Supplemental Material

Experiments 1, 2, and 3 were preregistered (https://osf.io/8n3kd) and materials and data are available at https://osf.io/pczb3.

Sensitivity Analysis

For graphical illustration, we fit the parameters of the model (β and θ) by minimizing the sum of squared errors between the model predictions and aggregate participant responses. However, all of the qualitative predictions that we emphasize in the main text are insensitive to the exact parameters. That is, for all "reasonable" (defined in the following paragraph) parameter values, the direction of the key predictions (also defined below) is the same as in the main plots. This is illustrated in Figure S2.

To demonstrate this, we first identified a subset of the parameter space that produced "reasonable" behavior in terms of the probability of choosing the higher valued option and range of likely response times. Specifically, this was defined as satisfying two conditions: first, the probability of choosing the higher valued option (marginalizing over each option's value) was required to be between 55% and 95%; second, the probability of a response time less than 50 milliseconds or more than 120 seconds both had to be less than 1%. We identified a set of reasonable parameters by grid search, expanding the size of the grid until further expansion yielded no additional reasonable parameters (Figure S1).

Then, for each of the reasonable parameters, we checked whether the following properties held.

- For Experiment 1 (Figure 3), the inferred preference decreased monotonically with response time.
- For Experiment 2 (Figure 6), all inferred likelihoods were greater than 50%.
- For Experiment 3 (Figure 9), inferred preference was higher for the short response time (across both thresholds), inferred preference was higher for the higher threshold (across

both response times), and the difference in inferred preference between short and long response time was greater for high threshold than for low threshold (a positive interaction).

We found that all properties held for every parameter value we checked (Figure S2) with one exception: 3 out of 530 parameter configurations did not show the interaction in Experiment 3. Thus overall, this analysis provides strong empirical evidence that our qualitative model predictions are truly *predictions* that result from the structure of the model, rather than post hoc explanations of patterns observed in the data.

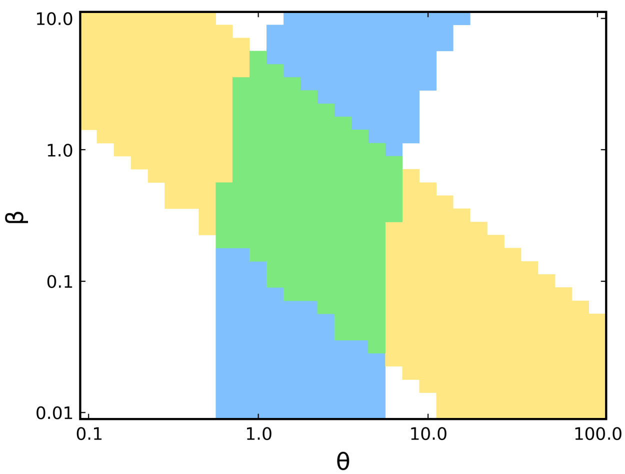

Figure S1. Illustration of reasonable parameters for the drift rate multiplier β and threshold θ. The yellow and blue regions indicate combinations that satisfy the accuracy and response time constraints, respectively. The sensitivity analysis was applied to all combinations in the green area.

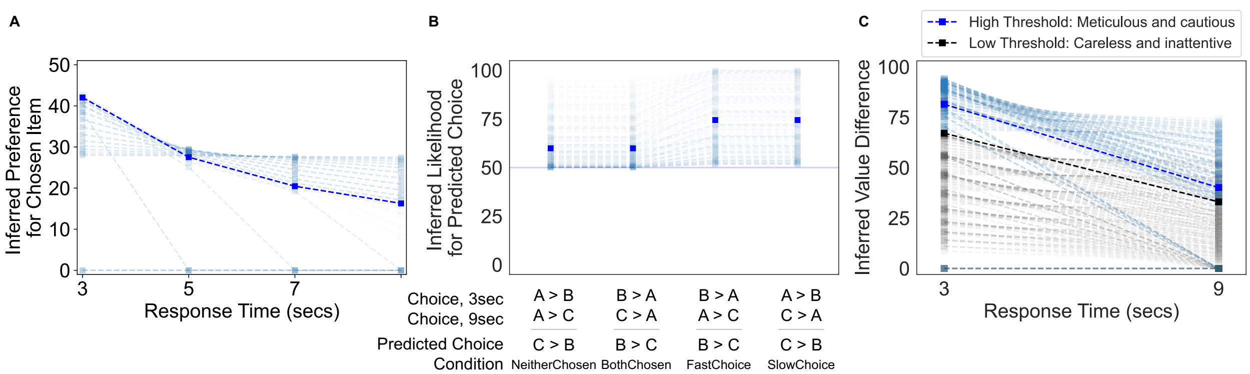

Figure S2. Model Predictions with Varied Parameters for (A) Experiment 1, (B) Experiment 2, and (C) Experiment 3. Light lines indicate the predictions made with varied parameters; dark points show the predictions that were fit to the empirical data (and are shown in the main text).

Random Starting Points

As alluded to in the General Discussion, one common extension of the so-called "vanilla" DDM that we used to produce our main results is the addition of a non-zero starting point that is randomly sampled on each trial/decision. Typically the starting point is assumed to follow a uniform distribution centered on zero. This results in one new free parameter, A, which is the maximum absolute value of the starting point. That is, on each trial the starting point is sampled uniformly from the range $(-A, A)$. We did not include such a parameter in our main model for two reasons. First, it was not necessary to explain our results. Second, including a random starting point poses a considerable computational challenge: the quadruple integral in Equation 9 is already very computationally intensive to evaluate. Adding an additional dimension (the starting point) makes the problem intractable without resorting to lower fidelity and more complex methods such as Markov chain Monte Carlo (MCMC).

To illustrate how the starting point can affect the predictions of our model, without addressing the aforementioned computational challenges, we focused on Experiment 1, which investigated the basic relationship between response time and inferred preference. Figure S3 shows the predictions of the DDM with no random starting point ("Vanilla DDM" in Figure S3), the maximal range of starting points ("DDM with fixed large starting point range"), and a range fit to the Experiment 1 results ("DDM with fitted starting point range"). The main results show that inferred preference is a monotonically decreasing function of response time, a pattern that holds across all reasonable parameter values (see Sensitivity Analysis). However, when the starting point range is sufficiently large ("DDM with fixed large starting point range"), we see that inferred preference is actually lower for very fast response times (less than a second) than for moderately fast times (two to three seconds). This is due to the "explaining away" effect mentioned in the General Discussion: Very fast decisions are best explained by the starting point being sampled close to a threshold, which makes a strong preference unnecessary to explain the observed data. When

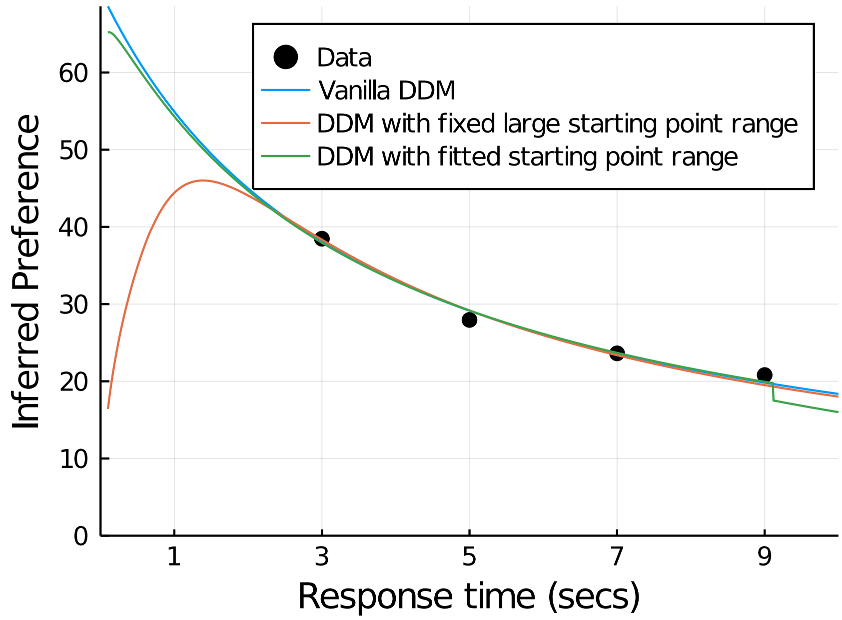

Figure S3. Illustration of the effect of starting points on model predictions. The blue line shows the predictions from the vanilla DDM with parameters fit to Experiment 1 only. The orange line shows a modification of the DDM with the same β and θ parameters, but with the starting point range fixed to $A = \theta$ (the highest allowable value). The green line shows the DDM with a random starting point, with all three parameters fit to Experiment 1 only.

we fit the starting point to the Experiment 1 data ("DDM with fitted starting point range"), we do not see this pattern because the best-fitting range is not sufficiently large ($\theta = 0.9614, \beta = 0.3706, A = 0.1088$). Moreover, it does not provide a substantially better fit to the data than the vanilla DDM fit to this experiment alone ($\theta = 1.589, \beta = 0.3769$). However, the predictions of the model with highly variable random starting points only diverge from the vanilla DDM for faster response times than we considered. Thus, further experimentation would be necessary to provide a strong test of whether people's inferences are more consistent with the model with or without random starting points.

Linear Ballistic Accumulator

To demonstrate that our findings are not specific to the DDM, and also to further investigate a possible role for random starting points in our data, we created an alternate version of the inference model in which the DDM is replaced by a linear ballistic accumulator (LBA). Briefly, the LBA is a sequential sampling model like the DDM, but it assumes that the evidence accumulates at a constant rate (i.e., linearly). We refer the reader to original paper (Brown & Heathcote, 2008) for a full description of this model. Here, we define

$$p_{\text{LBA}}(a \succ b, t \mid u_a - u_b) = \text{PDF}_1(t) \tag{1}$$

where $\text{PDF}_1(t)$ is defined in Equation 3 of Brown and Heathcote (2008). It is the probability of the accumulator for option a reaching the evidence threshold at time t, before the accumulator for b has reached the threshold. This function implicitly depends on several parameters. Most importantly, the drift rates for each accumulator are drawn from Normal distributions with means v_i. Following Trueblood et al. (2014), we define the average drift rate of each item relative to the other items and include both an intercept and slope. In the binary, single-attribute case, this reduces to

$$v_a = I_0 + \beta(u_a - u_b) \tag{2}$$

$$v_b = I_0 + \beta(u_b - u_a) \tag{3}$$

This has the convenient property of making the predictions of the model sensitive only to the difference $u_a - u_b$, allowing us to use the same equations for inference given the substitution of p_{LBA} for p_{DDM}. Note that the perhaps more intuitive definition of v_a based on absolute (rather than relative) values produces unreasonable predictions (e.g. similar response times for choosing between two very good options versus one very good and one very bad option). We implemented the model in Julia based on a publicly available MATLAB implementation: https://github.com/smfleming/LBA.

In total the LBA has four free parameters: the base accumulation rate I_0, the effect of relative value on accumulation β, the maximum starting point A, and the threshold b. We fit these parameters using the same procedure as for the DDM except that we sampled 100,000 points rather than 1,000 to account for the two additional free parameters. The best-fitting parameters were ($I_0 = 1.004, \beta = 0.7993, A = 13.13, b = 13.42$).

As illustrated in Figure S4, the LBA produces a very similar preference curve as the DDM across the range of response times used in the experiment (although it predicts weaker preferences for very fast reaction times, likely due to the random stopping rule as discussed below). Accordingly, Table S1 shows that the LBA and DDM provide a very similar account of the data in Experiments 1 and 2. The LBA was able to achieve a closer quantitative fit to the data in Experiment 3, primarily because it could capture the relative insensitivity to response time in this experiment. Importantly, all parameters except the two thresholds were separately fit to Experiments 1 and 2. Thus, the LBA and DDM had the same number of free parameters (two thresholds) to fit this specific behavioral pattern, and the ability of the LBA to better capture this data cannot be explained by overfitting. That being said, lacking a clear qualitative difference, we do not draw any strong conclusions about whether people's inferences from response times are more consistent with the DDM versus the LBA.

The LBA also allows us to conduct an additional test of random starting points because it can efficiently make predictions for all three experiments (unlike the DDM with random starting points, for which Equation 9 is intractable). As shown in Figure S4 the fitted LBA does indeed demonstrate the non-monotonic preference curve made possible by random starting points. However, overall, we found that the LBA was not able to capture any qualitative trends that the DDM could not, which suggests that random starting points cannot explain the key deviations we saw between people and the vanilla DDM model. A stronger test of the role of random starting points would of course be to compare the full LBA with a lesioned version without random starting points. Unfortunately, this is not possible because random starting points appear to be essential for the LBA to make reasonable predictions; we were not able to find any parameter configurations that were "reasonable" by our definition above without a non-zero value of A.

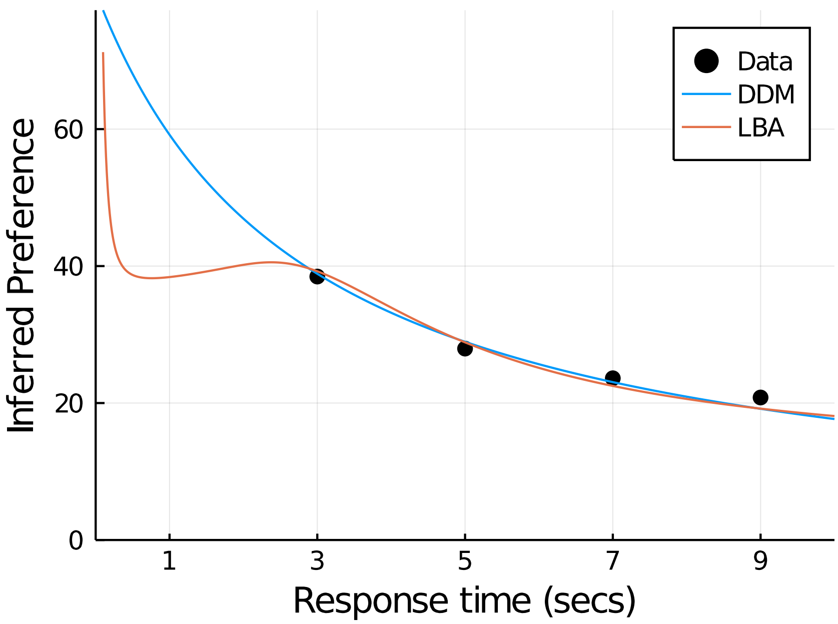

Figure S4. Illustration of the inferred preference curves of the LBA and DDM with parameters fit to Experiments 1 and 2.

Table S1

Comparing LBA and DDM model predictions.

	LBA	DDM	Mean Data
Experiment 1			
3 seconds	39.19	38.88	38.47
5 seconds	28.36	28.95	27.95
7 seconds	23.45	23.06	23.62
9 seconds	18.90	19.16	20.81
Experiment 2			
NeitherChosen	59.44	59.90	57.99
BothChosen	59.44	59.90	64.53
FastChoice	74.14	74.44	75.17
SlowChoice	74.14	74.44	72.34
Experiment 3			
Low Threshold, 3 seconds	60.17	67.09	56.77
Low Threshold, 9 seconds	34.98	33.06	54.01
High Threshold, 3 seconds	73.55	81.44	69.58
High Threshold, 9 seconds	37.11	40.13	64.20

Note. LBA and DDM model predictions, alongside the mean empirical data points from Experiments 1, 2, and 3.

Experiment 2, Additional Analysis

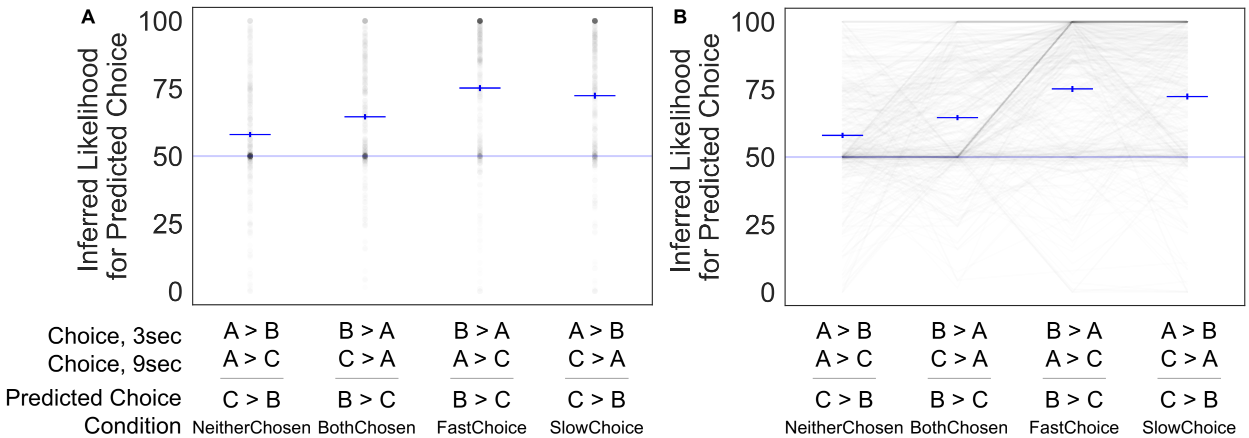

Figure S5. Experiment 2 Results, Additional Analysis. (A) Main Figure (see Figure 6). (B) Individual (Connected) Participant Responses. Note the prevalence of responses following a "50 - 50 - 100 - 100" pattern for the conditions "NeitherChosen - BothChosen - FastChoice - SlowChoice."

Future work may also explore some quirks of the empirical results. In Figure 6, the results indicate that while participants did not often select the non-predicted choice, very many participants chose 50 for the timing-inference based conditions, and 100 for the control conditions. (These numbers indicate participants' inferred likelihood for the predicted choice; the critical question participants were asked was: "This person is now offered a choice between items B and C. What choice do you think they'd make, and how likely do you think it is that they'd make that choice?") This pattern indicates that participants were not using response times when making their choices, and that choices alone could make them estimate 100% likelihoods for the decision-maker's choice on an unseen choice. Figure S5 suggests that a number of participants answered in this "50-50-100-100" sequence (for conditions "NeitherChosen-BothChosen-FastChoice-SlowChoice"), implying that if such responses had been discouraged the effect observed here may have been even stronger.

Experiment 3, Excluded Participants

The analysis was identical to Experiment 3 except that participants were excluded if they answered more than two out of eight manipulation check questions incorrectly, as was defined in the preregistration. Of 481 recruited participants, 212 met this exclusion criteria, so 269 participants were included. These results were similar to the Experiment 3 results with all participants included, and are described below (Figure S6).

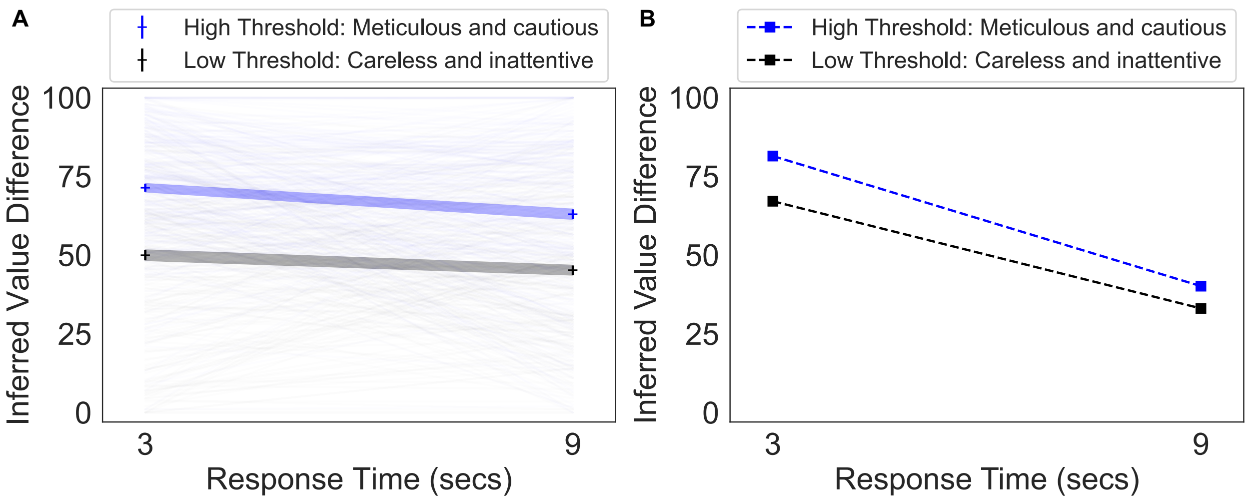

Figure S6. Experiment 3 Results, with Excluded Participants. (A) Experimental Results. The mean ± SE of the inferred value difference, n=269, is shown for the high threshold ("cautious," blue) and the low threshold ("careless," black) conditions, for each of the 3 and 9 second response time conditions. Responses are connected by threshold condition for emphasis. Individual participants' averaged values for each threshold/time pair are also shown and connected. 0 represents the participant feeling that the decision-maker was neutral between items, 50 that the decision-maker moderately preferred their chosen item, and 100 that the decision-maker strongly preferred their chosen item. (B) Model Predictions.

From our linear mixed model, the estimate (beta parameter) for the main effect of response time was 6.58 (95%CI=[4.53,8.62], $t(1880.0)$=6.31, p=4e-10, t-tests using Satterthwaite's method, η_p^2=0.02). The estimate for the main effect of threshold was 19.59 (95%CI=[17.54,21.63], $t(1880.0)$=18.70, p<2e-16, t-tests using Satterthwaite's method,

η_p^2=0.16). The estimate for the interaction effect of response time and threshold was 3.65 (95%CI=[-0.43,7.74], t(1880.0)=1.75, p=0.08, t-tests using Satterthwaite's method, η_p^2=0.002): not a significant effect. The estimated power of the interaction effect for α=0.01 was 20.00% (95%CI=[14.69,26.22]) using a Type-III F-test from the R package "car," generated via 200 simulations with the R package "simr." (The estimate for the fixed effect intercept was 57.37, 95%CI=[55.67,59.08], t(268.0)=66.15, p<2e-16, t-tests using Satterthwaite's method, η_p^2=0.94.)

References

Brown, S. D., & Heathcote, A. (2008). The simplest complete model of choice response time: Linear ballistic accumulation. *Cognitive Psychology*, *57*(3), 153–178. https://doi.org/10.1016/j.cogpsych.2007.12.002

Trueblood, J. S., Brown, S. D., & Heathcote, A. (2014). The multiattribute linear ballistic accumulator model of context effects in multialternative choice. *Psychological Review*, *121*(2), 179. https://doi.org/10.1037/a0036137