Choice-Induced Preference Change under a Sequential Sampling Model Framework

Douglas G. Lee[1,2] Giovanni Pezzulo[1]

[1] Institute of Cognitive Sciences and Technologies, National Research Council of Italy

[2] Tel Aviv University, School of Psychological Sciences

Email for correspondence:

DouglasGLee@gmail.com

Abstract

Sequential sampling models of choice, such as the drift-diffusion model (DDM), are frequently fit to empirical data to account for a variety of effects related to choice accuracy/consistency and response time (RT). Sometimes, these models include extensions that can also account for choice confidence. However, no model in this class is able to account for the phenomenon of *choice-induced preference change*. Studies have reported choice-induced preference change for many decades, and the principle findings are robust: decision-makers tend to rate options higher after they choose them and lower after they reject them. This *spreading of alternatives* (SoA) in terms of their rated values is fundamentally incompatible with traditional sequential sampling models, which consider the rated values of the options to be stationary throughout choice deliberation. Here, we propose a simple modification of the basic DDM that allows the drift rate to vary across deliberation time depending on which attributes are attended to at which points in time. Critically, the model assumes that initial ratings are based only on the more salient attributes of the individual options, and that more attributes will be considered when decision-makers must choose between options with different salient attributes. We show that this model can account for SoA (in addition to choice consistency and RT), as well as all previously reported relationships between SoA and choice difficulty, attribute disparity, and RT.

Keywords: spreading of alternatives, drift-diffusion model, multi-attribute choice

Introduction

Contemporary research on value-based decision-making typically focuses on choices between options, using individual ratings obtained separately as independent variables to help explain choice behavior. This research has exposed many robust patterns with respect to both response (i.e., which option the agent chose from the set of available choice options) and response time (RT; the temporal delay between the agent becoming aware of the choice options and declaring its response). First among such patterns is that so-called difficult decisions (between options of similar subjective value) are generally more stochastic (i.e., predictions about which option will be chosen are less accurate) and take longer to conclude, relative to so-called easy decisions (Lee & Coricelli, 2020; Lee & Daunizeau, 2020, 2021; Lee & Holyoak, 2021a, 2021b; Milosavljevic et al., 2010; Ratcliff & McKoon, 2008). Decisions are also thought to be more difficult when the value estimates of the choice options are less certain (Gwinn & Krajbich, 2020; Lee & Coricelli, 2020; Lee & Daunizeau, 2020, 2021; Lee & Usher, 2021). Another common pattern is that decisions for which the value of the choice set as a whole is larger are usually made more quickly, and often more consistently (Hunt et al., 2012; Lee & Usher, 2021; Polanía et al., 2014; Sepulveda et al., 2020; Shevlin et al., 2022; Shevlin & Krajbich, 2021; Smith & Krajbich, 2019; Teodorescu et al., 2016). Finally, choices between options with high *attribute disparity* (e.g., where each option dominates within a different attribute dimension) are correspond with faster RT and perhaps higher consistency (Lee & Holyoak, 2021b). Each of these independent variables (absolute value difference, value certainty, value sum, and disparity) also has a robust positive impact on choice confidence (the agent's subjective belief that it chose the best option from the choice set).

Various computational models of simple decisions (e.g., two-alternative forced-choice or 2AFC) exist, the most common of which come from the sequential sampling / accumulation-to-bound class (Brown & Heathcote, 2008; Busemeyer & Townsend, 1993; Ratcliff et al., 2016; Usher & Mcclelland, 2001). Under this type of model, the agent processes information about the choice options incrementally until the so-called evidence in favor of one of the options passes a threshold and that option is declared to be the winner (and is therefore chosen). Although multiple variations of this core model have been proposed, they all generally rely on a relative value accumulator that tallies information about the options based on random samples from underlying probability distributions (one for each option) until it reaches a preset evidence threshold / response boundary. This class of model has provided an elegant account of choice and RT patterns in a variety of domains, in particular with respect to the impact of choice difficulty (value proximity; see (Lee & Usher, 2021) for a version that also includes value uncertainty as a moderator of choice difficulty). With some minor modifications, this type of model has also simultaneously accounted for the impact of overall set value on choice behavior (Krajbich et al., 2010; Lee & Usher, 2021; Smith & Krajbich, 2019). Even choice confidence can be accounted for by certain variations of sequential sampling model, either as a secondary readout of the evidence accumulation process (see (Calder-Travis et al., 2021)) or as the primary readout itself (Lee, Daunizeau, et al., 2022).

The evidence accumulation-to-bound framework has proven to be a powerful tool, yet there remains an important gap in its explanatory power. In spite of the success that it has had in accounting for most relevant variables in the study of simple decisions (i.e., accuracy, RT, and sometimes confidence), there remains one variable that it is fundamentally incapable of accounting for, by virtue of the very assumptions that it is built on. Though largely ignored in the

decision-making literature, the phenomenon known as *choice-induced preference change* (CIPC) has been well documented in the social psychology literature (Bem, 1967; Brehm, 1956; Festinger, 1957; Harmon-Jones & Mills, 2019) and more recently demonstrated in value-based decision-making studies (Chammat et al., 2017; Izuma et al., 2010, 2015; Lee & Coricelli, 2020; Lee & Daunizeau, 2020, 2021; Lee & Holyoak, 2021a, 2021b; Sharot et al., 2009, 2010, 2012; Voigt et al., 2017, 2019); see (Enisman et al., 2021) for a meta-analysis). With CIPC, a preference (direction and/or magnitude) between options that is implicitly inferred via isolated value ratings differs before versus after the preference that is explicitly observed via choices between pairs of options. The key variable here is often referred to as the *spreading of alternatives* (SoA), because it typically entails the values of the choice options spreading apart as the decision maker reassesses them during deliberation (see also the literature on *coherence shifts* and *information distortion*: (Carlson & Russo, 2001; Holyoak & Simon, 1999; Russo et al., 1996, 2008; D. Simon et al., 2001, 2004, 2008)).

Contrary to some interpretations, SoA is likely to result from cognitive processes and not merely statistical noise (Lee & Pezzulo, 2022). Studies such as those listed above have shown that SoA is a robust phenomenon, and is particularly high when choices are difficult (e.g., the difference in option values and/or the certainty about option values is low) or when attribute disparity is high (Lee & Holyoak, 2021b). It has also been shown that SoA and RT are strongly negatively correlated, and that SoA and choice confidence are strongly positively correlated (Lee & Coricelli, 2020; Lee & Daunizeau, 2020, 2021; Lee & Holyoak, 2021a, 2021b). Figure 1 summarizes these key empirical findings reported in previous studies.

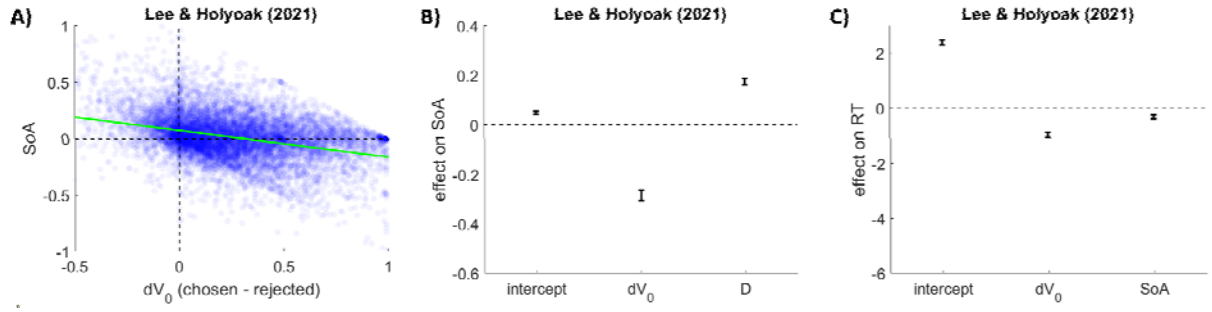

Figure 1: Empirical relationships between decision variables. A) Spreading of Alternatives (SoA) plotted against pre-choice value difference (value of the chosen option minus value of the rejected option); each dot represents one trial, pooled across participants. B) Fixed effects regression coefficients for the intercept, pre-choice value difference, and disparity on SoA; error bars represent 95% confidence intervals. C) Fixed effects regression coefficients for the intercept, pre-choice value difference, and SoA on RT; error bars represent 95% confidence intervals.

Sequential sampling models and the spreading of alternatives (SoA)

Lee & Daunizeau (Lee & Daunizeau, 2021) suggest that SoA occurs as a result of a tradeoff between mental effort and choice confidence. When faced with a difficult decision (e.g., if the options seem equally valuable at first and thus initial confidence about knowing the best option is low), a decision maker might prefer to invest mental effort towards processing additional information with the intention to distinguish the option values more confidently, rather than making an immediate choice with low confidence. SoA would therefore be instrumental to the decision as it unfolds. Sequential sampling models take an entirely different perspective, in which preferences are *retrieved* (e.g., from memory) and compared rather than being *constructed* during deliberation (see the literature on *constructed preferences*: (Ariely et al., 2006; DeKay et al., 2011; Lichtenstein & Slovic, 2006; Payne et al., 1999; D. Simon et al., 2008; Warren et al.,

2011)). Because the classical drift (evidence accumulation) rate is assumed to represent the (difference in the) true values of the options, there is no possibility for preferences to shift even were additional information to be considered.

The model of Tajima and colleagues (Tajima et al., 2016), which considers evidence accumulation to summarize a Bayesian updating of option values, could offer a different explanation of SoA (as prior estimates evolve to posterior estimates). This model assumes that both options are initialized with null (or at least equal) prior value estimates, but that the value estimates can come to differ as they are elucidated over time. From this perspective, preferences are not constructed during deliberation and do not change over time; they are simply more clearly recognized as their signals stabilize. However, SoA is typically calculated using subjective value ratings obtained both prior to and after the actual choice. It is unclear how the Tajima et al model might account for ratings that were systematically different when elicited after versus before the choice. Furthermore, this model would hold that value estimates could never shift closer together, and thus, importantly, decision makers could never change their minds during deliberation. Previous studies have shown such changes of mind (i.e., choices in favor of options that were initially rated lower, but later rated higher than the alternatives) to be common (Lee & Daunizeau, 2020, 2021; Lee & Holyoak, 2021a), suggesting that deliberation might sometimes cause an apparent preference reversal if the agent considers new information that it initially neglected. In standard accumulation-to-bound models, deliberation starts off with total ambivalence about preference, precluding the very possibility of a change of mind / preference reversal.

One way to potentially resolve this issue and enable this class of model to provide a measure of SoA would be to allow each option to enter deliberation with its own specific prior

value estimate and precision. Mathematically, this would entail initializing the evidence particle at a point of non-neutrality. Previous work has presented sequential sampling models with such a starting point bias, representing prior knowledge about which option is more likely *a priori* to be the better one, based on instruction (Mulder et al., 2012), choice history (Urai et al., 2019), category preference (Lopez-Persem et al., 2016), or personality trait (F. Chen & Krajbich, 2018). A starting point bias could also be used to capture option-specific priors. These priors would form immediately upon presentation of the choice options, during the time it takes for perceptual processing to inform the brain about the identities of the options, but before explicit and intentional value deliberation begins. This would be in line with research showing that value signals in the brain arise automatically (Lebreton et al., 2009), even when not relevant (e.g., during perception). The formation of priors in this way would not only explain the starting point bias, but also the non-decision time (NDT), which is sometimes included in sequential sampling models to represent the time required for stimulus encoding or visual pre-processing (Fontanesi et al., 2019; Mulder & Maanen, 2013; Nunez et al., 2017). Once the priors are established, the evidence accumulation process at the core of the model will begin, terminating when the desired evidence level is reached. The difference in the posterior value estimates of the options (i.e., those at the end of deliberation) could then be compared to the difference in the priors to calculate a measure of SoA. In this way, this type of model could also account for changes of mind, because the chosen option would not always be the option with the higher prior value estimate.[1]

[1] One further assumption that this approach would require is that the process resulting in individual value estimates for each option (during the pre-choice rating task in a typical experimental paradigm) is similar to that which establishes the initial imprecise value representation at the start of a decision (because SoA is measured experimentally using pre-choice ratings, as any latent value representation that forms at the beginning of a choice will not be observable). Although this assumption might seem unreasonable, it may nevertheless be reasonable to

The mechanisms presented above would result in SoA on average, and they would also account for the robust relationship between choice difficulty (absolute difference in certainty-adjusted value estimates) and SoA observed in the empirical data. But if deliberation is formalized as a process of recognition (or purification) of pre-existing value representations, the signals that inform the initial and final preferences would have identical means – it would only be the precision that would differ. So, any observed SoA would be purely due to noise. Previous studies have shown that mere statistical noise could indeed yield a positive correlation between choice difficulty and SoA (Alós-Ferrer & Shi, 2015; K. M. Chen & Risen, 2010; Izuma & Murayama, 2013; Lee & Pezzulo, 2022). However, empirical data across multiple previous studies supports a cognitive source of SoA rather than a purely artifactual one (see (Lee & Pezzulo, 2022)). A more satisfactory way to account for the SoA data would require a distinction between easy decisions (during which choices largely reflect prior value estimates and hence there is little or no SoA) and difficult decisions (during which novel sources of evidence are considered and hence there is often greater SoA). Choice difficulty might then be defined not based on some comprehensive measure of value, but rather on an incomplete measure of value based on whatever (presumably non-exhaustive) information was considered during the isolated rating task.

It seems plausible that the intra-decision dynamics of evidence accumulation might differ as a function of choice difficulty. During easy decisions, when the choice is between options with very different values, evidence based only on the priors might already surpass the target threshold, in which case the choice could be made immediately without the need for deliberation, or with little additional evidence (see (Alós-Ferrer, 2018; Caplin & Martin, 2016; Diederich &

assume that the process that yields isolated ratings lies somewhere between the very imprecise prior estimates and the relatively precise posterior estimates (i.e., it is a similar process but with an intermediate level of precision).

Trueblood, 2018; Pezzulo et al., 2013) for alternative models that allow for this sort of automatic choice)[2]. This would imply no significant differences between the initial and final preferences and hence no SoA for such easy decisions. Difficult decisions, in contrast, are purported to elicit higher levels of information processing (Lee & Daunizeau, 2020, 2021; Pezzulo et al., 2013), which could potentially perturb the means of the option value representations if newly-considered information is not fully congruent with previously-considered information. For example, it has been shown that information processing related to more salient or important attributes begins earlier than that related to less salient or important attributes (Lim et al., 2018; Maier et al., 2020; Sullivan et al., 2015). Furthermore, it has been proposed that the information considered in the formation of prior values and during deliberation is distinct and develops by means of passive retrieval from memory or active sampling via mental simulation, respectively (Pezzulo et al., 2013). Because the information considered at the start and at the end of a decision will not necessarily be the same, the value estimate for each option (not only its precision) could change over the course of deliberation. This suggests that the evidence accumulation rate in sequential sampling models should perhaps not be taken as static, but rather dynamic across deliberation time. We therefore propose an extension to a standard sequential sampling model in the form of a time-varying intra-trial drift rate. A dynamic drift rate combined with the notion that deliberation tends to last longer for difficult versus easy decisions could directly account for the observation that SoA positively correlates with choice difficulty.

The notion of an evidence accumulation rate that is dynamic across deliberation time is not new. It has been proposed that this rate should vary as a function of attention, biased in favor

[2] Alternatively, in tasks where participants are required to report a decision by reaching or clicking with the computer mouse one of two response buttons located far away from the "start" position, the prior could determine the initial direction of (finger or mouse) movement, which may eventually be revised online during deliberation (Barca & Pezzulo, 2012, 2012; Dotan et al., 2018; Lee, D'Alessandro, et al., 2022; Lepora & Pezzulo, 2015).

of the option that is visually fixated at any point in time (*attentional Drift-Diffusion Model* or aDDM; (Krajbich et al., 2010; Sepulveda et al., 2020; Yang & Krajbich, 2022)). Here, the drift rate oscillates between relatively high and relatively low, as gaze (and presumably, attention) shifts between the choice options. Other work has proposed a drift rate that changes mid-deliberation, as an additional attribute (related to both options) enters consideration (*start time DDM* or stDDM; (Maier et al., 2020; Sullivan et al., 2015)). Here, the drift rate changes at one specific point in time, as two attributes start to be compared rather than just one. Clearly, this could be expanded to include additional attributes, and thus additional changes in the intra-trial drift rate. This would also be in line with *Decision Field Theory* (DFT) models, where the drift rate fluctuates across deliberation time as (latent) attention shifts from one attribute to another (Busemeyer & Townsend, 1993; Diederich, 1997). Here, however, the momentary evidence accumulation rate is a function of only one attribute at a time, rather than an accumulation of attributes over time (as in the stDDM). Regardless of the differences between these models, they all demonstrate ways in which the preferences between options that are constructed during deliberation might differ depending on how attention (across options or across attributes) is dynamically allocated. Therefore, it is possible, in theory, that these models could account for the SoA phenomenon.

Here we advance a novel sequential sampling model designed to account for SoA, which combines key insights from three models that already include dynamic evidence accumulation (aDDM, DFT, stDDM). From these models, our new model would borrow the following notions: from the aDDM, evidence weights for specific options that diminish (but do not disappear entirely) when attention shifts away from those options; from DFT, attention that shifts between attributes (not only between options); and from the stDDM, additional attributes that come into

consideration as deliberation time advances. In addition, unique to our model, we extend the concept of attributes beyond the inherent features of the options themselves to include contextual aspects of the decision that might alter the mapping from attributes to value[3]. For example, when shopping for a new car, one might consider the size attribute. Initially, it could be that larger size is always deemed more valuable, as the most salient aspects of size are most easily recalled (e.g., more spacious seats, more room for luggage). However, upon further reflection (or mental simulation (Pezzulo et al., 2013)), larger size might come to be valued less, as other aspects are recalled (e.g., less room in the garage, more difficult to maneuver). Such contextual aspects might not be considered if the initial information is powerful enough to result in a fast choice. Note that these contextual aspects could be considered as sub-attributes of the higher-level attributes (i.e., size is merely a summary attribute that includes the true attributes involved in choosing the option: comfort, utility, capacity, flexibility). So, it could be said that deliberation involves not only the consideration of new attributes, but also the refinement of previous attributes. In our model, these could both be represented in the same mathematical format.

In summary, our proposed model would allow for momentary attention to focus on specific attributes of specific options, which would temporarily diminish information processing for all other attributes and all other options, with more (and more refined) attributes entering the competition for attention across deliberation time. In what follows, we present an analytical formulation for such a model and describe the key predictions that it makes with respect to choice behavior.

[3] Standard theory might posit that decisions are guided by the *retrieval* of value representations cached somewhere in the brain. In contrast, we support the notion that such value representations are rather *constructed* at the time of choice, possibly through the use of model-based computations. Note that this is similar to a key difference between model-free and model-based reinforcement learning (or active inference, or planning-as-inference).

Model specification

The core foundation of our model is a standard evidence accumulation process, with:

$$E_{t+1} = E_t + d(\mu_1 - \mu_2 + \varepsilon)$$

$$\varepsilon \sim N(0, \sigma^2)$$

where E is the momentary evidence for one option over the other, t is an index for deliberation time (arbitrary units), d is an agent-specific scalar representing evidence efficiency, μ_i is the value estimate of option i, and σ^2 is the strength of the white processing noise ε common to every decision. Choice probability and expected response time can be analytically calculated using the following equations (Cox & Miller, 2017):

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

$$E[RT] = \frac{\theta(2P - 1)}{d(\mu_1 - \mu_2)}$$

where θ is the magnitude of the evidence threshold for a response.

Expanding this so that the process is based on multiple attributes would be straightforward, replacing the overall value estimates μ_i with estimates of individual attributes μ_i^a and the drift scalar d with attribute-specific drift scalars d^a:

$$E_{t+1} = E_t + \sum_{a=1}^{n} d^a(\mu_1^a - \mu_2^a + \varepsilon)$$

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

$$E[RT] = \frac{\theta(2P - 1)}{\sum d^a(\mu_1^a - \mu_2^a)}$$

where n is the dimensionality of the attribute space for the decision at hand.

The next component of our model allows for the relative weights of the attribute dimensions to vary across time, as the attentional allocation of the agent fluctuates. The process then becomes:

$$E_{t+1} = E_t + \sum_{a=1}^{n} d_t^a (\mu_1^a - \mu_2^a + \varepsilon)$$

$$P = \frac{1}{1 + e^{\left(\frac{-2\theta \sum d_t^a (\mu_1^a - \mu_2^a)}{\sigma^2}\right)}}$$

$$E[RT] = \frac{\theta(2P - 1)}{\sum d_t^a (\mu_1^a - \mu_2^a)}$$

All that remains to complete the computational description of our model is to specify how the attribute weights d_t^a evolve across time. We will start by setting $d_0^a = 0 \; \forall \; a \in A = \{1,2,\ldots,n\}$. For clarity, we will assume that only one attribute will become active at a time (e.g., $d_1^i > 0 \; and \; d_1^j = 0 \; \forall \; j \in A \backslash i$). The attribute that becomes active first will be the most salient attribute, which we define as the attribute dimension that has the highest individual rating for any given set of choice options. More specifically, the most salient attribute of the first option that the agent gazes at will become active first, and the most salient attribute of the next option that the agent gazes at will become active next (if such attributes are distinct). While active, the magnitude of an attribute weight (i.e., the contribution of evidence regarding that attribute towards the eventual choice) will be determined by the amount of attention the agent allocates towards that dimension (subject to a scaling factor specific to the agent and common across attributes). Here, the attention need not be conscious in the sense that the agent could report it if queried, but rather in the more general sense of the focus of information processing. Attention will be attracted towards more inherently salient dimensions, and it can be willfully turned towards dimensions that are more important to the specific agent or decision context.

Whether or not additional attributes (beyond the initial, most salient ones) will become active, and the point in time at which this will occur, will depend upon how informative the attributes already under consideration are with respect to the decision at hand (or in other words, the value of information or VOI that they offer; see (Behrens et al., 2007; Kobayashi et al., 2021; Pezzulo et al., 2013)). If the subset of attributes currently attended to provides sufficient evidence that a specific option is the best, the agent might make its response without probing additional attribute dimensions. On the contrary, if the agent is unable to sufficiently distinguish the options based on the current active attribute subset, it might start to consider (i.e., allocate attention towards) an additional attribute – or proceed to simulate novel outcomes that may inform the decision. Furthermore, if the agent has already considered a particular attribute and concluded that that attribute dimension was unhelpful in distinguishing the option values, it might cease to consider that attribute for the remainder of the deliberation. This will be a function of the precision with which the agent has assessed the current attributes. For each estimate of the value of an attribute, an agent will also have a sense of precision about that estimate. Low precision will encourage further information processing, whereas high precision might discourage further information processing (as it would be redundant and thus inefficient to continue).

To better illustrate the process, take for example a simple choice between two snack food options: a chocolate bar and an orange. Let us assume that the most salient attribute of either option is its tastiness, so the agent will first consider that dimension. Imagine that after a brief deliberation, the agent realizes that both options are very similar in terms of tastiness and that it is difficult to choose which it prefers. The agent might then start to consider the healthiness of the options. If the agent were highly certain about its evaluation of the tastiness of the options,

and thus highly certain that it would be of no use to consider that dimension further, it might replace tastiness with healthiness as the sole attribute under consideration (at that point in time). Otherwise, it might start to consider both attributes in parallel, with attention split between them. Clearly, this process could iterate until either all attributes had been considered or a subset of attributes had successfully provided enough evidence that either the chocolate bar or the orange was better than the other. Alternatively (or additionally), the agent might be motivated to mentally simulate the experience of selecting either option, potentially adding context-dependent attributes (e.g., the chocolate bar might melt if kept in the agent's pocket, the orange might get moldy if not eaten soon) to the active set.

The above description provides an account for how the spreading of alternatives (SoA) between choice alternatives might arise *during* a decision. However, such intra-decisional changes in evaluations are typically not observable (but see (Lee & Holyoak, 2021a)). As such, experimental evidence for SoA generally comes in the form of isolated evaluations of the choice options, which are (on average) more separated after the choice task relative to before. Under the assumptions of our model described above, one would need to assume that pre-choice evaluations are made based on a smaller subset of attributes than the subset that determines the post-choice evaluations. In the example above, this might mean that the chocolate bar and the orange were both initially (pre-choice) rated very highly (in isolation), based entirely on their respective tastiness. During a decision in which these options were pitted against each other, however, both the tastiness and the healthiness attributes might have been considered before the orange was chosen as the best option. Then, when the items were later re-evaluated in isolation (i.e., for the post-choice ratings), the orange might be rated higher than the chocolate bar, because now the nutrition attribute (which favored the orange) is included in the overall rating.

We focus here on identifiable attributes, but in a more general sense, the idea is that earlier assessments (either in isolation or in tandem) are less information-rich than later assessments, provided that the agent has proper motivation to seek enrichening information (Lee & Daunizeau, 2020, 2021). Most likely, the posterior beliefs about option values from one task will inform the prior beliefs for a subsequent related task, following standard Bayesian computations.

Model predictions

Using simulation, we here illustrate the core predictions of our model and validate them by qualitative comparison with empirical findings. We first generated 100,000 choice trials, each containing a unique pair of options. Each option was composed of two attributes, each of which was assigned a random value from the interval: (0,1). The value for each option was calculated as the weighted sum of its attributes (with the first attribute arbitrarily weighing twice as much as the second), and the value difference (dV) for each choice trial was calculated as the magnitude of the difference between the option pair values. We then proceeded to simulate a drift-diffusion process under different scenarios, as described below. For each trial, we recorded the choice response and the response time (RT). Using the data obtained from each simulation run, we tested for correlations between the relevant decision variables: dV, RT, and spreading of alternatives (SoA). In our simulated data, we calculated SoA as dV at the end of each trial (dV_T) minus dV at the start of each trial (dV_0). For all of the formulations across all considered scenarios, we defined dV_T as the total amount of evidence accumulated by the end of the trial (i.e., at time T) divided by T. In this way, dV_T can be thought of as the "observed" evidence accumulation rate, whereas dV_0 can be thought of as the "anticipated" evidence accumulation rate. For trials on which the second option was chosen, the SoA was multiplied by -1 (since dV

was calculated as the value of the first option minus the value of the second option, but SoA is calculated with respect to the chosen option).

First, we considered a scenario where the decision was governed by a basic evidence accumulation-to-bound process (e.g., the classic DDM), where the evidence represents the difference in the aggregated subjective value of each option (i.e., $dV_0 = dV$). Previous studies have shown that SoA can arise simply due to noise in the valuation process, even if that process is static ((Alós-Ferrer & Shi, 2015; K. M. Chen & Risen, 2010; Izuma & Murayama, 2013; Lee & Pezzulo, 2022); note, however, that the noise in those studies was related to isolated ratings and not choices). We confirmed that in data simulated under a basic DDM, the average SoA across all trials is positive. We also found that this data shows a negative relationship between dV_0 and SoA, as seen in the empirical results (Figure 2A). However, GLM regression of SoA on dV_0 and disparity (D) shows that the predictive effect of D on SoA is negative, which does not match the empirical findings (Figure 2B). GLM regression of RT on dV_0 and SoA shows that both regressors have a negative predictive effect, which does match empirical findings (Figure 2C).

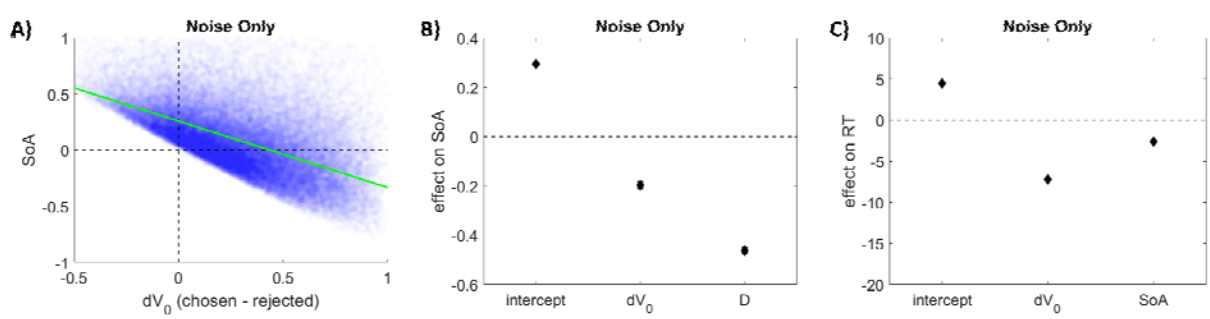

Figure 2: Simulated relationships between decision variables – Noise Only Model. A) Spreading of Alternatives (SoA) plotted against pre-choice value difference (value of the chosen option minus value of the rejected option); each dot represents one simulated trial. B) Regression

coefficients for the intercept, pre-choice value difference, and disparity on SoA. C) Regression coefficients for the intercept, pre-choice value difference, and SoA on RT.

Next, we considered a sequential sampling model where the evidence accumulation rate was not stochastic but rather deterministically driven by the attributes (cf, (Brown & Heathcote, 2008)). For this model, we assume that one attribute is more important (or salient) to the agent and thus the deliberation process begins by considering only that attribute (A1). We here introduce a new term: *relative start time* (RST). The RST term establishes the point in the deliberation process where the next attribute (A2) starts to be considered by the agent. This is in line with previous work showing that versions of multi-attribute DDM that include RST outperform those that do not (Lim et al., 2018; Maier et al., 2020). We set RST equal to 1000 "time step" iterations in our simulation. So, for $t < 1000$, evidence accumulated only for the difference in A1 (dA1) across the pair of options on each trial, but for $t >= 1000$, evidence accumulated for both dA1 and dA2. In this simulated data, the average SoA across all trials was negative (which does not match the empirical findings), although we did find the correct negative relationship between dV_0 and SoA (Figure 3A). Furthermore, this data showed a mismatch with the empirical results related to both the predictive effect of D on SoA (Figure 3B) and the predictive effect of SoA on RT (Figure 3C). It seems that noise in the process might be necessary to reproduce the observed empirical effects, but that noise alone cannot account for everything.

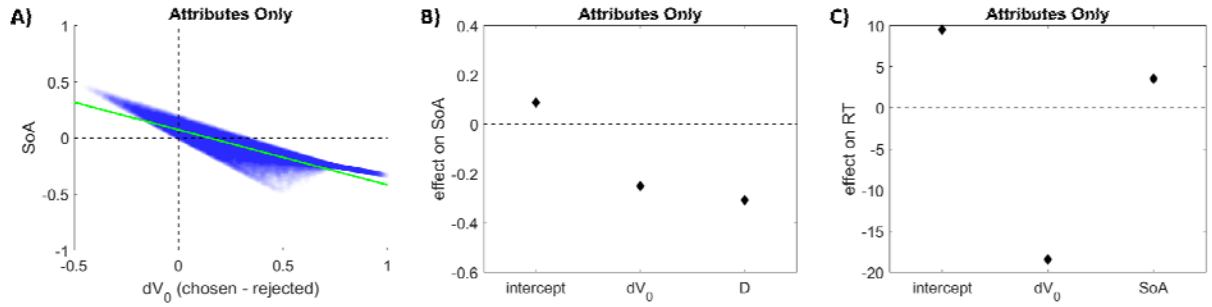

Figure 3: Simulated relationships between decision variables – Attributes Only Model. A) Spreading of Alternatives (SoA) plotted against pre-choice value difference (value of the chosen option minus value of the rejected option); each dot represents one simulated trial. B) Regression coefficients for the intercept, pre-choice value difference, and disparity on SoA. C) Regression coefficients for the intercept, pre-choice value difference, and SoA on RT.

We then combined the first two models, such that evidence accumulated for individual attributes depending on RST, and the process also included diffusion noise. The match between our simulated data and the empirical data seemed to improve, although it remained less than perfect (see Figure 4). Specifically, we found positive overall SoA, a negative relationship between dV_0 and SoA, and a negative predictive effect of SoA on RT (all in line with the empirical findings). However, we still found the unwanted negative predictive effect of D on SoA.

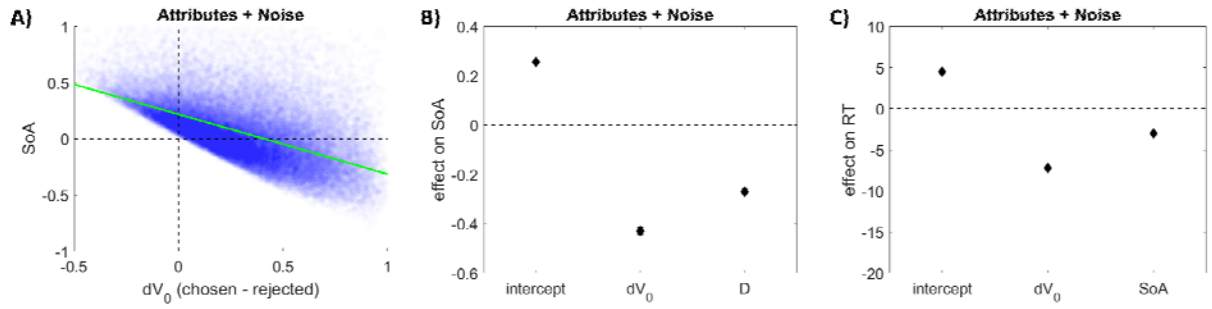

Figure 4: Simulated relationships between decision variables – Attributes + Noise Model. A) Spreading of Alternatives (SoA) plotted against pre-choice value difference (value of the chosen option minus value of the rejected option); each dot represents one simulated trial. B) Regression coefficients for the intercept, pre-choice value difference, and disparity on SoA. C) Regression coefficients for the intercept, pre-choice value difference, and SoA on RT.

It could be the case that when initially considering each option in isolation (e.g., during a pre-choice rating task), different attributes might be salient for different options. In that case, for some options, dV_0 might be related to A1, whereas for other options, dV_0 might be related to A2. But when choosing with a specific goal in mind, there could still be a default primary attribute that would be considered before the others. With this in mind, we designed an alternative version of the model described above (Attributes + Noise). Here, we defined dV_0 as the difference in the value of the salient attribute of each option (regardless of whether they were the same attribute or not). Then, the evidence accumulation process was as described above, but with one difference. On trials where the salient attribute for both options was A1, RST was set to 1000. On all other trials, RST was set to 0. This establishes a positive average RST that diminishes when the "non-default" attribute (i.e., A2 in our simulations) is more salient. Under this model formulation, our simulated data successfully reproduce all relevant patterns in the empirical data (see Figure 5).

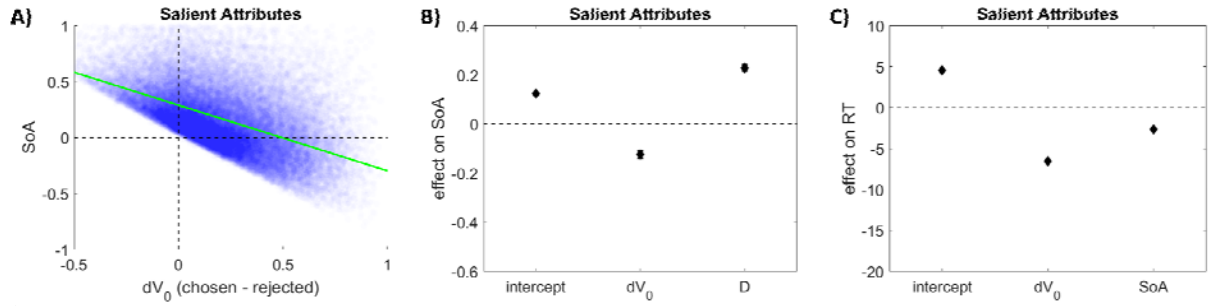

Figure 5: Simulated relationships between decision variables – Salient Attributes Model. A) Spreading of Alternatives (SoA) plotted against pre-choice value difference (value of the chosen option minus value of the rejected option); each dot represents one simulated trial. B) Regression coefficients for the intercept, pre-choice value difference, and disparity on SoA. C) Regression coefficients for the intercept, pre-choice value difference, and SoA on RT.

Discussion

In this work, we proposed a simple extension to classic evidence accumulation-to-bound models of value-based decision-making that can simultaneously account for response probability, response time (RT), and choice-induced preference change (as measured by the spreading of alternatives or SoA). We focused on the drift-diffusion model (DDM) because of its simplicity and analytical tractability, and also because of its widespread success in accounting for standard decision variables (response probability and RT). Ours is the first computational model of choice that provides a principled account for how SoA could arise during choice deliberation. Indeed, we are the first to explore such a possibility.

Our model assumes that the decision processes in the brain begin deliberation with imprecise representations of the values of each option, likely constructed from a subset of relevant attributes. Evidence starts to accumulate for only the most salient attributes at the beginning of deliberation, and evidence related to other attributes starts to accumulate at some later time point (Lim et al., 2018; Maier et al., 2020; Sullivan et al., 2015). We show via simulation that this is sufficient to give rise to SoA, complete with the correct (i.e., in line with the empirical findings) relationships between SoA and both choice difficulty and attribute disparity, as well as between RT and SoA. According to our theoretical model, the attention of

the agent can switch between attributes over the course of deliberation (Busemeyer & Townsend, 1993; Diederich, 1997), and attention will modulate the rates of evidence accumulation (Krajbich et al., 2010; Yang & Krajbich, 2022). These theoretical elements of our model yield dynamic intra-trial evidence accumulation rates that could often result in posterior value estimates (i.e., at the time of choice response or during a post-choice rating task) that differ from the prior value estimates (i.e., at the beginning of choice deliberation or during a pre-choice rating task). Thus, the well-established SoA phenomenon can be fully explained by a principled process-based model, removing the need to rely on more descriptive accounts such as the post-choice resolution of cognitive dissonance (Brehm, 1956; Festinger, 1957).

As a first step towards better understanding the elements that give rise to SoA during deliberation, we considered several simplified versions of the same core model. We first examined a model where all attributes always enter contemplation from the very beginning of deliberation. Under this model, SoA arises entirely as a result of stochasticity in the evidence accumulation process. Next, we considered a model where there is no noise, but different attributes enter contemplation at different points in time. We then considered a model that includes both noise and different contemplation starting times for each attribute. Each of these models was able to reproduce the general pattern of SoA with respect to mean levels and its relationship with choice difficulty (dV0), but none was able to account for all of the other key relationships reported in the literature: namely, the positive relationship between attribute disparity (D) and SoA, and the negative relationship between SoA and RT (Lee & Holyoak, 2021b). The final model that we considered is similar in form to the “attributes + noise” model, but it relies on the concept of attribute saliency to distinguish the trial-by-trial dynamics of

evidence accumulation. Only this model was able to recreate all of the patterns of relationships between the decision variables that we examined.

One feature of our model that sets it apart from other potential formulations is its ability to match the empirical relationship between attribute disparity and SoA. In brief, SoA is highest when the attribute composition of the options in a choice set is very different (high disparity), all else equal. In order to account for this, our model relies on an assumption that disparity relates to the relative salience of the individual attributes across options. With low disparity, attribute composition is similar across options. Therefore, the most salient attribute is likely to be the same for all options. This would allow decisions to be more likely to be made based on only that mutually-salient attribute. On the contrary, with high disparity, attribute composition is dissimilar across options. Therefore, the most salient attribute is likely to be different for different options. This would make it more likely that decisions would be made based on multiple attributes, which would give rise to SoA (because pre-choice ratings would have been made based on only the attributes that were salient for each isolated option). In sum, the *basis of valuation* will differ as a function of disparity. SoA will occur primarily during when the basis (mathematically, the set of attribute dimensions that comprise the value space where options are evaluated) changes between isolated ratings and comparison with alternative options, which is exactly the situation of high disparity.

Our model relies on the assumption that post-choice ratings will resemble the latent value estimates that are constructed during choice deliberation up until the time a choice is made. This idea would generally align with the idea that value estimates are updated in a Bayesian manner during deliberation (Tajima et al., 2016), where the posterior estimates would endure without decaying back to the priors. It has been shown that the SoA effect is maximal *during* choice

deliberation, although it then partially diminishes during a subsequent rating task (Lee & Holyoak, 2021a). This suggests that post-choice ratings might be based on some combination of the information on which pre-choice ratings were based and the information on which choices were based. It seems that during isolated evaluations after choices have recently been made, the agent will tend to focus on the attributes inherently salient for each option, but will also consider those attributes that contributed to the basis of comparison during the relevant choices. It is as if thinking about non-salient attributes during choice deliberation made those attributes more likely to be activated when considering the options again in the future (even in isolation). Future work should investigate the transient nature of SoA, as well as the neural underpinnings of bases of comparison during pre-choice rating, choice, and post-choice rating tasks.

Most studies related to SoA, including this one, consider only changes in ratings of overall value, but are agnostic about ratings of individual attributes. However, it has also been shown that SoA occurs even at the level of individual attributes (Lee & Holyoak, 2021b). In other words, attribute evaluations tend to favor the chosen options more after the choices are made, relative to before (Lee & Holyoak, 2021b). Perhaps SoA at the attribute level is caused by the same mechanisms that we have proposed for SoA at the level of overall value. This could be formalized as weights of sub-attributes (onto attributes) that change across deliberation time (or during choice versus rating tasks). Future work should explore sub-attributes and how their relationship to attributes does or does not mirror the relationship of attributes to overall value.

Finally, the theory advanced in this paper contributes to the broader ongoing debate amongst researchers of value-based decision-making as to how the notion of value is represented in the minds of decision-makers. The idea that value-related representations emerge at the time that they are called upon (e.g., when evaluating individual options or choosing between multiple

options) is itself not widely contested. The debate is rather about *how* the value representations emerge. Some believe that option values are stored in memory and retrieved when needed (Padoa-Schioppa, 2011). While this hypothesis is intuitive, it is difficult to reconcile with the frequent observation that values reported for the same options differ depending on context. Others believe that value is constructed anew with each new task, built by an agglomeration of assessments of individual attributes of the option under consideration weighted by the importance of each attribute within the context where it is being evaluated (Bettman et al., 1998; Bhatia & Stewart, 2018; Johnson & Payne, 1985; Levav et al., 2010; Payne et al., 1988, 1993; Russo et al., 1996; Shah & Oppenheimer, 2008). Still others hold that memories of previous interactions with the option under consideration are retrieved (or imaginary interactions are simulated) and compiled into summary statistics (mean and variance across the sample of memories; (Ronayne & Brown, 2017; Stewart et al., 2006)). While each of these perspectives holds some merit, there is another aspect that is not directly addressed by this particular debate: value representations need not be the identical across tasks (e.g., rating versus choice) – representations could have higher or lower fidelity depending on the task. In the ideal, there may be no reason for the value of an option (to a specific individual in a specific context) to be task-dependent. However, cognitive operations are not cost-free (H. A. Simon, 1957; Zénon et al., 2019). Therefore, it is reasonable to believe that the sophistication of value representations formed with the purpose of enabling different tasks should align with motivational factors associated with each task. Specifically, the agent should conserve cognitive resources when establishing a value representation, investing only enough to satisfactorily perform the task at hand (Lee & Daunizeau, 2021; Sims, 2016; Warren et al., 2011). In other words, the agent might intentionally neglect some information in order to balance against the costs of information

processing (Caplin & Dean, 2015; Fudenberg et al., 2018). What constitutes satisfactory performance is likely to differ between rating and choice tasks (for example). For an isolated individual rating, it might suffice simply to ascertain the valence of the option – the agent will want to approach the option if it has a value greater than zero, but to avoid the option if it has a value less than zero. Knowing that the rating will not be entirely isolated, in the sense that many different options need to be individually assessed, might induce a somewhat more refined value representation – the agent will want to accurately report its preferences, but will not be motivated to be highly precise. For a choice between options, on the contrary, the agent might not be able to perform the task satisfactorily (i.e., with a high enough feeling of confidence) unless it refines the value representations of the options by allocating additional cognitive resources. Of course, this should only be the case if decisions are relatively difficult (i.e., when low-fidelity value representations would not be easily discriminable). In sum, we propose that value representations for individual ratings will be lower fidelity than for choices (in general). We also propose that value representations that emerge during choice deliberation will initially be of lower fidelity (similar to that for ratings), but will convert to higher fidelity as necessary (i.e., if the decision is initially too difficult). At this point, our theory suggesting that decision-makers might use more- or less-refined value representations in a task-related way is speculative. This hypothesis remains to be tested in future studies.

Code Availability

The analysis code used in preparation of this manuscript is publicly available on the Open Science Framework at https://osf.io/2bs73/.

Funding

This research received funding from the European Union's Horizon 2020 Framework Programme for Research and Innovation under the Specific Grant Agreement Nos. 785907 and 945539 (Human Brain Project SGA2 and SGA3) to GP, the European Research Council under the Grant Agreement No. 820213 (ThinkAhead) to GP.

References

Alós-Ferrer, C. (2018). A Dual-Process Diffusion Model. *Journal of Behavioral Decision Making*, *31*(2), 203–218. https://doi.org/10.1002/bdm.1960

Alós-Ferrer, C., & Shi, F. (2015). Choice-induced preference change and the free-choice paradigm: A clarification. *Judgment and Decision Making*, *10*(1), 16.

Ariely, D., Loewenstein, G., & Prelec, D. (2006). Tom Sawyer and the construction of value. *Journal of Economic Behavior & Organization*, *60*(1), 1–10. https://doi.org/10.1016/j.jebo.2004.10.003

Barca, L., & Pezzulo, G. (2012). Unfolding visual lexical decision in time. *PloS One*, *7*(4), e35932.

Behrens, T. E. J., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. S. (2007). Learning the value of information in an uncertain world. *Nat Neurosci*, *10*(9), 1214–1221. https://doi.org/10.1038/nn1954

Bem, D. J. (1967). Self-perception: An alternative interpretation of cognitive dissonance phenomena. *Psychological Review*, *74*(3), 183–200. https://doi.org/10.1037/h0024835

Bettman, J. R., Luce, M. F., & Payne, J. W. (1998). Constructive Consumer Choice Processes. *Journal of Consumer Research*, *25*(3), 187–217. https://doi.org/10.1086/209535

Bhatia, S., & Stewart, N. (2018). Naturalistic Multiattribute Choice. *Cognition*, *179*, 71–88. https://doi.org/10.1016/j.cognition.2018.05.025

Brehm, J. W. (1956). Postdecision changes in the desirability of alternatives. *The Journal of Abnormal and Social Psychology*, *52*(3), 384–389. https://doi.org/10.1037/h0041006

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

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. https://doi.org/10.1037/0033-295X.100.3.432

Calder-Travis, J., Bogacz, R., & Yeung, N. (2021). Expressions for Bayesian confidence of drift diffusion observers in dynamic stimuli tasks. *BioRxiv*, 2020.02.25.965384. https://doi.org/10.1101/2020.02.25.965384

Caplin, A., & Dean, M. (2015). Revealed Preference, Rational Inattention, and Costly Information Acquisition. *American Economic Review*, *105*(7), 2183–2203. https://doi.org/10.1257/aer.20140117

Caplin, A., & Martin, D. (2016). The Dual-Process Drift Diffusion Model: Evidence from Response Times. *Economic Inquiry*, *54*(2), 1274–1282. https://doi.org/10.1111/ecin.12294

Carlson, K. A., & Russo, J. E. (2001). Biased interpretation of evidence by mock jurors. *Journal of Experimental Psychology: Applied*, *7*(2), 91–103. https://doi.org/10.1037/1076-898X.7.2.91

Chammat, M., Karoui, I. E., Allali, S., Hagège, J., Lehongre, K., Hasboun, D., Baulac, M., Epelbaum, S., Michon, A., Dubois, B., Navarro, V., Salti, M., & Naccache, L. (2017). Cognitive dissonance resolution depends on episodic memory. *Scientific Reports*, *7*. https://doi.org/10.1038/srep41320

Chen, F., & Krajbich, I. (2018). Biased Sequential Sampling Underlies the Effects of Time Pressure and Delay in Social Decision Making. *Nature Communications*, *9*(1), 1–10. https://doi.org/10.1038/s41467-018-05994-9

Chen, K. M., & Risen, J. L. (2010). How choice affects and reflects preferences: Revisiting the free-choice paradigm. *Journal of Personality and Social Psychology*, *99*(4), 573–594. https://doi.org/10.1037/a0020217

Cox, D. R., & Miller, H. D. (2017). *The theory of stochastic processes*. Routledge.

DeKay, M. L., Stone, E. R., & Miller, S. A. (2011). Leader-driven distortion of probability and payoff information affects choices between risky prospects. *Journal of Behavioral Decision Making*, *24*(4), 394–411. https://doi.org/10.1002/bdm.699

Diederich, A. (1997). Dynamic Stochastic Models for Decision Making under Time Constraints. *Journal of Mathematical Psychology*, *41*(3), 260–274. https://doi.org/10.1006/jmps.1997.1167

Diederich, A., & Trueblood, J. S. (2018). A dynamic dual process model of risky decision making. *Psychological Review*, *125*(2), 270–292. https://doi.org/10.1037/rev0000087

Dotan, D., Meyniel, F., & Dehaene, S. (2018). On-line confidence monitoring during decision making. *Cognition*, *171*, 112–121. https://doi.org/10.1016/j.cognition.2017.11.001

Enisman, M., Shpitzer, H., & Kleiman, T. (2021). Choice changes preferences, not merely reflects them: A meta-analysis of the artifact-free free-choice paradigm. *Journal of Personality and Social Psychology*, *120*(1), 16–29. https://doi.org/10.1037/pspa0000263

Festinger, L. (1957). *A Theory of Cognitive Dissonance*. Stanford University Press.

Fontanesi, L., Palminteri, S., & Lebreton, M. (2019). Decomposing the effects of context valence and feedback information on speed and accuracy during reinforcement learning: A meta-analytical approach using diffusion decision modeling. *Cognitive, Affective, & Behavioral Neuroscience*, *19*(3), 490–502. https://doi.org/10.3758/s13415-019-00723-1

Fudenberg, D., Strack, P., & Strzalecki, T. (2018). Speed, Accuracy, and the Optimal Timing of Choices. *American Economic Review*, *108*(12), 3651–3684. https://doi.org/10.1257/aer.20150742

Gwinn, R., & Krajbich, I. (2020). Attitudes and Attention. *Journal of Experimental Social Psychology*, *86*, 103892. https://doi.org/10.1016/j.jesp.2019.103892

Harmon-Jones, E., & Mills, J. (2019). An introduction to cognitive dissonance theory and an overview of current perspectives on the theory. In *Cognitive dissonance: Reexamining a pivotal theory in*

psychology, 2nd ed (pp. 3–24). American Psychological Association. https://doi.org/10.1037/0000135-001

Holyoak, K. J., & Simon, D. (1999). Bidirectional reasoning in decision making by constraint satisfaction. *Journal of Experimental Psychology: General*, *128*(1), 3–31. https://doi.org/10.1037/0096-3445.128.1.3

Hunt, L. T., Kolling, N., Soltani, A., Woolrich, M. W., Rushworth, M. F. S., & Behrens, T. E. J. (2012). Mechanisms underlying cortical activity during value-guided choice. *Nature Neuroscience*, *15*(3), 470–476. https://doi.org/10.1038/nn.3017

Izuma, K., Akula, S., Murayama, K., Wu, D.-A., Iacoboni, M., & Adolphs, R. (2015). A causal role for posterior medial frontal cortex in choice-induced preference change. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *35*(8), 3598–3606. https://doi.org/10.1523/JNEUROSCI.4591-14.2015

Izuma, K., Matsumoto, M., Murayama, K., Samejima, K., Sadato, N., & Matsumoto, K. (2010). Neural correlates of cognitive dissonance and choice-induced preference change. *Proceedings of the National Academy of Sciences*, *107*(51), 22014–22019. https://doi.org/10.1073/pnas.1011879108

Izuma, K., & Murayama, K. (2013). Choice-Induced Preference Change in the Free-Choice Paradigm: A Critical Methodological Review. *Frontiers in Psychology*, *4*. https://doi.org/10.3389/fpsyg.2013.00041

Johnson, E. J., & Payne, J. W. (1985). Effort and Accuracy in Choice. *Management Science*, *31*(4), 395–414. https://doi.org/10.1287/mnsc.31.4.395

Kobayashi, K., Lee, S., Filipowicz, A. L. S., McGaughey, K. D., Kable, J. W., & Nassar, M. R. (2021). Dynamic Representation of the Subjective Value of Information. *Journal of Neuroscience*, *41*(39), 8220–8232. https://doi.org/10.1523/JNEUROSCI.0423-21.2021

Krajbich, I., Armel, C., & Rangel, A. (2010). Visual Fixations and the Computation and Comparison of Value in Simple Choice. *Nature Neuroscience*, *13*(10), 1292–1298. https://doi.org/10.1038/nn.2635

Lebreton, M., Jorge, S., Michel, V., Thirion, B., & Pessiglione, M. (2009). An Automatic Valuation System in the Human Brain: Evidence from Functional Neuroimaging. *Neuron*, *64*(3), 431–439. https://doi.org/10.1016/j.neuron.2009.09.040

Lee, D. G., & Coricelli, G. (2020). An Empirical Test of the Role of Value Certainty in Decision Making. *Frontiers in Psychology*, *11*, 574473. https://doi.org/10.3389/fpsyg.2020.574473

Lee, D. G., D'Alessandro, M., Iodice, P., Calluso, C., Rustichini, A., & Pezzulo, G. (2022). *Risky Decisions Accumulate Information for Individual Attributes as a Function of Salience and Risk Preference*. PsyArXiv. https://doi.org/10.31234/osf.io/h3r6x

Lee, D. G., & Daunizeau, J. (2020). 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. https://doi.org/10.1371/journal.pone.0231081

Lee, D. G., & Daunizeau, J. (2021). Trading mental effort for confidence in the metacognitive control of value-based decision-making. *ELife*, *10*, e63282. https://doi.org/10.7554/eLife.63282

Lee, D. G., Daunizeau, J., & Pezzulo, G. (2022). *Evidence or Confidence: What Really Accumulates During a Decision?* (p. 2021.04.02.438275). bioRxiv. https://doi.org/10.1101/2021.04.02.438275

Lee, D. G., & Holyoak, K. (2021a). Transient Value Refinements during Deliberation Facilitate Choice. *PsyArXiv*. https://doi.org/10.31234/osf.io/7pmgh

Lee, D. G., & Holyoak, K. J. (2021b). Coherence shifts in attribute evaluations. *Decision*, *8*(4), 257–276. https://doi.org/10.1037/dec0000151

Lee, D. G., & Pezzulo, G. (2022). *Changes in Preferences Reported After Choices Are Informative, Not Merely Statistical Artifacts*. PsyArXiv. https://doi.org/10.31234/osf.io/btj95

Lee, D. G., & Usher, M. (2021). Value certainty in drift-diffusion models of preferential choice. *Psychological Review*, No Pagination Specified-No Pagination Specified. https://doi.org/10.1037/rev0000329

Lepora, N. F., & Pezzulo, G. (2015). Embodied Choice: How Action Influences Perceptual Decision Making. *PLoS Comput Biol*, *11*(4), e1004110. https://doi.org/10.1371/journal.pcbi.1004110

Levav, J., Heitmann, M., Herrmann, A., & Iyengar, S. S. (2010). Order in Product Customization Decisions: Evidence from Field Experiments. *Journal of Political Economy*, *118*(2), 274–299. https://doi.org/10.1086/652463

Lichtenstein, S., & Slovic, P. (2006). *The Construction of Preference* (y First edition edition). Cambridge University Press.

Lim, S.-L., Penrod, M. T., Ha, O.-R., Bruce, J. M., & Bruce, A. S. (2018). Calorie Labeling Promotes Dietary Self-Control by Shifting the Temporal Dynamics of Health- and Taste-Attribute Integration in Overweight Individuals. *Psychological Science*, *29*(3), 447–462. https://doi.org/10.1177/0956797617737871

Lopez-Persem, A., Domenech, P., & Pessiglione, M. (2016). How prior preferences determine decision-making frames and biases in the human brain. *ELife*, *5*, e20317. https://doi.org/10.7554/eLife.20317

Maier, S. U., Raja Beharelle, A., Polanía, R., Ruff, C. C., & Hare, T. A. (2020). Dissociable Mechanisms Govern When and How Strongly Reward Attributes Affect Decisions. *Nature Human Behaviour*, *4*(9), 949–963. https://doi.org/10.1038/s41562-020-0893-y

Milosavljevic, M., Malmaud, J., Huth, A., Koch, C., & Rangel, A. (2010). The drift diffusion model can account for value-based choice response times under high and low time pressure. *Judgment and Decision Making*, *5*, 437–449.

Mulder, M. J., & Maanen, L. van. (2013). Are Accuracy and Reaction Time Affected via Different Processes? *PLOS ONE*, *8*(11), e80222. https://doi.org/10.1371/journal.pone.0080222

Mulder, M. J., Wagenmakers, E.-J., Ratcliff, R., Boekel, W., & Forstmann, B. U. (2012). Bias in the Brain: A Diffusion Model Analysis of Prior Probability and Potential Payoff. *Journal of Neuroscience*, *32*(7), 2335–2343. https://doi.org/10.1523/JNEUROSCI.4156-11.2012

Nunez, M. D., Vandekerckhove, J., & Srinivasan, R. (2017). How attention influences perceptual decision making: Single-trial EEG correlates of drift-diffusion model parameters. *Journal of Mathematical Psychology*, *76*, 117–130. https://doi.org/10.1016/j.jmp.2016.03.003

Padoa-Schioppa, C. (2011). Neurobiology of Economic Choice: A Good-Based Model. *Annual Review of Neuroscience*, *34*(1), 333–359. https://doi.org/10.1146/annurev-neuro-061010-113648

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1988). Adaptive strategy selection in decision making. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *14*, 534–552.

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker*. Cambridge University Press. https://doi.org/10.1017/CBO9781139173933

Payne, J. W., Bettman, J. R., & Schkade, D. A. (1999). Measuring Constructed Preferences: Towards a Building Code. *Journal of Risk and Uncertainty*, *19*(1), 243–270. https://doi.org/10.1023/A:1007843931054

Pezzulo, G., Rigoli, F., & Chersi, F. (2013). The Mixed Instrumental Controller: Using Value of Information to combine habitual choice and mental simulation. *Frontiers in Cognition*, *4*, 92. https://doi.org/10.3389/fpsyg.2013.00092

Polanía, R., Krajbich, I., Grueschow, M., & Ruff, C. C. (2014). Neural Oscillations and Synchronization Differentially Support Evidence Accumulation in Perceptual and Value-Based Decision Making. *Neuron*, *82*(3), 709–720. https://doi.org/10.1016/j.neuron.2014.03.014

Ratcliff, R., & McKoon, G. (2008). The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks. *Neural Computation*, *20*(4), 873–922. https://doi.org/10.1162/neco.2008.12-06-420

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. https://doi.org/10.1016/j.tics.2016.01.007

Ronayne, D., & Brown, G. D. A. (2017). Multi-Attribute Decision by Sampling: An Account of the Attraction, Compromise and Similarity Effects. *Journal of Mathematical Psychology*, *81*, 11–27. https://doi.org/10.1016/j.jmp.2017.08.005

Russo, J. E., Carlson, K. A., Meloy, M. G., & Yong, K. (2008). The goal of consistency as a cause of information distortion. *Journal of Experimental Psychology: General*, *137*(3), 456–470. https://doi.org/10.1037/a0012786

Russo, J. E., Medvec, V. H., & Meloy, M. G. (1996). The Distortion of Information during Decisions. *Organizational Behavior and Human Decision Processes*, *66*(1), 102–110. https://doi.org/10.1006/obhd.1996.0041

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

Shah, A. K., & Oppenheimer, D. M. (2008). Heuristics made easy: An effort-reduction framework. *Psychological Bulletin*, *134*(2), 207–222. https://doi.org/10.1037/0033-2909.134.2.207

Sharot, T., De Martino, B., & Dolan, R. J. (2009). How choice reveals and shapes expected hedonic outcome. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *29*(12), 3760–3765. https://doi.org/10.1523/JNEUROSCI.4972-08.2009

Sharot, T., Fleming, S. M., Yu, X., Koster, R., & Dolan, R. J. (2012). Is Choice-Induced Preference Change Long Lasting? *Psychological Science*, *23*(10), 1123–1129. https://doi.org/10.1177/0956797612438733

Sharot, T., Velasquez, C. M., & Dolan, R. J. (2010). Do Decisions Shape Preference? Evidence From Blind Choice. *Psychological Science*, *21*(9), 1231–1235.

Shevlin, B. R. K., & Krajbich, I. (2021). Attention as a source of variability in decision-making: Accounting for overall-value effects with diffusion models. *Journal of Mathematical Psychology*, *105*, 102594. https://doi.org/10.1016/j.jmp.2021.102594

Shevlin, B. R. K., Smith, S. M., Hausfeld, J., & Krajbich, I. (2022). High-value decisions are fast and accurate, inconsistent with diminishing value sensitivity. *Proceedings of the National Academy of Sciences*, *119*(6). https://doi.org/10.1073/pnas.2101508119

Simon, D., Krawczyk, D. C., Bleicher, A., & Holyoak, K. J. (2008). The transience of constructed preferences. *Journal of Behavioral Decision Making*, *21*(1), 1–14. https://doi.org/10.1002/bdm.575

Simon, D., Krawczyk, D. C., & Holyoak, K. J. (2004). Construction of preferences by constraint satisfaction. *Psychological Science*, *15*(5), 331–336. https://doi.org/10.1111/j.0956-7976.2004.00678.x

Simon, D., Pham, L. B., Le, Q. A., & Holyoak, K. J. (2001). The emergence of coherence over the course of decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *27*(5), 1250–1260. https://doi.org/10.1037/0278-7393.27.5.1250

Simon, H. A. (1957). *Models of man—Social and rational.* John Wiley and Sons.

Sims, C. R. (2016). Rate-distortion theory and human perception. *Cognition*, *152*, 181–198. https://doi.org/10.1016/j.cognition.2016.03.020

Smith, S. M., & Krajbich, I. (2019). Gaze Amplifies Value in Decision Making. *Psychological Science*, *30*(1), 116–128. https://doi.org/10.1177/0956797618810521

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

Sullivan, N., Hutcherson, C., Harris, A., & Rangel, A. (2015). Dietary Self-Control Is Related to the Speed With Which Attributes of Healthfulness and Tastiness Are Processed. *Psychological Science*, *26*(2), 122–134. https://doi.org/10.1177/0956797614559543

Tajima, S., Drugowitsch, J., & Pouget, A. (2016). Optimal Policy for Value-Based Decision-Making. *Nature Communications*, *7*(1), 12400. https://doi.org/10.1038/ncomms12400

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. https://doi.org/10.3758/s13423-015-0858-8

Urai, A. E., de Gee, J. W., Tsetsos, K., & Donner, T. H. (2019). Choice history biases subsequent evidence accumulation. *ELife*, *8*, e46331. https://doi.org/10.7554/eLife.46331

Usher, M., & Mcclelland, J. L. (2001). On the time course of perceptual choice: The leaky competing accumulator model. *Psychological Review*, 550–592.

Voigt, K., Murawski, C., & Bode, S. (2017). Endogenous formation of preferences: Choices systematically change willingness-to-pay for goods. *Journal of Experimental Psychology Learning, Memory, and Cognition*, *43*(12), 1872–1882. https://doi.org/10.1037/xlm0000415

Voigt, K., Murawski, C., Speer, S., & Bode, S. (2019). Hard Decisions Shape the Neural Coding of Preferences. *Journal of Neuroscience*, *39*(4), 718–726. https://doi.org/10.1523/JNEUROSCI.1681-18.2018

Warren, C., McGraw, A. P., & Van Boven, L. (2011). Values and preferences: Defining preference construction. *Wiley Interdisciplinary Reviews. Cognitive Science*, *2*(2), 193–205. https://doi.org/10.1002/wcs.98

Yang, X., & Krajbich, I. (2022). A dynamic computational model of gaze and choice in multi-attribute decisions. *Psychological Review*, No Pagination Specified-No Pagination Specified. https://doi.org/10.1037/rev0000350

Zénon, A., Solopchuk, O., & Pezzulo, G. (2019). An information-theoretic perspective on the costs of cognition. *Neuropsychologia*, *123*, 5–18. https://doi.org/10.1016/j.neuropsychologia.2018.09.013