Value Certainty in Drift-Diffusion Models of Preferential Choice

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

Abstract

The drift-diffusion model (DDM) is widely used and broadly accepted for its ability to account for binary choices (in both the perceptual and preferential domains) and for their response times (RT), as a function of the stimulus or the option values. The DDM is built on an evidence accumulation to bound concept, where, in the value domain, a decision maker repeatedly samples the mental representations of the values of the options until satisfied that there is enough evidence in favor of one option over the other. As the signals that drive the evidence are derived from value estimates that are not known with certainty, repeated sequential samples are necessary to average out noise. The classic DDM does not allow for different options to have different levels of variability in their value representations. However, recent studies have shown that decision makers often report levels of certainty regarding value estimates that vary across choice options. There is therefore a need to extend the DDM to include an option-specific value certainty component. We present several such DDM extensions and validate them against empirical data from four previous studies. The data supports best a DDM version in which the drift of the accumulation is based on a sort of signal-to-noise ratio of value for each option (rather than a mere accumulation of samples from the corresponding value distributions). This DDM variant accounts for the positive impact of value certainty on choice consistency and for the negative impact of value certainty on RT.

The drift-diffusion model (***DDM***; Ratcliff & Rouder, 1998; Gold & Shadlen, 2007) 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 value comparisons (e.g., stochastic fluctuations in value signals that may result from variability in the firing of neural populations, from variability in attention or in memory retrieval of the values associated with the choice alternatives, or from variability intrinsic in the environment when the alternatives are characterized by variable/stochastic values) and implementing an optimal stopping rule (optimizing response time (RT) for a specified accuracy; Wald, 1948; Gold & Shadlen, 2007; Bogacz et al, 2006). This model accounts well for the dependency of the RT distribution on the values (or stimulus strength) of the choice options (Ratcliff & Rouder, 1998), and for the speed-accuracy tradeoff — the often-observed phenomenon that decision makers are able to improve their choice accuracy by taking longer to decide (on average; Wickelgren, 1977) — which corresponds to an expansion of the *DDM* response boundaries.

While initially used in the domain of perceptual (Ratcliff & Rouder, 1998) or memory-based decisions (Ratcliff, 1978), the *DDM* core principles of sequential sampling and integration to boundary have since become a central component of models of preference-based decisions (Busemeyer & Townsend, 1993; Busemeyer & Diederich, 2002; 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; Roe, Busemeyer & Townsend, 2001). For example, in the *Decision Field Theory* (***DFT***) model (Busemeyer & Townsend, 1993; Busemeyer & Diederich, 2002; Busemeyer et al, 2019; Roe et al., 2001) — one of the first models that applied sequential sampling principles to preferential choice — it is assumed that the variability in the value integration corresponds to fluctuations of attention between the relevant attributes that characterize the alternatives. In a more recent application of the *DDM* to preferential choice (between food items), the drift rate is assumed to be modulated by visual attention across the alternatives (i.e., looking at an alternative increases the momentary drift of the accumulator towards the corresponding decision boundary; Krajbich, Armel & Rangel, 2010). In another more recent study, Tajima et al (2016) carried out a normative analysis of the optimal decision policy in preferential choice. They converged on the conclusion that "similar to perceptual decisions, drift diffusion models implement the optimal strategy for value-based decisions." An important

assumption for the derivation of this optimal policy is that "the value of each option is represented by a probability distribution whose mean is the true (subjective) value of the option, and whose variance corresponds to sampling noise or uncertainty about the true value." This noise could be interpreted as imprecision in the value representations themselves. Thus, 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 value (the so-called drift rate, which corresponds to the difference in the means of the two value distributions). In order to average out this noise, the values of the alternatives are repeatedly sampled (possibly from memory; Bakkour et al, 2019) and accumulated over time until a sufficient amount of evidence has been acquired to allow for a choice to be made. Accordingly, the *DDM* includes response boundaries for each option (typically symmetric) that trigger a choice once reached by the evidence accumulator. Thus, the fundamental components of the drift-diffusion process are the drift rate (proportional to the difference in the option values), the diffusion coefficient (the degree of stochasticity of the system), and the choice boundaries (the threshold of minimum required evidence for a given level of caution, which controls the speed-accuracy tradeoff; see Figure 1).[1]

[1] According to Tajima et al (2016), the optimal policy requires a collapsing rather than fixed boundary. The collapse of the boundary results, however, from auxiliary assumptions about the cost of deliberation time to the subject. Since these assumptions can be debated, and since there is an intensive debate on whether experimental data best support DDM with fixed or collapsing boundaries (REFS), we consider both cases in this study. We will start from the simple fixed-boundary which is the most prominent DDM model in the literature, and we will then show that our conclusions are unchanged for the collapsing boundary variant.

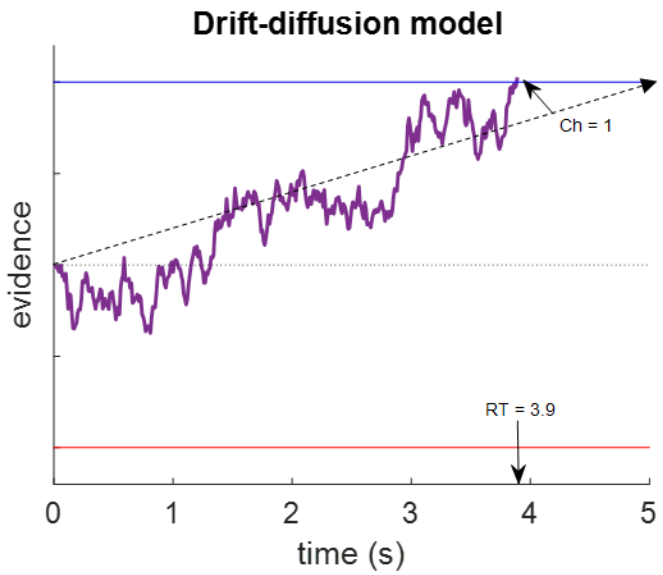

Figure 1: *An illustration of the classic* DDM. *Evidence accumulates across time, following a fixed drift trajectory (black dashed arrow, whose slope corresponds to the value difference of the choice options multiplied by a scalar coefficient), corrupted by white processing or sampling noise. Here, the accumulated evidence reaches the upper boundary after 3.9 seconds, and a choice for option 1 is recorded.*

The classic *DDM* implicitly assumes that processing (or sampling) noise is independent of the identity of the options contained in a particular choice set. That is to say, the sampling noise in the *DDM* is not option-specific (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. However, the brain is known to encode not only the subjective value of options, but also the subjective certainty about the value (Lebreton et al, 2015). It is thus reasonable to suggest that the representations of value that the brain uses to inform the decision process fluctuate (Tajima et al, 2016), 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, the judged value certainty, *C*, correlates positively with choice

consistency (the equivalent of choice accuracy for preferential choice) and negatively with RT[2] (see Figure 2). This provides a qualitative benchmark that any *DDM* variant that includes option-specific uncertainty should be able to account for.

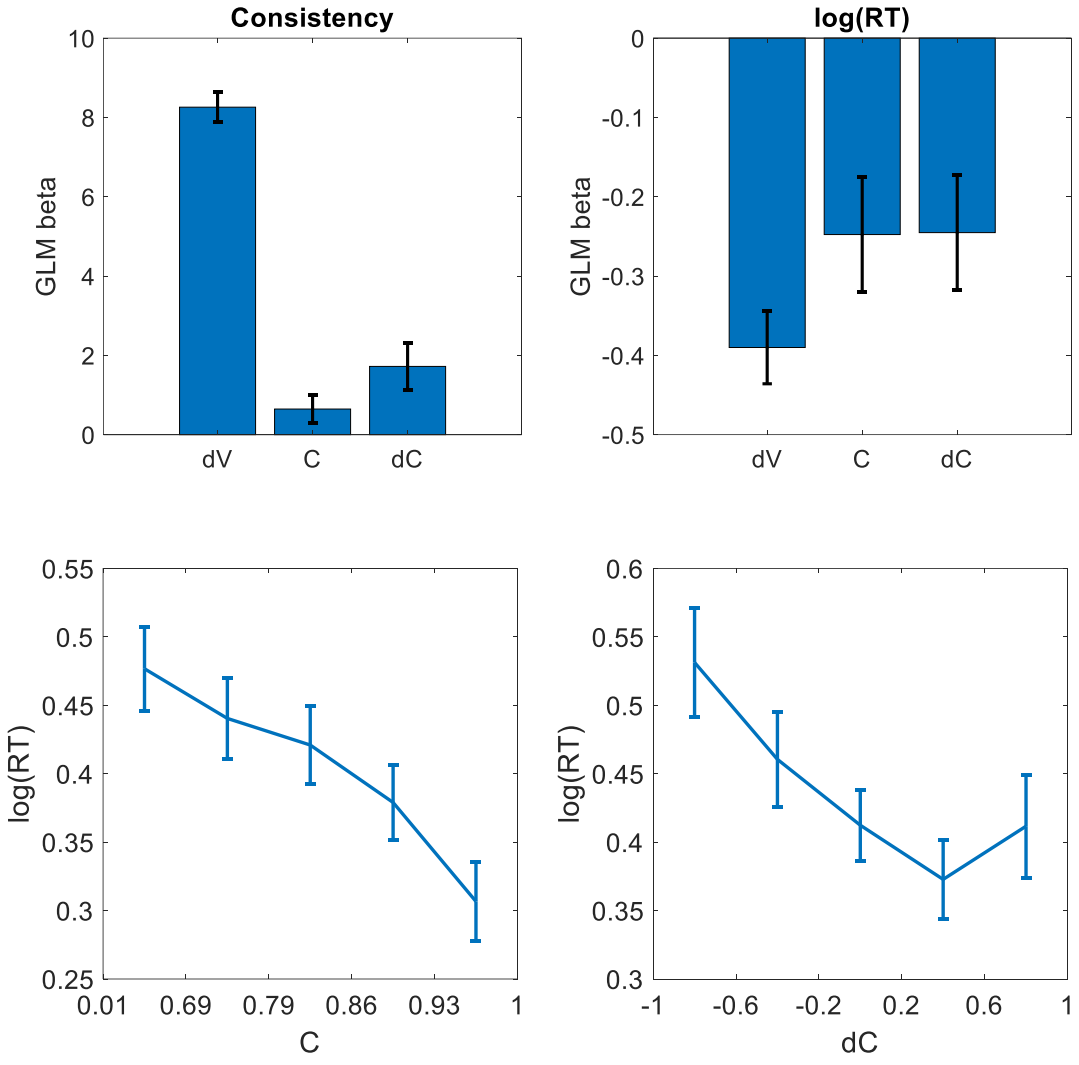

Figure 2: *Previous results (pooled across four studies, n=130) demonstrate the relationships between value difference (dV = v1 - v2), value certainty (C = (c1 + c2) / 2), and certainty difference (dC = c1 - c2) with choice consistency and log(RT). We define v1 (v2) as the higher (lower) value rating, and c1 (c2) as the certainty report for the option with the higher (lower)*

[2] As shown in Figure 2, we regressed the various measures of certainty on log(RT) rather than on RT. This is because whereas the RT variable is not normally distributed (it has a relatively long right tail), the transformed variable of log(RT) is normally distributed and is more adequate for statistical testing (see Glockner & Betch, 2008). Our conclusions hold if we use RT (without their log-transformation).

*value rating. The **upper** panels show the beta weights from a logistic regression on choice consistency (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. The **lower** panels show the isolated relationships between log(RT) and C (left) and between RT and dC (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)*

A straightforward way in which one could introduce option-specific noise in the preferential *DDM*, consistent with the assumptions of the optimal policy model (Tajima et al, 2016), would be to assume that the diffusion noise parameter increases with value uncertainty. This is because alternatives whose values are less certain will have greater variance in their representations, from which the evidence samples are drawn. However, this would lead to the wrong qualitative predictions, with regards to the RT x certainty benchmark. Specifically, under such a model, diffusion noise would decrease as certainty increases, resulting in longer 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 classic *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 & McClelland, 2001)), which also predict that higher noise in the system will result in *faster* responses, in direct contrast to the empirical data. An alternative way to include option-specific noise in the *DDM* could be to assume that the response boundary increases with value uncertainty. Such a model could account for the observed negative correlation between RT and certainty sum, but not certainty difference (between options). Furthermore, such a model would also predict that choice consistency (i.e., in favor of the higher rated option) decreases with certainty (resulting from a decrease in the response boundary, all else equal), which contradicts the empirical data. It is therefore most probable that option-specific value certainty affects the *DDM* process through the drift rate. (We provide additional support for this conclusion in the Supplementary Material, using median splits on certainty and comparing the fitted model parameters.)

The aim of this paper is to examine a number of *DDM* variants, which could account 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 classic *DDM*, each of which incorporates the concept of option-specific value certainty in a unique and realistic way. We start with the classic *DDM* (Model 1) that has no option-specific noise, and progress to models that include option-specific certainty in the sampling noise or in the response boundary (Models 2-3; although we know that these models are less likely to succeed, we believe that it is instructive to formally test them and to examine their potential failures). We then examine signal-to-noise ratio or uncertainty aversion concepts in the formulation of the drift rate (Models 4-6). We 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 use the results to 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 well 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, (with v1>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: the choice (ch = {0, 1}), determined by which boundary was reached; RT, determined by the number of integration time steps elapsed before that boundary was reached.

Model 1: Classic DDM

As a baseline default model for comparison (without any option-specific certainty term), we first consider the classic 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 sensitivity parameter, μ_i is the value estimate of option i, and σ^2 represents processing noise in the evidence accumulation and/or comparator systems. In the classic 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 boundary where a decision is triggered (shown here as the upper boundary, for a choice of option 1), p(ch=1) is the probability that the upper boundary will be reached (rather than the lower boundary), and RT is the expected time at which the accumulation process will end (by reaching one of the boundaries)[3]. Because this system of equations is over-parameterized, one of the free parameters must be fixed for a practical application of the equations. In this work, we will fix the boundary θ to a value of 1 when fitting the models, for simplicity. Choice probability and RT will thus be functions of the drift rate coefficient d and the noise parameter σ^2 (see Figure 3). As expected, accuracy increases and RT decreases with the drift rate. The noise, on the other hand, reduces both accuracy and RT.

[3] 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.

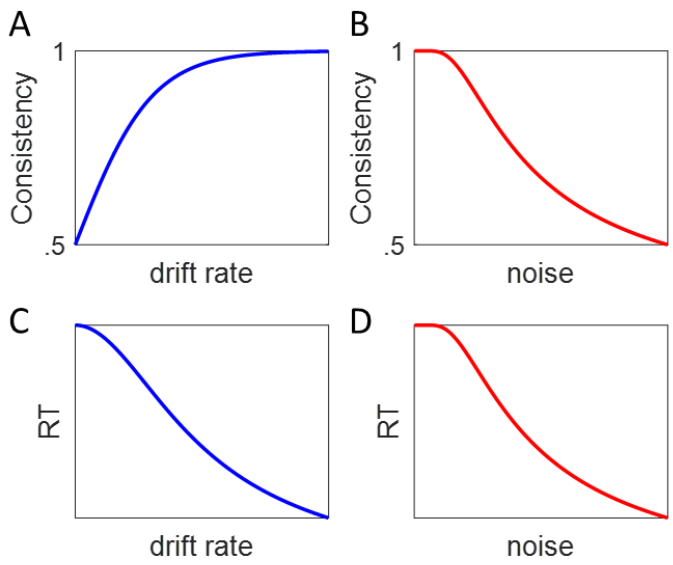

Figure 3: *With all other parameters fixed, an increase in drift rate leads to an increase in the probability of choosing option 1 from 0.5 (random guess) to 1, and a decrease in RT (left plots in blue). With all other parameters fixed, an increase in processing noise leads to a decrease in the probability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

Model 2: Certainty-Adjusted Diffusion (noise) Parameter

One simple potential solution to incorporate option-specific uncertainty into the DDM would be to model the process as:

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

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

This corresponds to the classic 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, and γ serving as a sensitivity parameter that controls the impact of value certainty on the level of diffusion noise.The only difference between this formulation and the classic one is that here the variance of Δ is specific to the options in the current choice set, whereas in the classic *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):

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

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

While the predicted relationship between certainty and choice consistency would match empirical observations, the relationship between certainty and RT would be inconsistent with the empirical data.

Model 3: Certainty-Adjusted Response Boundary

Another potential solution could be to allow the magnitude of the response boundary to vary as a function of option-specific (or trial-specific) value certainty. The evidence sampling process would be identical to that of the classic DDM:

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

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

Under this model, the height of the boundary would increase as the value certainty of the pair of options decreases, on a trial-by-trial basis. Choice probability and mean RT would thus be calculated using the following equations:

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

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

where γ is a sensitivity parameter that controls the impact of certainty on the magnitude of the response boundary. As can be seen in the equations, increasing value uncertainty (all else equal) will result in higher choice consistency and higher RT. While the predicted relationship between

certainty and RT would match empirical observations, the relationship between certainty and choice consistency would be inconsistent with the empirical data.

Model 4: Certainty-Adjusted Value Comparison

One possible 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 classic one is that here the mean of the option value difference is divided by its standard deviation ($\sqrt{\sigma_1^2 + \sigma_2^2}$). 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:

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

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

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 5: Certainty-Adjusted Option Values

An alternative 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 *i* is scaled by its own precision term σ_i^2, 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 4), but prior to comparison between the alternative options. Note that here the certainty weighting is truly specific to each option, whereas in Model 4 the certainty weighting is specific to the pair of options. Under this framework, the equation governing the DDM process would be:

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

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

The only difference between this formulation and the classic 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:

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

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

Here the impact of option-specific uncertainty on both choice and RT is more complex than in the previous models. If the value signal 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 the classic DDM (e.g., choices will be less stochastic and faster). On the other hand, if the option with the larger mean value also has a higher variance (due to its value uncertainty), the effective drift rate will be lower than in the classic DDM (e.g., choices will be more stochastic and slower).

Model 6: Certainty Difference Accumulator

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. Other ongoing research has proposed a model based on a similar principle (Li & Ma, 2020). Under this framework, the equation governing the DDM process would be:

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

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

$$\Delta_c \sim N\left(\frac{1}{\sigma_1} - \frac{1}{\sigma_2}, \sigma_c^2\right)$$

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:

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

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

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 $\frac{1}{\sigma_1}-\frac{1}{\sigma_2}$, 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, σ^2, γ, and α terms described in the model formulations, 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 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, and how this compares to the experimental data (see Figure 4). For the experimental data, we pooled the participants from all four studies, fit GLM regression models for each participant, and report the mean beta weights across participants (but see the Supplementary Material for results separated by study). For the synthetic data, we first simulated 10^7 trials for each model, based on random input (option values and certainties) with the same range as in our experimental data (0 to 1; simulated data uniformly distributed). The model parameters we used were the best fit parameters to the experimental data, with mean and variance equal to the mean and variance across all participants in all studies (simulated data normally distributed). 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 5 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 shown in Figure 4, Model 4 and Model 6 fail to account for the dependency of RT on either certainty difference (dC) or average certainty (C), respectively. However, it is worth noting that a joint model that calculates the drift rate as in Model 4 (d_v*dV/C) and also includes a secondary drift rate as in Model 6 (d_c*dC) would qualitatively capture the observed effects of C and dC in the empirical data (see Figure 4). We therefore expected that the winner of the model comparison would be either Model 5 or the hybrid Model 4+6.

We note that Models 2 and 3 cannot possibly account for the experimental data due to their qualitative predictions (see Figure 4). Model 2 necessitates a negative relationship between value certainty and choice consistency and a positive relationship between value certainty and RT, which is the opposite of what we observe in all of the empirical datasets. Model 3 necessitates a negative relationship between value certainty and choice consistency. Furthermore, as expected, neither the "noise model" nor the "boundary model" received any support when quantitatively compared with the classic DDM.

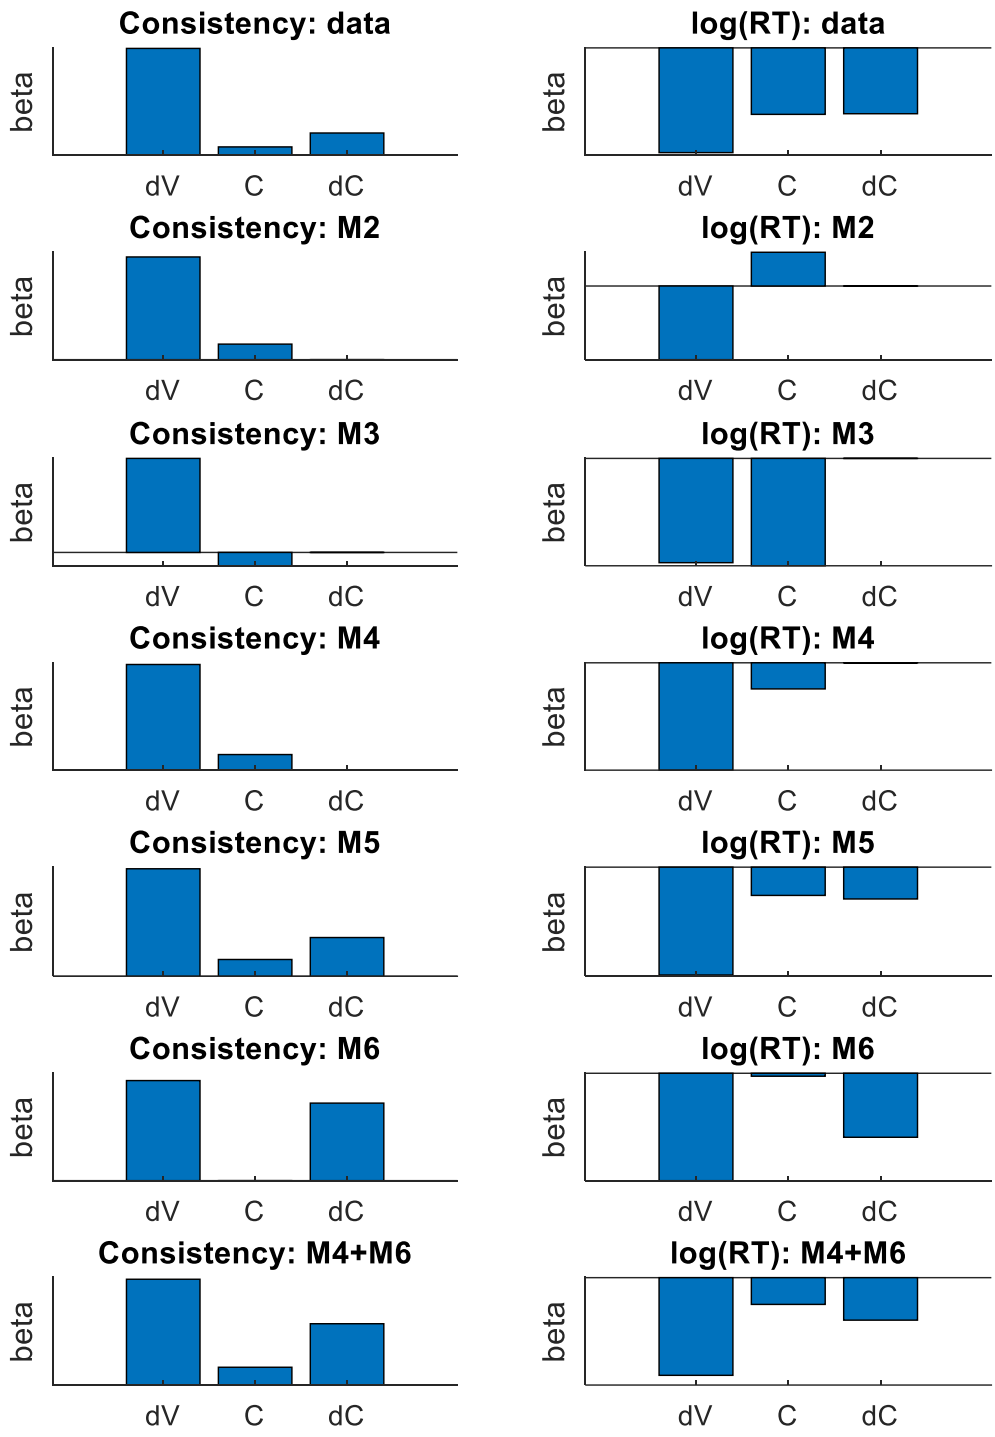

Figure 4: *Qualitative predictions of the effects of value difference, value certainty ((C1+C2)/2), and certainty difference (C1-C2) on choice consistency and log(RT) (shown for the best fit parameters of Models 2-6, plus a hybrid Model 4+6; top row shows experimental data). (Bar heights represent mean GLM beta weights based on 10^7 simulated trials for each model.)*

The classic 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 described above are new and have therefore never been tested with empirical data. Thus, we begin our model comparison exercise by comparing Model 1 with each of Models 2-6, 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 different manners described above, improves the fit of the classic DDM (while being penalized for utilizing additional free parameters). We then perform a comparison across all 10 models simultaneously (the four base models, the five non-linear versions of the base models, and the hybrid model described above), and test whether any of them dominates the others in terms of best fit to the data. We summarize the results of the model comparison below, present the main quantitative results in Figure 5, and describe them in more detail in the Supplementary Material.

Across all four studies, Models 2 and 3 were each dominated by Model 1 (the classic DDM) as anticipated, as was Model 4. Models 5 and 6 each dominated Model 1. A simultaneous comparison of all models showed that Model 5 dominated all others, with Model 6 coming in a distant second place (Model 5 fit best for 70% of participants, Model 6 fit best for 30%).

In recognition of the possibility that the impact of certainty on the drift rate might be non-linear, we considered variations of the surviving models (Models 5 and 6), where the only change is the inclusion of a (strictly positive) exponent α in computation of the signal to noise ratio for Model 5*:

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

and the inclusion of a (strictly positive) free parameter α in the mean of the evidence source distribution for Model 6*:

$$\Delta_c \sim N\left(sign\left(\frac{1}{\sigma_1} - \frac{1}{\sigma_2}\right) \cdot \left|\frac{1}{\sigma_1} - \frac{1}{\sigma_2}\right|^{\alpha}, \sigma^2\right)$$

In a formal comparison of the expanded set of models {1,2,3,4,5,5*,6,6*,4+6}, Model 5* was the clear winner (α median = 0.74, median average deviation = 0.38). This result can be seen in Figure 5, which shows the support that each model received within participants (left panels) and across participants (middle panels). The model attributions across participants and the exceedance

probability that each model was the best are shown in Figure 5. We report the same information separately for each study in the Supplementary Material.

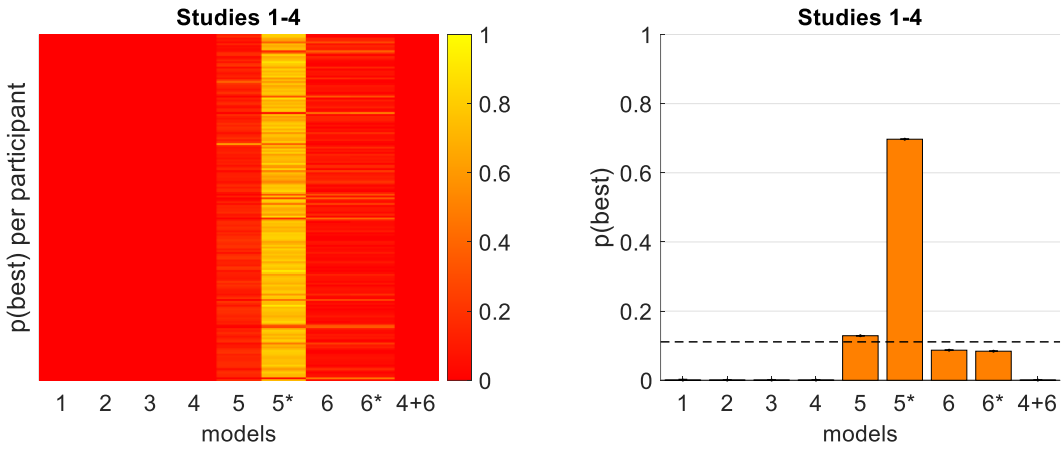

Figure 5: *Model comparison results. We show here the probability that each model best accounted for the data at the participant level (left panel), across the studies we examined; each cell represents the probability that the model (column) best represents the behavior of the participant (row). We also show the probability that each model best explains the data across the participant population (right panel), across all studies. The dashed line indicates chance level if all models were equally probable a priori.*

DISCUSSION

The aim of this study was to examine a number of extensions of the drift-diffusion model for preferential choice that include option-specific value certainty, and to probe them in their ability to account for benchmark data on the dependency of choice and RT on value certainty. As illustrated in Figure 2, the experimental data that we examined show that value certainty has a clear impact on both choice and RT, thereby motivating an extension of the classic DDM without option-specific noise and providing strong constraints on the way one can introduce option-specific noise into the models. As we have shown, perhaps the simplest DDM extension, in which the evidence accumulation noise decreases with value certainty, produces the wrong qualitative prediction: RT increases with certainty (certainty reduces the stochasticity of the system, which slows down RT; see Figure 3, right panel). Moreover, the problem with this method of introducing option-specific value certainty in modeling value-based decisions is not particular to the DDM,

but also applies to the broader class of evidence accumulation-to-bound models (e,g., LCA (Usher & McClleland, 2001) or independent accumulators (Vickers, 1970)), in which noise decreases RT.

An alternative way to introduce option-specific value uncertainty in the DDM could be to assume that the uncertainty affects the response boundary. Accordingly, decision makers would compensate for their uncertainty by increasing the height of the 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 certainty difference (dC). Moreover, such a model would predict that choices become more stochastic as value certainty increases (due to the lower boundary), which is both counterintuitive and in contrast to the experimental data. Those theoretical predictions are supported by our model comparison (see Figures 4 and 5). Thus, we believe that the way in which value uncertainty affects the decision process is via its impact on the drift rate. (We provide additional support for this conclusion in the Supplementary Material, using median splits on certainty and comparing the fitted model parameters.)

In order to understand the way in which value certainty affects the drift rate, we examined and tested three core DDM variants based on certainty-modulated drift rates. The first two (Models 4 and 5) were based on signal-to-noise principles, while the last one (Model 6) included an independent and additive diffusion process based on the certainty of the different option values. While each of these models was able to account for some of the qualitative relationships in the data, only Model 5 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 et al, 2016), but rather a difference between the ratios of the mean 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, the 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. Furthermore, 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); for 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).

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 more strongly fluctuating items. One interesting possibility is that the effective drift rate might be modulated by the temporal consistency of the evidence in successive samples (Glickman, Moran & Usher, 2020); the higher the value certainty, the lower the variance of the value signals, thus leading to a higher probability that successive samples will provide consistent choice evidence. Future research is also needed to understand the nonlinearity indicated in this mechanism (deviation of α exponent from 1), 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). Recent work has suggested that attention allocated toward a particular choice option will serve to reduce the uncertainty about the value of that option (Callaway, Rangel & Griffiths, 2020; Jang, Sharma & Drugowitsch, 2020). If true, this would provide a link between this current work and the so-called gaze bias that has already been well established (Krajbich, Armel & Rangel, 2010; Sepulveda et al, 2020).

Additional DDM Variants or Parameters

Contemporary versions of the DDM sometimes include response boundaries that collapse over time (see Palestro et al, 2018). Such collapsing bounds were first introduced in order to truncate the long tail of RT distributions that are predicted by a fixed-boundary DDM, especially for error trials (Milosavljevic et al, 2010). Some researchers have since included some form of collapsing bound parameter in their usage of the DDM, because it tends to provide better fits to the data (for preferential choice see Glickman et al, 2019). Furthermore, Tajima et al (2016) proved that the DDM was an optimal model only in the case that the response boundary decreased across deliberation time (although their work related to series of sequential decisions that created an

explicit opportunity cost of time, thereby mandating the collapsing bounds). In the current study, we remained agnostic as to the shape of the DDM response boundary (fixed versus collapsing). Whereas we relied on the fixed-boundary DDM in our model comparison, we verified that the inclusion of a collapsing boundary does not change any of our conclusions (see Figure S3A in the Supplementary Material), showing that, like for the fixed-boundary DDM (see Figure 3D), RT also decreases with sampling noise that could be associated to value-uncertainty (contrary to the data; see Figure 2) RT. This can easily be understood by examining the basic relationships between sampling noise and both response probability and RT. If there were no noise in the sampling/accumulation process, the evidence would simply accumulate along a straight trajectory until reaching a bound. Adding noise perturbs the trajectory (at any point in time) and thus increases the probability that the response boundary will be reached sooner. This is because the probability of reaching a boundary on the next time step is a decreasing function of the current distance from the accumulator to the boundary, and an increasing function of both the mean and the *variance* of the trajectory (which reflects the accumulation noise). Note that this will hold for any response boundary, regardless of the shape. Thus, models such as our Model 2 will also fail even when including collapsing boundaries (see Figure S3A in the Supplementary Material for a confirmation of this via simulations).

Some versions of the DDM (Ratcliff & Rouder, 1998) include additional parameters, which we did not examine in this study. Such parameters include non-decision time (time for encoding and response, or T_{er}), starting point bias (and starting point variability), and drift rate variability. While it might be mathematically possible to account for the effects of option-specific certainty by allowing T_{er} to vary as a function of certainty, we believe that this would not be a principled approach. Because T_{er} corresponds to time spent on perceptual (encoding) or motor (response) processes, we see no reason why this should vary with value certainty (determined by valuation processes). Similarly, we did not consider a starting point bias in the current study, as the behavioral effects we seek to explain were demonstrated in decision paradigms where there is no principled reason to include a starting point bias. Such a bias is typically meant to capture prior expectations about which option is more likely to be better. In the experimental paradigms that we considered, there are no such prior expectations. This is because the choice pairs were randomly created, each option was randomly assigned as either option 1 or option 2, and the sequence of trials was randomized. Therefore, as participants proceeded to make choices across trials, any

potential biasing effects from one trial should be obliterated by the effects from all other trials, leaving the prior expectations about which option will be better (on the next trial) centered on zero. It would be theoretically reasonable, however, to consider starting point variability (around a mean of zero). This could be due to a carry-over effect of the previous trial, where residual neural activation might lead the decision maker to be inclined to repeat the same choice on the next trial. This directly implies that any potential starting point variability will be caused by information from previous trials, and not from the current trial (for which deliberation will have not yet begun). Hence, the starting point variability is unlikely to be related to the option-specific value certainty on the current trial. If the starting point variability is unrelated to value certainty, it cannot mediate the impact of certainty on either choice or RT. Regardless, see Supplementary Material for a confirmation (via simulations) that even if one were to imagine that starting point variability increased as a function of value uncertainty (a type of noisier responding), it would have no impact on the mean RT (see Figure S3C), and it could definitely not account for the slowdown with value uncertainty shown in the data (see Figure 2).

Finally, we have not considered drift rate variability in the current study. In some versions of the DDM, the drift rate on a given trial is a random variable selected from a normal distribution whose mean equals the value difference of the options on that trial and with a variance that corresponds to either attentional or memory biases that are specific to that trial (so if the same choce is presented again, that drift may vary). In principle, it could be possible (but less likely, we think) that value uncertainty would affect the drift rate variability (across trials) rather than within-trial sampling noise. As we show in the Supplementary Material (Figure S3C), however, such an assumption would not help, because increasing the variability in the drift rate parameter d across trials with uncertainty, would have a similar effect to on choice and RT as increasing the sampling variability (i.e., the within-trial diffusion noise) – higher variability will yield lower choice consistency and faster RT, contrary to the data (see Figure S3C in the Supplementary Material for a confirmation of this via simulations).

Choice Confidence

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). 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). 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

This research was supported by a grant from the United States-Israel Binational Science Foundation/CNCRS (grant no 2014612 to MU).

Acknowledgments

We wish to thank Konstantinos 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.

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (2006). The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. Psychological review, 113(4), 700.

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., & Diederich, A. (2002). Survey of decision field theory. Mathematical Social Sciences, 43(3), 345-370.

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.

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.

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.

Callaway, F., Rangel, A., & Griffiths, T. L. (2020). Fixation patterns in simple choice are consistent with optimal use of cognitive resources. PsyArXiv preprint PsyArXiv: https://doi. org/10.31234/osf. io/57v6k.

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.

Glickman, M., Moran, R., & Usher, M. (2020). Evidence integration and decision-confidence are modulated by stimulus consistency. bioRxiv.

Glickman M, Sharoni O, Levy DJ, Niebur E, Stuphorn V, & Usher M (2019). The formation of preference in risky choice. PLOS Computational Biology.

Glöckner, A., & Betsch, T. (2008). Multiple-reason decision making based on automatic processing. Journal of experimental psychology: Learning, memory, and cognition, 34(5), 1055.

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.

Jang, A. I., Sharma, R., & Drugowitsch, J. (2020). Optimal policy for attention-modulated decisions explains human fixation behavior. bioRxiv. doi.org/10.1101/2020.08.04.237057

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.

Li, Z., & Ma, W. (2020). An uncertainty-based model of the effects of fixation on choice. PsyArXiv. doi.org/10.31234/osf.io/ajmwxMilosavljevic, 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.

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.

Palestro, J. J., Weichart, E., Sederberg, P. B., & Turner, B. M. (2018). Some task demands induce collapsing bounds: Evidence from a behavioral analysis. Psychonomic bulletin & review, 25(4), 1225-1248.

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. (1978). A theory of memory retrieval. Psychological review, 85(2), 59.

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., & Rouder, J. N. (1998). Modeling response times for two-choice decisions. Psychological science, 9(5), 347-356.

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.

Roe, R. M., Busemeyer, J. R., & Townsend, J. T. (2001). Multialternative decision field theory: A dynamic connectionist model of decision-making. Psychological Review, 108,370–392.

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.

Wickelgren, W. A. (1977). Speed-accuracy tradeoff and information processing dynamics. Acta psychologica, 41(1), 67-85.

Supplementary Material

Regression Results for Individual Datasets

In the Results section of the main manuscript, we presented GLM regression results for the impact of value difference (dV), certainty (C), and certainty difference (dC) on choice consistency and log(RT), for all datasets pooled together (see Figure 4). Here we present the same results, but separately for each dataset (Figure S1).

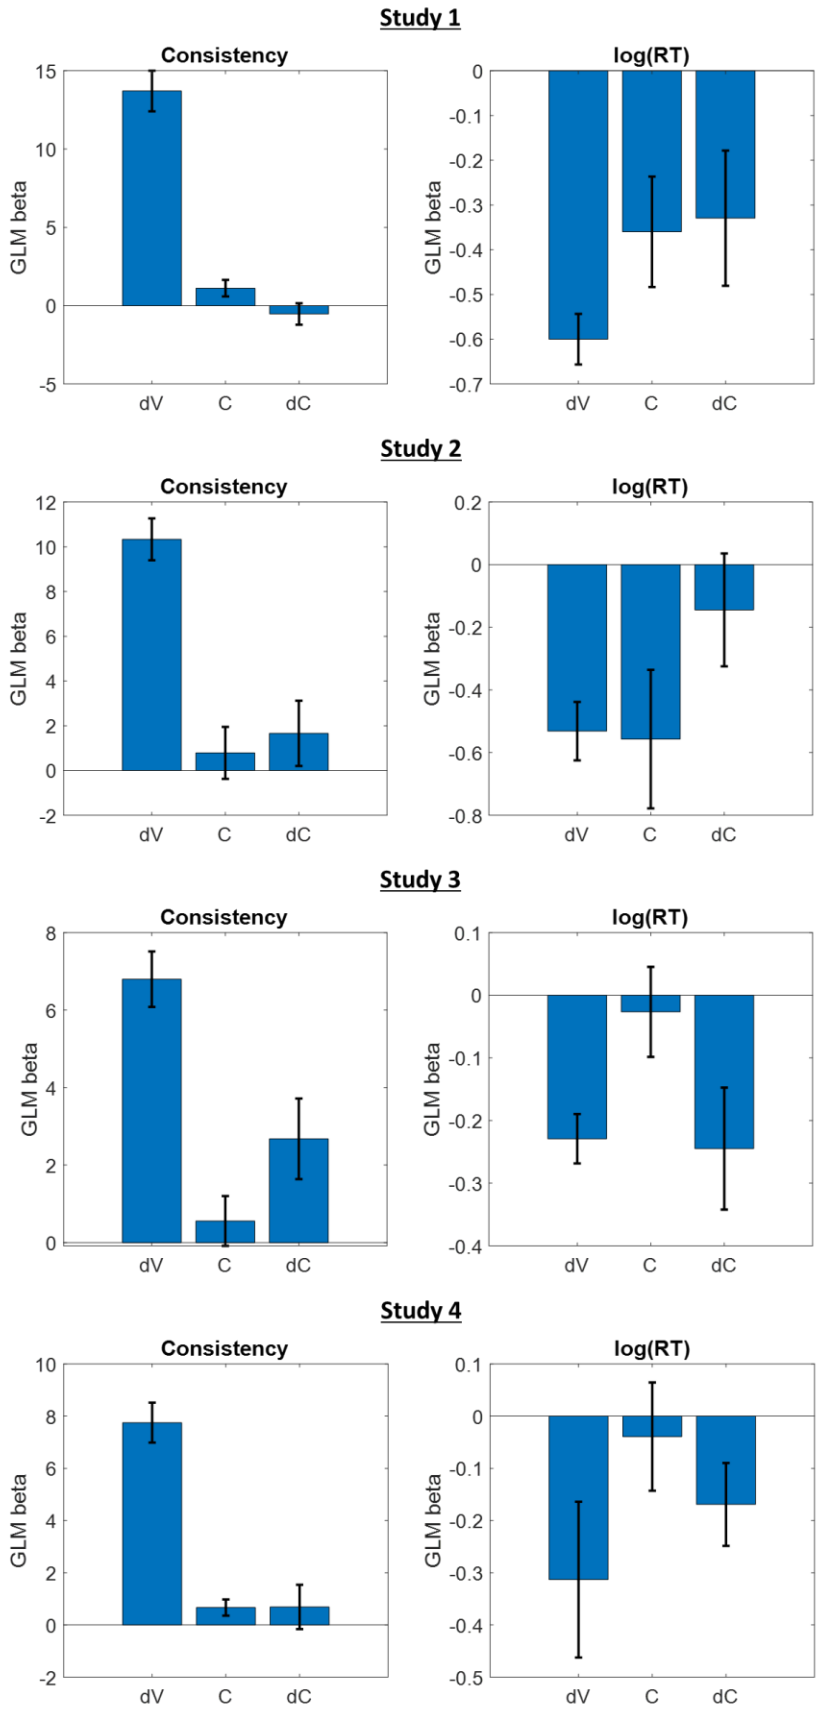
Study 1
Consistency
log(RT)
GLM beta
dV
C
dC
Study 2
Consistency
log(RT)
GLM beta
dV
C
dC
Study 3
Consistency
log(RT)
GLM beta
dV
C
dC
Study 4
Consistency
log(RT)
GLM beta
dV
C
dC

741

Figure S1: *Previous results demonstrate the relationships between value difference (dV = v1 - v2), value certainty (C = (c1 + c2) / 2), and certainty difference (dC = c1 - c2) with choice consistency and log(RT). We define v1 (v2) as the higher (lower) value rating, and c1 (c2) as the certainty report for the option with the higher (lower) value rating. Here we show the beta weights from a logistic regression on choice consistency and a linear regression on log(RT), separately for each study; 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.*

Model Comparison Results for Individual Datasets

In the Results section of the main manuscript, we presented the model comparison results for all studies pooled together (see Figure 5). Here we present the same results, but separately for each study (Figure S2). These results demonstrate that the victory of Model 5* over the other models was consistent across all four datasets, but that there was some variability in how well the losing models performed. In particular, Model 6 received some support in three of the four datasets, suggesting that it might be retained as a contender in future studies.

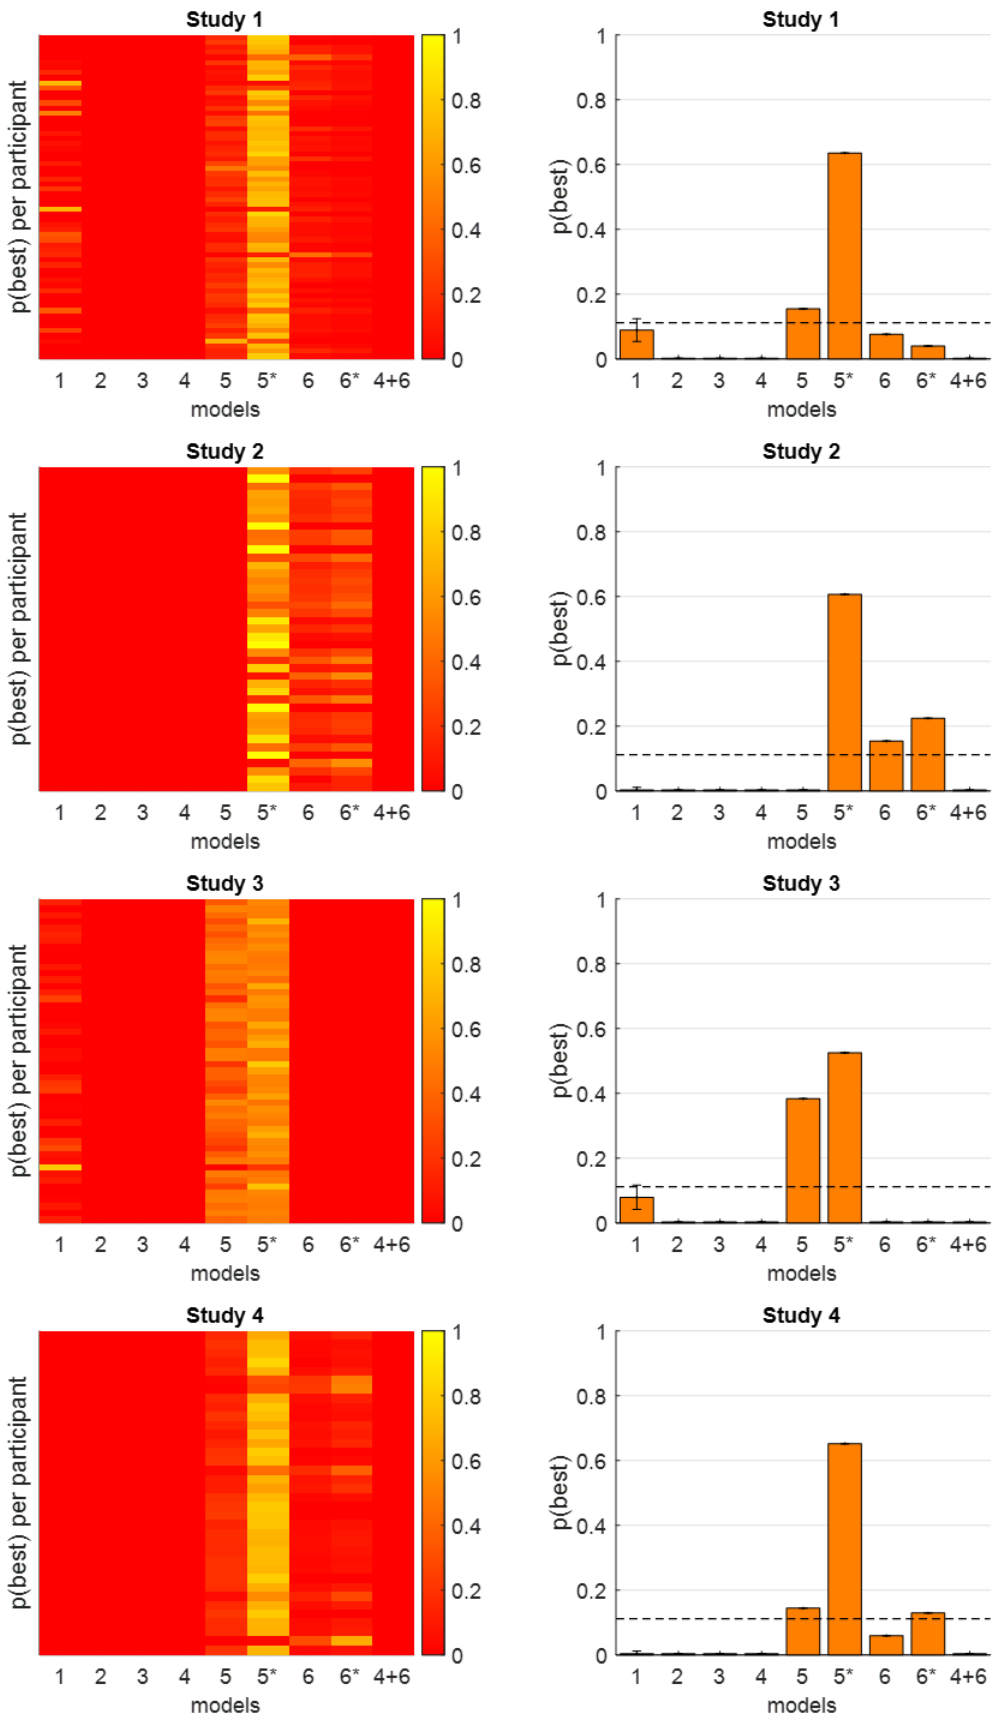
Study 1
p(best) per participant
models
1 2 3 4 5 5* 6 6* 4+6
0 0.2 0.4 0.6 0.8 1
Study 1
p(best)
Study 2
Study 2
Study 3
Study 3
Study 4
Study 4

757

Figure S2: *Model comparison results. We show here the probability that each model best accounted for the data at the participant level (left panel), separately for each study; each cell represents the probability that the model (column) best represents the behavior of the participant (row). We also show the probability that each model best explains the data across the participant population (right panel), separately for each study. The dashed line indicates chance level if all models were equally probable a priori.*

Simulation Results for Assumptions on the Impact of Value Certainty on Other DDM Parameters

In the Discussion section of the main manuscript, we explained why additional DDM variants (collapsing boundaries) or other assumptions about the impact of value certainty on the additional parameters in the standard DDM (drift rate variability, starting point variability) would not change the conclusions we presented here. Here we present simulation results that provide additional support to our conceptual explanations (Figure S3).

We start with the collapsing-boundary DDM variant. We show via simulations that if value uncertainty increased the sampling noise, this would decrease RT, just like in the fixed-boundary DDM (Figure S3A). We then consider the classic DDM with fixed boundaries and we examine whether value certainty could be introduced via an effect on either starting point variability or on drift rate variability (we consider the possibility that uncertainty might increase those variability parameters, so as to make to make the choice more random). We show via simulations that if value uncertainty increased the drift rate variability, this would decrease RT (Figure S3B). We also show via simulations that if value uncertainty increased the starting point variability, this would have no impact on mean RT (Figure S3C).

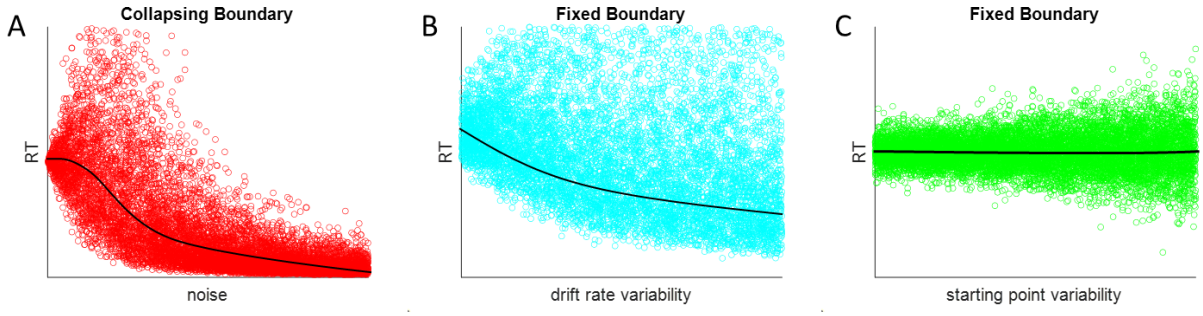

Figure S3: *Simulated results demonstrating the relationship between additional DDM parameters and RT. For each plot, we simulated 10^4 trials. Panel A show the relationships between*

diffusion noise and RT (see Figure 3), when the response boundaries collapse over time. Panel B shows the relationships between drift rate variability and RT. Panel C shows the relationship between starting point variability and RT. Each data point represents one simulated trial. The solid black lines indicate the mean RT as a function of the independent variable of interest.

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)

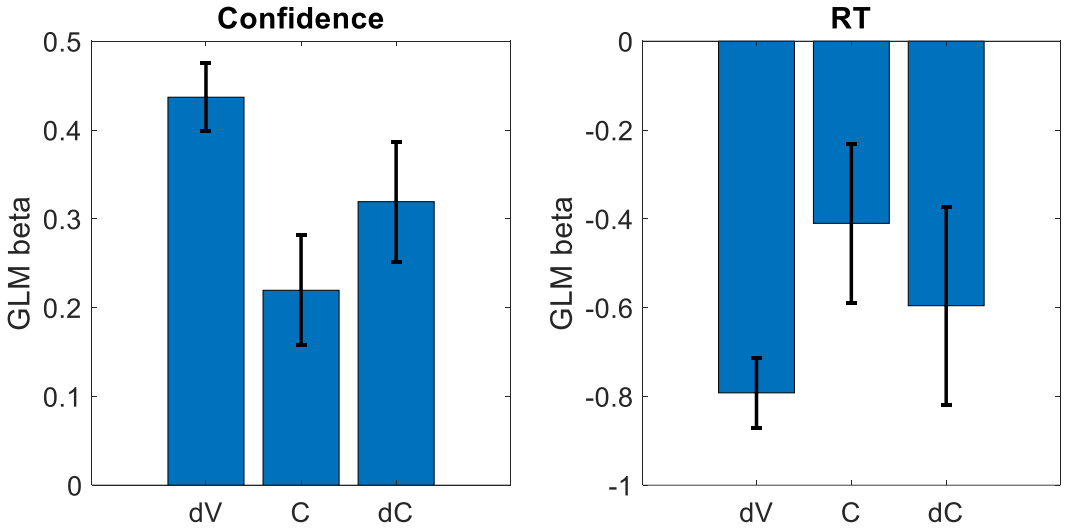

Figure S4: *Beta weights from a linear regression of dV, C, and dC on choice confidence (left) and RT (right; n=104; bar heights represent population means, error bars represent s.e.m.).*

DDM Median Split Parameter Comparison

To provide additional corroborating evidence that value certainty should relate to the drift rate and not the diffusion noise parameter, we performed classic median-split analyses on both the sum and the difference of the value certainties of the choice pair options. For each participant across all studies, we divided the trials based on a participant-specific median split of either certainty sum or certainty difference. We then fit the classic DDM separately for each of the

subsets of data and compared the fitted parameters across subsets. As shown in Figure S5, in both the certainty sum median split and the certainty difference median split, the drift rate parameter was larger in the high versus low subsets; the noise parameters did not differ across subsets. This shows that our conclusion that Model 5 (with drift rate adjusted by certainty) is valid, while Model 2 (with diffusion noise adjusted by certainty) is not.

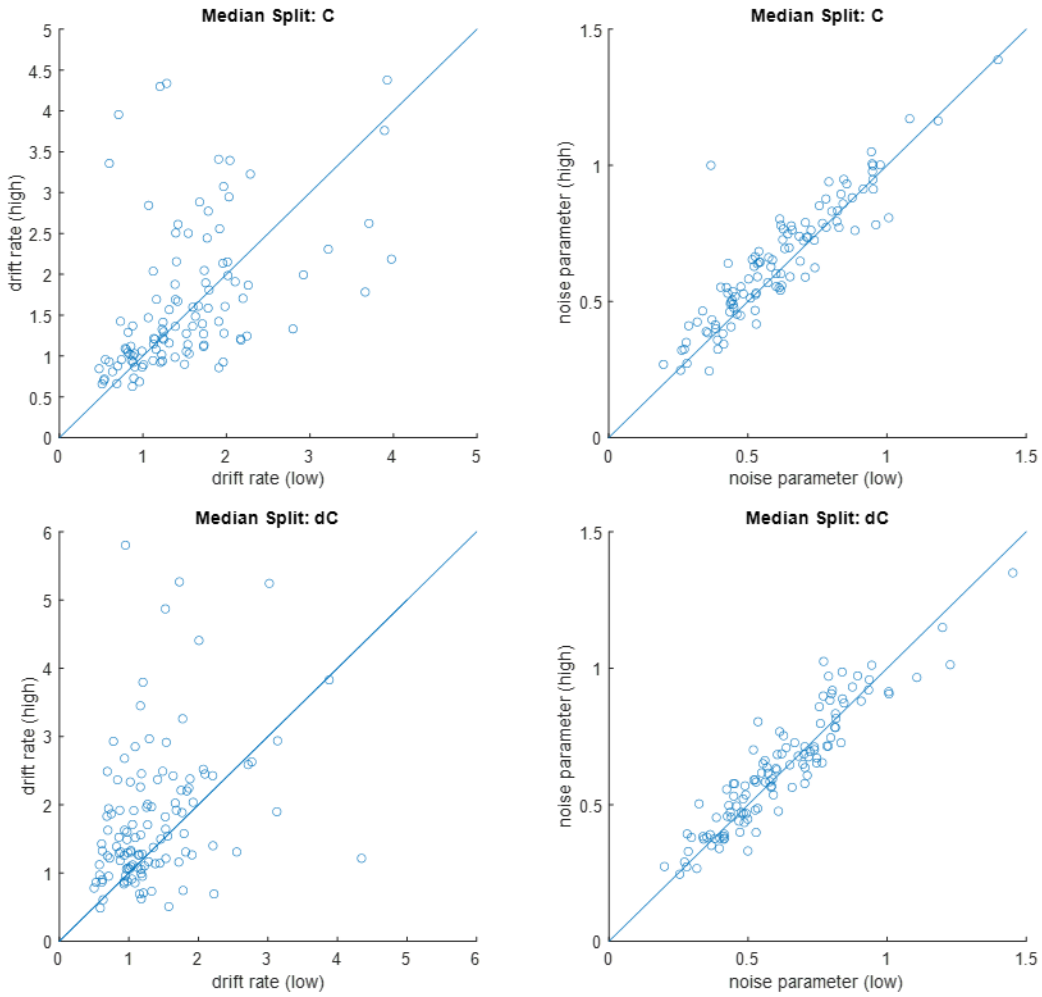

Figure S5: *DDM parameter comparison across median splits of certainty sum (top panels) and certainty difference (bottom panels). Each data point represents one participant. Across the population, the fitted drift rate parameters were larger for the high certainty trials as well as for the high certainty difference trials (left panels); the fitted diffusion noise parameters were not different across high and low certainty or certainty difference (right panels).*