; B, 5 0.044; B, 10 0.027; B, 20 0.040; P, 5 0.070; P, 10 0.062; P, 20 0.035; P, 30 0.038; P, 40 0.085; P, 50 0.234.

E Range Sensitivity Simulations

We simulate choices using four alternative models of value normalization to show that predicted range sensitivity is higher in the Narrow condition than in the Wide condition across all models. Across all models, $f(a_{ij})$ is the normalized value of an attribute value a_{ij}, $i \in \{L, R\}$, and $j \in \{1, 2\}$.

Model 1: Range Normalization. Each value is expressed as its position in the range of its attribute type, from 0 at the worst value to 1 at the best (Eum et al., 2026; Padoa-Schioppa, 2009; Rangel & Clithero, 2014; Soltani et al., 2012):

$$f(a_{ij}) = \frac{a_{ij} - a_{\min}}{a_{\max} - a_{\min}}$$

where $a_{\min}$ and $a_{\max}$ are the smallest and largest values that its attribute type takes in the participant's range condition. The model has no free parameters.

Model 2: Context Divisive Normalization. Each value is divided by the typical magnitude of the values its attribute type takes in the participant's range condition (Khaw et al., 2017; Louie et al., 2013):

$$f(a_{ij}) = \frac{a_{ij}}{\xi + \gamma\,\overline{|a|}}$$

where $\overline{|a|}$ is the mean absolute value of the values that the attribute type takes in the participant's range condition, ξ is a semi-saturation constant that keeps the denominator away from zero, and γ weights the context. We set $\xi = 5$ and $\gamma = 1$.

Model 3: Trial Divisive Normalization. Each value is divided by the summed magnitude of the two values that its attribute takes on the current trial, one per option (Landry & Webb, 2021; Louie et al., 2013):

$$f(a_{Lj}) = \frac{a_{Lj}}{\xi + \gamma\,(|a_{Lj}| + |a_{Rj}|)}, \qquad f(a_{Rj}) = \frac{a_{Rj}}{\xi + \gamma\,(|a_{Lj}| + |a_{Rj}|)}$$

where a_{Lj} and a_{Rj} are the values that the attribute takes in the left and right options on that trial, and ξ and γ are the same constants as above, $\xi = 5$ and $\gamma = 1$.

Model 4: Rank Normalization. Each value is replaced by its rank among the values its attribute type can take, from the worst to the best, as in the Decision by Sampling (DbS) framework (Bhui & Gershman, 2018; Stewart et al., 2006):

$$f(a_{ij}) = \frac{\mathrm{rank}(a_{ij}) - 1}{b - 1}$$

where $\mathrm{rank}(a_{ij})$ is the rank of a_{ij} among the b values that its attribute type takes in the participant's range condition, from 1 at the worst value to b at the best. Here $b = 6$ in both conditions, so f runs from 0 to 1 in steps of 0.2. The model has no free parameters.

Simulating Choices. We use the trials from Study 2 to simulate two studies, one in the Narrow condition and one in the Wide condition. For each simulated study, we simulate

choices from 260 participants. To simulate choices, we first define an option's subjective value U as the sum of the normalized values of its two attributes. Then we use a logistic choice rule to generate choices:

$$\Pr(\text{Left}) = \frac{1}{1 + \exp\left[-\left(\theta_s \left(U_L - U_R\right)\right)\right]}$$

where θ_s is participant s's sensitivity to subjective value differences (an inverse temperature). Each participant's temperature was drawn once and reused for all four models. $\theta_s \sim U(1, 5)$. Critically, we do not let this temperature depend on the range condition, so any Narrow versus Wide difference in the simulated choices comes from the normalization. Choices were then drawn as Bernoulli trials with these probabilities. On each simulated data set, we then ran the same choice regression from Table S2 and report the interaction between value difference and range condition.

Figure S2 shows the results. In every model, choice is less sensitive to a given difference in points in the Wide condition than in the Narrow condition, and the interaction is significant at $p < .001$.

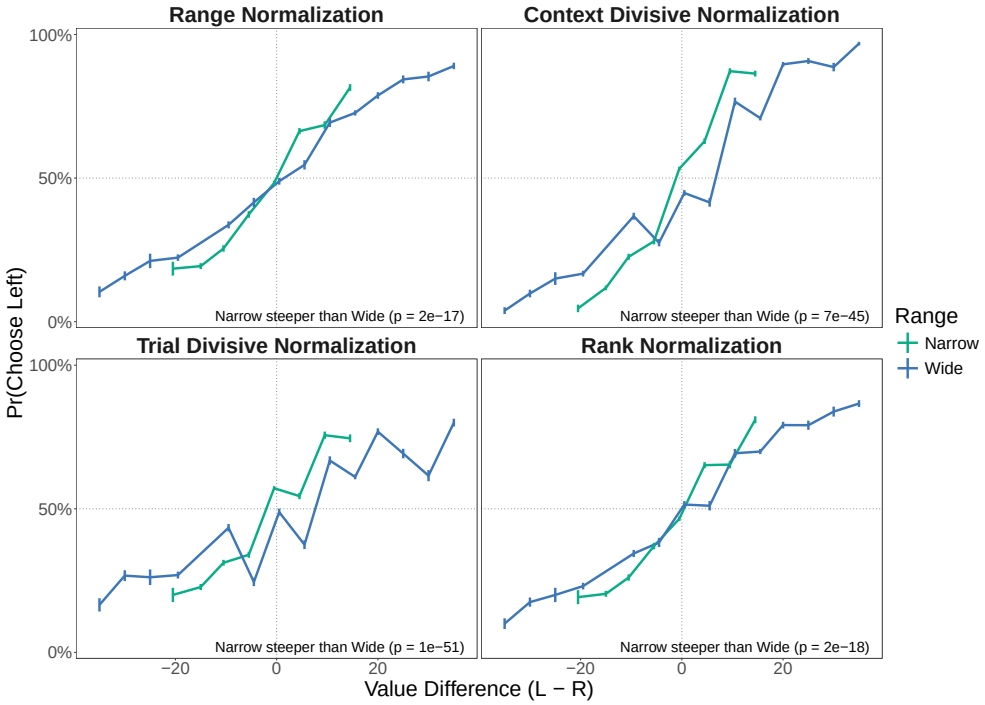


Figure S2: **Range Sensitivity Simulations.** Probability of choosing the left option, as a function of the left option value minus right option value, separately by range condition, for choices simulated from four models of value normalization (panels). Error bars depict standard error of the mean across simulated participant means.

F Range Sensitivity by Attribute Type

Figure S3 explores the sensitivity of attribute values to the range condition, without attentional effects.

To measure range sensitivity, we ran the same choice regression in Table S2, except that we added interactions between value difference, study, range condition, and attribute type. This gives us the slope of choosing the left option as a function of value difference, separately by study, range condition, and attribute type. Figure S3 plots these slopes, separately by study, range condition, and attribute type. Indeed, across all attribute types and studies, we find that value sensitivity was higher in the Narrow condition than in the Wide condition:

$$\Delta^{\text{Wide - Narrow},N}_{\text{Study 1}} = -0.460\ [-0.577, -0.342]$$
$$\Delta^{\text{Wide - Narrow},B}_{\text{Study 1}} = -0.286\ [-0.398, -0.174]$$
$$\Delta^{\text{Wide - Narrow},P}_{\text{Study 1}} = -0.546\ [-0.670, -0.421]$$
$$\Delta^{\text{Wide - Narrow},N}_{\text{Study 2}} = -0.645\ [-0.782, -0.507]$$
$$\Delta^{\text{Wide - Narrow},B}_{\text{Study 2}} = -0.454\ [-0.583, -0.325]$$
$$\Delta^{\text{Wide - Narrow},P}_{\text{Study 2}} = -0.767\ [-0.913, -0.621]$$

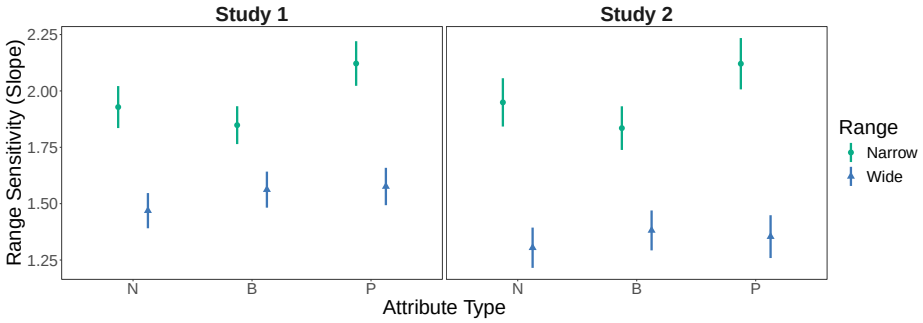


Figure S3: **Range Sensitivity by Range Condition and Attribute Type.** From a regression of choice on value difference × attribute type × study × range condition, we report the coefficient by attribute type, study, and range condition. Columns indicate study. Error bars depict standard error of the estimate.

G Bayesian Model Variants

Table S7 shows the group-level estimates and out-of-sample log-likelihoods of the main model and its variants. The main model reported in the paper, which has a Per-Condition Prior and Loss Aversion, is the best-fitting variant. It is statistically better than the Per-Type Prior Model, Zero Prior Model, and Flat Prior Model. While it is not significantly different from the Per-Condition Prior with No Loss Aversion Model at the 95% confidence level, it is at the 93% confidence level, despite the free loss aversion parameter λ being estimated close to one. The typical log-likelihood contributed by a single participant ranges from -14 to -30 ($M_{\text{Study 1}} = -14.72$, $SD_{\text{Study 1}} = 3.98$; $M_{\text{Study 2}} = -29.52$, $SD_{\text{Study 2}} = 5.87$), which means that the difference in log-likelihood between the main model and its variants are relatively small per-participant, but that over many participants these differences become statistically significant.

We also compute the out-of-sample mean absolute error (MAE) of the predicted attention weights, which serves as a more direct measure of how well the model variants capture the attentional patterns in the data. Similar to the out-of-sample log-likelihood results, the main model has the lowest MAE, though it is tied with the Per-Condition Prior with No Loss Aversion variant. Performance degrades for the Per-Type Prior Model, Zero Prior Model, and Flat Prior Model.

Table S7: Bayesian Model Variants.

Model	Range	σ_p^2(NP, B)	q	λ	exp(μ)	τ	OOS-LL	$\Delta^{\text{OOS-LL}}_{\text{Participants}}$ (p-value)	OOS-MAE Attn. Weights
Per-Condition Prior	Narrow	309	0.22	1.06	5.94	0.49	-29943	–	0.044
	Wide	712	0.19		8.78	0.49			
Per-Condition Prior No Loss Aversion	Narrow	309	0.23	1.00	6.00	0.50	-29963	-0.01 (0.070)	0.044
	Wide	712	0.22		8.76	0.49			
Per-Type Prior	Narrow	88, 140	0.17	1.12	6.02	0.50	-30184	-0.16 (3e-07)	0.074
	Wide	304, 210	0.15		8.76	0.50			
Zero Prior	Narrow	309	0.50	0.95	5.85	0.51	-30270	-0.22 (4e-08)	0.050
	Wide	712	0.50		8.56	0.50			
Flat Prior	Narrow	3.1×10^8	0.50	0.95	5.79	0.51	-30377	-0.29 (1e-09)	0.057
	Wide	7.1×10^8	0.50		8.53	0.50			

Note. σ_p^2(NP, B) jointly reports the prior variance for negative and positive attribute types, since by definition this was the same in these two attribute types. Then, it reports the prior variance for ambivalent attribute types, which had a wider range of possible values. For $\Delta^{\text{OOS-LL}}_{\text{Participants}}$, we took the out-of-sample log-likelihood of every participant for every variant, then compared the difference between a variant's log-likelihood and the best model's log-likelihood across all participants (the mean difference per participant, with the p-value of a paired t-test across participants in parentheses). OOS-MAE Attn. Weights is the mean, over the 36 cells (value × attribute type × range), of the absolute difference between the attentional weight estimated from a variant's out-of-sample simulated choices and the data's attentional weight.

Figure S4 shows the attentional weight predictions of the model variants. We plot the 95% HDI of the attentional weights from the data as vertical lines, which is the same across all

panels. Then we plot predictions of each model variant in a separate panel with lines.

The Per-Condition Prior with No Loss Aversion Model reproduces every feature that the main model does: the monotonic rise of the weights with value, the aversive crossover, the attentional repulsion at the lowest aversive values (to the same degree as the main model), the positive weights for ambivalent and appetitive attributes, and the lower value sensitivity of the Wide condition. This is because the main model estimates λ close to 1, so removing loss aversion leaves the reference point and the prior precision, which generate the attentional effects, essentially unchanged. The only qualitative difference is that the aversive crossover in the Wide condition sits slightly closer to zero than in the data.

The Per-Type Prior Model produces attentional weights that rise with value and cross zero inside each attribute's range. This model sees negative attentional weights at the lowest values for every attribute type, which we do not observe in the data except for aversive attributes. The model also overpredicts the sensitivity of attentional weights to attribute value for ambivalent attributes.

The Zero Prior Model reproduces the monotonic rise of the weights with value and the lower value sensitivity of the Wide condition, and it produces the deepest attentional repulsion effects at the lowest aversive values of any variant. This is because those values lie far below the reference. It misses two features of the data. First, attentional attraction to the highest aversive values is underpredicted. Second, the predicted attentional weights at the lowest ambivalent and appetitive values are negative, whereas the data's are at or near zero.

The Flat Prior Model (which essentially has no reference point) under-predicts the sensitivity of attentional weights to attribute value for all attribute types and fails to pick up on the strength of the different value sensitivities between range conditions. It also fails to reproduce the largely negative attentional weights seen in the data for extremely negative values within each range, despite allowing for loss aversion.

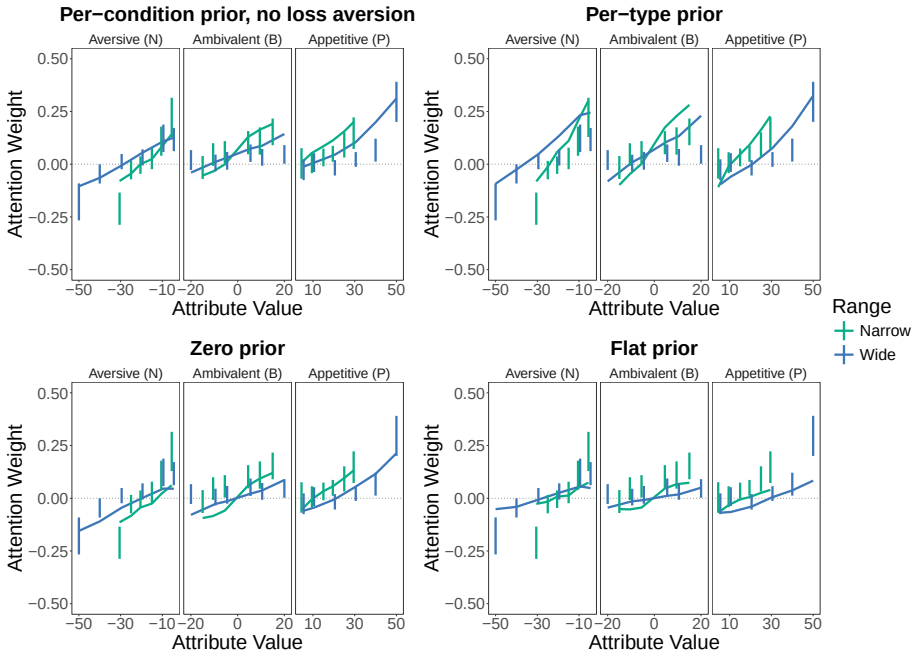


Figure S4: **Bayesian Model Variants.** Simulated, out-of-sample attention weights by attribute value, separately by model variants. Vertical bars indicate the 95% HDI of the attention weights from the data, which is the same across all panels. Lines indicate the predicted attention weights from each model variant. The variants are the Per-Condition Prior with No Loss Aversion variant (top left), Per-Type Prior variant (top right), the Zero Prior variant (bottom left), and the Flat Prior variant (bottom right). Colors indicate range condition.

H Attentional Weights by Study

The attentional weight regression in the main text pools Study 1 and 2. As a robustness check, we re-estimated it with every population-level coefficient (the intercept, the 15 utilities, and the 18 attentional weights) crossed with a Study 2 indicator (so that Study 1 is the reference level). Figure S5 shows the attentional weights of each study, and Tables S8 and S9 report the coefficients.

The Study 1 and Study 2 attentional weights largely agree. They exhibit a high correlation ($r = 0.77$, $p = 5e-08$). We find that only 5 of 36 attentional weights differ significantly between studies, but that all 5 are larger in Study 2 than in Study 1 ($\omega^{\text{Narrow}}_{\text{B,Study 2}-\text{Study 1}}(-15) = 0.187\ [0.071, 0.302]$, $\omega^{\text{Narrow}}_{\text{B,Study 2}-\text{Study 1}}(-10) = 0.114\ [0.004, 0.225]$, $\omega^{\text{Wide}}_{\text{N,Study 2}-\text{Study 1}}(-50) = 0.233\ [0.066, 0.435]$, $\omega^{\text{Wide}}_{\text{P,Study 2}-\text{Study 1}}(20) = 0.085\ [0.006, 0.164]$, $\omega^{\text{Wide}}_{\text{P,Study 2}-\text{Study 1}}(30) = 0.104\ [0.033, 0.188]$). This is slightly more than would be expected by chance (1.8 out of 36 using a 95% confidence interval), but still represents a high degree of agreement between the two studies. Averaged over all 18 cells, Study 2 weights are only slightly larger by about the same amount in both range conditions ($\bar{\omega}^{\text{Narrow}}_{\text{Study 2}-\text{Study 1}} = 0.031\ [-0.004, 0.066]$, $\bar{\omega}^{\text{Wide}}_{\text{Study 2}-\text{Study 1}} = 0.032\ [0.003, 0.062]$).

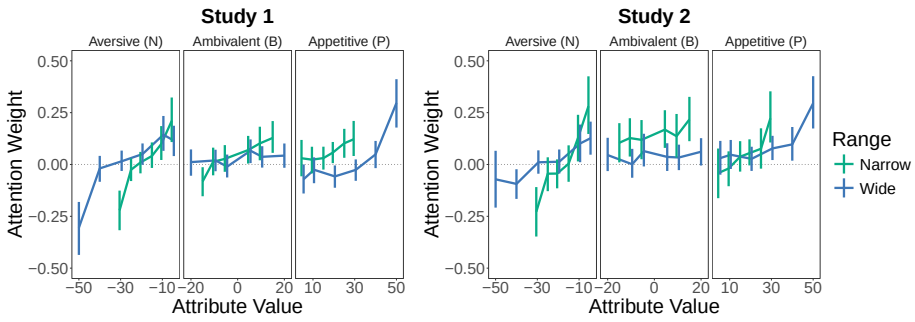


Figure S5: **Attentional Weights by Attribute Value and Study.** From the main-text Bayesian regression with every coefficient crossed with study. Columns indicate study. Colors indicate range condition. Error bars depict 95% highest density intervals.

Table S8: Attribute-Level Attention by Study, Narrow Condition.

	Study 1	× Study 2
Panel A: Utility of each value, relative to the lowest value of its type		
Intercept	-0.026 [-0.077, 0.025]	0.004 [-0.066, 0.069]
N, -30	0 (reference)	0 (reference)
N, -25	0.643 [0.452, 0.840]	-0.051 [-0.318, 0.224]
N, -20	1.943 [1.745, 2.139]	0.002 [-0.269, 0.265]
N, -15	2.917 [2.670, 3.134]	0.042 [-0.272, 0.347]
N, -10	3.813 [3.547, 4.058]	-0.159 [-0.494, 0.185]
N, -5	4.234 [3.940, 4.519]	-0.057 [-0.445, 0.331]
B, -15	0 (reference)	0 (reference)
B, -10	0.748 [0.560, 0.944]	0.089 [-0.189, 0.362]
B, -5	1.559 [1.360, 1.764]	0.089 [-0.209, 0.381]
B, 5	3.854 [3.617, 4.096]	0.001 [-0.319, 0.332]
B, 10	4.454 [4.199, 4.716]	0.034 [-0.306, 0.396]
B, 15	5.353 [5.045, 5.629]	0.034 [-0.367, 0.416]
P, 5	0 (reference)	0 (reference)
P, 10	0.746 [0.542, 0.954]	-0.065 [-0.360, 0.228]
P, 15	1.701 [1.490, 1.897]	-0.194 [-0.483, 0.081]
P, 20	2.797 [2.585, 3.012]	-0.011 [-0.285, 0.295]
P, 25	4.343 [4.077, 4.607]	-0.207 [-0.586, 0.125]
P, 30	5.162 [4.859, 5.454]	-0.362 [-0.770, 0.045]
Panel B: Attention weight per second of dwell		
N, -30	-0.221 [-0.317, -0.126]	-0.008 [-0.150, 0.143]
N, -25	-0.025 [-0.080, 0.030]	-0.019 [-0.118, 0.076]
N, -20	0.010 [-0.043, 0.061]	-0.054 [-0.138, 0.030]
N, -15	0.041 [-0.017, 0.106]	-0.037 [-0.136, 0.066]
N, -10	0.099 [0.022, 0.185]	0.029 [-0.098, 0.152]
N, -5	0.210 [0.108, 0.323]	0.070 [-0.084, 0.218]
B, -15	-0.082 [-0.153, -0.013]	0.187 [0.071, 0.302]
B, -10	0.013 [-0.052, 0.081]	0.114 [0.004, 0.225]
B, -5	0.027 [-0.035, 0.093]	0.092 [-0.016, 0.207]
B, 5	0.070 [0.004, 0.138]	0.098 [-0.012, 0.208]
B, 10	0.102 [0.020, 0.182]	0.033 [-0.092, 0.158]
B, 15	0.128 [0.056, 0.209]	0.089 [-0.035, 0.209]
P, 5	0.030 [-0.057, 0.118]	-0.071 [-0.222, 0.072]
P, 10	0.023 [-0.043, 0.086]	-0.043 [-0.146, 0.064]
P, 15	0.031 [-0.025, 0.086]	0.008 [-0.083, 0.096]
P, 20	0.057 [0.006, 0.109]	-0.002 [-0.080, 0.074]
P, 25	0.102 [0.032, 0.172]	-0.026 [-0.129, 0.087]
P, 30	0.122 [0.041, 0.210]	0.101 [-0.023, 0.234]
Num.Obs.	29100	
R^2 Marg.	0.379	
R^2 Cond.	0.464	
Estimation	Bayesian (brms)	

Note. Random-effect SDs across subjects (posterior medians). Utilities: N, -25 0.606; N, -20 0.301; N, -15 0.457; N, -10 0.557; N, -5 0.847; B, -10 0.397; B, -5 0.596; B, 5 0.558; B, 10 0.478; B, 15 0.772; P, 10 0.382; P, 15 0.136; P, 20 0.134; P, 25 0.684; P, 30 1.025. Weights: N, -30 0.294; N, -25 0.084; N, -20 0.070; N, -15 0.112; N, -10 0.184; N, -5 0.267; B, -15 0.167; B, -10 0.168; B, -5 0.125; B, 5 0.106; B, 10 0.208; B, 15 0.181; P, 5 0.233; P, 10 0.055; P, 15 0.068; P, 20 0.041; P, 25 0.130; P, 30 0.174.

Table S9: Attribute-Level Attention by Study, Wide Condition.

	Study 1	× Study 2
Panel A: Utility of each value, relative to the lowest value of its type		
Intercept	0.015 [-0.035, 0.066]	-0.085 [-0.153, -0.016]
N, -50	0 (reference)	0 (reference)
N, -40	0.846 [0.644, 1.064]	0.460 [0.199, 0.757]
N, -30	2.194 [1.968, 2.386]	0.302 [0.020, 0.586]
N, -20	3.466 [3.236, 3.711]	0.179 [-0.145, 0.490]
N, -10	4.582 [4.288, 4.866]	0.092 [-0.278, 0.475]
N, -5	5.089 [4.788, 5.385]	-0.079 [-0.482, 0.305]
B, -20	0 (reference)	0 (reference)
B, -10	1.359 [1.173, 1.548]	-0.056 [-0.316, 0.195]
B, -5	1.965 [1.754, 2.156]	-0.328 [-0.609, -0.066]
B, 5	3.475 [3.249, 3.702]	-0.224 [-0.523, 0.078]
B, 10	3.933 [3.695, 4.182]	-0.388 [-0.689, -0.061]
B, 20	5.488 [5.199, 5.779]	-0.407 [-0.779, -0.028]
P, 5	0 (reference)	0 (reference)
P, 10	0.314 [0.107, 0.523]	-0.033 [-0.322, 0.254]
P, 20	1.976 [1.774, 2.186]	-0.137 [-0.396, 0.146]
P, 30	3.312 [3.095, 3.545]	-0.357 [-0.663, -0.073]
P, 40	4.659 [4.396, 4.914]	-0.447 [-0.788, -0.102]
P, 50	5.459 [5.137, 5.786]	-0.491 [-0.918, -0.072]
Panel B: Attention weight per second of dwell		
N, -50	-0.305 [-0.436, -0.181]	0.233 [0.066, 0.435]
N, -40	-0.019 [-0.083, 0.041]	-0.076 [-0.164, 0.022]
N, -30	0.015 [-0.031, 0.067]	-0.005 [-0.078, 0.071]
N, -20	0.050 [-0.002, 0.101]	-0.037 [-0.115, 0.035]
N, -10	0.144 [0.066, 0.234]	-0.043 [-0.153, 0.073]
N, -5	0.115 [0.041, 0.187]	0.011 [-0.092, 0.115]
B, -20	0.011 [-0.054, 0.073]	0.034 [-0.066, 0.134]
B, -10	0.020 [-0.032, 0.071]	-0.018 [-0.103, 0.064]
B, -5	-0.010 [-0.063, 0.048]	0.074 [-0.011, 0.167]
B, 5	0.066 [0.007, 0.121]	-0.030 [-0.119, 0.055]
B, 10	0.037 [-0.015, 0.088]	-0.004 [-0.086, 0.077]
B, 20	0.043 [-0.016, 0.101]	0.019 [-0.067, 0.106]
P, 5	-0.071 [-0.140, 0.000]	0.102 [-0.005, 0.207]
P, 10	-0.025 [-0.090, 0.034]	0.073 [-0.017, 0.165]
P, 20	-0.058 [-0.112, -0.005]	0.085 [0.006, 0.164]
P, 30	-0.026 [-0.079, 0.024]	0.104 [0.033, 0.188]
P, 40	0.049 [-0.018, 0.114]	0.049 [-0.045, 0.143]
P, 50	0.296 [0.178, 0.411]	-0.001 [-0.152, 0.153]
Num.Obs.	29022	
R^2 Marg.	0.380	
R^2 Cond.	0.457	
Estimation	Bayesian (brms)	

Note. Random-effect SDs across subjects (posterior medians). Utilities: N, -40 0.402; N, -30 0.348; N, -20 0.132; N, -10 0.630; N, -5 0.807; B, -10 0.224; B, -5 0.425; B, 5 0.528; B, 10 0.523; B, 20 0.793; P, 10 0.558; P, 20 0.248; P, 30 0.333; P, 40 0.717; P, 50 1.097. Weights: N, -50 0.483; N, -40 0.126; N, -30 0.084; N, -20 0.039; N, -10 0.158; N, -5 0.060; B, -20 0.141; B, -10 0.090; B, -5 0.089; B, 5 0.049; B, 10 0.026; B, 20 0.044; P, 5 0.093; P, 10 0.068; P, 20 0.043; P, 30 0.057; P, 40 0.097; P, 50 0.233.

I Attention Weights from Pre-Registered Model

We pre-registered a different attention-weight regression from what is reported in the main text. It is a Bayesian hierarchical logistic regression of the choice on two kinds of regressors. The first is the signed value of each attribute type: the left minus the right value of the appetitive, aversive, and ambivalent attribute on that trial. The second is the signed dwell time to each of the 18 attribute-type-by-value cells: the dwell time to that box in seconds, positive if the box was in the left option and negative if it was in the right option. The coefficient on a cell's signed dwell time is its attention weight. The model was fit with Study 1 as the reference level and an interaction with Study 2 for every coefficient so that each study has its own weights, and every participant received random slopes on the three value channels and on all 18 cells. Figure S6 shows the attention weights by study. Tables S10 and S11 report the coefficients, where the columns are the Study 1 coefficient and its interaction with Study 2, Value Diff. (P), Value Diff. (N), and Value Diff. (B) are the left minus right value of each attribute type in units of 10 points, and cells show the posterior median and 95% highest density interval.

After pre-registration, we realized that this regression cannot separate the effect of dwell from re-encoded value effects because it is not allowed to assign different utilities to different raw values of that type. This is like saying a raw value of 5 must be worth exactly half of 10. If participants re-encode values, so that a raw value is worth more or less than what this fixed scale says, the model has no utility term to absorb the difference. The only term that could capture this re-encoding effect is the dwell-time predictor, so any re-encoding effect is attributed to dwell. This could result in an overestimation of attentional effects.

The regression in the main text therefore replaces the three value difference variables with one utility term per attribute value, so that re-encoding is captured by the utilities and each attention weight reflects only the effect of dwell.

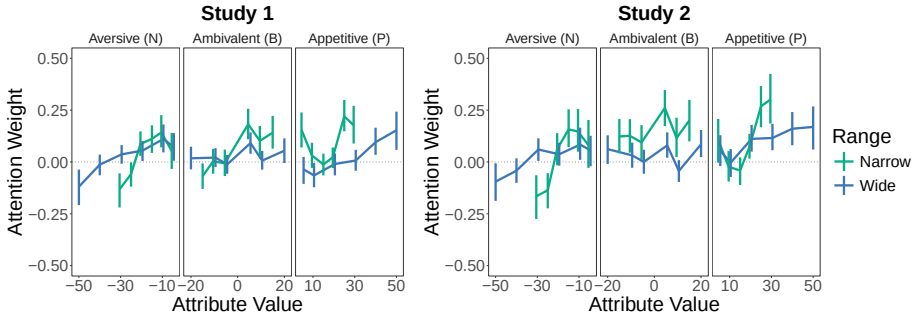


Figure S6: **Attention Weights by Attribute Value, from Pre-Registered Model.** Columns indicate study; colors indicate range condition. Error bars depict 95% highest density intervals.

Table S10: Attribute-Level Attention from the Pre-Registered Model, Narrow Condition.

	Study 1	× Study 2
Intercept	-0.029 [-0.079, 0.019]	-0.003 [-0.072, 0.065]
Value Diff. (P)	2.424 [2.281, 2.569]	-0.170 [-0.384, 0.036]
Value Diff. (N)	2.139 [2.000, 2.277]	-0.080 [-0.286, 0.114]
Value Diff. (B)	2.050 [1.936, 2.173]	-0.010 [-0.176, 0.165]
N, -30	-0.132 [-0.219, -0.055]	-0.034 [-0.157, 0.088]
N, -25	-0.058 [-0.119, -0.002]	-0.078 [-0.167, 0.022]
N, -20	0.086 [0.030, 0.147]	-0.025 [-0.113, 0.063]
N, -15	0.112 [0.044, 0.176]	0.046 [-0.054, 0.147]
N, -10	0.142 [0.071, 0.226]	0.003 [-0.113, 0.113]
N, -5	0.053 [-0.034, 0.141]	0.028 [-0.110, 0.153]
B, -15	-0.068 [-0.130, -0.007]	0.191 [0.087, 0.292]
B, -10	0.003 [-0.057, 0.064]	0.123 [0.028, 0.221]
B, -5	-0.006 [-0.068, 0.060]	0.101 [-0.003, 0.202]
B, 5	0.182 [0.110, 0.256]	0.079 [-0.028, 0.180]
B, 10	0.102 [0.025, 0.174]	0.013 [-0.097, 0.126]
B, 15	0.141 [0.064, 0.222]	0.059 [-0.057, 0.179]
P, 5	0.155 [0.070, 0.239]	-0.048 [-0.171, 0.087]
P, 10	0.032 [-0.028, 0.093]	-0.055 [-0.146, 0.041]
P, 15	-0.012 [-0.065, 0.040]	-0.030 [-0.110, 0.054]
P, 20	0.022 [-0.029, 0.070]	0.051 [-0.022, 0.126]
P, 25	0.220 [0.148, 0.299]	0.047 [-0.062, 0.157]
P, 30	0.177 [0.089, 0.271]	0.122 [-0.004, 0.257]
Num.Obs.	29100	
R^2 Marg.	0.392	
R^2 Cond.	0.470	
Estimation	Bayesian (brms)	

Note. Random-effect SDs across subjects (posterior medians). Value Diff. (P) 0.963; Value Diff. (N) 0.920; Value Diff. (B) 0.782; N, -30 0.187; N, -25 0.152; N, -20 0.160; N, -15 0.212; N, -10 0.201; N, -5 0.198; B, -15 0.088; B, -10 0.155; B, -5 0.216; B, 5 0.202; B, 10 0.215; B, 15 0.225; P, 5 0.201; P, 10 0.052; P, 15 0.089; P, 20 0.049; P, 25 0.243; P, 30 0.263.

Table S11: Attribute-Level Attention from the Pre-Registered Model, Wide Condition.

	Study 1	× Study 2
Intercept	0.024 [-0.024, 0.075]	-0.088 [-0.152, -0.019]
Value Diff. (P)	1.480 [1.399, 1.576]	-0.160 [-0.285, -0.036]
Value Diff. (N)	1.380 [1.298, 1.464]	-0.107 [-0.224, 0.004]
Value Diff. (B)	1.529 [1.439, 1.618]	-0.131 [-0.252, 0.000]
N, -50	-0.119 [-0.208, -0.037]	0.026 [-0.090, 0.146]
N, -40	-0.013 [-0.064, 0.036]	-0.029 [-0.108, 0.045]
N, -30	0.037 [-0.009, 0.082]	0.022 [-0.045, 0.094]
N, -20	0.055 [0.005, 0.101]	-0.018 [-0.088, 0.050]
N, -10	0.111 [0.046, 0.181]	-0.027 [-0.125, 0.068]
N, -5	0.072 [0.006, 0.139]	-0.017 [-0.112, 0.082]
B, -20	0.018 [-0.036, 0.073]	0.044 [-0.039, 0.130]
B, -10	0.021 [-0.025, 0.066]	0.010 [-0.056, 0.086]
B, -5	-0.012 [-0.057, 0.033]	0.012 [-0.060, 0.083]
B, 5	0.091 [0.040, 0.143]	-0.011 [-0.088, 0.066]
B, 10	0.007 [-0.037, 0.053]	-0.050 [-0.116, 0.022]
B, 20	0.054 [-0.005, 0.114]	0.031 [-0.051, 0.116]
P, 5	-0.036 [-0.106, 0.025]	0.089 [-0.009, 0.184]
P, 10	-0.065 [-0.122, -0.005]	0.062 [-0.023, 0.150]
P, 20	-0.010 [-0.065, 0.041]	0.122 [0.038, 0.198]
P, 30	0.007 [-0.042, 0.059]	0.109 [0.030, 0.180]
P, 40	0.095 [0.035, 0.165]	0.066 [-0.025, 0.159]
P, 50	0.153 [0.058, 0.243]	0.017 [-0.110, 0.153]
Num.Obs.	29022	
R^2 Marg.	0.395	
R^2 Cond.	0.458	
Estimation	Bayesian (brms)	

Note. Random-effect SDs across subjects (posterior medians). Value Diff. (P) 0.603; Value Diff. (N) 0.553; Value Diff. (B) 0.562; N, -50 0.204; N, -40 0.058; N, -30 0.102; N, -20 0.047; N, -10 0.151; N, -5 0.076; B, -20 0.074; B, -10 0.066; B, -5 0.076; B, 5 0.119; B, 10 0.028; B, 20 0.070; P, 5 0.069; P, 10 0.068; P, 20 0.112; P, 30 0.120; P, 40 0.172; P, 50 0.261.

Supplementary References

Bhui, R., & Gershman, S. J. (2018). Decision by sampling implements efficient coding of psychoeconomic functions. *Psychological Review*, *125*(6), 985–1001. https://doi.org/10.1037/rev0000123

Eum, B., Gonzalez, S., & Rangel, A. (2026, August 11). *Attention and Range-Normalized Sequential Integration in Aversive Value-Based Choices*. 7280000. https://doi.org/10.2139/ssrn.7280000

Khaw, M. W., Glimcher, P. W., & Louie, K. (2017). Normalized value coding explains dynamic adaptation in the human valuation process. *Proceedings of the National Academy of Sciences*, *114*(48), 12696–12701. https://doi.org/10.1073/pnas.1715293114

Landry, P., & Webb, R. (2021). Pairwise normalization: A neuroeconomic theory of multi-attribute choice. *Journal of Economic Theory*, *193*, 105221. https://doi.org/10.1016/j.jet.2021.105221

Louie, K., Khaw, M. W., & Glimcher, P. W. (2013). Normalization is a general neural mechanism for context-dependent decision making. *Proceedings of the National Academy of Sciences*, *110*(15), 6139–6144. https://doi.org/10.1073/pnas.1217854110

Padoa-Schioppa, C. (2009). Range-Adapting Representation of Economic Value in the Orbitofrontal Cortex. *Journal of Neuroscience*, *29*(44), 14004–14014. https://doi.org/10.1523/JNEUROSCI.3751-09.2009

Press, W. H. (Ed.). (2007). *Numerical recipes: The art of scientific computing* (3rd ed.). Cambridge University Press.

Rangel, A., & Clithero, J. A. (2014). The Computation of Stimulus Values in Simple Choice. In *Neuroeconomics* (pp. 125–148). Elsevier. https://doi.org/10.1016/B978-0-12-416008-8.00008-5

Soltani, A., De Martino, B., & Camerer, C. (2012). A Range-Normalization Model of Context-Dependent Choice: A New Model and Evidence (L. T. Maloney, Ed.). *PLoS Computational Biology*, *8*(7), e1002607. https://doi.org/10.1371/journal.pcbi.1002607

Stewart, N., Chater, N., & Brown, G. D. (2006). Decision by sampling. *Cognitive Psychology*, *53*(1), 1–26. https://doi.org/10.1016/j.cogpsych.2005.10.003

Tavares, G., Perona, P., & Rangel, A. (2017). The Attentional Drift Diffusion Model of Simple Perceptual Decision-Making. *Frontiers in Neuroscience*, *11*, 468. https://doi.org/10.3389/fnins.2017.00468