Value Certainty in Diffusion Decision Models

Douglas Lee[1] Marius Usher[2]

[1] California Institute of Technology, Division of the Humanities and Social Sciences

[2] Tel Aviv University, School of Psychological Sciences and Sagol School of Neuroscience

E-mail for correspondence:

DouglasGLee@gmail.com

Keywords:

Word count:

Abstract

Decision models such as the drift-diffusion model (DDM) are a widely used and broadly accepted tool that accounts remarkably well for binary choices and their response time distributions, as a function of the option values. The DDM is built on an evidence accumulation to bound concept, where a decision maker repeatedly samples a mental representation of the values of the options on offer until satisfied that there is enough evidence in favor of one option over the other. The value estimates that drive the DDM evidence are derived from the relative strength of value signals that are not stable across time, so that repeated sequential samples are necessary to average out noise. The standard DDM, however, typically does not allow for different options to have different levels of variability in their value representations. However, recent value-based decision studies have shown that a decision maker often reports levels of certainty regarding value estimates that vary across options. We thus propose that future versions of DDM should include an option-specific value certainty component. We present four different versions of such a model and validate them against empirical data from four previous studies. The data show that a model built around a sort of signal-to-noise ratio for each option (rather than a pure signal that randomly fluctuates) performs best, accounting for the positive impact of value certainty on choice consistency and the negative impact of value certainty on response time.

The drift-diffusion model (DDM) is ubiquitous in the contemporary literature on decision making, including research spanning the fields of psychology, neuroscience, economics, and consumer behavior. The DDM is a parsimonious mechanism that explains normative decisions by averaging out noise in information processing and implementing an optimal stopping rule (optimizing response time (RT) for a specified accuracy). This model accounts well for the dependency of the RT distribution on the values of the choice options, and for the speed-accuracy tradeoff—the often observed phenomenon that faster choices are typically less accurate (see Ratcliff et al, 2016; Gold and Shadlen, 2007 for reviews). While initially used in the domain of perceptual decision-making, the DDM has since become a widespread tool in studies of preferential (i.e., value-based) decision-making as well (Busemeyer et al, 2019; Tajima, Drugowitsch, & Pouget, 2016; Polania et al, 2015; Philiastides & Ratcliff, 2013; Milosavljevic et al, 2010; Krajbich, Armel, & Rangel, 2010; Basten et al, 2010). In brief, in its application to preferential choice, the DDM assumes that the value of each option is represented by a probability distribution whose mean is the true value of the option, and whose variance corresponds to processing noise. This noise can be interpreted as imprecision in the value representations themselves, or as a stochastic distortion of the value signals as they are relayed through the decision system by populations of neurons (whose firing patterns are known to be stochastic). Either way, the momentary signal about the relative value of the options (the so-called evidence for one option over the other) fluctuates randomly around a fixed trajectory (the so-called drift). In order to average out this processing noise, evidence signals are thought to accumulate over time until a sufficient amount of evidence has been acquired to allow for a choice to be made. The DDM thus includes thresholds for each option (typically symmetric) that trigger a choice once reached by the evidence accumulator. So, the fundamental components of the drift-diffusion process are the drift rate (the difference in the option values), the diffusion coefficient (the degree of stochasticity of the system), and the choice boundaries (the minimum required evidence threshold; see Figure 1).

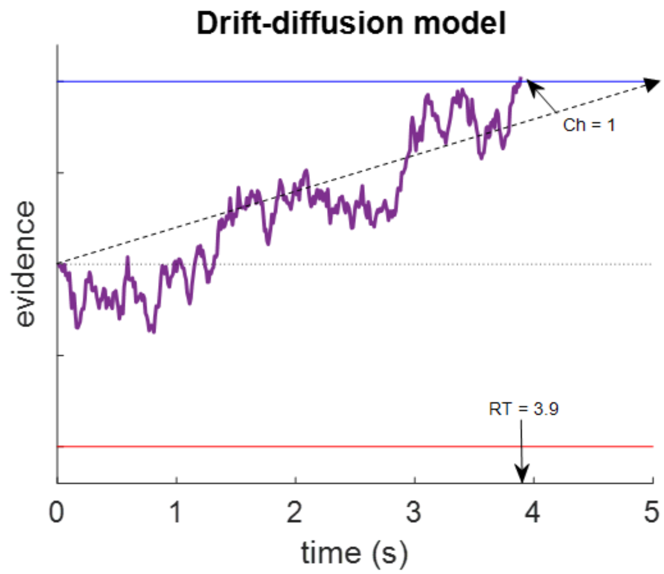

Figure 1: An illustration of the classic basic DDM. Evidence accumulates across time, following a fixed drift trajectory (black dashed arrow, corresponding to the value difference of choice options), corrupted by white processing noise. Here, the accumulated evidence reaches the upper threshold after 3.9 seconds, and a choice for option 1 is recorded.

The standard DDM implicitly assumes that processing noise is independent of the identity of the options contained in a particular choice set. That is to say, there is no option-specific noise in the DDM (but see Ratcliff, Voskuilen, & Teodorescu, 2018; Teodorescu, Moran, & Usher, 2016 for DDM variants in which the noise increases with task difficulty). Most of the DDM applications to preference-based decisions have assumed no option-specific noise. Since the brain encodes not only the subjective value of options, but also the subjective certainty about the value (Lebreton et al, 2015), it is possible to suggest that the representations of value that the brain uses to inform the decision process fluctuate, and that the degree of imprecision (or uncertainty) is not the same for all choice options. Indeed, it has been shown that decision makers hold highly variable beliefs about the certainty of their value estimates for different options, and that those beliefs are relatively stable within individuals (Lee & Coricelli, 2020; Lee & Daunizeau, 2020a, 2020b; Gwinn & Krajbich, 2020; Polania, Woodford, & Ruff, 2019). It has further been shown that the variability in value (un)certainty has a clear impact on both choice and RT (Lee & Coricelli, 2020; Lee & Daunizeau, 2020a, 2020b). Specifically, judged value certainty, C, correlates positively with choice consistency and negatively with RT (see Figure 2). This provides a qualitative benchmark that any DDM variant that includes option-specific uncertainty should be able to account for.

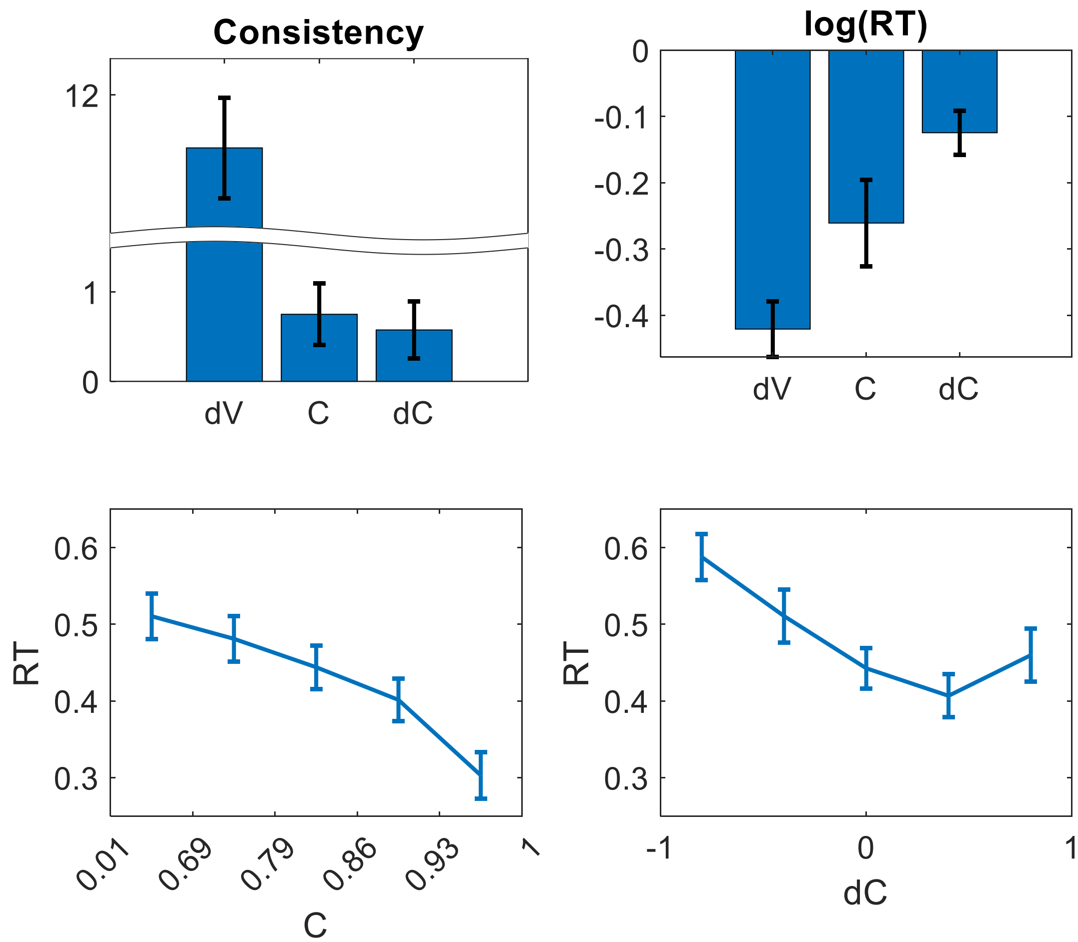

Figure 2: Previous results (pooled across four studies, n=152) demonstrate the relationships between value difference (dV), value certainty (C= .5 (c1+c2)), and certainty difference (dC= c1 –c2) with choice consistency and log(RT). Here we show the beta weights from a logistic regression on choice (upper left) and a linear regression on log(RT) (upper right; bar heights represent population means, error bars represent s.e.m.). Note: dV, C, and dC were all simultaneously included as independent variables in each regression model. Lower plots show the isolated relationships between log(RT) and C (left) and C1 - C2 (right). Note: the distribution of C was highly skewed, so we show the data binned by quantiles. (Lee & Daunizeau, 2020a, 2020b; Lee & Coricelli, 2020; Gwinn & Krajbich, 2020)

As we show below, the most straightforward way to include option-specific noise in the preferential DDM – by assuming that noise increases with value uncertainty – leads to the wrong qualitative predictions, with regards to the RT x certainty benchmark. In particular, as certainty increases, noise decreases, resulting in lower RT. Thus, unlike what we see in experimental data, such a model would predict that people would speed up when they are less certain of the options' values (all else equal). Moreover, this prediction is not specific to the standard DDM, but applies to the broader class of evidence accumulation-to-bound models (e.g., independent accumulators (Vickers, 1970; Brown & Heathcote, 2006); leaky competing accumulator (LCA; Usher & McLelland, 2001)), which also predict that higher noise in the system will result in *faster* responses, in direct contrast to the empirical data.

The aim of this paper is to examine a number of DDM variants, which could potentially rise to the challenge of accounting for the impact of value certainty on choice and RT in behavioral data. In particular, we will first present a variety of derivations of the standard DDM, each of which incorporates the concept of option-specific value certainty in a unique and realistic way, starting from the default DDM that has no option-specific noise, and progressing to signal-to-noise type of models. We then fit each of these models to experimental data from a variety of empirical datasets. Finally, we quantitatively compare the performance of each model across the different datasets and suggest the best recommended approach for future studies to incorporate option-specific value certainty in models derived from the DDM. To anticipate our results, we find that a signal-to-noise DDM variant provides the best fit to the data and accounts for all qualitative benchmarks.

METHODS

Computational models

In each of the models described below, we consider decisions between two alternatives, with values, v1 and v2, and with uncertainties, σ1, σ2, respectively. Evidence of these values is integrated over deliberation time, subject to noise. The evidence accumulator for each decision is initialized at a neutral starting point (i.e., zero evidence or default bias in favor of either option), and evolves until reaching one of two symmetric boundaries (see Figure 1 above). For each decision, the output of the model is: which boundary is reached, or the choice (ch = {0, 1}) and the number of integration time steps elapsed when that boundary is reached (RT).

Model 1

As a baseline default model for comparison (without any option-specific certainty term), we first consider the classic basic DDM. In this model, the equations that govern the evidence accumulation process are:

$$E_{t+1} = E_t + d\Delta$$

$$\Delta \sim N(\mu_1 - \mu_2, \sigma^2)$$

where E represents the cumulative balance of evidence that one option is better than the other, t represents deliberation time, Δ represents the incremental evidence in favor of one option over the other, d is a scalar, μ_i is the value estimate of option i, and σ^2 represents processing noise in the evidence accumulation and/or comparator systems. In a standard basic DDM, choice probability and expected response time can be analytically calculated using the following equations (Alos-Ferrer, 2018):

$$p(ch = 1) = \frac{1}{1 + e^{\left(\frac{-2\theta d}{\sigma^2}(\mu_1 - \mu_2)\right)}}$$

$$RT = \frac{\theta(2 * p(ch = 1) - 1)}{d(\mu_1 - \mu_2)}$$

where θ is the height of the evidence accumulation threshold where a decision is triggered (shown here as the upper threshold, for a choice of option 1), p(ch=1) is the probability that the upper threshold will be reached (rather than the lower threshold), and RT is the expected time at which the accumulation process will end (by reaching one of the thresholds)[1]. Because this system of equations is over-parameterized, one of the free parameters must be fixed for a

[1] In the standard version of the DDM, the RT distribution for correct and incorrect responses is identical. In a more complex version, additional variability parameters are introduced that allow to account for asymmetries between the RT distributions of correct and incorrect responses (see Ratcliff & McKoon, 2008 for review). We only consider the standard DDM without the variability parameters, as those cannot change the impact of value certainty on accuracy and RT illustrated in Figures 2 and 3.

practical application of the equations. In this work, we will fix the threshold θ to a value of 1 when fitting the models, for simplicity. Choice probability and RT will thus be functions of the drift rate d and the noise σ^2 (see Figure 3). As expected, accuracy increases and RT decreases with the drift. The noise, on the other hand, reduces both accuracy and RT.

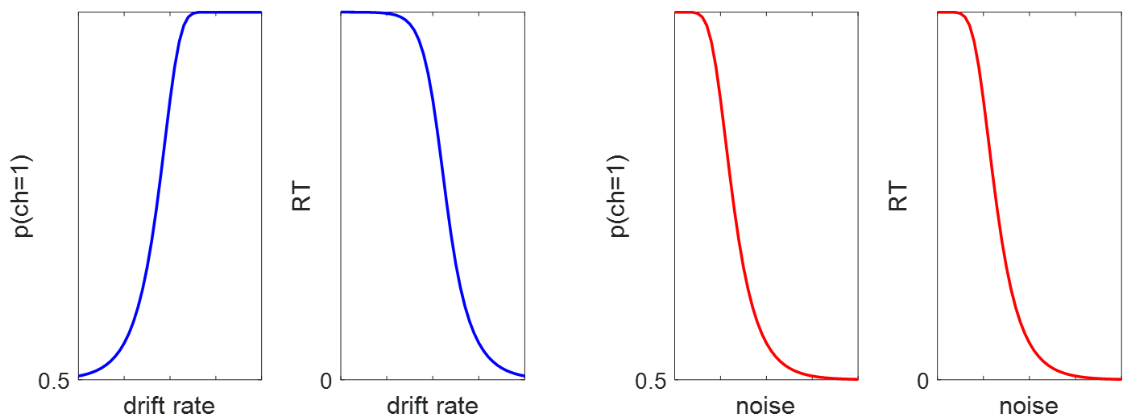

Figure 3: With all other parameters fixed, an increase in drift rate leads to a sigmoidal increase in the probability of choosing option 1 from 0.5 (random guess) to 1, and a sigmoidal decrease in RT (left plots in blue). With all other parameters fixed, an increase in processing noise will lead to a sigmoidal decrease in the probability of choosing option 1, and a parallel sigmoidal decrease in RT (right plots in red).

Model 2

The simplest and most obvious solution to incorporate option-specific uncertainty into the DDM would be to model the process as:

$$E_{t+1} = E_t + d\Delta$$

$$\Delta \sim N(\mu_1 - \mu_2, \sigma^2 + \sigma_1^2 + \sigma_2^2)$$

which is simply the standard DDM equation capturing the evolution of the accumulated evidence across time, but with σ_i^2 representing the uncertainty about the value estimate of option i. The only difference between this formulation and the standard one is that here the variance of Δ is specific to the options in the current choice set, whereas in the standard DDM it is fixed across choices (for an individual decision maker). A direct result of this reformulation is that choices between options with greater value uncertainty (lower C) will be more stochastic and take less time (on average), as can be seen by examining the (revised) DDM equations for choice probability and expected response time (see also Fig 3, red lines):

$$p(ch = 1) = \frac{1}{1 + e^{\left(\frac{-2\theta d(\mu_1 - \mu_2)}{\sigma^2 + \sigma_1^2 + \sigma_2^2}\right)}}$$

$$RT = \frac{\theta(2 * p(ch = 1) - 1)}{d(\mu_1 - \mu_2)}$$

Model 3

An alternative way in which the concept of option-specific value (un)certainty could be incorporated into the DDM would be through a signal-to-noise dependency in the drift rate. The drift rate in the DDM symbolizes the accumulation of evidence for one option over the other, equal to the value estimate of one option minus the value estimate of the other option (scaled by a fixed term). The accumulator variable is referred to as "evidence" because the probability distributions controlling it (or the neural activity driving it) are thought to provide reliable signal that will accurately inform the decision. If the value representations of different options can have different levels of uncertainty, it stands to reason that the reliability of the "evidence" that these signals provide about the correct decision will also be different. As such, evidence drawn from a more reliable source (i.e., one with higher certainty) should be weighted more heavily. Under this framework, the equation governing the DDM process would be:

$$E_{t+1} = E_t + d\Delta$$

$$\Delta \sim N\left(\frac{\mu_1 - \mu_2}{\sqrt{\sigma_1^2 + \sigma_2^2}}, \sigma^2\right)$$

where σ^2 (without a subscript) is the noise in the system unrelated to specific choice options. The only difference between this formulation and the standard one is that here the mean of the option value difference is divided by its standard deviation. A direct result of this reformulation is that choices between options with greater value uncertainty will be more stochastic and also take *more* time (on average), as can be seen by examining the (revised) DDM equations for choice probability and expected response time:

$$p(ch = 1) = \frac{1}{1 + e^{\left(\frac{-2\theta d(\mu_1 - \mu_2)}{\sigma^2\left(\sqrt{\sigma_1^2 + \sigma_2^2}\right)}\right)}}$$

$$RT = \frac{\theta(2 * p(ch = 1) - 1)}{\frac{d(\mu_1 - \mu_2)}{\sqrt{\sigma_1^2 + \sigma_2^2}}}$$

Here the impact of option-specific uncertainty on RT is more complex. First, greater uncertainty decreases RT through its effect on choice stochasticity (as before). Second, greater uncertainty directly increases RT by diminishing the slope of the drift rate. The second effect dominates.

Model 4

A variant of the DDM in which the drift rate is altered by the option-specific value certainty could be one in which the evidence in favor of each option is scaled by its own precision term, as is the case, for example, in multi-sensory integration (Drugowitsch et al, 2014; Fetsch et al, 2012). The drift rate would thus become the accumulation of adjusted evidence for one option over the other, equal to the precision-weighted value estimate of one option minus the precision-weighted value estimate of the other option (scaled by a fixed term). Here, the evidence drawn from a more reliable source (i.e., one with higher certainty) will be weighted more heavily (as in Model 3), but prior to comparison between the alternative options. Note that here the certainty weighting is truly specific to each option, whereas in Model 3 the certainty weighting is specific to the pair of options. Under this framework, the equation governing the DDM process would be:

$$E_{t+1} = E_t + d\Delta$$

$$\Delta \sim N\left(\frac{\mu_1}{\sqrt{\sigma_1^2}} - \frac{\mu_2}{\sqrt{\sigma_2^2}}, \sigma^2\right)$$

The only difference between this formulation and the standard one is that here the mean of the option value difference is adjusted by the standard deviations of the individual choice options. Because the evidence in favor of each option (prior to comparison) will be scaled by its own specific (and importantly, potentially different) precision term, the impact on both choice stochasticity and response time could go in either direction. This can be seen by examining the (revised) DDM equations for choice probability and expected response time:

$$p(ch=1)=\frac{1}{1+e^{\left(\frac{-2\theta d}{\sigma^2}\left(\frac{\mu_1}{\sqrt{\sigma_1^2}}-\frac{\mu_2}{\sqrt{\sigma_2^2}}\right)\right)}}$$

$$RT=\frac{\theta(2*p(ch=1)-1)}{d\left(\frac{\mu_1}{\sqrt{\sigma_1^2}}-\frac{\mu_2}{\sqrt{\sigma_2^2}}\right)}$$

Here the impact of option-specific uncertainty on both choice and RT is more complex than in the other models presented above. If the evidence stream for one option has both a larger mean and a smaller variance, relative to the other option, the effective drift rate will be higher than in a standard DDM (e.g., choices will be less stochastic and faster). On the other hand, if the option with the larger mean evidence stream is different from the option with the more reliable evidence stream, the effective drift rate will be lower than in a standard DDM (e.g., choices will be more stochastic and slower).

Model 5

Yet another alternative way in which the DDM could include option-specific value certainty is in the form of an independent evidence accumulator -- a secondary drift, with a rate proportional to the difference in certainty between the choice options. In this way, the decision about which option to choose would be influenced both by which value estimate was higher (via the primary, standard drift) and by which value estimate was more certain (via the secondary, novel drift). The secondary drift might represent an aversion to risk or ambiguity, where the deliberation process would be both slowed by the risk/ambiguity and pulled towards the more certain option. This would imply that the decision maker prefers options that are more valuable, but also options for which the value is more certain. Under this framework, the equation governing the DDM process would be:

$$E_{t+1}=E_t+d_v\Delta_v+d_c\Delta_c$$

$$\Delta_v \sim N(\mu_1 - \mu_2, \sigma^2)$$

$$\Delta_c \sim N(\sigma_1^2 - \sigma_2^2, \sigma^2)$$

where Δ_v and Δ_c are the incremental evidence for value and certainty, respectively, and d_v and d_c are scalars (it is assumed that d_v will always be positive, but d_c could take either sign). The inclusion of a secondary drift rate (in essence, a parallel evidence stream that monitors value certainty rather than value itself) can have either a positive or a negative impact on both choice and RT, depending on whether the option with the higher mean evidence stream is the same as or different than the option with the higher evidence reliability (as well as on the sign of d_c). This can be seen by examining the (revised) DDM equations for choice probability and expected response time:

$$p(ch = 1) = \frac{1}{1 + e^{\left(\frac{-2\theta}{\sigma^2}\left[d_v(\mu_1 - \mu_2) + d_c(\sigma_1^2 - \sigma_2^2)\right]\right)}}$$

$$RT = \frac{\theta(2 * p(ch = 1) - 1)}{d_v(\mu_1 - \mu_2) + d_c(\sigma_1^2 - \sigma_2^2)}$$

Here, the secondary drift rate will result in more consistent and faster choices if the sign of μ_1-μ_2 is the same as that of σ_1^2-σ_2^2, but it will result in less consistent and slower choices otherwise. This, of course, is under the assumption that the decision maker is risk/ambiguity averse (i.e., $d_c>0$). If the decision maker were risk/ambiguity seeking (i.e., $d_c<0$), the opposite predictions would hold.

Materials and Design

Using a variety of different datasets from previous studies, one at a time, we fit experimental data to each of the models that we described above. We then performed Bayesian model comparison to determine which of the models (if any) performed better than the others across the population of participants. For this model fitting and comparison exercise, we relied on the Variational Bayesian Analysis toolbox (VBA, available freely at https://mbb-team.github.io/VBA-toolbox/; Daunizeau, Adam, & Rigoux, 2014) for Matlab R2020a. We used the VBA_NLStateSpaceModel function to fit the data for each participant individually, followed by the VBA_groupBMC function to compare the results of the model fitting across models for the full group of participants. The input for the model inversion was a series of two-alternative forced choice (2AFC) data, including measures of value estimate and value estimate certainty for each choice option, the chosen option (left or right) for each trial, and response time (RT) for each trial. Some datasets also included choice confidence reports for each trial, which were included in supplementary analyses as will be described below. The parameters to be fitted included all of the d and σ^2 terms described above, plus additional parameters for an affine transformation of experimental certainty measures into theoretical ones. This is necessary because in the experimental data, the measures of value and those of value certainty reside on the same scale, but this is likely untrue for the cognitive variables that are meant to drive the models.

Dataset 1

The first dataset we examined was from Lee & Daunizeau (2020a). In this study, participants made choices between various snack food options based on their personal subjective preferences. Value estimates for each option were provided in a separate rating task prior to the choice task. Participants used a slider scale to respond to the question, "Does this please you?" After each rating, participants used a separate slider scale to respond to the question, "Are you sure?" This provided a measure of value estimate certainty for each item. During the choice task, participants were presented with pairs of snack food images and asked, "What do you prefer?" After each choice, participants used a slider scale to respond to the question, "Are you sure about your choice?" to provide a subjective report of choice confidence. This dataset contained 51 subjects, each of whom were faced with 54 choice trials.

Dataset 2

The second dataset we examined was from Lee & Daunizeau (2020b). In this study, participants made choices between various snack food options based on their personal subjective preferences. Value estimates for each option were provided in a separate rating task prior to the choice task. Participants used a slider scale to respond to the question, "How much do you like this item?" After each rating, participants used the same slider scale to respond to the question, "How certain are you about the item's value?" by indicating a zone in which they believed the value of the item surely fell. This provided a measure of value estimate certainty for each item.

During the choice task, participants were presented with pairs of snack food images and asked, “Which do you prefer?” After each choice, participants used a slider scale to respond to the question, “Are you sure about your choice?” to provide a subjective report of choice confidence. This dataset contained 32 subjects, each of whom were faced with 74 choice trials.

Dataset 3

The third dataset we examined was from Lee & Coricelli (2020). In this study, participants made choices between various snack food options based on their personal subjective preferences. Value estimates for each option were provided in a separate rating task prior to the choice task. Participants used a slider scale to respond to the question, “How pleased would you be to eat this?” After each rating, participants used a six-point descriptive scale to respond to the question, “How sure are you about that?” This provided a measure of value estimate certainty for each item. During the choice task, participants were presented with pairs of snack food images and asked, “Which would you prefer to eat?” After each choice, participants used a slider scale to respond to the question, “How sure are you about your choice?” to provide a subjective report of choice confidence. This dataset contained 47 subjects, each of whom were faced with 55 choice trials.

Dataset 4

The fourth dataset we examined was from Gwinn & Krajbich (2020). In this study, participants made choices between various snack food options based on their personal subjective preferences. Value estimates for each option were provided in a separate rating task prior to the choice task. Participants used a 10-point numerical scale to respond to the prompt, “Please indicate how much you want to eat this item.” After each rating, participants used a seven-point numerical scale to respond to the prompt, “Please indicate how confident you are in your rating of this item.” This provided a measure of value estimate certainty for each item. During the choice task, participants were presented with pairs of snack food images and instructed to choose the one that they preferred to eat. Choice confidence was not measured in this study. This dataset contained 36 subjects, each of whom were faced with 200 choice trials.

RESULTS

Model comparison

Before we present the quantitative model comparison results, we first show the qualitative predictions that each model (fitted to its optimal parameters) makes with respect to the effects of value difference, value certainty, and certainty difference on choice consistency, RT, and choice confidence (see Figure 4). We first simulated 10^7 trials for each model, based on random input values with similar distributions as in our experimental data. The model-parameters we used were in line with the best fit parameters to the experimental data. We then performed GLM regressions: dV, C, and dC on choice (binomial) and on RT (linear). (Note: we coded the data such that option 1 always had the higher value.) A preliminary inspection of the results suggests that Model 4 is the only model that accounts for all of the qualitative benchmarks of the certainty-RT correlations, in particular the decrease in RT with both average certainty and certainty difference. As expected, Model 2 makes the wrong qualitative prediction (higher RT with value certainty), while Model 3 and Model 5 fail to account for the dependency of RT on either certainty difference (dC) or average certainty (C), respectively.

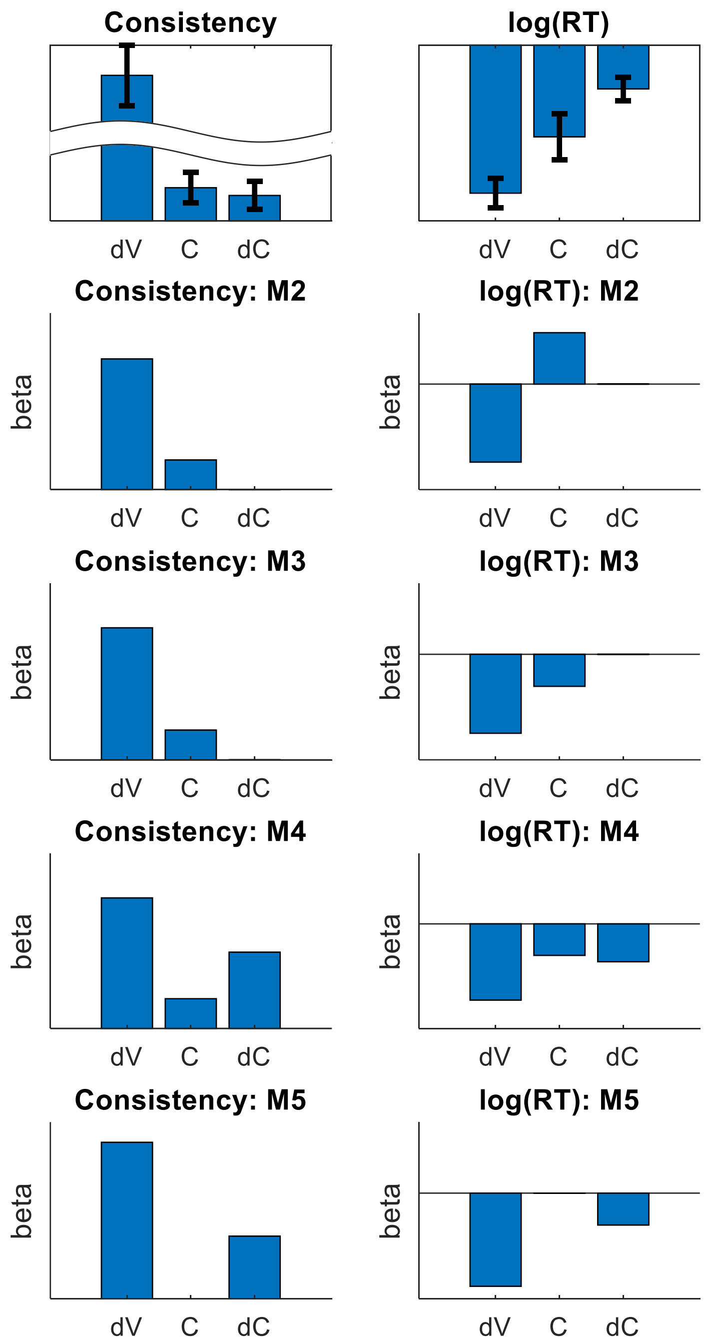

Figure 4: Qualitative predictions of the effects of value difference, value certainty, and certainty difference on choice consistency and RT (shown for models 2-5; top row shows experimental data). (Bar heights represent mean GLM beta weights based on 10^7 simulated trials for each model.)

The classic basic DDM, our Model 1, has been validated countless times for its ability to account for two-alternative forced choice responses and mean response times. The other models we described above, Models 2-5, are new and have therefore never been tested with empirical data. Thus, we start our model comparison exercise with one-on-one competitions between Model 1 and each of Models 2-5, separately. This serves as a simple test of whether the addition of the option-specific value estimate certainty term, as suggested in each of the four different manners described in Models 2-5, improves the fit of the classic DDM. We then perform a comparison across all five models simultaneously, and test whether any of them dominates the others in terms of best fit to the data. We present the quantitative results of the model-fit comparison in Figure 5, and describe them below.

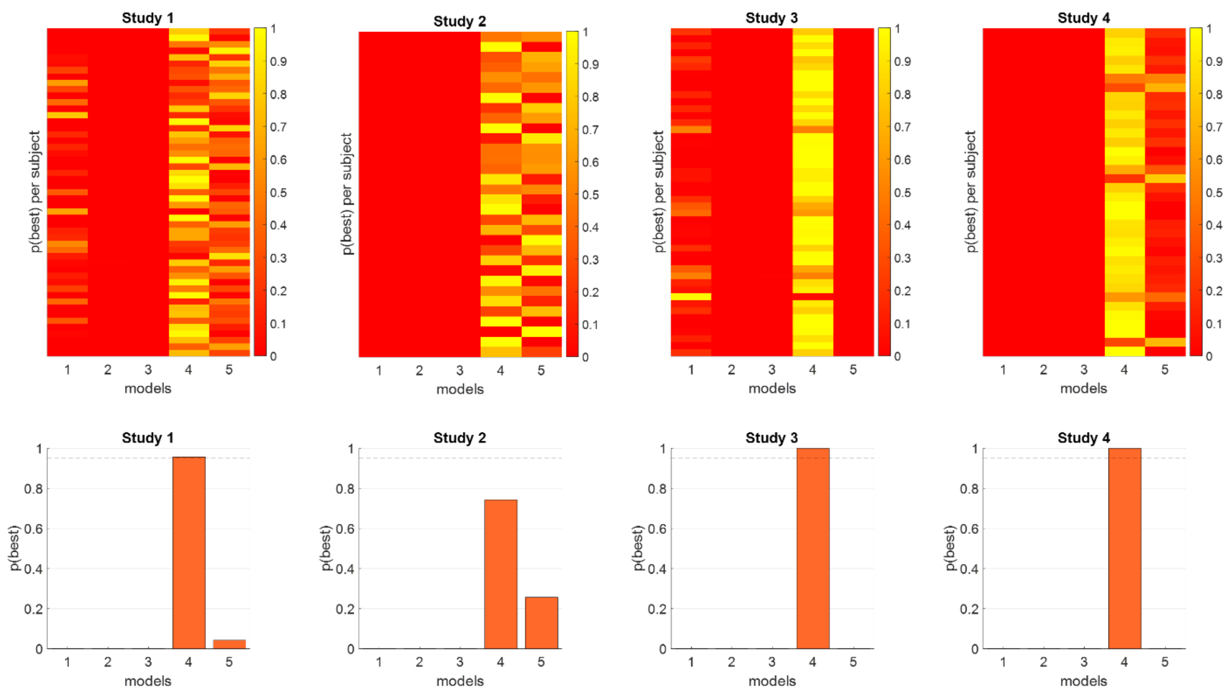

Figure 5: Model comparison results. For the five models we examined, we show here the model attributions (top row) for each of the studies we examined, at the subject level; each cell represents the probability that the model (column) best represents the behavior of the subject (row). We also show the exceedance probability for each model being the best in each study (bottom row).

For Dataset 1, Model 1 dominated Model 2, with an exceedance probability of 1 and an estimated model frequency (across the participant population) of 0.972. Model 1 dominated Model 3, with an exceedance probability of 1 and an estimated model frequency of 0.990. Model 4 dominated Model 1, with an exceedance probability of 1 and an estimated model frequency of 0.779. Model 5 dominated Model 1, with an exceedance probability of 1 and an estimated model frequency of 0.757. When comparing all models simultaneously, Model 4 dominated, with an exceedance probability of 0.956 and an estimated model frequency of 0.541. Models 1, 2, 3, and 5 had estimated model frequencies of 0.127, 0.004, 0.004, and 0.324, respectively. Because Models 4 and 5 each performed better than Model 1, we ran the comparison again including Model 6, which was a combination of Models 4 and 5. Model 4 again dominated, with Model 6 receiving no support.

For Dataset 2, Model 1 dominated Model 2, with an exceedance probability of 1 and an estimated model frequency of 0.984. Model 1 dominated Model 3, with an exceedance probability of 1 and an estimated model frequency of 0.983. Model 4 dominated Model 1, with an exceedance probability of 1 and an estimated model frequency of 0.897. Model 5 dominated Model 1, with an exceedance probability of 1 and an estimated model frequency of 0.889. When comparing all models simultaneously, Model 4 outperformed the others, with an exceedance probability of 0.742 and an estimated model frequency of 0.546. Models 1, 2, 3, and 5 had estimated model frequencies of 0.127, 0.004, 0.004, and 0.324, respectively. Because Models 4 and 5 each performed better than Model 1, we ran the comparison again including Model 6, which was a combination of Models 4 and 5. Model 4 again outperformed the others, with Model 6 receiving no support.

For Dataset 3, Model 1 dominated Model 2, with an exceedance probability of 1 and an estimated model frequency of 0.987. Model 1 dominated Model 3, with an exceedance probability of 1 and an estimated model frequency of 0.985. Model 4 dominated Model 1, with an exceedance probability of 1 and an estimated model frequency of 0.860. Model 5 slightly outperformed Model 1, with an exceedance probability of 0.561 and an estimated model frequency of 0.511. When comparing all models simultaneously, Model 4 dominated, with an exceedance probability of 1 and an estimated model frequency of 0.865. Models 1, 2, 3, and 5 had estimated model frequencies of 0.122, 0.004, 0.004, and 0.004, respectively. Because Models 4 and 5 each performed better than Model 1, we ran the comparison again including Model 6, which was a combination of Models 4 and 5. Model 4 again dominated, with Model 6 receiving no support.

For Dataset 4, Model 1 dominated Model 2, with an exceedance probability of 1 and an estimated model frequency of 0.955. Model 1 dominated Model 3, with an exceedance probability of 1 and an estimated model frequency 0.982. Model 4 dominated Model 1, with an exceedance probability of 1 and an estimated model frequency of 0.986. Model 5 dominated Model 1, with an exceedance probability of 1 and an estimated model frequency of 0.986. When comparing all models simultaneously, Model 4 dominated, with an exceedance probability of 1 and an estimated model frequency of 0.808. Models 1, 2, 3, and 5 had estimated model frequencies of 0.005, 0.005, 0.005, and 0.175, respectively. Because Models 4 and 5 each performed better than Model 1, we ran the comparison again including Model 6, which was a combination of Models 4 and 5. Model 4 again dominated, with Model 6 receiving no support.

DISCUSSION

The aim of this study was to examine a number of variants of drift-diffusion model for preferential choice, and to probe them in their ability to account for benchmark data on the dependency of choice and RT on value uncertainty. As illustrated in Figure 2, the experimental data that we examined show that value certainty has a clear impact on both choice and RT (thus extending beyond the default DDM without option-specific noise), and also provides strong constraints on the way one can introduce option-specific noise into the model. As we have shown, the simplest DDM extension, in which the noise increases with value uncertainty, produces the wrong qualitative prediction: RT increases with certainty (certainty reduces the noise in the system, which slows down RT; see Figure 3, right panel). Moreover, this problem with the introduction of option-specific value uncertainty in modeling value-based decisions is not particular to the DDM, but also applies to the broader class of evidence accumulation-to-bound models, in which noise speeds up RT.

We have examined and tested three additional DDM variants. The first two (Models 3-4) were based on signal-to-noise principles, while the last one (Model 5) included an independent and additive diffusion process based on certainty. While each of these models was able to account for some of the relationships in the data, only Model 4 accounted for all of them. In this model, the drift rate of the diffusion process is not simply the fluctuating difference in the values of the options (Tajima, Drugowitsch, & Pouget, 2016), but rather a difference between the ratios of the values and their corresponding value uncertainties. This mechanism has a normative flavor, as it penalizes values that are associated with uncertain alternatives. Some similar type of signal-to-noise models have also been supported by data in perceptual choice tasks. For example, de Gardelle and Summerfield (2011) examined choices in which an array of eight visual patches of variable color or shape are compared (in binary choice) to a reference (color or shape). By independently varying the set mean distance (in the relevant dimension) from the reference as well as the set variance, they found that both independently affect choice accuracy and RT. In particular, set variance (which is the analog of our value uncertainty) reduces choice accuracy and increases RT. As shown by de Gardelle and Summerfield (2011), a signal-to-noise model can account for this dependency. Indeed, the random dot motion task that is widely used alongside the DDM in perceptual decision making studies provides a signal-to-noise ratio as input for the drift rate (e.g., Gold & Shadlen, 2007). With this task, drift rate is typically determined by the motion coherence, which is composed of the number of dots moving in the same direction (signal) as well as the number of dots moving randomly (noise).

An alternative way to introduce option-specific value uncertainty in the DDM could be to assume that the uncertainty affects the response boundary rather than the drift rate. Accordingly, decision makers would compensate for their uncertainty by increasing the response boundary. While such a model could account for the negative correlation between RT and certainty (C) shown in Figure 2, it would not be able to account for the negative correlation between RT and dC. Moreover, such a model would predict that choices become more stochastic as value certainty increases, which is both counterintuitive and in contrast to the data. As we show in the Supplementary Materials, this model also fails in term of quantitative model comparison. Thus, we believe that the way in which value uncertainty affects the decision process is via its impact on the drift rate. Future work is needed to examine the neural mechanism that extracts the drift rate from fluctuating values (sampled from memory or prospective imagination; Bakkour et al, 2019; Poldrack et al, 2001; Schacter, Addis, & Buckner, 2007) and that reduces the drift rate of strongly fluctuating items. Future research is also needed to examine if the effects of value uncertainty on choice correlate with risk-aversion at the level of individual participants, and to integrate this type of model with dynamical attentional affects as in the attentional drift-diffusion model (aDDM; Krajbich et al, 2010; Sepulveda et al, 2020).

While we have focused here on how value certainty affects choice and RT, the experimental data also importantly show a marked and systematic effect of value certainty on choice confidence. In particular, higher average value certainty (C) and certainty difference (dC) both lead to higher choice confidence. This pattern raises a further challenge for most accumulation-to-bound style choice models that aim to account for both RT and choice confidence. For example, in the balance of evidence (BOE) type models (Vickers & Packer, 1982; De Martino et al, 2013), confidence corresponds to the difference in the activation of two accumulators that independently race to a decision boundary. If we were to naively introduce option-specific noise in such models, they would predict, contrary to the data, that the confidence becomes larger for options with more value uncertainty (as the noise increases the BOE; see Lee & Daunizeau, 2020b). Similarly, if we were to model confidence using a DDM with collapsing boundaries (e.g., Tajima et al, 2016), with confidence corresponding to the height of the boundary at the time the choice is made, naively introducing option-specific noise would once again provide us with a prediction opposite from what we see in the data. For uncertain alternatives, there would be more noise in the evidence accumulation process, resulting in faster choices and therefore higher boundaries, and thus higher confidence (in fact, this would be true for any model that assumes that confidence decreases with RT; Kiani & Shadlen, 2009).

There are very few value-based choice studies that simultaneously examined value certainty and choice confidence (but see Lee & Daunizeau, 2020a, 2020b; Lee & Coricelli, 2020; De Martino et al, 2013). We have not modeled choice confidence here, as there are many potential ways to do this, with substantial divergence among them (Vickers & Packer, 1982; Kiani & Shadlen, 2009; Pleskac & Busemeyer, 2010; De Martino et al, 2013; Moran, Teodorescu, &

Usher, 2015; Calder-Travis, Bogacz, & Yeung, 2020; see Calder-Travis et al, 2020). Nevertheless, all of these models strive to predict a strong negative correlation between RT and choice confidence, as has been demonstrated in a plethora of experimental data. We note that in the data we examined, the impact of value certainty on choice confidence was essentially the reverse of its effect on RT (see Supplementary Material, Figure S1). While we did not explore this further, it suggests that a signal-to-noise DDM can also capture the dependency of choice confidence on value certainty. Future work is needed to determine how signal detection style DDM variants might be extended towards an optimal unified account of choice, RT, and choice confidence.

Funding

MU was funded by supported by the United States-Israel Binational Science Foundation (CNCRS = 2014612).

Acknowledgments

We wish to thank Konstarinos Tsestsos and Giovanni Pezzulo for a critical reading of the manuscript and for helpful discussions. We thank Ian Krajbich for generously sharing his data with us. We also thank Antonio Rangel for encouraging support on this project.

BIBLIOGRAPHY

Alós-Ferrer, C. (2018). A Dual-Process Diffusion Model. Journal of Behavioral Decision Making, 31(2), 203-218.

Bakkour, A., Palombo, D. J., Zylberberg, A., Kang, Y. H., Reid, A., Verfaellie, M., ... & Shohamy, D. (2019). The hippocampus supports deliberation during value-based decisions. elife, 8, e46080.

Basten, U., Biele, G., Heekeren, H. R., & Fiebach, C. J. (2010). How the brain integrates costs and benefits during decision making. Proceedings of the National Academy of Sciences, 107(50), 21767-21772.

Boldt, A., & Yeung, N. (2015). Shared neural markers of decision confidence and error detection. *Journal of Neuroscience*, *35*(8), 3478-3484.

Brown, S. D., & Heathcote, A. (2008). The simplest complete model of choice response time: Linear ballistic accumulation. Cognitive psychology, 57(3), 153-178.

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.

Calder-Travis, J., Bogacz, R., & Yeung, N. (2020). Bayesian confidence for drift diffusion observers in dynamic stimuli tasks. bioRxiv.

Calder-Travis, J., Charles, L., Bogacz, R., & Yeung, N. (2020). Bayesian confidence in optimal decisions. psyArXiv.

Daunizeau, J., Adam, V., & Rigoux, L. (2014). VBA: a probabilistic treatment of nonlinear models for neurobiological and behavioural data. *PLoS Comput Biol*, *10*(1), e1003441.

De Gardelle, V., & Summerfield, C. (2011). Robust averaging during perceptual judgment. Proceedings of the National Academy of Sciences, 108(32), 13341-13346.

De Martino, B., Fleming, S. M., Garrett, N., & Dolan, R. J. (2013). Confidence in value-based choice. *Nature neuroscience*, *16*(1), 105-110.

Drugowitsch, J., DeAngelis, G. C., Klier, E. M., Angelaki, D. E., & Pouget, A. (2014). Optimal multisensory decision-making in a reaction-time task. Elife, 3, e03005.

Fetsch, C. R., Pouget, A., DeAngelis, G. C., & Angelaki, D. E. (2012). Neural correlates of reliability-based cue weighting during multisensory integration. Nature neuroscience, 15(1), 146-154.

Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual review of neuroscience*, *30*.

Gwinn, R., & Krajbich, I. (2020). Attitudes and attention. *Journal of Experimental Social Psychology*, *86*, 103892.

Kiani, R., & Shadlen, M. N. (2009). Representation of confidence associated with a decision by neurons in the parietal cortex. science, 324(5928), 759-764.

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.

Lebreton, M., Abitbol, R., Daunizeau, J., & Pessiglione, M. (2015). Automatic integration of confidence in the brain valuation signal. *Nature neuroscience*, *18*(8), 1159-1167.

Lee, D., & Coricelli, G. (2020). An empirical test of the role of value certainty in decision making. *bioRxiv,* 2020.06.16.155234.

Lee, D., & Daunizeau, J. (2020a). Choosing what we like vs liking what we choose: How choice-induced preference change might actually be instrumental to decision-making. *PloS one*, *15*(5), e0231081.

Lee, D., & Daunizeau, J. (2020b). Trading Mental Effort for Confidence: The Metacognitive Control of Value-Based Decision-Making. *bioRxiv*, 837054.

Milosavljevic, M., Malmaud, J., Huth, A., Koch, C., & Rangel, A. (2010). The drift diffusion model can account for the accuracy and reaction time of value-based choices under high and low time pressure. Judgment and Decision Making, 5(6), 437.

Moran, R., Teodorescu, A. R., & Usher, M. (2015). Post choice information integration as a causal determinant of confidence: Novel data and a computational account. Cognitive psychology, 78, 99-147.

Philiastides, M. G., & Ratcliff, R. (2013). Influence of branding on preference-based decision making. Psychological science, 24(7), 1208-1215.

Pleskac, T. J., & Busemeyer, J. R. (2010). Two-stage dynamic signal detection: a theory of choice, decision time, and confidence. Psychological review, 117(3), 864.

Polania, R., Woodford, M., & Ruff, C. C. (2019). Efficient coding of subjective value. *Nature neuroscience*, *22*(1), 134-142.

Polanía, R., Moisa, M., Opitz, A., Grueschow, M., & Ruff, C. C. (2015). The precision of value-based choices depends causally on fronto-parietal phase coupling. Nature communications, 6(1), 1-10.

Poldrack, R. A., Clark, J., Paré-Blagoev, E. A., Shohamy, D., Moyano, J. C., Myers, C., & Gluck, M. A. (2001). Interactive memory systems in the human brain. Nature, 414(6863), 546-550.

Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. Neural computation, 20(4), 873-922.

Ratcliff, R., Smith, P. L., Brown, S. D., & McKoon, G. (2016). Diffusion decision model: Current issues and history. *Trends in cognitive sciences*, *20*(4), 260-281.

Ratcliff, R., Voskuilen, C., & Teodorescu, A. (2018). Modeling 2-alternative forced-choice tasks: Accounting for both magnitude and difference effects. Cognitive psychology, 103, 1-22.

Schacter, D. L., Addis, D. R., & Buckner, R. L. (2007). Remembering the past to imagine the future: the prospective brain. Nature reviews neuroscience, 8(9), 657-661.

Sepulveda, P., Usher, M., Davies, N., Benson, A., Ortoleva, P., & De Martino, B. (2020). Visual attention modulates the integration of goal-relevant evidence and not value. bioRxiv.

Tajima, S., Drugowitsch, J., & Pouget, A. (2016). Optimal policy for value-based decision-making. *Nature communications*, *7*(1), 1-12.

Teodorescu, A. R., Moran, R., & Usher, M. (2016). Absolutely relative or relatively absolute: violations of value invariance in human decision making. Psychonomic bulletin & review, 23(1), 22-38.

Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: the leaky, competing accumulator model. Psychological review, 108(3), 550.

Vickers, D. (1970). Evidence for an accumulator model of psychophysical discrimination. Ergonomics, 13(1), 37-58.

Vickers, D., & Packer, J. (1982). Effects of alternating set for speed or accuracy on response time, accuracy and confidence in a unidimensional discrimination task. Acta psychologica, 50(2), 179-197.

Supplementary Materials

Choice Confidence

In this study, we chose not to include choice confidence in our model predictions, as there is not currently an agreed upon standard for doing so. Nevertheless, we did briefly examine this variable in those datasets that contained it (Studies 1-3; Lee & Daunizeau, 2020a, 2020b; Lee & Coricelli, 2020). In general, choice confidence exhibited patterns qualitatively opposite to those exhibited by RT. Specifically, regression beta weights for dV, C, and dC were of similar magnitude as those for RT, but were all positive (whereas for RT, they were all negative). (see Figure S1)

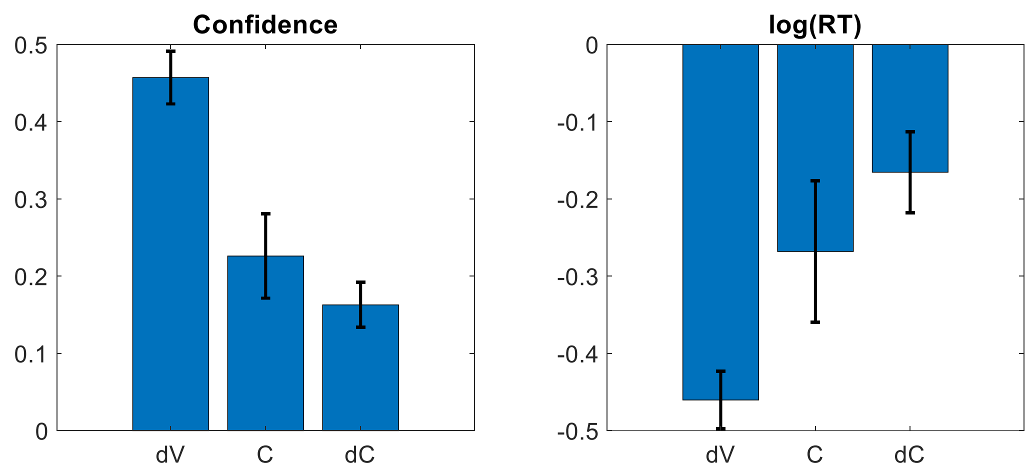

Figure S1: Beta weights from a linear regression of dV, C, and dC on choice confidence (left) and log(RT) (right; n=124; bar heights represent population means, error bars represent s.e.m.).

Certainty-Adjusted Response Threshold

We considered a model that was a standard DDM, but with the response threshold determined as a function of option-specific (or more accurately, trial-specific) value certainty. Under this model, the height of the threshold increases as the value certainty of the pair of options decreases, on a trial-by-trial basis. Choice probability and mean RT are thus calculated using the following equations:

$$p(ch = 1) = \frac{1}{1 + e^{\left(\frac{-2\theta(\sigma_1^2+\sigma_2^2)d}{\sigma^2}(\mu_1-\mu_2)\right)}}$$

$$RT = \frac{\theta(\sigma_1^2 + \sigma_2^2)(2 * p(ch = 1) - 1)}{d(\mu_1 - \mu_2)}$$

As can be seen in the equations, increasing value uncertainty will result in higher choice consistency and higher RT. This is inconsistent with the experimental data. Furthermore, as expected, this model received no support when included in a quantitative comparison with the other models.