

# Psychological Review

## **A Spiking Neural Model of Decision Making and the Speed-Accuracy Trade-Off**

Peter Duggins and Chris Eliasmith

Online First Publication, December 12, 2024. <https://dx.doi.org/10.1037/rev0000520>

### CITATION

Duggins, P., & Eliasmith, C. (2024). A spiking neural model of decision making and the speed-accuracy trade-off. *Psychological Review*. Advance online publication. <https://dx.doi.org/10.1037/rev0000520>

# A Spiking Neural Model of Decision Making and the Speed–Accuracy Trade-Off

Peter Duggins and Chris Eliasmith  
Centre for Theoretical Neuroscience, University of Waterloo

The speed–accuracy trade-off (SAT) is the tendency for fast decisions to come at the expense of accurate performance. Evidence accumulation models such as the drift diffusion model can reproduce a variety of behavioral data related to the SAT, and their parameters have been linked to neural activities in the brain. However, our understanding of how biological neural networks realize the associated cognitive operations remains incomplete, limiting our ability to unify neurological and computational accounts of the SAT. We address this gap by developing and analyzing a biologically plausible spiking neural network that extends the drift diffusion approach. We apply our model to both perceptual and nonperceptual tasks, investigate several contextual manipulations, and validate model performance using neural and behavioral data. Behaviorally, we find that our model (a) reproduces individual response time distributions; (b) generalizes across experimental contexts, including the number of choice alternatives, speed- or accuracy-emphasis, and task difficulty; and (c) predicts accuracy data, despite being fit only to response time data. Neurally, we show that our model (a) recreates observed patterns of spiking neural activity and (b) captures age-related deficits that are consistent with the behavioral data. More broadly, our model exhibits the SAT across a variety of tasks and contexts and explains how individual differences in speed and accuracy arise from synaptic weights within a spiking neural network. Our work showcases a method for translating mathematical models into functional neural networks and demonstrates that simulating such networks permits analyses and predictions that are outside the scope of purely mathematical models.

*Keywords:* drift diffusion, random dot motion task, Neural Engineering Framework, empirical validation

The speed–accuracy trade-off (SAT) is the well-documented tendency for individuals who make quick decisions to make more errors, compared to individuals who gather more information and deliberate longer. The SAT occurs in both natural and artificial contexts and has been widely studied by researchers interested in the cognitive mechanisms of decision making (DM). Neural and behavioral results indicate that a variety of individual differences and contextual factors influence how humans (Drugowitsch et al., 2015; Kanai & Rees, 2011) and animals (Chittka et al., 2009) accommodate the SAT (Sih & Del Giudice, 2012). Various models have been developed to explain the cognitive mechanisms that give rise these differences (Brown & Heathcote, 2008; Rao, 2010;

Ratcliff & McKoon, 2008). One model in particular, the drift diffusion (DD) model, explains how DM arises from the noisy accumulation of evidence toward one or more decision bounds (Ratcliff & McKoon, 2008). Over the past few decades, DD models have been quite successful in describing patterns of speed (or response time, RT) and accuracy across individuals, contexts, and tasks (Huang & Rao, 2013; Ratcliff et al., 2016). Furthermore, neural activity in specific brain regions has been shown to correlate with the decision variables (DVs) and parameters proposed by DD (Hanks et al., 2014; Turner et al., 2015), suggesting that DD may be consistent with DM processes in the brain. However, our understanding of the relationship between the cognitive mechanisms

Andrew Heathcote served as action editor.

Peter Duggins  <https://orcid.org/0000-0002-8948-0586>

Simulation code for the model and for all the simulated experiments in the remainder of the article are available on GitHub. This repository contains abbreviated and anonymized versions of the empirical data sets provided by Hanks et al. (2014), Churchland et al. (2008), Forstmann et al. (2011), and Fiedler et al. (2021). This study was not preregistered. Previous versions of the model were designed and analyzed in part by Dominik Krzemiński and Szymon Wichary as a conference article (Duggins et al., 2020), and subsequently as a thesis chapter (Duggins, 2023), but were substantially altered for the current work.

This work was supported by the Canadian Foundation for Innovation (52479-10006), the Ontario Innovation Trust (35768), the Natural Sciences and Engineering Research Council of Canada (261453), and the Air Force Office of Scientific Research (FA9550-17-1-0644) awarded to Chris

Eliasmith. The funders had no role in study design, data analysis, decision to publish, or preparation of the article.

The authors thank Anne Churchland, Timothy Hanks, Birte Forstmann, and Klaus Fiedler for providing them with the empirical data they used for model validation. The authors also thank Michael Shadlen for assistance with the literature review.

Peter Duggins played a lead role in conceptualization, formal analysis, investigation, methodology, software, visualization, and writing—original draft. Chris Eliasmith played a lead role in funding acquisition and supervision and a supporting role in conceptualization, formal analysis, investigation, methodology, and writing—review and editing.

Correspondence concerning this article should be addressed to Peter Duggins, Centre for Theoretical Neuroscience, University of Waterloo, 200 University Avenue West, Waterloo, ON N2L 3G1, Canada. Email: [pspeter@gmail.com](mailto:pspeter@gmail.com)

of DM (as proposed by DD and other mathematical models) and the neural mechanisms of DM (as realized by networks of biological neurons in the brain) remains incomplete.

In this article, we aim to integrate neural and cognitive explanations of the SAT using a spiking neural network (SNN) model. Our goals are twofold. First, we show that the cognitive mechanisms proposed by DD can be directly implemented in biologically plausible neural networks, and that the mathematical parameters proposed by DD (especially those responsible for individual differences and the SAT) can be instantiated within synaptic connection weights. Successfully achieving this goal demonstrates that neural mechanisms within the brain are capable of supporting the cognitive mechanisms laid out by DD, bridging algorithmic and implementational descriptions. Second, we show that a neural model of DM can explain (and make predictions about) classes of phenomena that are outside the scope of purely mathematical models. Successfully achieving this goal demonstrates that biologically plausible neural network models have utility beyond the original DD model, improving our understanding of cognition in the brain.

To this end, we develop a SNN model of DM with two key properties: Its architecture is based on the functional neuroanatomy of the brain; and its dynamics realize the mechanisms and parameters of the DD model. While our model realizes the processes of evidence accumulation and thresholding that are central to DD, it departs from (and extends) DD in several respects. First, our model posits an additional cognitive operation that we call “valuation,” in which accumulated evidence may be considered in either absolute or relative terms. We argue that parameterized valuation captures an important feature of DM that is sometimes obscured in evidence accumulation models and show that including valuation in our model improves behavioral fits. Second, while many DD models posit trial-to-trial variance in parameters like drift rate, starting point, or threshold, we treat our model parameters as invariant properties for each individual. In our model, variance instead emerges as a result of spike noise and neural communication, allowing us to recreate observed distributions of speeds and accuracies without positing parametric variation at the cognitive level. Third, our model can make predictions that span neural and cognitive levels of analysis: For instance, we can predict the cognitive deficits that arise from simulated neural degeneration or neurotransmitter deficits, a capacity that is important for studying and eventually treating age-related and mental disorders. Last, our model offers two key advantages over other neural network realizations of DD: (a) It can be described in terms of high-level DVs and cognitive operations, facilitating explainability and direct comparison with mathematical models; and (b) it leverages a modular, functional organization that can scale to accommodate more complex operations and more difficult tasks, which we demonstrate by using the same architecture across multiple tasks and experimental manipulations. Overall, the ways in which our network implements DD should be of interest to both psychologists and neuroscientists who study DM and the SAT.

We begin by describing our network in detail, discussing how it relates to the DD model and showing its basic behavior. We then evaluate the model’s performance by comparing its outputs to neural and behavioral data from four separate empirical experiments. In Experiment 1, we apply the model to the canonical random dot motion (RDM) task, investigating two experimental manipulations:

the degree of motion coherence (task difficulty), and contextual task instructions (speed vs. accuracy). We show that our model exhibits the well-established SAT curve relating task accuracy to RT across the spectrum of motion coherences, subject to speed versus accuracy instructions (Hanks et al., 2014). We also identify neurons in our model that exhibit ramping activities similar to those observed in cortical accumulators. We demonstrate that their buildup rate is higher in the speed-emphasis condition, but that their activities converge immediately before a decision is made; both these results are consistent with neural data from the same experiment (Hanks et al., 2014).

In Experiment 2, we illustrate the scalability of our model by applying it to a RDM task that features either two or four directions of motion (Churchland et al., 2008). We show that our model reproduces the SAT curves for both two- and four-choice variant without any structural changes. We also show that the model has longer RTs on error trials with high motion coherence, a result that is consistent with the behavioral data (Churchland et al., 2008).

In Experiment 3, we showcase the utility of neural modeling by simulating the effects of aging in our SNN. Compared to young individuals, elderly individuals have longer RTs and reduced accuracy, both in speed- and accuracy-emphasis conditions on RDM tasks (Forstmann et al., 2011). Similar speed–accuracy deficits have been observed in individuals with attention-deficit/hyperactivity disorder (Mulder et al., 2010), schizophrenia (Fish et al., 2018; Schweitzer & Lee, 1992), and prefrontal tumors (Campanella et al., 2016). These functional deficits have been attributed to numerous biological sources, including impaired working memory resulting from degraded connectivity in prefrontal cortex (Nissim et al., 2017; Park & Holzman, 1992) and impaired threshold setting resulting from degraded connectivity between subcortical and cortical regions (Bogacz et al., 2010; Forstmann et al., 2008). In this experiment, we train our SNN to match the performance of young individuals, then simulate the process of aging by degrading particular synaptic connections within the network. We show that some, but not all, of these hypothesized biological factors induce cognitive deficits reminiscent of elderly individuals, lending support to the corresponding cognitive theories and demonstrating the utility of our neural implementation of DD.

In Experiment 4, we move beyond the RDM task to study the SAT in a nonperceptual (i.e., cognitive), temporally extended setting. In this task, participants sample noisy evidence about stock prices over time, and must decide which stock to invest in after a self-truncated sampling period. Participants complete as many trials of the task as possible within the allotted timeframe and are rewarded for the total number of correct choices. Participants who sample fewer cues before deciding will complete more trials overall, but will be less accurate than participants who sample more cues: This task thus establishes an explicit SAT with respect to the rewards obtained. Experimental evidence shows that individuals adopt different strategies when performing this task, but that most humans are biased toward strategies that overemphasize accuracy (Fiedler et al., 2021). We apply our model to this task with the goal of recreating individual differences between these strategies and their corresponding behaviors. We fit models to each of the 55 participants in the empirical data set, attempting to reproduce a particular “training” subset of the behavioral data while reserving the remainder as a “validation” subset. We show that our fitted models reproduce (a) the behaviors characteristic of strategies across

SAT spectrum (a fast-inaccurate, a middle-of-the-road, and a slow-accurate strategy) and (b) the SAT curve evident across the entire population (mean accuracy vs. mean number of sampled cues for all individuals). Furthermore, our model (c) recreates the observed accuracy bias and (d) generalizes to data outside the training set, making accurate predictions about the number of cues sampled and resulting accuracy for novel task difficulties.

In order to highlight the advantages of our SNN model for understanding the SAT, we also simulate two mathematical models, the DD model and the extrema detection model, and analyze their behavior in these four experiments. These two simple models represent traditional approaches to computational modeling of the SAT, and their behavior serves as a baseline against which we compare the behavior our relatively complex model. We find that, while DD and extrema detection can fit the observed RT distributions, these models struggle to predict data outside the training set, often make incorrect accuracy predictions, and perform poorly in the nonperceptual task.

We conclude with a discussion of how our model relates to other computational models, to empirical data, and to more complex tasks. We first review the biological and cognitive plausibility of our model: We lay out our theoretical assumptions and enumerate the unique predictions made by our model, then analyze how well our model captures the brain's cognitive abilities in this limited domain. In doing so, we discuss how our model derives from DD and how it extends DD through neural implementation; we also compare our network to existing neural network models of DM. We argue that our model helps integrate multiple levels of analysis by describing how neural mechanisms in the brain implement key cognitive mechanisms of DM, and that such an understanding helps improve our understanding of individual differences with respect to the SAT. We finish by proposing extensions for the model and identifying directions for future work.

## Background

### Decision-Making Task

To study the neural and cognitive mechanisms of DM, researchers use cognitive tasks that require individuals to sample information from the world, evaluate that information with respect to some goal, and choose between alternative actions. Different tasks involve different processes for sampling, evaluation, and decision, and therefore inform our understanding of DM in different ways. In this article, we study two classes of DM task: perceptual and nonperceptual.

In perceptual tasks, participants are presented with sensory input and must classify the input as quickly as possible, given the presence of distractors. Speed is measured as the time between stimulus presentation and choice selection (RT), while accuracy is measured as the mean number of correct responses across many trials. Task difficulty is determined by the number of distractor stimuli; a greater fraction of distractors from the incorrect category implies a more difficult trial. Perhaps the most widely studied perceptual discrimination task is the RDM task (Brown & Heathcote, 2008; Churchland et al., 2008; Hanks et al., 2014; Lo et al., 2015; Palmer et al., 2005; Ratcliff et al., 2016). In this task, participants must identify the predominant direction of motion among a cloud of moving dots. Some of the dots move in random directions (distractors), and the

remainder move in a coherent direction (signal). The percentage of dots moving in the coherent direction (or motion coherence,  $C$ ) determines the task difficulty: when  $C \approx 100\%$ , participants complete the task quickly and accurately, but when  $C \approx 0\%$ , participants take longer to respond and are less accurate.

In nonperceptual or cognitive tasks, participants are presented with symbolic information and must choose between alternative actions, given some abstract goal. Speed is measured as the amount of information collected, while accuracy is measured as the mean number of correct responses across many trials. Task difficulty is determined by the relative uncertainty of the presented information with respect to the goal. Uncertainty and speed in nonperceptual tasks are intended to be independent of the brain's perceptual system: uncertainty is defined in abstract or symbolic terms, and speed should depend on deliberative cognitive processes rather than perceptual ones. For example, in the task developed by Fiedler et al. (2021), participants are presented with sequential evidence about the changing financial value of two hypothetical stocks ("A" and "B") and are asked to select the stock with greater value. Participants view a screen that displays information about Stock A on the left and Stock B on the right. Every 500 ms, the screen displays new evidence about one stock, indicating whether it increased or decreased in value. Samples alternate between Stocks A and B. Participants may press a button at any time to stop sampling, then another button to select Stock A or Stock B. Once participants complete a single trial, the next trial begins immediately. In this task, speed is the number of samples taken before a decision, and accuracy is the mean number of correct selections across multiple trials. Behind the scenes, each stock has a hidden, unique probability  $P$  of increasing in value. The difference in probabilities between A and B is fixed for each trial and determines the task difficulty; larger  $\Delta P$  implies an easier task (less uncertainty), which produces faster and more accurate decisions.

There are important similarities and differences between perceptual and nonperceptual tasks with respect to DM and the SAT. In both tasks, the brain must continuously reevaluate potential actions as more information is gathered (Domenech et al., 2018; Kayser et al., 2010; Turner et al., 2015), and studies have shown that multiple factors, such the quality of incoming evidence and the need for urgency, influence people's behavior in these tasks (Yau et al., 2020). Similarly, both tasks involve the SAT: Some individuals perform the task quickly but inaccurately, while others do so slowly but accurately; and participants can adjust their strategies to accommodate task instructions to emphasize speed or accuracy (Fiedler et al., 2021; Hanks et al., 2014). However, some researchers have questioned whether perceptual tasks are well-suited to investigate high level, deliberative DM (Fiedler et al., 2021). Perceptual tasks rarely require conscious deliberation, extended risk assessment, or the voluntary seeking of additional information, three features that define high-level DM in humans. Furthermore, the SAT must often be externally incorporated into perceptual tasks: For instance, experimenters must instruct participants to favor speed or accuracy (Hanks et al., 2014; Palmer et al., 2005), create payoff structures that reward fast versus accurate decisions, or impose external deadlines for responses (Heitz, 2014). The extent to which these explicit instructions influence perceptual processes is a matter of debate. Finally, differences in speed and accuracy in perceptual tasks are often determined by perceptual abilities and by the allocation of attention, rather than by individual differences in DM

(Fiedler et al., 2021). In contrast, nonperceptual tasks, such as those presented in Gluth et al. (2014) and Liu et al. (2022), naturally address many of these concerns. Even without explicit instructions or payoff structures, participants in nonperceptual tasks conceptually understand they are faced with a SAT, and attempt to behave accordingly, with varying levels of success (Fiedler et al., 2021). Sampling information in nonperceptual tasks is also largely divorced from sensory processing, meaning that individual differences in perceptual abilities do not confound the analysis of the SAT. That said, perceptual tasks have a rich history in the DM literature (Heekeren et al., 2008), supported by neural and behavioral data (Churchland et al., 2008; Hanks et al., 2014; Palmer et al., 2005) and mathematical theory (Brown & Heathcote, 2008; Ratcliff & McKoon, 2008).

In this article, we investigate DM and the SAT using both perceptual and nonperceptual tasks. By simulating perceptual tasks, we can (a) compare our model with existing theories of DM and (b) validate our results against well-known behavioral and neural data. By simulating nonperceptual tasks, we investigate whether our model (c) applies to temporally extended, deliberative DM and (d) captures the variety of high-level cognitive strategies exhibited by individual humans. Finally, by applying our model to both tasks, we show the generality, scalability, and predictive power of our approach.

## Neuroanatomy

Extensive evidence supports the idea that the brain maintains value estimates for candidate actions via neural accumulators (Luo, 2018; Pärnamets et al., 2020; Rusch et al., 2020; Suzuki & O’Doherty, 2020), and that accumulators for different modalities are realized in domain-specific working memory buffers throughout cortex (Christophel et al., 2017; Sreenivasan et al., 2014). In perceptual tasks, evidence accumulation occurs primarily in visual areas like the lateral intraparietal area (LIP), frontal eye fields, medial temporal area (MT), and the fusiform gyrus (Churchland et al., 2008; Hanks et al., 2014; Heitz, 2014; Kayser et al., 2010; Yau et al., 2020). In nonperceptual tasks, evidence is abstract and is often processed in cortical areas such as ventromedial prefrontal cortex, dorsolateral prefrontal cortex (dlPFC), orbitofrontal cortex, and anterior cingulate cortex (Domenech et al., 2018; Gluth et al., 2014; Gupta et al., 2021; Liu et al., 2022; Mulder et al., 2014). Interestingly, value accumulation is also apparent in parts of motor cortex, such as the supplementary motor area and the premotor cortex, for both perceptual and nonperceptual tasks (Gupta et al., 2021; Kayser et al., 2010), indicating that internal representations of the action plans themselves might accumulate evidence.

Various brain areas have been implicated in flexibly controlling the SAT. Areas like dlPFC may influence the weighting of incoming information, based on the reliability of evidence or the urgency of acquiring more information (Domenech et al., 2018). Various areas of the basal ganglia (BG), including the caudate nucleus, substantia nigra pars compacta, and the subthalamic nucleus, are also thought to effectively shift decision thresholds through projections to cortex (Gupta et al., 2021; Yau et al., 2020). The involvement of these areas in thresholding decisions is further supported by data from humans with Parkinson’s, schizophrenia, and autism: disrupted substantia nigra pars compacta and subthalamic nucleus activity in these disorders correlates with higher decision thresholds and slower evidence accumulation (Gupta et al., 2021). Other researchers have argued that changes in

decision threshold and/or the rate of evidence accumulation are realized directly in cortical areas, where sensory processing evaluates, and working memory integrates, evidence (Domenech et al., 2018; Simen, 2012; Standage et al., 2011). Finally, action selection is often thought to involve winner-take-all competition, either through inhibitory connections within the cortex (particularly motor regions) or via mutual connections between cortex and BG (Guthrie et al., 2013; Redgrave et al., 1999).

## Computational Models

Many mathematical and computational models have explored the neural and cognitive basis of DM. Evidence accumulation models describe DM by positing one or more DVs, which encode the dynamic likelihood of choosing one or more actions. DVs grow and shrink as the system *samples* information from the environment. This growth is termed *accumulation*. When a DV reaches a *threshold*, a decision is triggered: The time between this event and the onset of the stimulus (i.e., the start of the simulation) is taken to be the RT. While most evidence accumulation models share these core features, they make different assumptions about the processes of sampling, accumulation, and thresholding.

## Drift Diffusion

The DD model is arguably the most widely adopted model of DM: It has been very successful in modeling animal behavior across many domains (Ratcliff & McKoon, 2008), and its variables and parameters have been correlated with neural activity in brain areas associated with DM (Ratcliff et al., 2016; Steinemann et al., 2022; Stine et al., 2023; Turner et al., 2015). In DD, the DV begins at some starting value  $S$ , then “drifts” toward the decision threshold as evidence is sampled and accumulated. This drift is a continuous random walk process, whose slope is determined by the mean drift rate  $R$  and whose variance is determined by the drift variance parameter  $V$ . When the DV crosses the decision threshold  $T$ , the model chooses the action associated with that DV. The model’s total RT is given by the decision time ( $t_d$ , the time between when drift begins and ends), plus a nondecision time (NDT or  $t_{nd}$ , the time between when the stimuli is presented and when drift begins). With respect to the SAT, DD models have been widely successful in capturing contextual differences in experiments and in describing individual differences between participants. For instance, numerous studies have shown that differences in RT and accuracy, in experiments that emphasize speed versus accuracy, can be described by differences in  $R$  and  $T$ , and that these differences in model parameters correlate with differences in neural activities between individuals or between experimental contexts (Karalunas et al., 2012; Kofler et al., 2013; Palmer et al., 2005; Ratcliff & Childers, 2015; Trueblood et al., 2014; Turner et al., 2015; Vallesi et al., 2015; Zhang & Rowe, 2014).

Conceptually, the parameters and processes in DD represent various DM processes in the brain. The starting point  $S$  captures bias toward a decision that is formed before a trial begins; these biases reflect prior beliefs held by an individual and may arise from numerous sources, including decision statistics from trials in a previous experimental block (Hanks et al., 2011), or a failure to properly reset the DVs represented in working memory (Simen, 2012). Incorporating  $S$  helps DD capture behavioral phenomenon related to perseveration (Urai & Donner, 2022) and RTs on error

trials (Ratcliff et al., 2016). The ramp rate  $R$  (and its variance  $V$ ) capture the processes of sampling and accumulation.  $R$  is set according to the motion coherence: The direction of motion determines the sign of  $R$  and the value of  $C$  determines the magnitude of  $R$ . Larger values of  $V$  drive greater variance in the drift of the DV; this may arise from noise in either the perceptual sampling process (e.g., more jitter in the random dots) or the accumulation process (e.g., imprecise representations of the DV in working memory). The decision threshold  $T$  dictates the amount of evidence that must be accumulated before the model is ready to make a decision. Finally,  $t_{nd}$  captures the time delay between the presentation of sensory stimuli and the initiation of the accumulation process; including  $t_{nd}$  allows DD to account for unknown perceptual processes and improves the ability of DD models to fit RT data.

The DD model can be modified in many ways. Researchers have extended the basic model described above with the goal of increasing the model's cognitive realism and fitting more behavioral data. Some models assume that decision thresholds shrink over time as the pressure to make a decision steadily increases (Drugowitsch et al., 2012; Kira et al., 2024): The existence of a ramping urgency signal is consistent with behavioral data (Ditterich, 2006) and neural data (Cisek et al., 2009), though the strength of this effect has been called into question (Evans et al., 2017). In other models, the drift rate increases as time pressures increase, leading to faster decisions (Murphy et al., 2016). Finally, confidence may also stimulate changes in model parameters, speeding or slowing decisions based on past performance, current evidence estimates, or elapsed time (Kiani et al., 2014; Kira et al., 2015; Lee & Usher, 2021). In contrast, other models have sought to simplify the mechanisms of DD while preserving the ability to fit neural or behavioral data. For instance, in the linear ballistic accumulator (LBA) model (Brown & Heathcote, 2008), the DV grows linearly, rather than according to a continuous random walk. This accumulation can be contrasted with other ballistic accumulator models (Brown & Heathcote, 2005), which place fewer constraints on the mathematical form of accumulation; or with competitive accumulator models (Usher & McClelland, 2001), which assume that DVs compete with (i.e., inhibit) one another as they grow. Despite their relative simplicity, LBA models have managed to reproduce many of the behavioral phenomenon explained by DD models (Brown & Heathcote, 2008; Donkin et al., 2011).

### **Extrema Detection**

Many other mathematical models of the SAT are based on cognitive heuristics for DM. These models involve various methods for sampling, evaluating, and comparing alternative actions that are typically biased or irrational in some manner, in that they discard or discount some information in service of making decisions quickly or efficiently. A comprehensive review of these models is beyond the scope of this work, but we refer interested readers to Pietsch and Vickers (1997), which tests a number of sampling and memory heuristics for DM in perceptual tasks; to Canellas and Feigh (2016), which compares heuristic models within the “fast-and-frugal” family on naturalistic decision tasks; and to Gigerenzer and Gaissmaier (2011), which reviews heuristic DM in applied settings such as business and health care. Here, we highlight one family of models, called “nonintegration” models, which do not involve accumulating evidence over time. Surprisingly, studies have shown

that simple nonintegration models can sometimes reproduce empirical speed and accuracy data in many perceptual tasks (although not always, Kira et al., 2015). For instance, the “extrema detection” model chooses an option if it observes any sample that exceeds a decision threshold  $T$ ; despite its simplicity, this model behaves similarly to DD models in the RDM task (Stine et al., 2020). Interestingly, Stine et al. (2020) found that several commonly accepted signs of evidence accumulation were also predicted by nonintegration strategies. These results showcase the importance of studying DM using models other than DD, which currently dominates the field, as well as the explanatory power of simple models. In this work, we compare our model, which resembles DD, to an extrema detection model in order to highlight their similarities and differences.

### **Neural Models**

Although the DD model has been instrumental in advancing our understanding of DM and the SAT and provides an intuitive high-level description of the cognitive mechanisms of DM, it does not (directly) provide a low-level description of the neural mechanisms that support these cognitive mechanisms. We believe, in conjunction with other researchers (Frank, 2006; Lo et al., 2015; Shen et al., 2023; Standage et al., 2011), that SNN models are critical for explaining these neural mechanisms, and have the potential to explain how the mathematical operations described by DD might be realized in the brain. By implementing DD in an SNN, researchers can show how the brain might realize the proposed cognitive mechanisms, thereby supporting the claim that DD explains DM in biological systems. Conversely, if an SNN implementation of DD cannot explain animal behavior, this raises questions about whether the biological brain actually performs the cognitive operations described by DD. Such a failure might motivate a search for more biologically compatible explanations, either by modifying the cognitive mechanisms of DD or by adjusting the neural mechanisms of the SNN.

More practically, neural models make predictions about phenomena that may be outside the scope of behavioral models. For example, SNNs generate patterns of neural activity that can be directly compared with single-unit activities in the brain. Previous SNNs have produced spiking activities that align closely with the activity in visual areas such as LIP and superior colliculus (Lo et al., 2015; Shen et al., 2023), supporting theories about the neural codes underlying DVs (Steinemann et al., 2022; Stine et al., 2023). Furthermore, SNNs can potentially predict the deficits induced by mental disorders (schizophrenia, Parkinson's, etc.), neurotransmitter deficits (dopamine, serotonin, etc.), or physical damage to the brain (Duggins & Elias-Smith, 2022; Frank, 2006). Cognitive models explain these deficits by fitting model parameters to healthy versus impaired groups: For example, DD models have been applied to explain DM deficits in older adults (Forstmann et al., 2011) and individuals with attention-deficit hyperactivity disorder (Mulder et al., 2010) or schizophrenia (Moustafa et al., 2015). However, the resulting explanations are often descriptive (top-down) in nature: They summarize the effects of biological deficits on the dynamics of DM, but do not predict how a novel biological perturbation would change these dynamics. In contrast, SNNs have the potential to explain the neural origins of cognitive deficits in a mechanistic (bottom-up) manner. For example, in this article, we model the cognitive effects

of aging by simulating the biological deficits induced by aging. We test three different biological hypotheses related to aging and the SAT (Bogacz et al., 2010) by degrading the synaptic connectivity of our SNN in specific ways. We then simulate the model performing a RDM task and generate bottom-up predictions about the resulting behavioral deficits. Our results support the conclusion that working memory deficits induced by cortical degeneration are the best explanation for age-related decline in accuracy and speed.

### Summary

Both the cognitive explanations provided by mathematical models (such as DD and extrema detection) and the neural explanations provided by SNNs (such as our model) are valuable for understanding DM and the SAT: They are concerned with different levels of analysis and appeal to different scientific audiences. We believe that a robust understanding of these phenomena requires a diversity of such explanations, and that models that integrate multiple perspectives are especially valuable. In the remainder of this article, we aim to develop just such a model. We use a well-established theoretical framework to develop a SNN that implements evidence accumulation in a biologically plausible manner. We apply our model to both perceptual and nonperceptual tasks and validate our predictions using a variety of neural and behavioral data. We also compare our simulated results to the behavior of DD and extrema detection models in these tasks. Later, we review the neural and cognitive realism of our model, arguing that it extends existing SNNs in several respects. We conclude by summarizing the ways in which our model expands our current understanding of DM and the SAT and by suggesting directions for future work.

## Materials and Method

### Neural Engineering Framework

Our spiking neural DM model is built using the Neural Engineering Framework (NEF) developed by Eliasmith and Anderson (2003). The NEF characterizes spiking activity within populations of neurons as encoding information in a latent *state space*. While spikes are the physical means of communication between neurons, cognition can be analyzed as transformations of these states, permitting a more compact description of what brains do (Boerlin et al., 2013; Gallego et al., 2017; Recanatesi et al., 2022). We assume that states in this state space can be represented by a vector-valued signal  $\mathbf{x}(t)$ , and that the cognitive operations performed in the brain may be described as dynamical transformations of  $\mathbf{x}(t)$ .

The NEF defines methods for encoding and decoding between neural activity and the state space. A neuron will fire most frequently when presented with its particular “preferred stimulus” and will respond less strongly to increasingly dissimilar stimuli. Each simulated neuron  $i$  is accordingly assigned a preferred direction vector or *encoder*. When driven with an external signal  $\mathbf{x}(t)$ , the firing rate of the neuron is given by

$$a_i(t) = G[\alpha_i \mathbf{e}_i \cdot \mathbf{x}(t) + \beta_i], \quad (1)$$

where  $a_i(t)$  is the spiking activity of neuron  $i$ ,  $G$  is the neuron model with electrical current inputs  $[\cdot]$ ,  $\alpha_i$  is the gain,  $\beta_i$  is the bias current, and  $\mathbf{e}_i \cdot \mathbf{x}(t)$  is the dot product between the state space inputs and the neuron’s encoder. A distributed encoding extends the notion of

representation: If  $\mathbf{x}(t)$  is fed into multiple neurons, each with a unique tuning curve defined by  $\mathbf{e}_i$ ,  $\alpha_i$  and  $\beta_i$ , then each neuron will respond with a unique spiking pattern  $a_i(t)$ , and the collection of all neural activities will robustly encode the signal.

In order to recover, or decode, the state space information encoded in neural spike trains, the NEF also defines *decoders*  $\mathbf{d}_i$ , which either perform this recovery or compute arbitrary functions,  $f(\mathbf{x})$ , of the represented vector. A functional decoding with  $\mathbf{d}_i^f$  allows networks of neurons to *transform* the signal into a new state, which is essential for performing cognitive operations. To compute these transformations, a linear decoding is applied to the neural activities:

$$\hat{f}(\mathbf{x}(t)) = \sum_{i=0}^n a_i(t) \mathbf{d}_i^f, \quad (2)$$

where  $a_i(t)$  is the activity of neuron  $i$ ,  $n$  is the number of neurons, and the hat notation indicates that the computed quantity is an estimate of the target function. Connection *weights* between each presynaptic neuron  $i$  and each postsynaptic neuron  $j$  are composed of encoders and decoders:

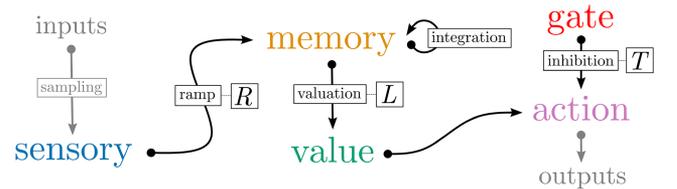
$$\mathbf{w} = \mathbf{e} \times \mathbf{d}^f. \quad (3)$$

In our model, encoders are chosen by sampling vectors from a  $D$ -dimensional hypersphere, ensuring that each neural population effectively represents a state space signal  $\mathbf{x} \in \mathbb{R}^D$ . Decoders are computed by specifying the target function  $f(\mathbf{x})$  for each connection in the model, then performing an offline least-squares optimization to solve for decoders that minimize the error between the neural estimate and the target function. For further details on encoder selection and decoder optimization, see Eliasmith and Anderson (2003) and Bekolay et al. (2014).

### A Spiking Neural Model of Decision Making

Our model, summarized in Figure 1, contains five spiking neural populations and one nonneural component. We begin by describing the nonneural perceptual system, which samples information from the environment. In general, we are more interested in investigating the cognitive processes of DM than in characterizing the perceptual

**Figure 1**  
Network Architecture for the Decision-Making Model



*Note.* The populations (sensory, memory, value, gate, and action) each contain 500 spiking LIF neurons. Grey arrows are external inputs and outputs, and black arrows are current-based exponential synapses connecting neural populations. Boxes indicate cognitive operations being computed via synaptic weights (black) or in a preprocessing step (gray). The model’s three free parameters ( $R$ ,  $L$ , and  $T$ ) are indicated next to the corresponding operation. LIF = leaky integrate-and-fire. See the online article for the color version of this figure.

processes of encoding external stimuli. However, we recognize that, especially in perceptual tasks, these processes may substantially affect the quality of the representations that are manipulated by subsequent cognitive processes. We hypothesize that perceptual processes introduce noise into DVs that, in conjunction with the noise inherent to computation in SNNs, contribute to trial-to-trial variance in individual behavior. This hypothesis is supported by evidence that both external noise and noisy neural communication increase the internal noise of DV representations, reducing task performance (Chen et al., 2014; Lu & Doshier, 1998). In this article, we use a nonneural system for *sensory sampling* that generates noisy estimates of percepts in the RDM task; these representations are then fed to the SNN as inputs. In the Discussion section, we note that previous NEF models have simulated this sensory sampling system using a more complete biologically plausible model of the visual system and discuss how such a system could be incorporated into our model in future work.

In the RDM task, our *sampling* module takes as input the motion coherence value of the current trial and returns an estimate of the perceived motion in each direction (i.e., fraction of dots moving in each direction) for the current sampling period. To do so, we first calculate the true percentage of motion in each direction from coherence: For example, in a two-choice task with a coherence value of  $C = 0$ , these percentages are  $P = [0.5, 0.5]$ ; and in a four-choice task with  $C = 100$ , these fractions are  $\mathbf{P} = [0, 0, 0, 1]$ . Next, we assume that our system cannot perfectly perceive these percentages and instead takes periodic, noisy samples based on  $\mathbf{P}$  to estimate the current fraction of dots moving in each direction. Specifically, our system draws perceptual samples from a normal distribution with mean  $\mathbf{P}$  and variance  $\sigma$  every  $t_s$  seconds. The input to our model is therefore given by the motion input vector  $\mathbf{i} = N(\mathbf{P}(C), \sigma)$ , and  $\mathbf{i}$  is resampled every  $t_s$  seconds. Note that  $\sigma$  and  $t_s$  are free parameters of the model:  $\sigma$  controls the magnitude of perceptual error (how far perceived motion in each direction differs from the ground truth), while  $t_s$  controls how long these errors are propagated before resampling occurs. In the nonperceptual stock market task, inputs are presented for a sustained duration without any distractors, so we simplify the *sampling* module to remove noise and resampling: inputs alternate between  $\mathbf{i} = [c, 0]$  and  $[0, c]$  every  $t_s = 0.1$  s. Negative cues are assigned a value of  $c = 0$ , and positive cues are valued at  $c = 1$ .

The input vector  $\mathbf{i}(t)$  is sent to our first neural population, labeled *sensory*, which encodes  $\mathbf{i}(t)$  in spiking neural activities according to Equation 1. Activity in this population thus represents the currently perceived motion strength in each direction. This information is then passed to the *memory* population, which is recurrently connected to itself: This combination of feedforward and recurrent connectivity implements a neural integrator, which continuously adds the currently perceived motion strength into a running estimate of the cumulative motion in each direction. As noisy percepts are accumulated in *memory*, the model’s cumulative motion estimates drift up or down: The drift rate is proportional to motion coherence,  $C$  determines  $\mathbf{P}$  and hence the mean of  $\mathbf{i}(t)$ , and the extent of random walk is proportional to  $\sigma$  and  $t_s$ . However, we also introduce a ramp rate parameter  $R$ , which is used to scale the magnitude of synaptic weights between *sensory* and *memory*. Higher  $R$  directly scales the rate of accumulation in the model and is the first free parameter that we assume differs between individual participants.

Importantly, *memory* tracks the cumulative motion in each direction independently, but every neuron in *memory* is sensitive to every direction of motion. In the Discussion section, we compare this method of neural coding with other neural models of DM and argue that permits a greater degree of flexibility and scalability. One consequence of this encoding scheme is that each DV can be separately decoded from the population activities in *memory*, but these activities can also be used to compute functions over multiple DVs. We make use of this property to flexibly evaluate each action based on the history of accumulated information. The NEF allows us to compute arbitrary functions of this information; in this article, we use the connection between *memory* and the downstream *value* population to realize a simple form of “valuation.” Specifically, this connection computes a function that considers accumulated evidence in either absolute or relative terms,

$$V_i(t) = \frac{1}{A-1} \sum_{j \neq i}^A M_j(t) - LM_j(t), \quad (4)$$

where  $V$  is the value assigned to each candidate action  $i$ ,  $A$  is the number of possible actions, and  $j$  indexes alternative actions.  $M$  is the cumulative motion estimate encoded in *memory*, and  $L$  is a free parameter of the model that changes whether value computation is relative or absolute. In a two-choice task, Equation 4 reduces to

$$V_0(t) = M_0(t) - LM_1(t), \quad (5)$$

$$V_1(t) = M_1(t) - LM_0(t). \quad (6)$$

When  $L = 0$ ,  $\mathbf{V}(t)$  simply reduces to the cumulative motion estimate  $\mathbf{M}(t)$ , and we can interpret the DV represented in *value* as the absolute sensory evidence for each action. On the other hand, when  $L = 1$ ,  $\mathbf{V}(t)$  reduces to the relative difference between the cumulative estimates for leftward and rightward motions; positive values indicate one estimate is greater than another. Finally, values of  $L$  between 0 and 1 lead to mixed evaluation, which accounts for both the absolute amount of evidence that has accumulated for a given choice and the relative difference between each choice and its alternatives. The parameter  $L$  governs how the model evaluates its estimates of cumulative motion and is the second free parameter of the model that we assume differs between individual participants. In the Discussion section, we argue that this parameterization helps clarify a distinction that is often ambiguous in mathematical models.

Finally, the model must use the DV represented in *value* to make a decision. Neurons in the *action* population are tuned such that they exhibit minimal activity if the combined inputs to *action* are less than zero. This population receives two inputs: The DV estimates from *value*, and an inhibitory hold from the *gate* population. Consequently, the excitatory signal from *value* will only produce a decision when it overcomes the inhibitory signal from *gate*. The amount of inhibition applied by *gate* is determined by the third (and final) free parameter of the model, the decision threshold  $T$ . When any element of  $\mathbf{V}(t)$  exceeds  $T$ , neurons in *action* become active, and we decode an action from the model. Note that this thresholding mechanism is agnostic about how the value function is computed: When  $L = 0$ , actions are thresholded based on the absolute cumulative motion, whereas when  $L = 1$ , actions are thresholded based on the relative dominance of one motion direction.

To summarize, our model has three parameters: the decision threshold  $T$ , the ramp rate for motion accumulation  $R$ , and the relative valuation parameter  $L$ . An individual model is instantiated by choosing values for these three parameters, randomly selecting gains and biases for each neural population, and using the NEF to train a neural network whose synaptic weights realize the necessary cognitive operations. It is important to note that these three parameters are theoretical abstractions: They functionally describe how the network manipulates the represented information when performing the task and influence the selection of synaptic weights, but do not govern network dynamics during the simulation itself. Instead, the network operates solely via the communication of spikes through weighted synaptic connections. Each population in our model contains 500 spiking leaky integrate-and-fire (LIF) neurons, which typically fire at a rate below 80 Hz. Neurons communicate via spikes transmitted through current-based exponential synapses with time constant  $\tau = 30$  ms (the recurrent connection in *memory* has a longer time constant,  $\tau = 100$  ms).

### Drift Diffusion Model

We simulate a DD model to provide a point of comparison for our SNN model and to contrast the assumptions and behaviors of our relatively complex model with a traditional computational model of DM. The model begins each trial by waiting  $t_{nd}$  seconds, after which the DV begins to drift.  $t_{nd}$  is drawn from a normal distribution on each trial:  $t_{nd} = N(\mu_{nd}, \sigma_{nd})$ . However, we choose to eliminate intertrial variance in the NDT by setting  $\sigma_{nd} = 0$  in all four experiments; our preliminary experiments suggested that optimizing  $\sigma_{nd}$  did not improve fits to the RT training data, but that doing so reduced the quality of fits to RT data in novel contexts, and lead to poorer accuracy predictions. Next, the DV is set to its starting point  $S$ , which is also drawn from a normal distribution on each trial:  $S = N(\mu_S, \sigma_S)$ . However, we simplify our DD model slightly by setting  $\mu_S = 0$  for all experiments. During drift, we update the DV every  $\Delta t = 0.001$  s by sampling a new percept from the distribution  $N(R, V)$  and adding the result to the current DV.  $R$  is the trial-specific drift rate and  $V$  is the trial-independent drift variance. In order to ensure that  $R$  is systematically related to the coherence of RDM, we draw  $R_0$  from a normal distribution on each trial,  $R_0 = N(\mu_R, \sigma_R)$ , then compute  $R$  by multiplying by the motion coherence:  $R = R_0 C$ . This method for setting  $R$  differs somewhat from other approaches, which typically fit a unique value of  $R$  to each value of  $C$ . Finally, when the DV crosses the trial-independent decision threshold  $T$ , a choice is made. For the two-choice tasks in this article, a correct choice occurs when the DV crosses  $+T$  and an incorrect choice occurs when the DV exceeds  $-T$ . In summary, our DD model contains six free parameters:  $\mu_{nd}$ ,  $\sigma_S$ ,  $\mu_R$ ,  $\sigma_R$ ,  $V$ , and  $T$ . Below, we discuss how these parameters are fit to the behavioral data.

### Extrema Detection Model

Although integration models like DD capture many of the neural and behavioral phenomena associated with the SAT, they are not the only class of models capable of doing so. Recent work by Stine et al. (2020) has shown that nonintegration models fit data from RDM tasks as well as DD in many experimental setups. This finding adds to a long history of work with simple mathematical models, dating as far back as Pietsch and Vickers (1997), Smith and Vickers (1989),

and Vickers et al. (1971), that have also been applied to nonperceptual tasks (or “expanded judgment” tasks, Ratcliff et al., 2016). The exact criteria and computations required for decisions range widely between these models (Pietsch & Vickers, 1997), but most are considerably simpler than DD.

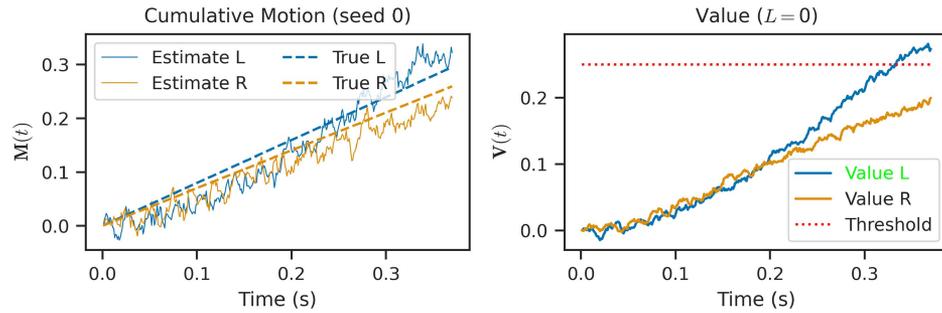
We simulate and analyze a nonintegration model to provide another point of comparison for our SNN model. However, choosing which simple model to simulate is problematic, because not all models are applicable to all DM tasks. For instance, some nonintegration models use clusters (runs) of consistent information to drive decisions: If  $T$  consecutive samples all point to the same decision, the model will choose that action. This model can be applied to nonperceptual tasks that involve sequential sampling, but not easily to perceptual tasks with continuous sampling (such as RDM, where the consistency of each sample is poorly defined). On the other hand, other nonintegration models use extreme samples to drive decisions: If the magnitude of some sample is greater than  $T$ , the model will choose that action. This can be applied to perceptual tasks with noisy, real-valued information, but not easily to nonperceptual tasks where samples have equal magnitude (such as the task in Fiedler et al., 2021, where each sample is  $\pm 1$ ).

We adopt an extrema detection model in this article for several reasons. First, Stine et al. (2020) showed that this model behaves similarly to DD in many experiments, raising questions about whether a simple heuristic may explain DM as well as an extended reasoning process. Second, extrema detection is mechanistically simple (samples are taken, compared to a threshold, and immediately rejected or accepted) and parametrically simple, the decision threshold  $T$  is the critical parameter, although we also include a normally distributed NDT  $t_{nd} = N(\mu_{nd}, \sigma_{nd})$ . Third, extrema detection is flexible enough to simulate various tasks with the appropriate definition of a “sample.” For the RDM, we adopt the same sampling procedure for extrema detection that is used in the SNN model: a new sample  $\mathbf{i} = N(\mathbf{P}(C), \sigma)$  is taken every  $t_s$ , and its magnitude is compared with the decision threshold  $T$ . For the stock market task, we add a simple memory: The model stores the length of the current “run” of consecutive, consistent samples. If a new sample arrives that points in the same direction as the last sample, this memory is incremented by one; otherwise, it is reset to one. If a run longer than  $T$  is detected, the model chooses that option. These mechanisms (memory and cluster detection) are reminiscent of several models previously applied to nonperceptual tasks (Pietsch & Vickers, 1997), while preserving the core one-shot heuristic of extrema detection.

## Results

### Dynamics and Parameters

We begin by inspecting the dynamics of the model in a two-choice RDM task. Figure 2 plots the cumulative motion estimates (left panel) and the value estimates (right panel) as percepts are dynamically sampled from the environment. The *memory* population integrates external signals represented in the *sensory* population, causing cumulative motion estimates to grow at a rate proportional to  $R$  (here,  $R = 1.5$ ). For comparison, we plot the ground truth cumulative motion, which we calculate by multiplying the true motion probabilities  $P$  by  $R$ . In this trial, motion coherence is set to  $C = 6.4\%$ , leading to  $P = [0.436, 0.564]$ . The *memory* population

**Figure 2***Network Dynamics in a Two Choice, Random Dot Motion Task (Correct Choice)*

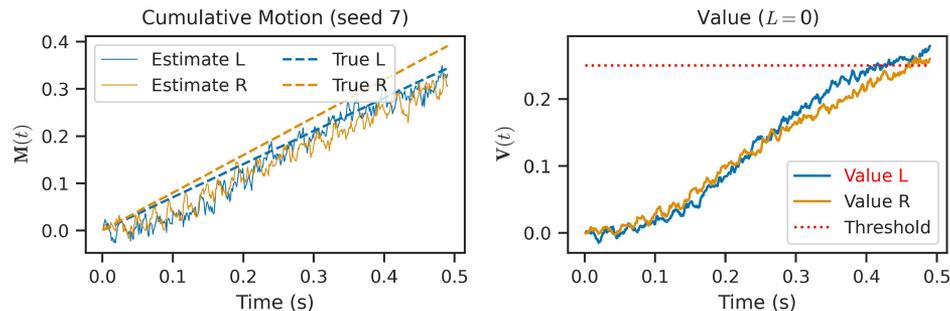
*Note.* The left panel shows the cumulative motion in each direction (blue for leftward motion, orange for rightward motion): The dotted lines are the ground truth, and the solid lines are the estimates of  $\mathbf{M}(t)$  decoded from the spiking activity of memory. The right panel shows the computed DVs: For an absolute value calculation ( $L = 0$ ),  $\mathbf{V}(t)$  mirrors  $\mathbf{M}(t)$ , although spike noise is reduced due to additional synaptic filtering. The red dashed line plots the decision threshold  $T$ ; when either DV crosses this line, excitation from value overcomes inhibition from gate, and action neurons begin to fire, producing a decision. On this trial, the model makes the correct decision after 0.37 s. DVs = decision variables; L = left; R = right. See the online article for the color version of this figure.

maintains a reasonable estimate  $\mathbf{M}(t)$  of this true value, but sample noise and spike noise lead to drift over time (as we would expect in any biological system). The right panel shows the DVs represented in the *value* population, after the synaptic connection between *memory* and *value* computes  $\mathbf{V}(t)$  according to Equation 4. Here, we show an absolute value computation ( $L = 0$ ). For the given decision threshold ( $T = 0.25$ ), the DV for leftward motion overcomes inhibition at  $t = 0.37$  s, leading to a correct decision with a RT of 0.37 s on this trial.

Figure 3 shows another trial with the same experimental setup (identical model and motion coherence) but a different random seed for perceptual sampling. In this trial, noisy sampling and neural representation lead to inaccurate representations of  $\mathbf{M}(t)$ , causing the model to make the incorrect choice after 0.49 s. While these inaccuracies could be solved by increasing the number of neurons in our model (improving the signal-to-noise ratio), or by reducing sampling noise by lowering  $\sigma$  and  $t_s$ , these inaccuracies are an important source of incorrect decisions in our model. In fact, we

found that they are necessary to reproduce the behavioral variance in RT and accuracy observed in the empirical data. To reduce model complexity, we hold the number of neurons, and their maximum firing rates, constant across all our experiments.

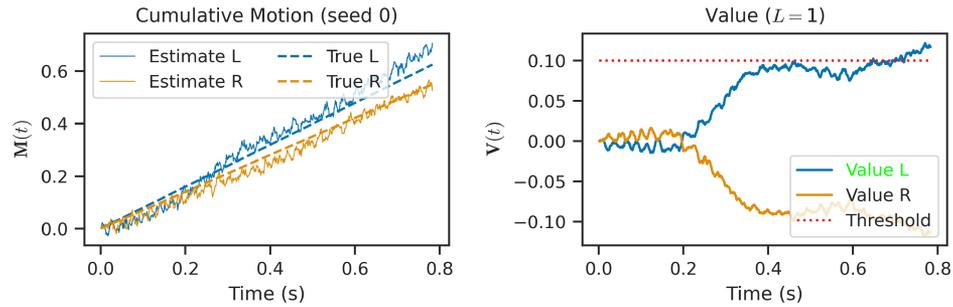
An alternative strategy for handling noisy motion estimates is to compute value based on relative evidence rather than absolute evidence. Figure 4 show the same trial (perception seed) as in Figure 3, but with the network computing relative value ( $L = 1$ ) using a lower decision threshold  $T = 0.1$  (the *sensory* and *memory* populations are identical to the previous simulations). When computing relative value, the network estimates minimal difference between the two DVs during the first 0.2 s of simulation: Only when  $M_0(t) \gg M_1(t)$  does relative value cross threshold, resulting in a slower RT (0.78 s) but a correct choice (unlike the  $L = 0$  network in Figure 3). Of course, relative valuation does not guarantee correct choices in the face of noise: using a lower threshold (e.g.,  $T = 0.01$ ) may still result in an incorrect choice, while using a higher threshold may prevent decisions altogether. In

**Figure 3***Network Dynamics in a Two Choice, Random Dot Motion Task (Incorrect Choice)*

*Note.* On this trial, spike and sampling noise lead to inaccurate representations of  $\mathbf{M}(t)$  and poor separation between the DVs represented by  $\mathbf{V}(t)$ , causing the model to make the incorrect choice after 0.49 s. DVs = decision variables; L = left; R = right. See the online article for the color version of this figure.

**Figure 4**

*Network Dynamics in a Two Choice, Random Dot Motion Task (Relative Valuation)*



*Note.* When computing relative value from cumulative motion strength ( $L = 1$ ), the network can avoid errors caused by small DV fluctuations near the decision threshold. The downside is a slower RT (0.78 s) and greater sensitivity to  $T$ . RT = response time; DVs = decision variables; L = left; R = right. See the online article for the color version of this figure.

many cases, an intermediate value of  $L$  is the ideal compromise: Including an absolute component in valuation ensures that the DVs always grow toward threshold as time passes, while including a relative component in valuation reduces the likelihood that small, momentary differences between the DVs will prompt a hasty decision. Essentially, choosing  $0 < L < 1$  ensures that the model simulates an urgency signal, but requires a degree of certainty before choosing between two comparable options. An example is shown in Figure 5, which demonstrates that a mixed-valuation network

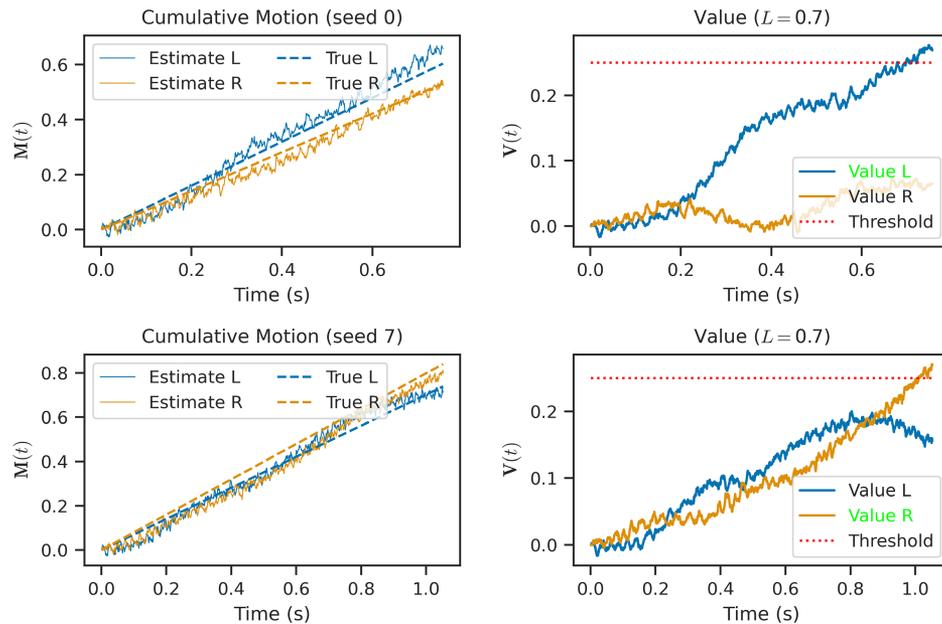
( $L = 0.7$ ) can make decisions quickly (faster than a  $L = 1$  network) and accurately (more correct choices than a  $L = 0$  network).

### Variance and Noise

To estimate the effects of noise in our model and to visualize the distributions of RTs that result, we created a model with middle-of-the-road parameters ( $R = 1.5$ ,  $T = 0.25$ ,  $L = 0.7$ ,  $\sigma = 0.3$ ,  $t_s = 0.02$ ), simulated 500 trials with different perception seeds, and analyzed

**Figure 5**

*Network Dynamics in a Two Choice, Random Dot Motion Task (Mixed Valuation)*



*Note.* When valuation includes both an absolute and a relative component, temporal urgency guarantees an eventual decision as the “effective relative threshold” shrinks over time, but decisions are still driven by differences between the cumulative motion estimates. Here, we show a mixed-valuation network ( $L = 0.7$ ) choosing correctly on both the previous trials, with intermediate RTs. RT = response time; L = left; R = right. See the online article for the color version of this figure.

the variance. Recall that there are two sources of randomness in our model: noisy perception of the external world, which occurs in our *sampling* module; and different realizations of spike noise on each trial. We do *not* simulate trial-to-trial variance in any model parameters. The left panel of Figure 6 depicts model dynamics: Each line represents the DV associated with the correct choice on each trial (decoded from the *value* population) and is colored according to whether the model chose correctly (blue) or incorrectly (orange) on that trial. The right panel plots the distribution of RTs across all trials, again color-coded by accuracy.

Several patterns are apparent in the simulated data. First, the decoded DVs exhibit significant variation, both over time and across trials. This variation qualitatively resembles the DVs decoded from monkey LIP while performing the RDM task (Steinemann et al., 2022). Second, DV variation produces a distribution of RTs whose mean and variance depend on the model parameters. Third, it appears that error trials have longer RTs on average; this result (slow errors) is qualitatively consistent with behavioral results in the RDM task (Ratcliff et al., 2016). Together, these results suggest that variance within and across trials may be explained by sample and spike noise in a neural network, rather than by variance in model parameters across trials (as is common in DD models, see the Discussion section). In the following sections, we compare simulated neural and behavioral variance with empirical data more quantitatively.

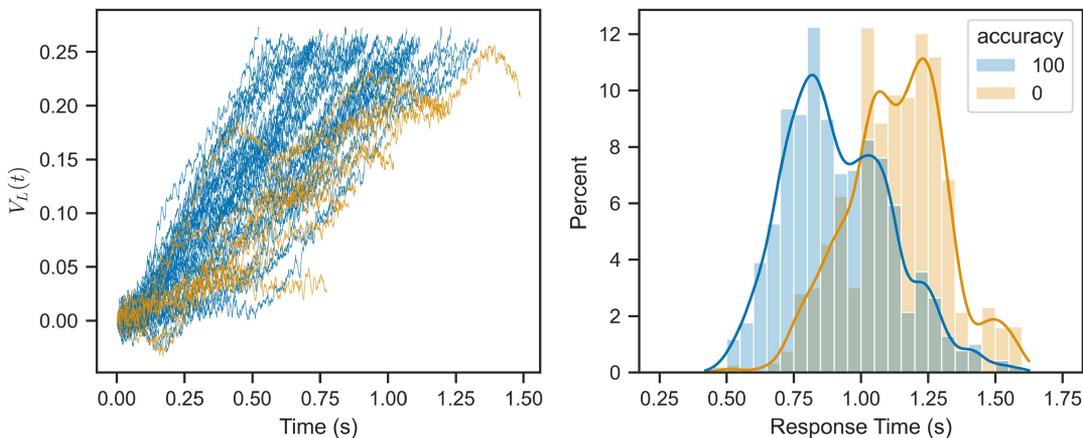
### Parameter Fitting Procedure

To validate our model against empirical data in the following experiments, we fit the SNN parameters ( $T$ ,  $R$ ,  $L$ ,  $\sigma$ ,  $t_s$ ), the DD

parameters ( $\mu_{nd}$ ,  $\sigma_S$ ,  $\mu_R$ ,  $\sigma_R$ ,  $V$ ,  $T$ ), and the extrema parameters ( $T$ ,  $\mu_{nd}$ ,  $\sigma_{nd}$ ,  $\sigma$ ,  $t_s$ ), to behavioral data using an iterative procedure with the Optuna optimization framework (Akiba et al., 2019). In this procedure, we randomly choose values for each parameter, build the model with these parameters, simulate the model on a large number of trials using different sampling and/or network seeds, and record the RT and accuracy on each trial. Depending on the experiment, we hold some of these parameters constant across experimental conditions. For instance, in the RDM experiments, we assume that all model parameters stay the same for different values of the motion coherence  $C$ ; and in the speed- versus accuracy-emphasis experiments (Experiments 1 and 3), we only allow one model parameter (the decision threshold  $T$ ) to vary between conditions. We explain how these choices support our claims about model generalization and prediction in the corresponding section for each experiment.

Following simulation, we create histograms of RTs for the simulated and empirical data, then calculate the difference between these histograms by computing a kernel density estimate, evaluating each kernel over the range of RTs, and computing the root-mean-square error between them. The resulting loss measures how different the distribution of simulated RTs is from the empirical RT distribution. Optuna uses this loss metric, along with a tree-structured Parzen estimator algorithm, to generate a new set of parameters to test. This resampling and testing continues for a fixed number of optimization trials, with the algorithm generally finding better parameters as the search continues. Note that this fitting procedure is qualitatively different from the parameter estimation techniques used to fit DD models to empirical data (Ratcliff & Childers, 2015).

**Figure 6**  
*Variance in Decision Variables (Left) and Response Times (Right) Arise From Sample Noise and Spike Noise Within the Network*



*Note.* We simulated 500 trials with unique seeds for sensory sampling, then analyzed how DVs and RTs varied within and across trials. We held all model parameters ( $R$ ,  $T$ ,  $L$ , etc.) constant across trials: All intertrial variance emerged from noise within the network. Plotting the distribution of RTs across trials, we observed that (a) the model exhibited significant behavioral variance due to this noise and (b) error trials tended to have longer RTs than correct trials, despite the absence of parametric variance or a starting point parameter  $S$  in the model. DD = drift diffusion; RT = response time; DVs = decision variables; LIP = lateral intraparietal area. As expected from Figures 2–5, the DVs decoded from our spiking neural activities drifted over time, consistent with the DD model and with DVs decoded from neurons in LIP (Steinemann et al., 2022). See the online article for the color version of this figure.

Importantly, we do *not* provide accuracy data or neural data to the fitting procedure: We treat these data as a separate “validation set” that the model must predict, given that it has been calibrated to the “training set” of RT distributions. We review these topics further in the Discussion section. The fitted parameter values for all simulations can be found in Table A1.

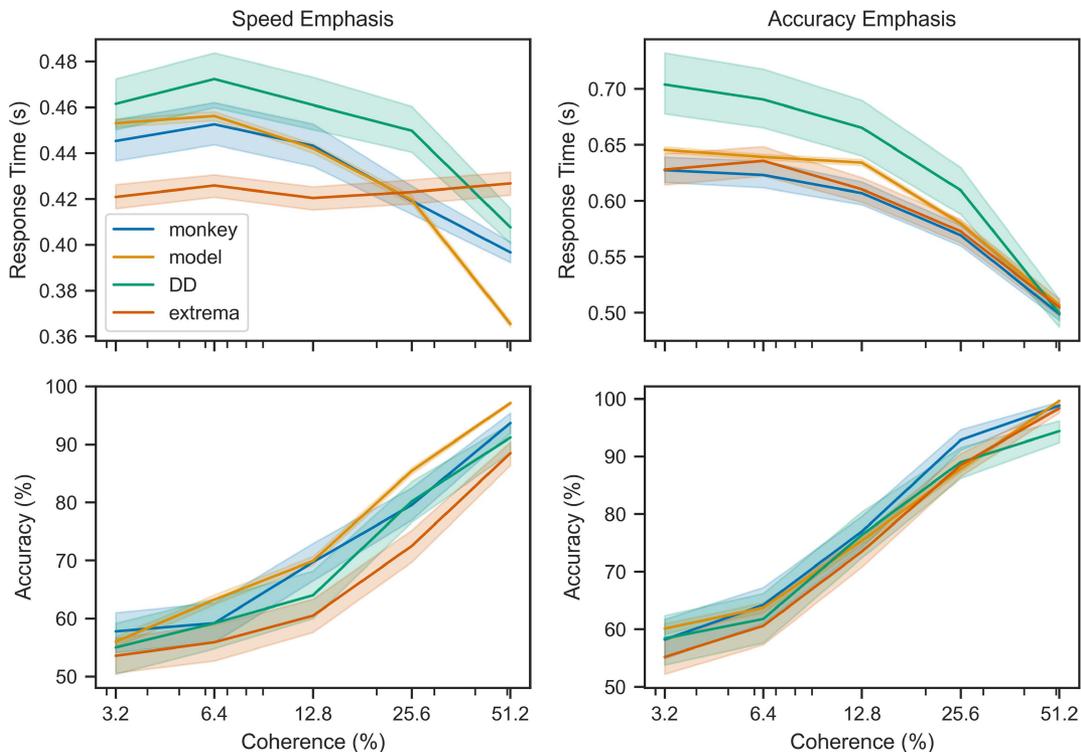
### Experiment 1: Speed Versus Accuracy Emphasis

Having established the basic behavior of our model, we now seek to validate it against various classes of empirical data related to the SAT. Studies have shown that instructing participants to favor speed or accuracy leads to systematic behavioral changes in RDM tasks: Emphasizing speed leads to shorter RTs but lower accuracy, while emphasizing accuracy leads to longer RTs and higher accuracy (Palmer et al., 2005; Rae et al., 2014; Vandekerckhove et al., 2011; Voss et al., 2004; Wagenmakers, 2009). Other studies have shown that these behavioral changes are associated with specific patterns of

neural activity: The firing rate of LIP neurons steadily increases during dot motion but converges to a consistent firing threshold immediately before decision (Churchland et al., 2008; Hanks et al., 2014; Steinemann et al., 2022). In this experiment, we first show that our model can perform the RDM task with speed and accuracy comparable to monkeys, then recreate the effects of speed- and accuracy-emphasis by adjusting a single model parameter, the decision threshold  $T$ . In doing so, we demonstrate that executive control may adjust the cognitive strategy realized by our model, shifting its behavior in a manner consistent with neural and behavioral data.

We begin by fitting our model to recreate the RTs of an individual monkey from Hanks et al. (2014). Figure 7 plots RT and accuracy as a function of motion coherence for this individual, our fitted model, and the fitted DD and extrema models. Qualitatively, all three models captured how RT decreases as coherence increases for the accuracy-emphasis condition (upper right) and were able to predict the corresponding changes in accuracy (lower right). We then

**Figure 7**  
*Response Time (Top) and Accuracy (Bottom) as a Function of Experimental Instructions to Favor Speed (Left) or Accuracy (Right) in a Two Choice RDM Task*



*Note.* We fit model parameters to reproduce the empirical RT distributions for a single monkey (see Figure 8), keeping all parameters the same between conditions except for the decision threshold  $T$ . Our model (orange) reproduced the monkey’s RT curves with reasonable accuracy in both the speed and accuracy conditions. While the DD model (green) and extrema model (red) qualitatively reproduced the shape of the RT curves, their RT estimates were consistently worse than our SNN: The DD model produced longer and more variable RTs, while the extrema model failed to capture the relationship between RT and  $C$  in the speed condition. All three models successfully predicted the monkey’s accuracy curves, despite not being trained on accuracy data. Plotted data are the mean and 95% bootstrapped confidence intervals across all trials. DD = drift diffusion; RT = response time; RDM = random dot motion; SNN = spiking neural network; LIP = lateral intraparietal area. Empirical data are from Hanks et al. (2014). See the online article for the color version of this figure.

lowered the  $T$  parameter in all three models to emulate the speed-emphasis condition. As expected, this change decreased RT and accuracy while preserving the characteristic shape of coherence curves. Overall, our SNN, the DD model, and the extrema detection model all reproduced the SAT in this task, predicted the changes in accuracy accompanying the changes in RT, and accommodated contextual changes (speed-emphasis) through changes in a single model parameter, the decision threshold.

To compare the quantitative performance of these three models, we computed the mean absolute error between the empirical data and the simulated data for each value of  $C$  in Figure 7, for both experimental contexts (speed- vs. accuracy-emphasis) and dependent variables (RT and accuracy). We found that our model had the lowest average error in two conditions, RT error in speed condition ( $0.009 \pm 0.01$ ) and accuracy error in accuracy condition ( $1.96 \pm 1.62$ ), while the extrema model had the lowest average RT error in the speed ( $0.005 \pm 0.004$ ), and the DD model had the lowest average accuracy error in the speed condition ( $2.32 \pm 2.0$ ). However, despite these successes, the mathematical models fell short in several respects. For example, the extrema model did not predict a significant decrease in mean RT with  $C$ , and the DD model produced a few slow responses for each  $C$ , leading it to overestimate the mean and variance of RTs.

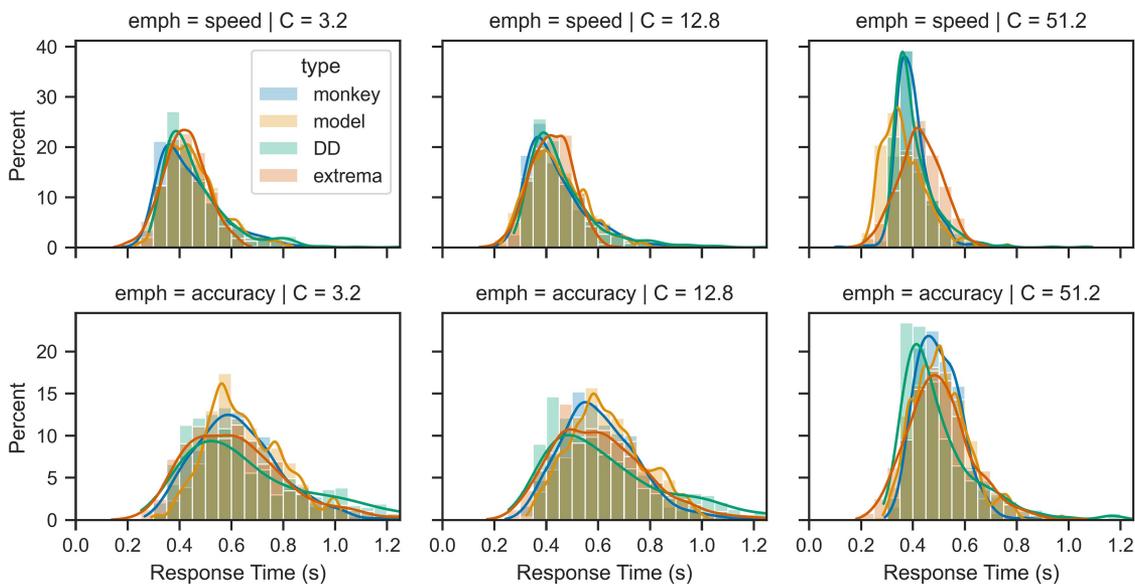
There are several reasons why the DD and extrema models struggle to recreate the RT data. First, the effect size is quite small in the speed condition (a difference of about 20 ms between  $C = 3.2$  and  $C = 51.2$ ), compared to the accuracy condition (a difference of about 100 ms). In the speed condition, individuals will sometimes answer as quickly as possible, regardless of coherence, making an analysis of RT versus coherence unreliable. In fact, the other monkey in the Hanks et al. (2014) data set had an even smaller effect

size. Second, we placed an important constraint on model fitting: All parameters except  $T$  were identical in the two experimental conditions. While this choice enforced a cognitive constraint that we argue increases model realism (see the Discussion section), it limited the flexibility of our fitting procedure. Indeed, we found that when we allowed all model parameters to vary between the speed- and accuracy-emphasis conditions, our model fits improved significantly (for all three models). Notably, our constraint that drift rate in the DD model be directly proportional to motion coherence ( $R = R_0C$ ) and identical across speed-emphasis conditions significantly limited the performance of the DD model, compared to the traditional approach in which  $R$  is fit to every value of  $C$  and/or experimental context separately. Third, and perhaps most important, we argue that focusing on changes in the mean RT obscures an important aspect of the model fit: The ability to capture variance in RTs across trials. Figure 8 plots the distribution of RTs for each  $C$  and provides a richer picture of how RTs vary across both trials and motion coherence. For instance, this plot shows that, in both simulated and empirical data, the RT distributions generated by all three models have long tails at low  $C$  but short tails at high  $C$ . Overall, these results demonstrate that our model can capture (and predict) the behavioral differences induced by explicit instructions to favor speed or accuracy.

Next, we recorded the spiking activity from neurons in the *value* population of the fitted models and analyzed their dynamics. As mentioned above, previous studies have measured ramping activities in purported neural accumulators in sensory and motor cortex (Heitz, 2014; X.-J. Wang, 2008); many of these studies suggest that higher motion coherence is associated with faster ramping, but that firing rates converge to a set firing rate threshold immediately before a decision is made (Churchland et al., 2008;

**Figure 8**

*Response Time Distributions as a Function of Motion Coherence and Speed/Accuracy Emphasis*



*Note.* These data expand on the top panel of Figure 7, showing that, while the model may not exactly reproduce the mean RT, it reproduces the entire distribution with surprising accuracy. DD = drift diffusion; RT = response time. See the online article for the color version of this figure.

Hanks et al., 2014). To select model neurons that had the appropriate tuning curves, we recorded spikes from simulated neurons in the *value* population that were strongly selective for motion in the direction of coherent motion (i.e., those that had encoders with large, positive components in the appropriate direction). We then filtered our simulated spikes using a low-pass filter ( $\tau = 0.03$  s) and plotted the mean activity (with confidence intervals across neurons and trials) over time. Figure 9 shows the firing rate of model neurons in the 0.4 s following stimulus presentation (left panel) and the 0.3 s preceding action selection (right panel).

Several trends are apparent in the neural data. First, both simulated neurons exhibit ramping activities that start approximately 0.05 s after stimulus presentation: starting from a baseline firing rate, neurons monotonically increase their firing rate, and the “buildup rate” is higher in trials with greater motion coherence. Second, if the spikes across trials are aligned to the RT (instead of the start time), simulated neurons ramp to a final firing rate that seems to be constant across motion coherence. Third, ramp rate appears to be greater in the speed-emphasis condition, but the firing rate threshold is not strongly affected by speed-emphasis. These trends are all consistent with neural data in previous studies (Churchland et al., 2008; Hanks et al., 2014; Steinemann et al., 2022). However, a closer look at firing patterns in the empirical data reveals that the activities preceding stimulus presentation and following action selection are not well-approximated by our model. We attribute these differences to startup transients and resetting

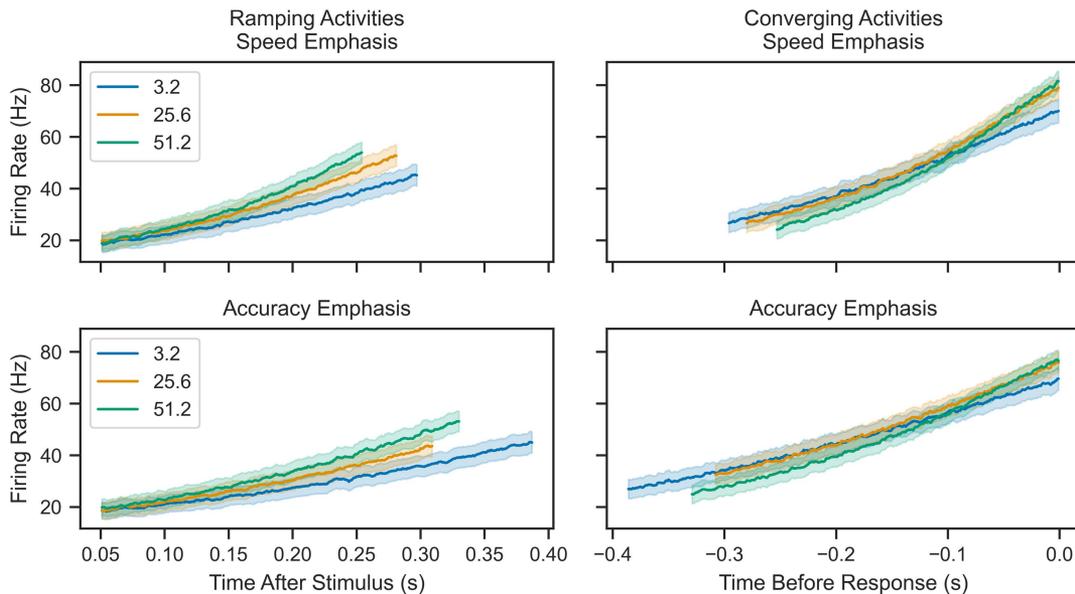
dynamics dictated by parts of the brain outside the scope of the current project (see the Discussion section).

We also recreated the analysis from Hanks et al. (2014) in which the authors calculate the mean buildup rate and threshold for neural activities at multiple coherence values, then plot these two metrics as a function of coherence. We calculated the buildup rate  $B$  as

$$B_i = \frac{A(t_1) - A(t_0)}{t_1 - t_0}, \quad (7)$$

where  $i$  is the neuron index and  $A(t)$  is the firing rate at time  $t$ , which is calculated by smoothing  $i$ 's spikes with a low-pass filter. We choose  $t_0$  to be 0.1 s after stimulus presentation and  $t_1$  to be 0.1 s before the recorded RT: discarding the initial and final 100 ms eliminates irrelevant spiking activity at the beginning and end of the simulation. Figure 10 plots the mean  $B$  and mean  $A(t_1)$  as a function of motion coherence and speed-emphasis.  $B$  increases with coherence, and the rate of this increase is similar for both the speed- and accuracy-oriented models. However, the speed-oriented model exhibits greater  $B$  for each value of  $C$ , reproducing the offset apparent in the empirical data. In contrast, the final firing rates change very little as coherence increases, regardless of speed-emphasis. This result is also consistent with the neural data: see Figure 4 of Hanks et al. (2014). We again note that the model was not built with these neural constraints in mind, nor were model parameters fit to reproduce these neural data.

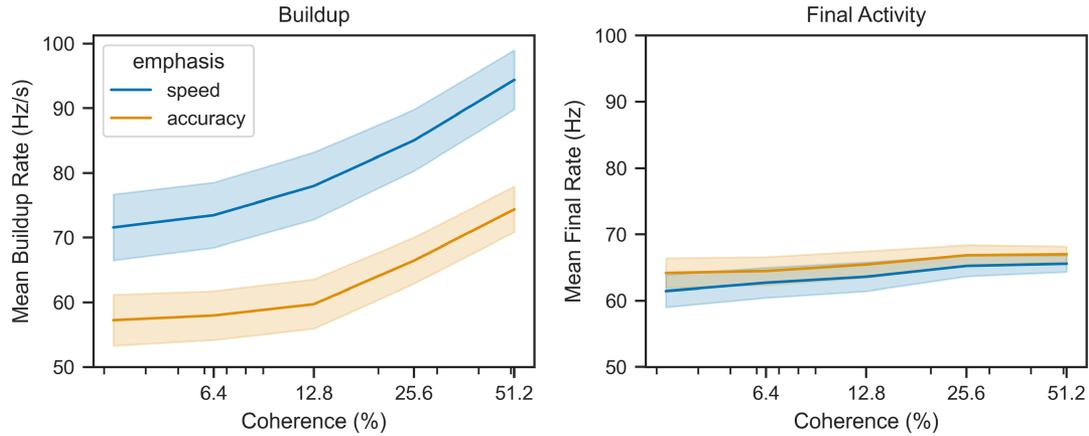
**Figure 9**  
*Neural Activities Following Stimulus Presentation (Left) and Preceding Action Selection (Right) as a Function of Motion Coherence and Speed/Accuracy Emphasis*



*Note.* We recorded the spikes from direction-selective neurons in the *value* population while the fitted model performed the task, smoothed them using a low-pass filter, and plotted the mean and variance (95% bootstrapped confidence intervals) over time. The left panel clearly exhibits ramping activity as motion evidence from the sensory population accumulates in memory and drives activity in *value*, while the right panel shows that neurons in *value* reach approximately the same firing rate about 0.1 s before the model makes a decision. Higher motion coherences are associated with a faster buildup in firing rate, but do not affect the observed final firing rate. See the online article for the color version of this figure.

**Figure 10**

*Firing Rate Buildup (Left) and Final Firing Rate (Right) of Direction-Selective Value Neurons as a Function of Motion Coherence and Experimental Instructions*



*Note.* Buildup rate and final rate are calculated from neural spikes according to Equation 7, once for each neuron (70) and each trial (100). We plot the mean values of  $B$  and  $A(t_i)$  for each  $C$ ; error bars depict data between the 40th and 60th percentile, giving a coarse indication of the variance in firing rate (buildups) across neurons and trials. Buildup rates increase linearly with motion coherence regardless of speed- or accuracy-emphasis, but these rates are consistently higher when participants are instructed to favor speed over accuracy. In contrast, final firing rates are largely independent of motion coherence or accuracy instructions: Any differences in mean rates are small in absolute terms and are dwarfed by the variance across neurons and trials. Both these model trends are consistent with empirical trends reported in Figures 3 and 4 of Hanks et al. (2014). See the online article for the color version of this figure.

Overall, these results confirm that the firing properties of our model are consistent with empirical data from monkeys performing the RDM task. In particular, we have recreated the differences (or lack thereof) in firing rates that are induced by instructing monkeys to favor speed versus accuracy. One interesting conclusion is that, while the decision threshold parameter  $T$  is the only parameter that changes between these conditions, we do not observe significant differences in the firing rates of neurons immediately before a decision. In the Discussion section, we explain how this is possible given our distinction between “threshold” in decision variable space and neural activity space, and argue that, compared to other neural models of the SAT, our model is uniquely capable of exploring their relationship.

## Experiment 2: Multiple Alternative Actions

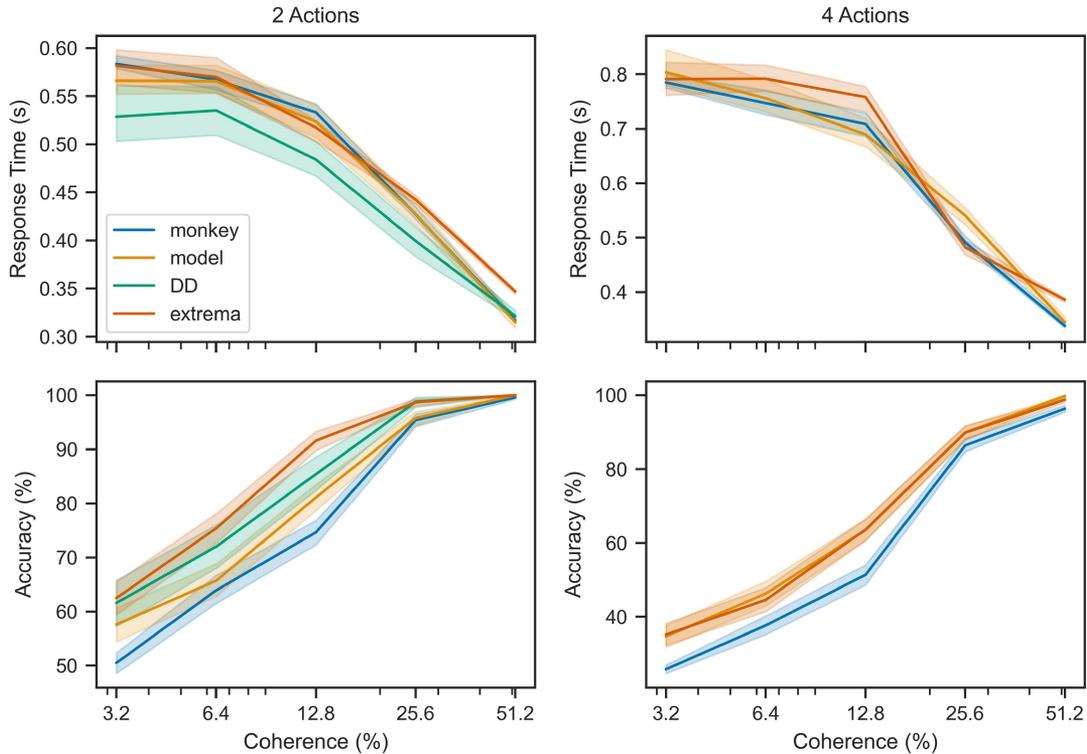
In this experiment, we apply our model to a RDM from Churchland et al. (2008), in which dots may move toward either two targets (the upper right or lower left of a square screen, 180° separation) or four targets (each corner of a square screen, 90° separation). Our goal is to show that the model naturally scales to tasks that feature more than two decisions, and that the same neural network architecture can complete both the two- and four-choice variants. As before, we validate our model by comparison to both behavioral and neural data from monkeys in Churchland et al. (2008). To fit the model parameters, we use the same approach as in Experiment 1, but we relax the constraint that model parameters are the same in the two- and four-choice conditions. However, we reiterate that our model architecture (the number of neural populations and their structural connectivity) does *not* change between these conditions; the only difference between the models

performing the two- and four-choice variant of the task are (a) the amount of sampling noise ( $\sigma$  and  $t_s$ ) and (b) the synaptic weights, which realize the cognitive parameters ( $T$ ,  $R$ ,  $L$ ). We did not apply the DD model to the four-choice task: while some studies have extended the DD model to multiple alternative actions (Krajbich & Rangel, 2011; Roxin, 2019), doing so requires significant changes to the model that are outside the scope of the current work.

Figure 11 plots RT and accuracy as a function of motion coherence for an individual participant, our fitted model, and fitted DD and extrema models. The simulated data from all three models closely match the RT data in both task conditions. In addition, all three models predict the correct shape for the accuracy curves in both experimental contexts, but overestimate accuracy to varying degrees. As in Experiment 1, the simulated RT distributions have long tails at low  $C$  but short tails at high  $C$ , matching the empirical distributions. Similarly, we observe coherence-dependent ramping of neural activities, and coherence-independent firing thresholds, that are consistent with the neural data from Churchland et al. (2008), and which capture the differences between the two- and four-choice contexts. We omit these two plots for brevity.

We also analyzed how the distribution of RTs differed between correct and incorrect trials. As mentioned above, individuals performing RDM tasks will often have longer reaction times on trials in which they choose incorrectly. To reproduce this trend, DD models typically employ trial-to-trial variance in  $S$ ,  $R$ , or  $T$  (Ratcliff et al., 2016). We were curious whether our model would exhibit “slow errors” without these mechanisms. Figure 12 plots the distribution of RTs for correct and incorrect trials at each value of  $C$  for the two-choice context. Monkeys frequently produce slow errors for  $C > 3.2\%$ . Our model predicts the same trend, with an effect size comparable to the empirical data. In contrast, the DD model only

**Figure 11**  
*Response Time (Top) and Accuracy (Bottom) as a Function of Motion Coherence and Number of Candidate Actions in a RDM Task With Two (Left) or Four (Right) Choice Alternatives*



*Note.* As in Figure 7, we fit the parameters of our model (orange), the drift diffusion model (green), and the extrema model (red) to reproduce the empirical RT data (blue), then had the models predict accuracy in the two- and four-choice contexts. All three models successfully reproduced the RT curves in both contexts, but our model better predicted accuracy in the two-choice context. RT = response time; RDM = random dot motion. Empirical data are from Churchland et al. (2008). See the online article for the color version of this figure.

exhibited slow errors for  $C > 25.6\%$ , and the extrema model did not exhibit slow errors at all. We observed similar trends in the four-choice context. We also investigated whether removing our constraints on  $R$  and  $\mu_S$  for the DD model would produce more slow errors. However, when we fit  $\mu_R$ ,  $\sigma_R$ , and  $\mu_S$  separately for each value of  $C$ , we did not observe an increase in slow errors and instead observed poorer accuracy predictions.

Overall, these results show that our model naturally scales to multiattribute decision tasks without architectural changes, indicate that our model better predicts performance of monkeys in these tasks than does the DD or extrema model, and suggest that sampling and spike noise may be an alternative explanation for the slow errors observed in RDM tasks.

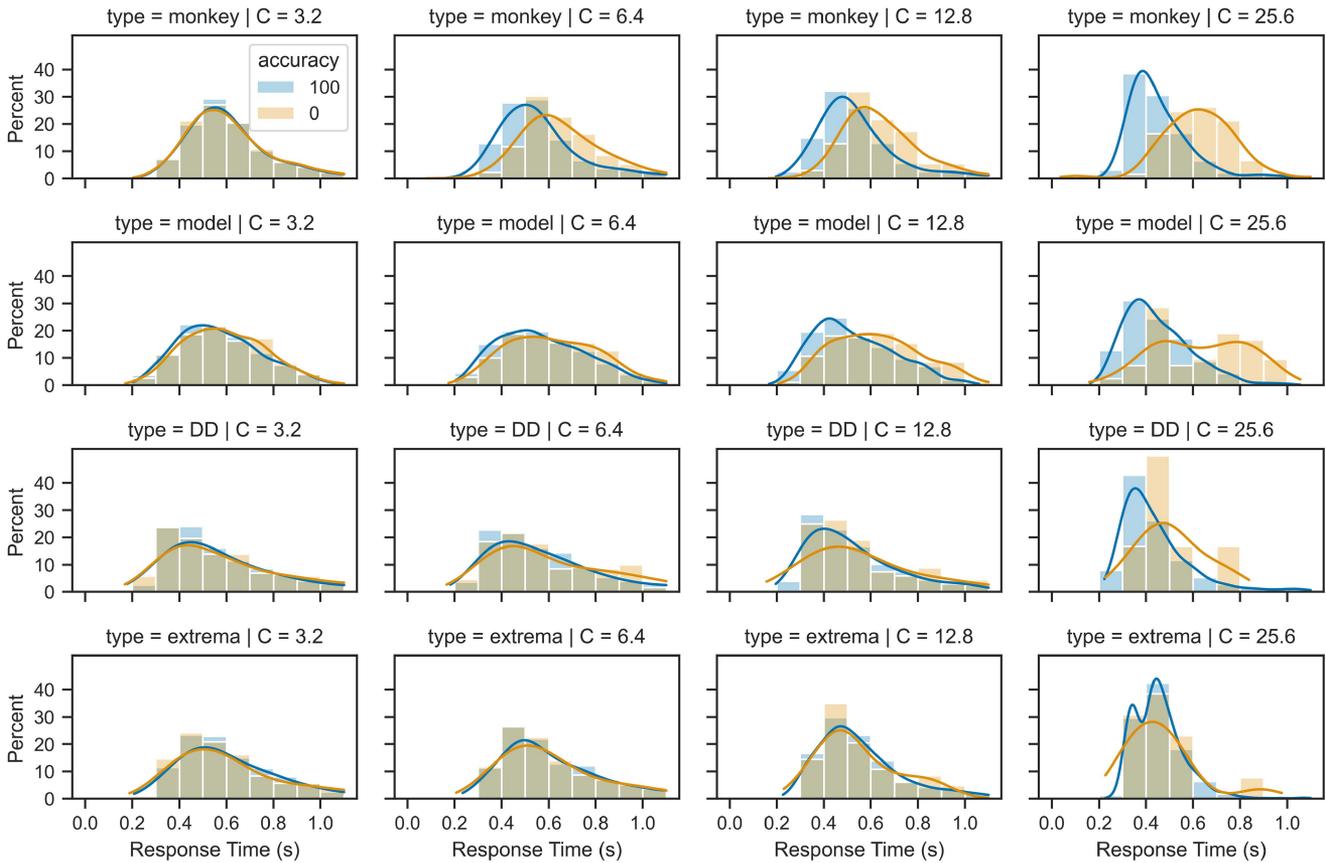
### Experiment 3: Aging and Cognitive Decline

By investigating DM and the SAT with spiking neuron models, we gain the ability to explore how low-level biological changes impact high-level cognitive performance. In this experiment, we recreate the changes in speed and accuracy that accompany aging by perturbing the synaptic connectivity in our model. As mentioned in the

introduction, different cognitive theories attribute these age-induced changes to different biological deficits. We investigated three potential biological origins of age-related decline: perceptual inaccuracy, visual working memory deficits, and impaired threshold setting. According to the first theory (Faubert, 2002; Hutchinson et al., 2012), elderly individuals have an impaired ability to perceive or encode visual motion signals, leading to deficits in perceptual DM tasks. According to a second theory, aging degrades connectivity within cortical areas that realize visual working memory (such as neural integrators in LIP), introducing noise into the neural estimates of accumulated DVs, reducing task performance, and increasing RTs. Given that working memory deficits may underlie decreased RDM performance in patients with attention-deficit/hyperactivity disorder (Mulder et al., 2010) and schizophrenia (Fish et al., 2018; Schweitzer & Lee, 1992), we hypothesize that comparable working memory deficits in elderly individuals (Nissim et al., 2017; Reinhart & Nguyen, 2019; M. Wang et al., 2011) could also explain elderly speed-accuracy deficits. Finally, according to the third theory, elderly individuals have an impaired ability to switch between prioritizing speed and accuracy according to task instructions (Forstmann et al., 2011). Neural and anatomical data show that elderly individuals have

**Figure 12**

Response Time Distributions for Correct Versus Incorrect Trials as a Function of Model Type and Motion Coherence



*Note.* RTs from trials in which the individual monkey gave the correct response are shown in blue, and incorrect trials are shown in orange. Monkeys exhibit slower RTs on error trials (slow errors), especially for higher motion coherences  $C$ . Our model (second row) exhibited a similar trend, as does the DD model (third row) to a lesser extent, but the extrema model (fourth row) does not. DD = drift diffusion; RT = response time; RDM = random dot motion. Empirical data are from the two-choice RDM task in Churchland et al. (2008). See the online article for the color version of this figure.

degenerated white matter connectivity between cortex and striatum, which may impair their ability to set the decision threshold, especially when switching between speed- and accuracy-emphasis conditions (Bogacz et al., 2010; Forstmann et al., 2008).

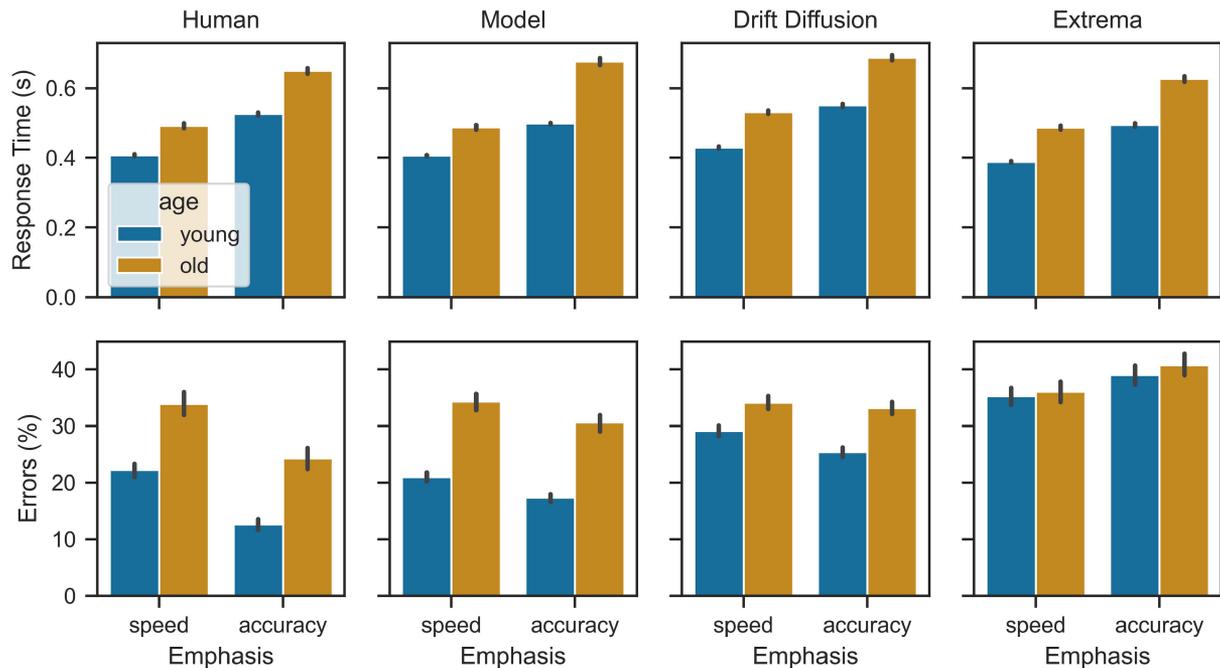
To test these three theories, we began by fitting the parameters of our model to the RTs from young participants in Forstmann et al. (2011). We fit one model to the RTs of each of the 18 young participants in this data set, then predicted their accuracy. Each individual model was characterized by a unique combination of the parameters ( $R$ ,  $L$ ,  $T_1$ ,  $T_2$ ,  $\sigma$ , and  $t_s$ ), where  $T_1$  and  $T_2$  are the decision thresholds applied during the speed- and accuracy-emphasis conditions, respectively. The results are shown in blue in Figure 13; as expected, both models reproduce the mean RTs (and the shape of the RT distributions, not shown) and predict the effects of speed-emphasis on accuracy. To simulate aging, we then degraded the synaptic connectivity between specific neural populations in the model by setting some fraction of the connection weights to zero, effectively simulating the partial loss of connectivity between those populations. To test the perceptual theory, we applied this degradation to the connection between *sensory* and *memory*. To test

the working memory theory, we weakened the recurrent connections in the *memory* population, which correspond to LIP. To test the “striatal” theory of threshold setting, we created a new population labeled *threshold*, which receives external input about the desired  $T$  for the current context, and which modulates the *gate* population. This new *threshold* population corresponds to cortical populations that modulate the decision network according to task instructions (Bogacz et al., 2010), while the *gate* population corresponds to striatum, inhibiting *action* until the decision threshold is exceeded. To simulate aging, we then weakened the “cortico-striatal” connections between *threshold* and *gate*. We implemented these three degradations for each of the previously fitted models, then simulated them to predict the resulting changes in RT and accuracy in the elderly cohort.

The results are shown in the center panel of Figure 13. Consistent with the empirical data, the elderly models have higher error rates than their young counterparts. However, only the accuracy deficits produced by degrading neural accumulators in LIP (shown in orange) captured the magnitude of the accuracy deficits observed in the elderly population. Furthermore, this manipulation produced

**Figure 13**

Mean Response Time (Top) and Accuracy (Bottom) as a Function of Experimental Instructions and Age



*Note.* After fitting our model to each young individual in Forstmann et al. (2011), we artificially aged our model by degrading the synaptic connections between specific model populations. We then compared the speed and accuracy of our young models (blue) with our elderly models (orange). We found that degrading the recurrent connections in the memory population produces changes in RT and accuracy that most closely resemble the empirical data from elderly individuals. In contrast, the DD and extrema models cannot fully reproduce the accuracy deficits of elderly individuals, even when all model parameters are allowed to vary between the young and elderly populations. DD = drift diffusion; RT = response time. See the online article for the color version of this figure.

longer RTs that resemble the slower responses of elderly individuals. In contrast, the other manipulations did not fit the behavioral data as well. Weakening connectivity between the perceptual system and neural accumulators did increase the RT in the expected way, but did not degrade accuracy by a substantial amount. Surprisingly, we found that weakening the model's corticostriatal connections actually reduced model RTs, contrary to the proposed theory. Our model therefore predicts that the empirical trends in Forstmann et al. (2011) are best explained by age-induced deficits in visual working memory related to weakened recurrent connectivity within the cortex. We discuss these results further in the Discussion section and hope that future neuroanatomical studies will directly test this prediction by examining the relationship between RDM performance and cortical connectivity.

We also fit the DD and extrema model to the human data and generated predictions about changes in task performance. In these mathematical models, there was no way to simulate degraded synaptic connectivity, so we instead performed separate parameter fits for the young and elderly cohorts independently. As before, we only allowed the decision threshold  $T$  to change between the speed-emphasis and accuracy-emphasis conditions. The results are shown in the right panels of Figure 13. Both the DD and extrema models do a good job fitting the mean RT across ages and emphases, capturing the trend that all individuals have longer RTs in the accuracy condition, and that older individuals have longer RTs overall.

However, neither model made good predictions about the resulting accuracy. The DD model correctly predicted that accuracy-emphasis leads to lower error rates, and that elderly individuals have higher error rates overall, but the magnitude of these trends was small compared to the empirical data. The extrema model predicted high error rates in all conditions, an insignificant increase in errors among the elderly cohort, and an increase in error rates in the accuracy condition. In the Discussion section, we argue that purely mathematical models, and especially the extrema model, are too simple to correctly predict how age independently influences RT and accuracy, but that our neural network model is able to dissociate these two trends and provide novel insights about the relationship between biological aging and cognitive performance.

#### Experiment 4: Speed–Accuracy Trade-Off in a Nonperceptual Task

In the final experiment, we apply our model to the stock market task from Fiedler et al. (2021). Recall that in this nonperceptual task, participants are presented with information about the rising or falling value of two stocks, each of which has a hidden probability of rising on each sample. Participants view samples sequentially and may choose either option at any time. After choosing one stock, the next trial begins immediately. Participants continue this process for a total of 20 min and are rewarded based on the total number of

correct responses they give. This task is distinct from RDM tasks in three respects: The SAT is implied by the reward structure of the task, rather than verbal instructions to emphasize speed or accuracy; sampling is divorced from sensory processing, while memory and deliberation occur over cognitive timescales (seconds, rather than milliseconds); and individuals frequently adopt dramatically different strategies to complete the task. These features allow us to investigate temporally extended, deliberative DM, as well as the high-level cognitive strategies exhibited by individuals.

One key feature of this task is that the correct choice is not immediately available to the viewer, and a momentary dominance in cumulative evidence for one choice over another does not imply that choice is correct. Even with perfect perception and noiseless internal representations, an individual cannot be assured of a correct choice until they sample multiple cues, and only when the maximum number of cues per trial (24) have been displayed does the cumulative evidence dictate the correct answer. Thus, a speed-oriented decision maker may choose an action after viewing several cues that objectively favor one action, but an accuracy-oriented individual may, after waiting to view all the available cues, discover that the other action is the correct choice.

Fiedler et al. (2021) argued that the best strategy in this task is to sample very few cues per trial, but complete many trials in the 20 min allotted by the experiment. This is because, while accuracy increases with the number of cues sampled, these gains are subject to diminishing returns; that is, accuracy begins at chance (50%) and increases rapidly for the first few samples, but plateaus as more samples are gathered (asymptotically approaching 100% after the final cue). In contrast, the number of trials that an individual can complete in 20 min decreases linearly with the number of cues sampled. Given certain assumptions (e.g., the losses for an incorrect decision are symmetric with the gains for a correct decision), the optimal strategy heavily favors speedy decisions over accurate ones (Fiedler et al., 2021). However, when the authors asked humans to perform the task, they found that people exhibited a consistent “accuracy bias”: They sampled far more cues per trial than the optimal decision strategy recommends. This trend held across multiple experimental conditions designed to provide instructive feedback to participant about their performance and the SAT. We were curious whether our model could reproduce the various strategies that humans exhibited in these experiments and to see whether our models exhibited a comparable accuracy bias.

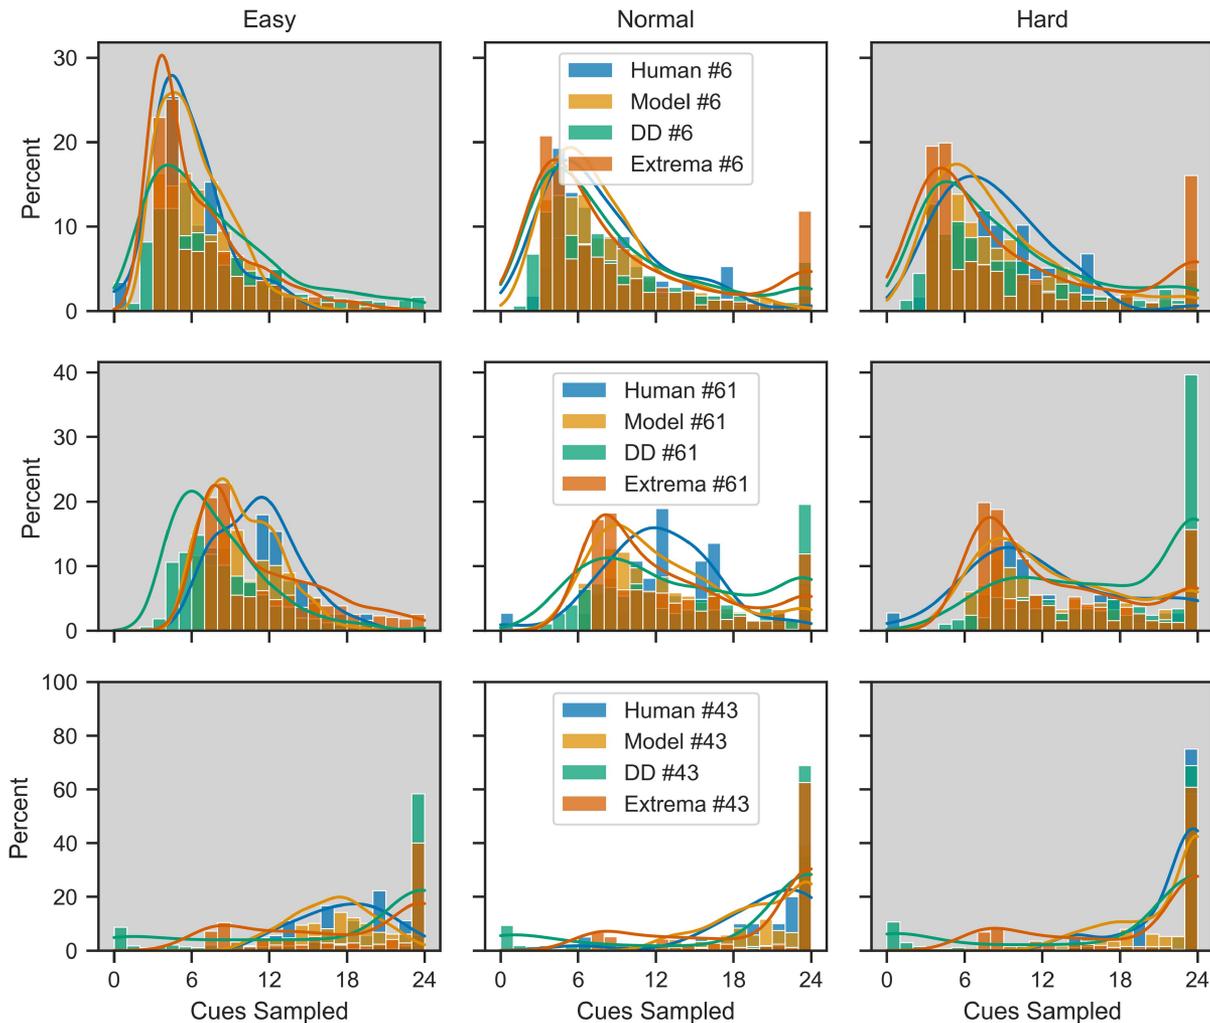
Our model remains the same as in the previous experiments, except that we simplify the *sampling* module to remove noise and resampling, as described previously. Noisy visual sampling and NDT were also removed from the extrema model: NDT is not relevant in this task because real-valued timing is replaced by integer-valued sampling. The extrema model instead made decisions using a simple memory: The model remembers the length of the current run of consecutive samples, and if a run longer than  $T$  is detected, the model chooses that option. We tested two definitions of “consecutive samples.” In the first, the model only tracks the longest run of positive samples for each stock, and chooses when one run exceeds  $T$ . In the second, the model tracks the longest run of positive samples and the longest run of negative samples for each stock; the model chooses a stock if its positive run exceeds  $T$  or if the negative run of the alternative stock exceeds  $T$ . These mechanisms of memory and cluster detection are intended to capture the sampling and memory heuristics of successful models from Pietsch and Vickers (1997)

while preserving the simplicity and flexibility of extrema detection models from Stine et al. (2020). Finally, to apply our DD model to this nonperceptual task, we assumed that sampling cues was analogous to sampling percepts in the RDM task. Recall that, in this task, each stock has a probability  $P$  of increasing in value for each sample, such that larger  $\Delta P$  between the two stocks implies an easier task. We accordingly set the drift rate  $R$  proportional to the task difficulty ( $R = R_0 \Delta P$ ) and had the model update its DV at every timestep by drawing samples from the distribution  $N(R, V)$ . To ensure DD dynamics were comparable to our SNN model, we used a simulation timestep of  $\Delta t = 0.001$  s and defined one cue to be equivalent to 0.1 of simulation; that is, if the DD model chooses after 1.2 s of simulation, we codify this as a RT of 12 cues. However, note that while both our model and the DD model update their DVs every  $\Delta t = 0.001$ , our model samples a new cue value every 0.1 s, while the DD model necessarily samples a new cue value every  $\Delta t$ .

We began by selecting participants from the human data set that exhibited three different strategies: a slow-but-accurate strategy, a middle-of-the-road strategy, and a fast-but-inaccurate strategy. The strategic behavior of each participant was consistent across the three difficulty conditions ( $\Delta P$ ) tested in the experiment. Each participant also tended to sample more cues as the task became more difficult, confirming that these strategies were context-sensitive. As before, we fit model parameters to RT data and predicted the resulting accuracy data for each individual. In this task, RT is defined as the number of cues sampled before a decision is made. To demonstrate the generality and predictive power of our model, we trained the model using only RT data from the moderate difficulty condition; we then tested the model in the easy and hard difficulty conditions, predicting both RT and accuracy in these novel contexts. We trained the extrema model in the same way. Figure 14 plots RT as a function of task difficulty for these three behavioral strategies. We found that our model, as well as the DD extrema models, fit the training data well, but that our model gave the best RT predictions across all conditions, see Table 1. Furthermore, our model predicted human accuracy as well as, or better than, the other models. These initial results suggest that our model can capture diverse strategies for navigating the SAT, but that simpler models may fail to do so.

Next, we investigated whether our model exhibited the same population-level speed-accuracy trends as the humans in Fiedler et al. (2021). We fit one model to each participant in the data set, again using only RTs in the moderate difficulty condition to fit the model. To investigate the SAT across the population, we recorded, for each individual, (a) the mean number of cues they sampled, as well as (b) the total number of trials they completed correctly within 20 min and (c) their mean accuracy. We then repeated the analysis from Fiedler et al. (2021). Figure 15 plots mean accuracy versus mean sampled cues, and Figure 16 plots the total number of correct trials versus mean sampled cues. In these figures, each data point corresponds to one individual in the data set. Several features of the human SAT are apparent. First, mean accuracy increases as a function of mean cues sampled; this trend appears linear in the moderate and hard conditions, but asymptotic in the easy condition, in which individuals regularly reach 100% accuracy. Second, the number of correct trials completed within the allotted 20 min decreases as a function of mean cues sampled; this trend appears exponential in all three conditions, with slower individuals correctly completing less than half the trials as the fastest individuals. Our model predicts both these features.

**Figure 14**  
*Response Time Distributions in the Nonperceptual Stock Market Task*



*Note.* Each row depicts the behavior of one individual who demonstrates a different SAT strategy: fast-but-inaccurate (top, ID 6), middle-of-the-road (middle, ID 61), and slow-but-accurate (bottom, ID 43). Each column depicts a different difficulty condition: easy ( $\Delta P = 0.4$ , left), easy ( $\Delta P = 0.2$ , center), and hard ( $\Delta P = 0.1$ , right). Colors indicate participant type: blue (human) our SNN model (orange), the DD model (green), and the extra model (red). As in the human experiment, we allowed each model to complete as many trials in the allotted time window as possible and recorded its RT and accuracy on each trial. We then drew histograms, plotted kernel density estimates, and computed the mean accuracy in each condition. Differences between human and model RTs and accuracies are reported in Table 1. Grey backgrounds indicate model predictions in novel contexts: Empirical data from these plots were not used to fit model parameters. SAT = speed-accuracy trade-off; DD = drift diffusion; SNN = spiking neural network; RT = response time. See the online article for the color version of this figure.

To verify the similarity of the simulated and empirical data, we performed an analysis of covariance on the data from Figure 15, with mean accuracy as the dependent variable, mean cues sampled as the covariate, and participant type (human or model) as the factor variable. We found no statistically significant difference between the model and human trends in any of the three conditions ( $p > .01$ ). In fact, the only noticeable differences between simulated and empirical data in Figures 15 and 16 are that (a) the model completes more trials and achieves higher accuracy in the easy difficulty condition, and (b) the human data contains more extreme outliers, whose accuracy departs significantly from the population trend. To further quantify

model performance, we computed a loss metric between simulated data and empirical data for both RT and accuracy. The former was identical to the loss function computed during optimization and quantified the dissimilarity between RT distributions by calculating a kernel density estimate over the RTs using *SciPy's* gaussian-kde function (Virtanen et al., 2020), evaluating the kernel for all possible values of sampled cues, then computing the root-mean-square error between the evaluated kernels. The latter was simply the absolute difference between mean accuracy across all trials. Across all individuals and difficulties, our model had an RT error of  $\Delta RT = 0.488 \pm 0.662$  and an accuracy error of  $\Delta Acc = 4.73 \pm 5.54$ . Overall,

**Table 1**

*Differences Between Human and Model RTs and Accuracies for the Individuals in Figure 14*

Type	ID	$\Delta RT$ (Easy)	$\Delta RT$ (Moderate)	$\Delta RT$ (Hard)
Model	6	<b>0.24</b>	<b>0.14</b>	<b>0.25</b>
Model	61	<b>0.51</b>	<b>0.41</b>	<b>0.14</b>
Model	43	<b>0.29</b>	<b>0.26</b>	<b>0.41</b>
DD	6	0.59	0.22	0.37
DD	61	1.08	0.64	0.65
DD	43	1.03	0.60	0.77
Extrema	6	0.44	0.33	0.56
Extrema	61	0.67	0.60	0.31
Extrema	43	0.94	0.77	0.90

Type	ID	$\Delta Acc$ (Easy)	$\Delta Acc$ (Moderate)	$\Delta Acc$ (Hard)
Model	6	<b>1.96</b>	2.25	5.20
Model	61	3.11	8.03	7.73
Model	43	<b>0.00</b>	<b>0.83</b>	<b>7.61</b>
DD	6	15.87	12.31	<b>1.39</b>
DD	61	<b>0.00</b>	<b>0.02</b>	<b>0.33</b>
DD	43	16.08	28.33	17.78
Extrema	6	3.59	<b>0.83</b>	1.62
Extrema	61	2.97	16.79	14.76
Extrema	43	1.74	11.54	17.31

*Note.* The model that best predicts the human data for each condition is bolded. For these three individuals, our model made the best RT predictions for all three difficulties and did as well as the other models in predicting accuracies across all conditions. DD = drift diffusion; RT = response time.

our model reproduces multiple aspects of the SAT that are evident in the human data, including the accuracy bias reported in Fiedler et al. (2021) and makes accurate predictions about the behavior of individual humans in novel difficulty conditions.

The DD and extrema models, in contrast, did not provide good fits to the empirical data across the population. While these models did an acceptable job capturing the mean number of cues sampled for each individual in the moderate difficulty condition, and in qualitatively predicting the changes in accuracy associated with changes in difficulty, their quantitative fits to the SAT data were lacking. The extrema model predicted a trend between mean accuracy and mean cues sampled, but significantly underestimated human accuracy across all conditions. Similarly, the DD model predicted a weak trend between speed and accuracy, but it consistently gave poor predictions for both RT and accuracy, and included many outliers that departed significantly from the overall trends. An analysis of covariance analysis revealed a statistically significant difference between the human trends and the DD trends in all three difficulty conditions ( $p < .01$ ), and between the human and extrema trends in all three conditions ( $p < .01$ ). Across all individuals and difficulties, the DD model had a RT error of  $\Delta RT = 0.74 \pm 0.6$  and an accuracy error of  $\Delta Acc = 10.57 \pm 10.24$ , while the extrema model had a RT error of  $\Delta RT = 0.74 \pm 0.56$  and an accuracy error of  $\Delta Acc = 9.25 \pm 7.68$ . In the Discussion section, we speculate as to why the DD and extrema models did poorly in this nonperceptual task.

## The Role of Valuation

One novel feature of our model is the computation of absolute and relative value from sensory information. We hypothesized that a

neural integrator (*memory*) tracks accumulated information and that a separate *value* population computes action values from the representations in *memory*. We also assumed that this value computation may differ between individuals and/or contexts. In our model, value is a mixture of the absolute evidence for a given action and the relative surplus of evidence for that action compared to other actions (Equation 4). Does our hypothesized valuation mechanism help explain the empirical data?

To answer this question, we removed complex valuation from our model and repeated Experiments 3 and 4. Specifically, we fixed valuation to be absolute ( $L = 0$ ), fit our model to the empirical data as before, and computed the validation error metrics for RT ( $\Delta RT$ ) and accuracy ( $\Delta A$ ) described in the previous section. Table 2 reports these errors across tasks and experimental conditions, comparing the reduced models (fixed  $L$ ) with the full model (fitted  $L$ ). The model that produced results closer to the empirical data (smaller error, averaged across individuals) is highlighted for each condition. The full model produced the best fits to the RT data in four of five cases and predicted empirical accuracy better in four of five cases. These results support the idea that valuation has both a relative and absolute component, which may manifest differently across individuals and contexts.

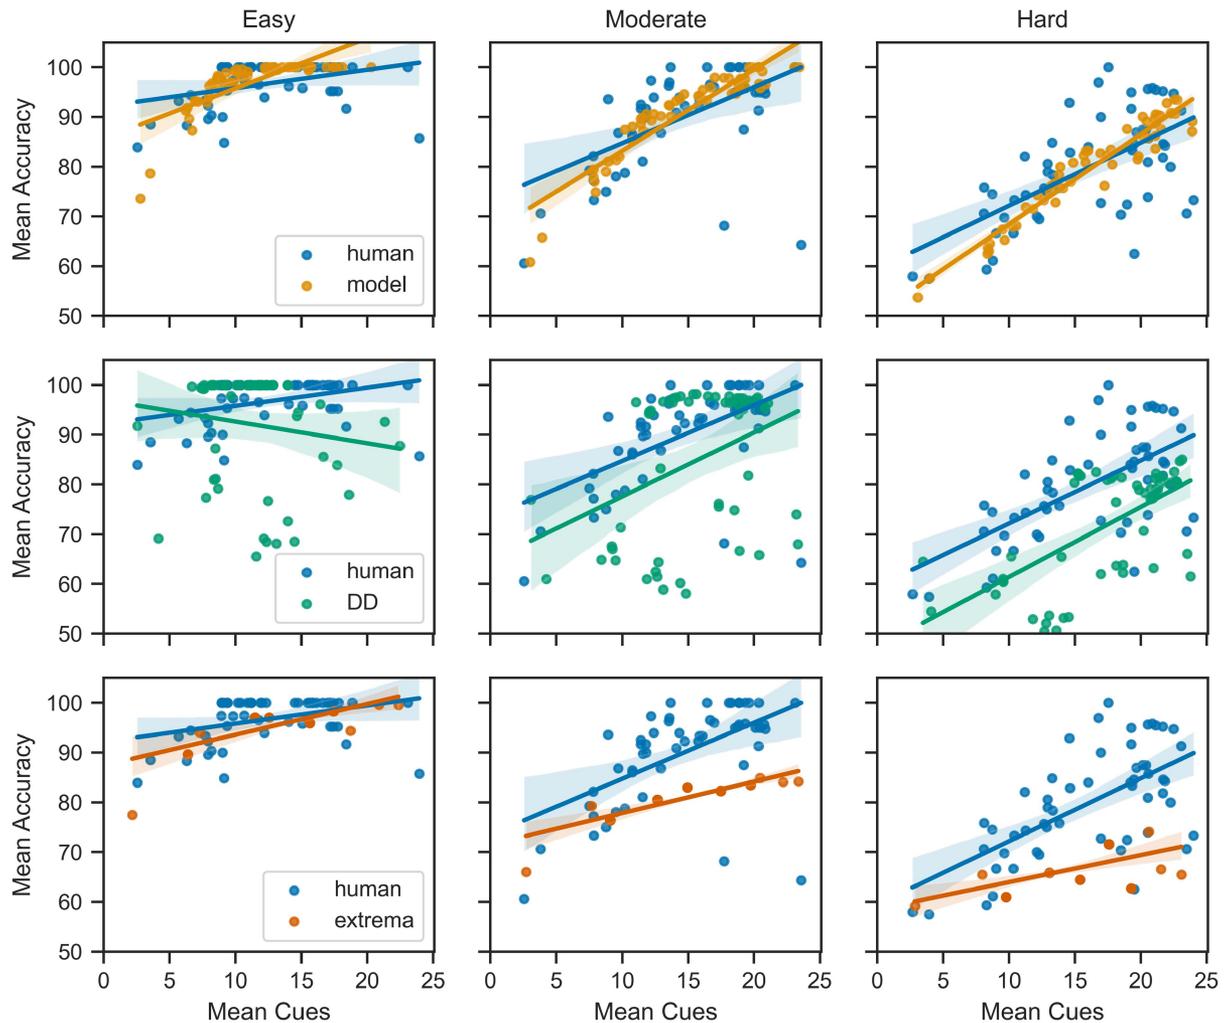
## Discussion

In this section, we begin by reviewing the central goals, methods, and results for our model, including a detailed summary of model predictions and empirical validation. We then examine the biological and cognitive plausibility of the model: We compare our neural network to the architecture and neural responses of the brain and reiterate how our functional connections and parameters realize (and extend) the DD model. We also point out the novel cognitive capabilities of our model, including the capacity to simulate both perceptual and nonperceptual tasks, capture multiple experimental contexts such as speed-accuracy emphasis and multiple alternative choices, recreate trial-to-trial variance using spike noise, and predict several classes of behavioral data that are not used during training. Next, we compare our model to existing mathematical, computational, and neural models of the SAT and argue that our model provides a novel understanding that brings together biological and cognitive explanations. First, we justify our decision to use SNN models by explaining the unique insights they provide in comparison to mathematical models. Second, we compare the results from our model with the results from the DD and extrema detection models and analyze the pros and cons of simple versus complex models. Third, we compare our model with other neural models of DM, ranging from brain-inspired symbolic models, to DD models constrained by neural data, to other SNN models. This allows us to review the novel features and unique predictions made by our model, while also pointing out its shortcomings. Last, we conclude by presenting directions for future work, including empirical experiments that can be used to further validate our model predictions, and theoretical extensions that would increase the realism and explanatory power of our model.

## Goals, Methods, Results, and Predictions

The overarching goal of this article was to integrate neural and cognitive explanations of the SAT using a SNN model of DM (DM).

**Figure 15**  
*Speed–Accuracy Trade-Off in the Stock Market Task*



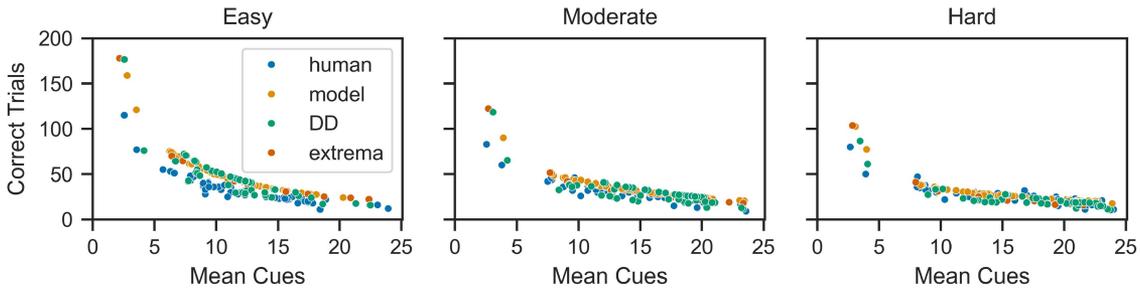
*Note.* We fit one model to each of the 55 participants in the human data set, calculated the mean cues sampled and the mean accuracy for each individual, and plotted each one as a single data point to identify trends across the population. As expected, sampling more cues leads to improved accuracy; this trend is approximately linear, although many individuals reach 100% accuracy in the easy condition. We found no statistically significant difference between the trendlines for the human and SNN model data in any of the three difficulty conditions. In contrast, the extrema model was significantly less accurate, and the DD model was significantly more variable; both mathematical models did a poor job in recreating the human SAT trend. DD = drift diffusion; SNN = spiking neural network; SAT = speed–accuracy trade-off. See the online article for the color version of this figure.

We sought to show that the cognitive mechanisms of evidence accumulation and thresholding can be directly implemented in a biologically plausible neural network, and that this network can be parameterized to capture differences in how individuals approach the SAT. We also sought to show that such a neural model can explain (and make predictions about) classes of phenomena that are outside the scope of purely mathematical models. In doing so, we hoped to demonstrate that neural mechanisms within the brain are capable of supporting the cognitive mechanisms laid out by DD (bridging algorithmic and implementational descriptions) and that biologically plausible neural network models can provide novel insights into our understanding of DM in the brain.

To this end, we designed a SNN model with two key features. First, the model's architecture captures the functional neuroanatomy of DM circuits in the brain. Second, the neural representations and computations realized by the model capture the dynamics of mathematical models such as DD. In other words, our model makes decisions according to a well-validated cognitive theory while respecting an array of biological facts. To put our results in context and to showcase the advantages of neural modeling, we also trained, simulated, and analyzed the behavior of a simple mathematical model, the extrema detection model.

We ran four simulated experiments: Experiments 1–3 investigated the RDM task, while Experiment 4 investigated a nonperceptual

**Figure 16**  
Accuracy Bias in the Stock Market Task



*Note.* As in Figure 15, each data point represents one individual, but we instead plot, as the dependent variable, the total number of trials completed in the allotted 20 min. Participants who decide quickly complete many more trials, but only have slightly less accuracy; as a result, individuals who sample fewer cues complete more trials correctly (and hence earned more total rewards) over the course of the experiment. Both simulated and empirical data were well-characterized by an exponential decrease in completed trials with cues sampled, exhibiting the accuracy bias reported in Fiedler et al. (2021). Interestingly, the DD and extrema models also exhibited this trend, despite having significantly lower accuracy than humans. DD = drift diffusion. See the online article for the color version of this figure.

“stock market” task from Fiedler et al. (2021). In Experiment 1, we reproduced the RT distributions of monkeys (a) across a spectrum of motion coherences and (b) under speed- or accuracy-emphasis instructions. Consistent with the theory that speed-accuracy instructions primarily impact participant’s decision threshold (see the Discussion section below), only the parameter  $T$  varied between the speed and accuracy conditions. We predicted the corresponding accuracy data and recreated the SAT curves of RT versus coherence and accuracy versus coherence (Hanks et al., 2014). We also identified model neurons that exhibited ramping activities similar to those observed in cortical accumulators and demonstrated that (a) their buildup rate is higher in the speed-emphasis condition but (b) their activities converge immediately before a decision is made (Hanks et al., 2014). In Experiment 2, we illustrated the scalability of our model by applying it to both a two- and four-choice RDM task (Churchland et al., 2008) without any structural changes. We reproduced the SAT curves from monkeys for both tasks, once again

predicting accuracy data across motion coherences. We also showed that the model has empirically consistent slower responses on error trials (Churchland et al., 2008), despite the model lacking any mechanism for starting point bias. In Experiment 3, we showcased the utility of neural modeling by simulating the effects of aging via biological perturbation. We trained our model to match the RTs of young individuals, then simulated the process of aging by weakening particular synaptic connections with the network. We predicted the behavioral signatures of aging (Forstmann et al., 2011) by degrading cortical tissue and connectivity, lending support to the theory that impaired working memory explains deficits in the elderly cohort. We also showed that changing  $T$  captures how people respond to speed-accuracy emphases across a population of individuals. In Experiment 4, we studied the SAT in a nonperceptual, temporally extended setting. Once again, we showed that our fitted models reproduces (a) a spectrum of individual strategies in this task and (b) the SAT curve of the entire population (Fiedler et al., 2021). Furthermore, our model (c) recreated the observed accuracy bias and (d) generalized to data outside the training set, making accurate predictions about speed and accuracy for novel difficulty conditions. In the following sections, we expand our analyses of these results in relation to other models, and describe our predictions in detail.

**Table 2**  
Comparison of Response Time Errors and Accuracy Errors Between the Reduced Model and the Full Model

Experiment	Condition	Reduced model ( $L = 0$ )		Full model ( $L \in [0, 1]$ )	
		$\Delta RT$	$\Delta A$	$\Delta RT$	$\Delta A$
Random dot motion	Speed emphasis	0.534	5.03	<b>0.497</b>	<b>3.24</b>
Random dot motion	Accuracy emphasis	0.588	12.3	<b>0.503</b>	<b>6.49</b>
Stock market	Easy difficulty	0.732	2.83	<b>0.663</b>	<b>2.44</b>
Stock market	Moderate difficulty	0.364	<b>4.45</b>	<b>0.288</b>	4.65
Stock market	Hard difficulty	<b>0.463</b>	7.60	0.512	<b>7.11</b>

*Note.* We reran Experiments 3 and 4 with a reduced model, in which relative valuation was removed ( $L = 0$ ). We then computed a RT error (RMSE between evaluated RT kernels) and accuracy error (absolute difference in mean accuracy) between the model and the empirical data. The smaller error is bolded for each row. In four of five experiments, the full model had better fits to the RT data; and in four of five experiments, the full model gave better accuracy predictions. RT = response time; RMSE = root-mean-square error.

## Biological Plausibility

We labeled the neural populations in our model according to their functional properties (see Figure 1), but its structure and connectivity can be mapped onto various parts of cortex and BG, as we described in the Background section. Processing of *sensory* information occurs in modality-specific regions, such as V1–V4 and MT for visual processing, or cortical regions like dIPFC, orbitofrontal cortex, and anterior cingulate cortex for abstract processing. Working *memory* buffers are found throughout the brain and are also modality specific: Visual memory may be realized in areas such as LIP, frontal eye fields, and MT, while memory for abstract information may require dIPFC, orbitofrontal cortex, and ACC. Projections from these areas to regions like ventromedial prefrontal cortex and vSTR may compute the *value*

of information in working memory for the current task, given the goals and context of the experiment. Value estimates for candidate actions are sent to areas in motor cortex, including supplementary motor area and premotor cortex, that may help plan and execute *actions*. These motor areas are under the inhibitory control of nuclei within the BG, which receive value estimates from cortex and *gate* action selection by projecting back to cortex. The dynamic interaction of cortical accumulators, motor cortex, and BG is therefore anatomically consistent with the functionality we have described in our model. However, functional networks resembling Figure 1 likely occur in multiple places throughout the brain, with each specialized to a different sensory or action modality. While the exact anatomical mapping between our model and the brain is therefore up for debate, it is appropriate to claim that the structure of our model is broadly anatomically plausible.

At the neural level, our model was constructed using populations of spiking LIF neurons connected via weighted current-based synapses. DVs, including percepts, accumulated evidence, action value, and decision threshold, were represented via the dynamic activity of these populations. Furthermore, cognitive operations, including sensory accumulation, valuation, and inhibitory thresholding, were realized in the synaptic connections between these populations. Most importantly, our free parameters ( $R$ ,  $T$ , and  $L$ ), which capture individual differences in decision criteria between different model instances, were all implemented directly within this biological substrate, either as synaptic weights on particular connections, or as the background activation of the inhibitory population. All the cognitive mechanisms within our model were therefore realized in a biologically plausible manner.

One advantage of using SNNs is that trial-to-trial variability may naturally arise from the dynamic activity of the network. In our preliminary experiments, we showed (Figure 6) that intratrial variance in DVs and intertrial variance in reaction times (RTs) and task performance (accuracy) arises from two components in our model: our *sampling* module, which emulates the visual system's noisy perceptual processes; and noisy spiking representations throughout the network, which lead to dynamically inaccurate estimates of accumulated evidence, value, and threshold. In our four main experiments, we showed that these biological mechanisms naturally produce the behavioral variance in RTs and accuracies observed across several tasks, experimental contexts, and species. Below, we argue that this biologically grounded explanation for randomness (within and across trials) provides a novel explanation of behavioral variance.

Another advantage of simulating SNNs is the ability to directly compare simulated neural activities with neural data from behaving brains. We performed model validation against neural data and found that, while the model was not designed to recreate activity patterns from particular experiments, our model neurons reproduced several empirical regularities. In particular, we observed ramping activities within our *value* population that (a) grew faster on trials with greater motion coherence, (b) converged to a consistent final firing rate before making a decision, and (c) were systematically greater when experimental instructions emphasized speed over accuracy (Hanks et al., 2014). However, we recognize that some neural dynamics, especially those related to resetting and response mechanisms in the brain, are not captured by our model; we discuss these shortcomings further when comparing our model to other neural implementations below.

## Cognitive Plausibility

Our SNN is an evidence accumulation model: It perceives external information about choice alternatives, accumulates this information over time, compares DVs to a decision threshold, and releases action inhibition to make a choice. Here, we briefly review these components and discuss their cognitive plausibility. For RDM, the model samples information from the environment using a *sampling* procedure that periodically retrieves a noisy estimate of the fraction of dots moving in each direction. This process coarsely captures the ways in which the eyes and visual system perceive inputs in RDM tasks. While this system is admittedly simplistic, we argue that it is more plausible than directly using the motion coherence values as inputs, as is common in other neural models of DM (see below). We also note that previous NEF networks have modeled the visual stream for RDM tasks in great biological detail (Hurzook et al., 2013), and that our *sampling* module could be replaced with this network to further enhance realism. Next, the *memory* population in our model performs evidence accumulation. While our use of a recurrently connected neural integrator for this operation is not novel, the scalability and flexibility of this component contribute to the cognitive plausibility of our model. As we discuss below, many existing neural models simulate a separate accumulator for each DV. In contrast, neurons in our *memory* population have graded sensitivities to each DV. This distributed representation allows our model to accommodate tasks that include multiple action alternatives (including continuous action spaces) or shared features across DVs, expanding the space of DM phenomena that our model can potentially explain.

To further expand the cognitive realism of our model, we introduced a novel mechanism: valuation. In simple DM tasks such as RDM, the accumulated evidence for each candidate action offers sufficient information to make a decision. However, in more complicated DM tasks and in naturalistic settings, accumulated information must often be compared to other quantities (such as goals, motivations, context, and alternatives) before a meaningful decision can be reached. We refer to this additional cognitive step as “valuation” of evidence. The notion of valuation is commonplace in neural models of reinforcement learning, where the state of the world is evaluated relative to an agent's goals, allowing the agent to plan ahead and make decisions that maximize its task performance (Duggins et al., 2022; Rasmussen et al., 2017). Given that the brain areas involved in biological reinforcement learning closely overlap with those involved in DM (Gešiarz & Crockett, 2015; Glimcher, 2011; Lee et al., 2012), it is reasonable to suppose that many forms of DM involve some form of valuation. In this article, we incorporated valuation into our model by having the connection between the *memory* and *value* population compute a value function.

We chose to investigate one common valuation criteria: comparison between competing choices. We introduced a parameter  $L$  to determine the extent to which valuation was absolute versus relative: when  $L = 0$ , each DV in our model captures the absolute amount of accumulated evidence for the corresponding choice; and when  $L = 1$ , each DV captures the relative gains in evidence for that choice compared to all others. We believe that this parameterization captures a meaningful dimension of variance that is often obscured in evidence accumulation models. Specifically, some mathematical models of two-choice RDM tasks simulate two DVs (one for each choice alternative) and feed each accumulator the instantaneous amount of

motion in that direction (the “absolute” case); while other models simulate one DV and feed to the accumulator the instantaneous *difference* in motion between the two directions (the “relative” case). In an extensive review, Teodorescu and Usher (2013) noted the theoretical differences between these approaches and developed a detailed taxonomy for models that feature competition at various processing stages (stimulus, input, response, etc.). They showed that these models make different predictions in DM tasks, validating the study of models that feature DV competition. We note that the valuation mechanism in our model (Equation 4) can be expressed as either the average difference between evidence for option  $i$  and each alternative  $j$ , or as the evidence for  $i$  minus the average evidence for all alternatives  $j$ ; this latter formulation closely resembles the “input competition” mechanism in their taxonomy (Equation 2 in Teodorescu & Usher, 2013) and has also appeared in mathematical models such as the “advantage” LBA (van Ravenzwaaij et al., 2020). However, we argue that our model is able to capture both independent and competitive accumulation without any structural changes, promoting cognitive flexibility. Specifically, because  $L$  is a free parameter in our model, we account for the possibility that valuation (or the degree of DV competition) varies between individuals, potentially capturing an important dimension of strategic variability with respect to the SAT. In future work, it would be interesting to examine different forms of competition, either by changing which neural populations feature competition, or by changing the mathematical form of the competitive interaction.

Interestingly, when valuation is neither strictly absolute nor strictly relative ( $0 < L < 1$ ), our model will wait to make a decision until the evidence favors one alternative over the other(s), but its threshold for the required disparity will effectively shrink over time. This is because the absolute value of both alternatives gradually rises over time, bringing both closer to the decision threshold. In other words, our valuation mechanism effectively implements relative DM in the context of a rising urgency signal, which is thought to be an important decision criteria, both in empirical studies (Cisek et al., 2009; Ditterich, 2006) and computational models (Lee & Usher, 2021; Murphy et al., 2016). Indeed, our results showed that mixed valuation improves the behavioral performance of our model in Experiments 3 and 4. Overall, we believe the addition of valuation is a significant theoretical development that can improve our understanding of the cognitive mechanisms of DM, our ability to predict animal behavior, and the flexibility of our models in more sophisticated cognitive tasks.

The modular, functional structure of our model accommodates a wide range of cognitive tasks and experimental manipulations. We demonstrated its flexibility by (a) applying it to both perceptual and nonperceptual tasks; and (b) investigating numerous contextual manipulations, including instructions to favor speed versus accuracy, the number of choice alternatives, the age of simulated individuals, and the task difficulty. Given previous successes in building large-scale brain models using the NEF and Semantic Pointer Architecture (Eliasmith, 2013), further extensions and applications of our model should be straightforward. For instance, we can replace the nonneural *sampling* component with a neural network that is trained to process particular kinds of visual inputs (Hurzook et al., 2013), experiment with more complex valuation functions and thresholding rules using Bayesian inference (Furlong & Eliasmith, 2022; Sharma, 2018), incorporate sophisticated inhibitory

competition using a detailed BG model (Stewart et al., 2010), and expand the complexity of the underlying neural representations to accommodate symbollike cognitive manipulations (Dumont et al., 2023).

We extensively validated the cognitive capacity of our model by comparing its outputs with behavioral data. We reproduced the RT distributions of humans and monkeys in four separate experiments and predicted the resulting accuracies. Beyond these successes, it is insightful to compare our process of model fitting, and the nature of our model’s predictions, to existing mathematical models. Recall that our model has five free parameters that determine an individual’s behavior: the ramp rate  $R$ , the decision threshold  $T$ , and the relative valuation  $L$ ; plus two sensory sampling parameters for the RDM experiments, the sampling period  $t_s$  and variance  $\sigma$ . These parameters remain constant across trials and experimental conditions. The only exception is the parameter  $T$ , which varies to accommodate speed instructions; this manipulation is consistent with previous work using DD models, which have shown that varying  $T$  can reproduce behavioral differences when participants are asked to favor speed versus accuracy (Voss et al., 2004; Wagenmakers, 2009). While DD models use similar parameters, they often specify a mean and variance for these parameters and draw new values from these distributions when simulating each new trial. This variance, in combination with mechanisms for starting point variability  $S$  and NDT  $t_{nd}$ , allow DD models to robustly fit empirical RT distributions, especially in error trials (Evans, 2020; Jones & Dzhafarov, 2014). However, varying such high-level parameters requires varying neural connection weights in most neural implementations, which is not consistent with typically assumed timescales of synaptic plasticity (Bi & Poo, 2001; Pulvermüller et al., 2021). In contrast, we hypothesized that intertrial variability arises from sampling noise and spike noise in the brain, rather than from fluctuation in high-level cognitive parameters. This hypothesis is consistent with evidence that both external noise (from the random motion of dots) and noisy neural communication (from a lack of attention, mental disorders, etc.) introduce internal noise into the brain’s representations of motion direction, which subsequently affects task performance (Chen et al., 2014; Lu & Doshier, 1998). Our experiments confirm that embedded noise is sufficient to recreate the observed behavioral variance, eliminating the need for intertrial parameter variation or fitting distributions of parameters to each individual. These simplifications allow us to interpret our model parameters as characterizing the persistent cognitive strategy (and sampling characteristics) of an individual.

With regards to prediction, we adopt an approach that is common in the field of machine learning, but is relatively rare in DM studies using mathematical models (although see Shoostari et al., 2019; Vabalas et al., 2020). In each of our experiments, we split the available empirical data into a training set and a validation set: The former was used to train the parameters of the model, and the latter was used to assess its performance. This approach can be contrasted with the prototypical DD study, in which all available data are used to fit the model, and the subsequent analyses identify the parameter(s) that explain the most variance. We use term “prediction” to mean any data generated by the model that was not used to fit model parameters. It is worth emphasizing that our validation data is sometimes of the same type as the training data but in a novel context (e.g., predicting RT in the elderly cohort for Experiment 3, or predicting RT for the easy/hard difficulty in

Experiment 4), but it is sometimes of a completely different type (e.g., accuracy data in Experiments 1–4, neural activities in Experiment 1, and slow errors in Experiment 2). Below, we argue that these bottom-up predictions are a novel contribution of our approach, one which complements the top-down predictions that are typically made using DD models. More broadly, the result that our model anticipates the SAT (in multiple tasks, across multiple species, for multiple experimental contexts, and between individuals) when provided only with speed data suggests that our theoretical account of the underlying cognitive processes are sound.

## Comparison to Other Decision-Making Models

### *Drift Diffusion*

Throughout this article, we have pointed to the differences between our model and traditional DD models and shown that our model outperforms DD in our simulated experiments. We summarize those differences and results here, then more broadly discuss the advantages of studying the SAT using SNNs. Fundamentally, our model implements the same core cognitive processes as DD: sampling from the environment, accumulation of evidence, and thresholding for a decision. In contrast to mathematical models, we realize these operations in a biologically plausible SNN. In doing so, we (a) formalize how high-level processes such as accumulation and thresholding are composed of specific cognitive operations that arise from low-level biophysical processes and (b) map these operations onto distinct brain areas. We also introduce valuation as a distinct cognitive processes and show that it helps capture individual differences in the cognitive strategies for the SAT. We leverage the biological aspect of our model to (a) predict simulated neural responses that align with observed responses in the corresponding brain areas, (b) simulate internal noise that generates intertrial variability and slow errors, and (c) predict how age-induced degeneration of synaptic connectivity induces behavioral deficits. We leverage our fitting and validation procedure to (a) predict accuracy data when trained only on speed data and (b) generalize these predictions across multiple tasks and experimental contexts without changing the model structure or refitting parameters to each condition.

Biological modeling with SNNs complements mathematical modeling with DD in multiple ways, and should be of interest to both psychologists and neuroscientists. First, while DD describes the cognitive mechanisms of DM, SNNs explain how these mechanism might be implemented in the brain. As an existence proof, if a phenomenon can be explained mathematically but not generated using a biologically plausible model of the brain, we should be concerned and seek more biologically compatible explanations. This process is critical for developing rigorous theories of how the brain works. For example, in the field of machine learning, neural networks trained using backpropagation are extremely powerful and can reproduce sophisticated human behavior, but many researchers argue that the backpropagation algorithm is not biologically plausible. This concern has driven the development of biological learning rules, such as Hebbian learning and spike-timing dependent plasticity, and greatly advanced our cognitive understanding of learning in the brain. We argue, alongside others (Frank, 2006; Lo et al., 2015; Shen et al., 2023;

Standage et al., 2011), that SNNs drive an analogous improvement for cognitive theories of DM.

Second, SNN models are able to predict some phenomena that are outside the scope of mathematical models and to explain other phenomena at different levels of analysis. In Experiment 3, we sought to showcase this capacity by studying the differences between young and old participants in the RDM task. Empirically, elderly individuals typically have slower RTs and make more errors than young individuals. Forstmann et al. (2011) showed that RT differences could be described by the LBA model, which is closely related to DD. They were able to explain the slower speed of elderly individuals through changes in the model's decision boundary and showed that these changes were correlated with degraded white matter connectivity between cortex and striatum. This result is consistent with the striatal theory, which posits that the brain flexibly adjusts decision thresholds by sending contextual signals from cortex to striatum and controlling the amount of inhibition applied to the decision circuit. The approach taken by Forstmann et al. (2011) showcases the descriptive power of mathematical models: They captured behavioral differences using a parameter with a clear cognitive interpretation and related variance in this parameter with a measured quantity from the behaving brain.

Our SNN model complements this work by implementing the proposed circuit and testing whether the hypothesized biological mechanism induces the expected deficits. We found that simply degrading connections between cortex and striatum did not produce the desired effect. Instead, we found that degrading the recurrent cortical connections responsible for accumulating sensory evidence better explained the observed data: Only this manipulation explained increases in both RT and error rate. The biological nature of this deficit cannot easily be captured by a pure mathematical model, as it cannot be directly described through changes in model parameters (or their distributions). Our manipulation relied on perturbing the (simulated) physical system that implements the dynamics and parameters proposed by DD and thus explains the phenomenon of aging at a different level of analysis. Note that we are not claiming to have disproven the striatal hypothesis: what we have shown is that, given certain assumptions about the functional role of neural populations and connections, the proposed mechanism does not explain the observed deficit, but another mechanism does. It is entirely possible that an alternative account of how striatum contributes to threshold setting would better explain elderly behavior. For example, if BG instead modulates how cortex evaluates sensory data, or conveys an urgency signal that dynamically controls threshold (Frank, 2006; Thura & Cisek, 2017), degrading inputs to striatum may have a different effect on behavior. Our point is that we need SNNs to specify these theories in enough detail to rigorously test them and to directly validate them against both neural and behavioral data. Our model and subsequent analyses exemplify this approach, and our network can be adjusted to accommodate alternate cognitive theories.

Third, SNNs explain SAT phenomena in a mechanistic, bottom-up manner rather than a descriptive, top-down manner. Our explanations are mechanistic in the neuroscientific sense (Machamer et al., 2000), in that they explain how the coordinated activity of neural processes give rise to higher level cognitive phenomenon (Bechtel, 2009). To return to our example from Experiment 3, our model explained the cognitive deficits of aging by (a) identifying the biological changes associated with aging, (b) recreating those changes within the network, (c) simulating the elderly model, and

(d) analyzing its behavior relative to the young model. In contrast, Forstmann et al. (2011) explained the cognitive deficits of aging by (a) fitting mathematical models (the LBA) to the data from the young and elderly cohort, (b) analyzing the differences in fitted parameters, (c) identifying the biological changes associated with aging, and (d) looking for statistical relationships between biological changes and parameter changes. This analysis exemplifies the approach taken by many studies using mathematical models. We believe that both mechanistic and descriptive explanations, as well as the theories that connect them, are essential to understanding DM and the SAT, and that our model is a step in this direction.

In order to assess the performance of our model, we simulated a simple DD model in our four main experiments. To ensure fair comparisons to our SNN, we used the same optimization procedure to fit DD parameters as we employed to fit our SNN parameters. This meant that our DD model was only trained to fit RT data and had to predict accuracy data without prior knowledge. We also enforced similar parametric constraints in specific experiments on both DD and on our model. For example, in all four experiments, we fit a single drift rate parameter  $R_0$  to the data, then set the mean drift rate on each trial to be proportional to the motion coherence  $C$ ,  $\mu_R = R_0 C$  (or  $\Delta P$  in the case of the nonperceptual task). In Experiments 1–3, we only allowed one parameter, the decision threshold  $T$ , to vary between experimental contexts, forcing the model to explain differences in speed- versus accuracy-emphasis (for each individual) through a single mechanism. However, unlike in our model, we did allow both NDT and intertrial variability in DD parameters such as  $R$  and  $S$ , enhancing its capacity to fully capture variance in RT and recreate slow errors.

Given these constraints, the DD model performed fairly well, recreating many of the behavioral trends we observed in the empirical data. In Experiments 1 and 2, the DD model recreated the SAT curves, fitting the RT distributions for all motion coherences and predicting the corresponding error rates with reasonable accuracy. However, we did notice a slight propensity for the DD model to generate more outlier trials than did our model, which caused over- or underestimation of the mean RTs in many cases. We did not apply the DD model to the four-choice decision task in Experiment 2, as doing so would have required significant revisions of the two-choice model. While the previous studies have shown that multiple alternative choice models of DD can explain multiattribute DM (Krajcich & Rangel, 2011; Roxin, 2019), we felt that this investigation was beyond the scope of the present study. In Experiment 3, we again found that the DD model fit the RT distributions of 18 individuals in the RDM task and that the DD model even predicted the corresponding accuracy deficits in the elderly population. Still, the quantitative fits of our model were superior, and (as discussed above) our SNN offers a more mechanistic description of how aging induces the observed deficits.

However, in Experiment 4, we found that the DD model struggled to fit both the RT and accuracy data of the 55 participants performing the stock market task. While the DD model was able to fit the RT of selected individuals in the moderate difficulty condition, we observed that most fitted models included numerous trials where the model did not choose until the cutoff time, after 24 cues had been presented. This was true even for individuals that adopted fast or middle-of-the-road strategies, as shown in Figure 14. These behaviors not only made for poor RT fits, but were also associated

with significantly lower (and more variable) accuracy. This was surprising, given that we applied the same choice criteria to the DD model upon reaching 24 cues as we did in our SNN model: We inspected the DV in the final simulation step and had the model choose whichever option was favored at that time. Even for the training data, the accuracy predictions of fitted DD models were often so poor that, in many parameter optimization passes with different constraints or parameter ranges, we did not observe any statistically significant trend between speed and accuracy. Furthermore, even when DD models achieved reasonable fits to the training data, they often gave extreme predictions for the easy and difficult conditions: For instance, most DD models predicted that RTs would become much shorter and that accuracies would become much higher, in the easy condition. In contrast, the human data showed more modest changes, which were correctly predicted by our model. We suspect that these failures were related to the challenges of applying DD to a nonperceptual task, and in extrapolating DD performance to data outside the training set. However, given that DD models have successfully explained behavior in certain nonperceptual tasks (Dutilh & Rieskamp, 2016), further analysis is needed to establish the exact causes of these failures.

From these results, we conclude that our SNN model is at least as good at explaining behavioral DM data as a simple DD model in the experiments we investigated. However, we acknowledge that more sophisticated versions of the DD model may provide better fits to the data, and that more advanced tools for fitting DD parameters to the data may improve DD performance. For instance, recent tools for fitting DD parameters like the HDDM python toolbox (Wiecki et al., 2013) have significantly boosted the ability of DD models to explain behavioral (and even neural) data. Furthermore, we suspect that the constraints we placed on DD parameter fitting (shared parameters, no access to accuracy data, etc.) significantly impaired the performance of the DD model. While we felt these constraints were important for a fair comparison to our model and indeed provide an interesting set of results on their own, we suspect that an unconstrained DD model would perform better. Thus, we do not claim that our model outperforms DD in the general case; we leave a more rigorous comparison between our model and state-of-the-art DD models to future work.

### *Extrema Detection*

In our four experiments, we also simulate a nonintegration model to provide a point of comparison for our SNN model. We adopt an extrema detection model in this article because (a) it is based on an intuitive cognitive heuristic for DM, (b) it behaves similarly to DD in many perceptual tasks (Stine et al., 2020), (c) it is mechanistically and parametrically simple, and (d) it can be applied to nonperceptual tasks with modest adjustments. We found that the extrema model performed well in some experiments, and poorly in others; its successes and failures offer insights into the explanatory power of simple versus complex models. In Experiments 1 and 2, which investigated the RDM task, the extrema model fit the RT data almost as well as our model: It captured the distribution of RTs at each coherence value and reproduced the shape of the speed-accuracy curves. Although its accuracy predictions were slightly worse than our model, it still did surprisingly well in predicting accuracy data, given that it was only trained on speed data. Furthermore, the extrema

model could be flexibly applied to different experimental contexts: Like our model, it explained changes in how individuals respond to speed- or accuracy-emphasis using only its decision threshold  $T$  (both in Experiment 1 and 3); and it performed the four-choice RDM task without requiring any extra components or decision rules.

In other contexts, the extrema model fell short. In Experiment 1, the extrema model could not make any predictions about neural activities, whereas our model correctly predicted the relationships between firing rate building, firing threshold, motion coherence, and speed-emphasis. This comes as no surprise, as there are no components in the extrema model that resemble neurons. In Experiment 2, our model correctly predicted that error trials at high coherence would have longer RTs, but the extrema model did not. It is possible that, with more model mechanisms and parameters, the extrema model might predict slow errors or perseveration; this seems to be the case for DD models (Ratcliff et al., 2016; Urai & Donner, 2022). In Experiment 3, the extrema model captured the RT distributions of individuals in the young cohort, but not as well as our model ( $\Delta RT = 0.500$  for our model and  $\Delta RT = 0.876$  for the extrema model). This was somewhat surprising, given that the extrema model includes parameters for NDT ( $\mu_{nd}$  and  $\sigma_{nd}$ , whose sole function is to increase fits to RT data), which our model forgoes. Moreover, its accuracy predictions were quite far off ( $\Delta A = 4.86$  for our model and  $\Delta A = 19.9$  for the extrema model).

One significant failure of the extrema model is its inability to explain or predict age-related phenomena. The simplicity of the extrema model makes it difficult to plausibly simulate the bottom-up effects of biological degradation, as we did with our model: Only by changing model parameters can the extrema model capture age-related deficits. When we allowed all these parameters to vary between the young and old cohorts, we did recreate the RTs of most elderly individuals. However, when we analyzed the corresponding accuracy of elderly extrema models, we found that they were *more* accurate than their younger counterparts. This directly contradicts the empirical data, which show that elderly individuals are slower but also *less* accurate. Why is this the case? In the extrema model, longer responses imply more information being sampled from the environment. Assuming that the sampled evidence reflects the ground truth (even with large variance), this implies that longer RTs should, on average, lead to more accurate decisions; this is the observed trend. In contrast, the (simulated) physical implementation of our system allows accuracy and/or RT to be partially decoupled from sampling, given the presence of noise in particular parts of the system. We showed that degrading visual working memory had exactly this effect; by removing a fraction of connections in the neural integrator responsible for accumulating evidence, we perturbed the represented DVs in such a way that both RT and error increased. We doubt that this kind of perturbation is possible in simple mathematical models, but this could be a valuable direction for future research. Regardless, our ability to (a) directly implement the biological degradation associated with aging, (b) to correctly predict the resulting deficits, and (c) to advocate for one cognitive theory of aging over another showcases the utility of SNNs relative to the extrema model.

The extrema model (plus memory) also performed poorly in the nonperceptual task: It had approximately twice the RT and accuracy error as our model. While it still reproduced the accuracy bias present in the human data, it had significantly lower accuracy, leading to a poor representation of the SAT. These failures can be

largely explained by the removal of NDT from the model for the nonperceptual task. This simplification was motivated by the timescales present in the stock market task: New cues were presented every 0.5 s, leading to trial times of 1–12 s. Given that NDT is intended to capture the cognitive processes associated with early visual processing, and given that RT in this task was quantified by the number of cues sampled, simulating NDT did not make sense for this task. Interestingly, the removal of this model component, alongside the removal of visual sampling for this nonperceptual task, seems to significantly reduce the power of the extrema model. Because the only remaining free parameter was the decision threshold  $T$ , and because extrema detection in this task occurs over clusters of inputs with integer length, the space of possible behaviors from the extrema model is quite limited. This was reflected in the limited number of unique behaviors observed in Figure 15 and contributed to poor RT fits across the human population, whose strategies were quite diverse. Once again, it is possible that an extension of the extrema model may resolve some of these issues, but doing so would introduce more cognitive complexity into the model, defeating the purpose of simulating this simple baseline.

In summary, comparing our SNN to the extrema model highlighted the advantages of complex models for understanding the SAT. While the heuristic of extrema detection may explain the RT and accuracy of individuals in simple perceptual tasks, they fail to capture features such as slow errors and accuracy deficits in elderly individuals. Moreover, this simple model relies heavily on NDT to explain the distribution of RTs and struggles to explain data from cognitive tasks that involve extended, strategic decisions. Indeed, in preliminary simulations, we found that removing NDT from the model in Experiments 1–3 drastically degraded its performance. While mathematical models based on heuristics for sampling, memory, and decision have a rich history in the DM literature (Pietsch & Vickers, 1997; Vickers et al., 1971), our results support the conclusion that more complex models are needed to explain phenomena other than RT and accuracy, such as neural activities, deficits induced by biological changes, and behavior in nonperceptual tasks.

### Neural Models

Others studies have implemented mathematical models of DM in existing cognitive frameworks, explored the relationship between DD models and neural data, and developed SNN models that resemble DD. For example, several articles have related the parameters in DD models to parameters in the Adaptive Character of Thought–Rational (ACT-R) DM system (Fisher et al., 2015; Grennan & Stocco, n.d.). While implementing these models in ACT-R helps formalize the underlying cognitive mechanisms and promotes generalization across task contexts, ACT-R is not a neural framework: Unlike our model, these models cannot be directly compared to neural data, and any comparison to the biological brain will require significant abstraction.

In contrast, numerous studies have related DD parameters to large-scale activation in regions of the brain, then argued that these brain areas might realize cognitive processes such as accumulation or thresholding (Forstmann et al., 2016; Gupta et al., 2021; Purcell & Palmeri, 2017). Other models have used neural data to further parameterize the DD model, improving its ability to explain

behavioral differences between individuals (Turner et al., 2015). While these approaches relate to neural data more closely than ACT-R models, they still rely on statistical descriptions of neural activities, such as blood oxygenation level dependent or electroencephalography responses within brain regions. Unfortunately, DD models are limited in their ability to distinguish sensory, integrative, and comparative processes, making it hard to pin down which brain areas are responsible for each of these cognitive operations. Furthermore, it is difficult to perform certain classes of simulated experiments with these models: For instance, once a model has been fit to an individual's data, it would be problematic to change some aspect of the model (to simulate, for instance, a new experimental context) and predict the neural or behavioral consequences for that individual.

Our model is best compared with other SNNs that implement evidence accumulation and thresholding to perform the RDM task. Lo et al. (2015) developed a model that implements accumulation of evidence via recurrent connections in a population of spiking LIF neurons. Like our model, this model is constructed using a framework that optimizes synaptic weights to achieve evidence accumulation. The authors applied this model to a two-choice RDM task and reproduced many of the results in Experiments 1 and 2. For instance, they (a) observed ramping neural activities with slopes proportional to motion coherence, (b) recreated the RT distributions associated with speed- versus accuracy-emphasis instructions, and (c) captured the SAT curves (RT and accuracy vs. coherence) for individuals performing the task with speed- or accuracy-emphasis. Similarly, Shen et al. (2023) developed a model that implements evidence accumulation via recurrent connections and thresholding via disinhibition. This model uses a framework that optimizes synaptic weights via divisive normalization, is applied to a two- and four-choice RDM task, and reproduces many of the neural and behavioral results from Experiments 1 and 2. For instance, they (a) observed that firing rate buildup depends on motion coherence and number of choice alternatives, (b) showed that error trials have longer RTs at high motion coherences, and (c) captured the SAT curves. Given that these two models, which we refer to as the balanced synaptic input model (BSI) model and the disinhibition and divisive normalization (DDN) model, have a similar structure and produce similar results as our model, it is worth comparing and contrasting them in greater detail.

The structure of the neural networks in BSI, DDN, and our model share several features, but our model differs in a few important respects. All three models feature populations of excitatory and inhibitory neurons, which are recurrently connected such that they integrate an input signal over time. In BSI and DDN, each DV has a dedicated accumulator, which receives an input conveying the magnitude of motion in the corresponding direction. These accumulators connect to a pool of inhibitory neurons, which in turn connect back to the accumulators, realizing competition that facilitates a decision. Both models also receive top-down control signals that flexibly adjust the models' SAT. Finally, both models monitor the activity of each accumulator and define a decision as one accumulator crossing a predefined firing rate threshold. In contrast, our model features only a single accumulator, which includes neurons that are partially sensitive to each DV. This population receives inputs about both motion directions and includes both excitatory and inhibitory neurons, which are recurrently connected in an all-to-all manner. This *memory* population does not implement mutual

competition, but rather connects feedforward to a *value* population: this connection can realize competition in the decision space ( $L > 0$ ), such that the DV for one action decreases when another increases. The *gate* population inhibits the downstream *action* population and receives a top-down signal that influences the SAT by controlling the magnitude of this inhibition. Rather than defining a firing rate threshold for the accumulator, our model makes a decision when the downstream *action* population becomes active, which happens when the DV in *value* overcomes the inhibition from *gate* (i.e.,  $DV_i > T$ ).

These structural differences imply different biological and cognitive assumptions. First, our model predicts that neurons in the brain should be sensitive to each direction of motion (and each alternative action) to some degree, even if some neurons are primarily tuned to a single direction. Most experiments that measure neural activities in LIP identify highly selective neurons and discard the rest; we predict that a different experimental design would reveal a spectrum of tuning curves in neural accumulators that are better described by our model than by BSI or DDN. This distinction may appear trivial for two-choice RDM tasks, but we believe that simulating one accumulator per action alternative will not scale appropriately to more naturalistic decision tasks. Our model realizes a more plausible population code: neurons that remember information or evaluate actions are sensitive to a variety of actions according to a multidimensional tuning curve, but will respond more strongly to some actions than others, according to Equation 1. Second, thresholding and inhibition play a different role in our model: in BSI and DDN, inhibition facilitates competition within the accumulator, and thresholding is applied directly to this population; while in our model, inhibition implements a threshold via a downstream *gate*. In general, our model has a more modular, functional design that captures the interaction of many brain areas: each population and connection is responsible for one cognitive operation, leading to a larger network with sparser anatomical connectivity; whereas in the BSI and DDN models, accumulation, competition, and thresholding all occur within each cortical neural integrator. Our model thus implements a different hypothesis about the content and target of inhibitory control signals; our model is closer to the striatal theory, while BSI and DDN are closer to the cortical theory (Bogacz et al., 2010). Third, an action is triggered when one accumulator crosses a predefined firing rate threshold in BSI and DDN, whereas an action is triggered in our model when the downstream *action* population is disinhibited. The former method of decoding actions relies on the problematic accumulator separation defined above, whereas our method leverages population (de)coding in the respective populations. Interestingly, in Experiment 1, we found that, despite differences in the decision threshold parameter  $T$ , neurons in the *value* population converged to a similar firing rate immediately before decision. The conclusion is that, in our model, the state space decision threshold  $T$  is not equivalent to the firing rate threshold: It is possible to adjust  $T$  in response to experimental conditions while still respecting the biological observation that firing rate thresholds remain constant (Churchland et al., 2008; Hanks et al., 2014).

Speaking more generally, our model offers several advantages over the BSI and DDN models. While all three models have similar degrees of biological realism, our ability to decode simulated activities and analyze the latent state space greatly enhances the scalability and explainability of our model. Many of these differences are derived from the theoretical frameworks used to optimize synaptic weights within the network. BSI and DDN place

a greater emphasis on using inhibition to carefully balance neural dynamics, while the NEF places greater emphasis on using mixed-weight connections to realize arbitrary cognitive dynamics within the neural network. Our focus on describing dynamics in the state space of DVs and cognitive operations is preferable in several respects. First, we can quantitatively relate the dynamic neural activities in our model to DVs: That is, we can decode a state space estimate of the information in each population from spike data, facilitating cognitive analyses and permitting direct comparison to modern methods of decoding DVs in the RDM task (Steinemann et al., 2022). Second, we can build networks that implement functions specified in the DV space: For example, we can build a neural accumulator that integrates the input signal over time: That is, decoding the spike data from *memory* at any given time will return the mathematical sum of the input motions, modulo any drift introduced by spike noise. Similarly, we can define mathematical operations for valuation or thresholding, with the expectation that our network will realize those dynamics during simulation: This allows us to compute the real-time difference between accumulated motion estimates during relative valuation, and to directly compare the represented DVs to a decision threshold. Third, we can leverage the above two properties to specify model parameters in the DV space, ensuring that these parameters have a clear cognitive interpretation. For example, when we define a decision threshold  $T$  and a ramp rate  $R$ , we guarantee that the model will make a decision when the sum of inputs, multiplied by  $R$ , exceeds  $T$ . This property ensures a close alignment between the functional dynamics of our model and pure mathematical models like DD, and lets us map our model parameters directly onto DD parameters (as appropriate). Fourth, we can easily expand our model to simulate cognitive control from other brain areas or investigate different rules for valuation and thresholding. This property arises from knowing what state space variables are represented by each population’s spiking activities and from the ability to specify changes to these state space representations using additional connections. For example, it is possible to extend our model to compute choice confidence over the history of inputs according to Bayesian inference (Dumont et al., 2023; Furlong et al., 2022; Sharma, 2018) or to implement an urgency signal by adding a dynamically valued signal to the current value estimate (as hypothesized by the cortical theory and recent theoretical and empirical work, Bogacz et al., 2010; Murphy et al., 2016). As we have previously mentioned, this property also allows us to substitute existing model components with more biologically or cognitively plausible components, such as a dedicated network for visual sampling (Hurzook et al., 2013) or inhibitory competition for action selection (Stewart et al., 2010).

While we are not as familiar with the theoretical guarantees of the BSI and Divisive Normalization Frameworks, our impression is that these tools implement biologically plausible patterns of excitatory and inhibitory connectivity within a population, but have not been shown to be able to implement large-scale networks that realize complex cognitive dynamics. For instance, Shen et al. (2023) recreated the dip in neural activities that follow stimulus presentation but precedes ramping activity, as well as the nonlinear ramping of neural activities during deliberation. Our model does not currently recreate these features, which likely arise from the specific patterns of recurrent connectivity specified by the DDN model (but see Stöckel et al., 2021, for an example of how to specify such connectivity using the NEF). However, we believe that our model

offers greater insights into DM and the SAT at the cognitive level of analysis. The BSI and DDN models do not dynamically decode the DVs represented by spiking activities and instead rely on mean activities to approximate these quantities. Similarly, these models do not precisely define  $R$  or  $T$  during model training and must observe the behavior of the networks in order to characterize the ramp rate and threshold. These shortcomings make it difficult to directly compare BSI and DDN to mathematical models. More importantly, we suspect that the optimization frameworks underlying BSI and DDN scale poorly to tasks involving many alternative actions, multiple interacting networks, or symbolic DM in multi-dimensional problem spaces. In contrast, the NEF and Semantic Pointer Architecture have been used to build large-scale brain models capable of solving complex cognitive tasks, such as the Tower of Hanoi and Raven’s Progressive Matrices (Eliasmith, 2013), as well as control robots online in naturalistic tasks, such as exploration and reaching (DeWolf et al., 2020). Indeed, our methods have been used to build a large-scale model of the brain that contains 6 million neurons and performs 12 different cognitive tasks (Choo, 2018). Overall, we feel that our model provides a better theoretical understanding of the neural mechanisms underlying the SAT because it (a) better relates the proposed neural dynamics to high-level DVs and cognitive operations and (b) is more extensible and scalable to complex DM tasks.

## Future Work

In future work, we plan to extend our model to study other cognitive tasks and to explain other DM phenomena. In the DD model, starting point bias  $S$  is often included to capture perseveration, which is the bias toward choosing the same action in the subsequent trial. We hypothesize that perseveration arises from two processes: The brain’s attempt to model the underlying statistics of the environment (i.e., estimating the prior probability that one action is favorable); and imperfect resetting of the DM system between trials (i.e., not zeroing the DVs before the next trial). While modeling the former process requires extensive revisions to the model (see below), the latter can easily be investigated by simulating trials in a continuous fashion (in a block, rather than one-at-a-time). To do this, we would simply add a neural mechanism for resetting the DV representation in *memory* following a decision, such as an inhibitory connection between *action* and *memory*, or simulating the biologically detailed resetting mechanism described in Stine et al. (2023). We hypothesize that such a mechanism would cause the model to make perseveration errors *and* recreate the postdecision reduction in neural activities observed in RDM tasks (Churchland et al., 2008; Hanks et al., 2014). We would also like to add more modulatory control systems to the network, allowing executive cortical populations to modify  $R$ ,  $T$ , or  $L$  in real-time; we expect these additions would allow the model to change its behavior over the course of an experimental block, responding more quickly or more accurately to meet task demands. This would permit the study of DM tasks where participants adjust their strategies based on feedback about their performance (Kira et al., 2024). We could even incorporate reinforcement learning into the model (Rasmussen et al., 2017), allowing it to learn behaviors that maximize its rewards through trial-and-error.

In the longer term, we would like to expand the cognitive complexity of the model in order to simulate more sophisticated

decision tasks and recreate more sophisticated choice behavior. For instance, we plan to explicitly model urgency as a cognitive mechanism. Numerous DD models assume that, as the pressure to make a decision steadily increases, the decision criteria applied by the model also shifts: This tendency has been ascribed to various mechanisms (Bogacz et al., 2010), including a positive bias applied to neural accumulators (Hanks et al., 2011), a shrinking decision threshold (Drugowitsch et al., 2012; Kira et al., 2024), and an increase in drift rates (Murphy et al., 2016). The existence of a ramping urgency signal is consistent with behavioral data (Ditterich, 2006) and neural data (Cisek et al., 2009). Furthermore, simulating a dynamic decision threshold  $T(t)$  can dramatically improve the fits of mathematical models to behavioral data in tasks featuring the SAT (Drugowitsch et al., 2012; Kira et al., 2024). Previous work using the NEF has already investigated urgency signals in the context of working memory (Singh & Eliasmith, 2006). In future work, we will follow a protocol similar to Experiment 3: We will implement different hypotheses about the neural mechanisms of urgency modulation, simulate our network to make neural and behavioral predictions, and compare these predictions with empirical results.

Finally, we plan to add Bayesian inference to our model. Rather than explicitly tracking DVs, the network would maintain beliefs about the external world using probability distributions, then use those beliefs to calculate the expected value of (and confidence in) each candidate action. Previous work has shown that computational models of probabilistic inference and belief updating may explain many aspects of DM under uncertainty, both in perceptual tasks (Huang & Rao, 2013; Rao, 2010) and social tasks (Khalvati, Mirbagheri, et al., 2019; Khalvati, Park, et al., 2019). These results align with empirical studies showing that individuals estimate their confidence in decision alternatives alongside their expected value, and weight incoming evidence accordingly (Hanks et al., 2011; Kiani et al., 2014). Recent work using the NEF has developed neural networks that implement several of the complex cognitive operations required for Bayesian inference: For instance, the Legendre delay network can record an entire history of sensory inputs (Furlong et al., 2022), and spatial semantic pointers can be used to represent probability distributions (Furlong & Eliasmith, 2022). We are interested in building networks capable of Bayesian reasoning using biologically plausible models (Sharma, 2018) and applying them to better understand the neural and cognitive mechanisms of DM.

## Conclusion

In this article, we presented a spiking neuron model of DM that implements and extends the DD model while respecting the functional neuroanatomy of the brain. Our goal was to integrate neurological and computational accounts of the SAT by showing that a biologically plausible neural network can realize evidence accumulation, reproduce empirical results, and make novel predictions. Our model extends previous work in several respects: It attributes intertrial variance in reaction time and accuracy to sensory sampling and to noisy neural representations, rather than to variance in model parameters for accumulation and starting point; and it introduces valuation of accumulated evidence as a core cognitive component. Over the course of four experiments, we applied our model to perceptual and nonperceptual tasks, investigated several contextual manipulations, and validated model performance using

neural and behavioral data. Behaviorally, our model (a) reproduced RT distributions in all experiments; (b) generalized across experimental contexts, including the number of choice alternatives, speed- or accuracy-emphasis, and task difficulty; and (c) predicted accuracy data, slower RTs in error trials, and RTs in novel contexts. Neurally, our model (a) recreated observed patterns of ramping and converging activities and (b) predicted the deficits observed in elderly individuals following targeted synaptic degradation. More broadly, our model explains how individual differences in speed and accuracy arise from synaptic weights within a neural network, is applicable to variety of tasks and contexts, and can be extended to simulate more complex forms of DM. Our work showcases a method for translating mathematical models into functional neural networks and demonstrates that simulating such networks permits analyses and predictions that are outside the scope of the original mathematical models.

## References

- Akiba, T., Sano, S., Yanase, T., Ohta, T., & Koyama, M. (2019). Optuna: A next-generation hyperparameter optimization framework. *Proceedings of the 25th international conference on knowledge discovery and data mining* (pp. 2623–2631). Association for Computing Machinery.
- Bechtel, W. (2009). Looking down, around, and up: Mechanistic explanation in psychology. *Philosophical Psychology*, 22(5), 543–564. <https://doi.org/10.1080/09515080903238948>
- Bekolay, T., Bergstra, J., Hunsberger, E., DeWolf, T., Stewart, T. C., Rasmussen, D., Choo, X., Voelker, A. R., & Eliasmith, C. (2014). Nengo: A python tool for building large-scale functional brain models. *Frontiers in Neuroinformatics*, 7, Article 48. <https://doi.org/10.3389/fninf.2013.00048>
- Bi, G.-Q., & Poo, M.-M. (2001). Synaptic modification by correlated activity: Hebb's postulate revisited. *Annual Review of Neuroscience*, 24(1), 139–166. <https://doi.org/10.1146/annurev.neuro.24.1.139>
- Boerlin, M., Machens, C. K., & Denève, S. (2013). Predictive coding of dynamical variables in balanced spiking networks. *PLOS Computational Biology*, 9(11), Article e1003258. <https://doi.org/10.1371/journal.pcbi.1003258>
- Bogacz, R., Wagenmakers, E.-J., Forstmann, B. U., & Nieuwenhuis, S. (2010). The neural basis of the speed-accuracy tradeoff. *Trends in Neurosciences*, 33(1), 10–16. <https://doi.org/10.1016/j.tins.2009.09.002>
- Brown, S., & Heathcote, A. (2005). A ballistic model of choice response time. *Psychological Review*, 112(1), 117–128. <https://doi.org/10.1037/0033-295X.112.1.117>
- Brown, S., & 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>
- Campanella, F., Skrap, M., & Vallesi, A. (2016). Speed-accuracy strategy regulations in prefrontal tumor patients. *Neuropsychologia*, 82, 1–10. <https://doi.org/10.1016/j.neuropsychologia.2016.01.008>
- Canellas, M. C., & Feigh, K. M. (2016). Toward simple representative mathematical models of naturalistic decision making through fast-and-frugal heuristics. *Journal of Cognitive Engineering and Decision Making*, 10(3), 255–267. <https://doi.org/10.1177/1555343416656103>
- Chen, Y., Norton, D., & McBain, R. (2014). Effects of domain-specific noise on visual motion processing in schizophrenia. *PLOS ONE*, 9(6), Article e99031. <https://doi.org/10.1371/journal.pone.0099031>
- Chittka, L., Skorupski, P., & Raine, N. E. (2009). Speed-accuracy tradeoffs in animal decision making. *Trends in Ecology & Evolution*, 24(7), 400–407. <https://doi.org/10.1016/j.tree.2009.02.010>
- Choo, F.-X. (2018). *Spaun 2.0: Extending the world's largest functional brain model* [PhD thesis]. University of Waterloo.

- Christophel, T. B., Klink, P. C., Spitzer, B., Roelfsema, P. R., & Haynes, J.-D. (2017). The distributed nature of working memory. *Trends in Cognitive Sciences*, 21(2), 111–124. <https://doi.org/10.1016/j.tics.2016.12.007>
- Churchland, A. K., Kiani, R., & Shadlen, M. N. (2008). Decision-making with multiple alternatives. *Nature Neuroscience*, 11(6), 693–702. <https://doi.org/10.1038/nn.2123>
- Cisek, P., Puskas, G. A., & El-Murr, S. (2009). Decisions in changing conditions: The urgency-gating model. *Journal of Neuroscience*, 29(37), 11560–11571. <https://doi.org/10.1523/JNEUROSCI.1844-09.2009>
- DeWolf, T., Jaworski, P., & Eliasmith, C. (2020). Nengo and low-power ai hardware for robust, embedded neurorobotics. *Frontiers in Neurobotics*, 14, Article 568359. <https://doi.org/10.3389/fnbot.2020.568359>
- Ditterich, J. (2006). Evidence for time-variant decision making. *European Journal of Neuroscience*, 24(12), 3628–3641. <https://doi.org/10.1111/j.1460-9568.2006.05221.x>
- Domenech, P., Redouté, J., Koehlin, E., & Dreher, J.-C. (2018). The neuro-computational architecture of value-based selection in the human brain. *Cerebral Cortex*, 28(2), 585–601. <https://doi.org/10.1093/cercor/bhw396>
- Donkin, C., Brown, S., & Heathcote, A. (2011). Drawing conclusions from choice response time models: A tutorial using the linear ballistic accumulator. *Journal of Mathematical Psychology*, 55(2), 140–151. <https://doi.org/10.1016/j.jmp.2010.10.001>
- Drugowitsch, J., DeAngelis, G. C., Angelaki, D. E., & Pouget, A. (2015). Tuning the speed–accuracy trade-off to maximize reward rate in multisensory decision-making. *Elife*, 4, Article e06678. <https://doi.org/10.7554/eLife.06678>
- Drugowitsch, J., Moreno-Bote, R., Churchland, A. K., Shadlen, M. N., & Pouget, A. (2012). The cost of accumulating evidence in perceptual decision making. *Journal of Neuroscience*, 32(11), 3612–3628. <https://doi.org/10.1523/JNEUROSCI.4010-11.2012>
- Duggins, P. (2023). *Learning and decision making in social contexts: Neural and computational models* [PhD thesis]. University of Waterloo.
- Duggins, P., & Eliasmith, C. (2022). Constructing functional models from biophysically-detailed neurons. *PLOS Computational Biology*, 18(9), Article e1010461. <https://doi.org/10.1371/journal.pcbi.1010461>
- Duggins, P., Krzemiński, D., Eliasmith, C., & Wichary, S. (2020). A spiking neuron model of inferential decision making: Urgency, uncertainty, and the speed–accuracy tradeoff. *Proceedings of the annual meeting of the cognitive science society*. <https://escholarship.org/uc/item/0sb290h8>
- Duggins, P., Stewart, T. C., & Eliasmith, C. (2022). Reinforcement learning, social value orientation, and decision making: Computational models and empirical validation. *Proceedings of the annual meeting of the cognitive science society*. <https://escholarship.org/uc/item/1b30p4sw>
- Dumont, N. S.-Y., Stöckel, A., Furlong, P. M., Bartlett, M., Eliasmith, C., & Stewart, T. C. (2023). Biologically-based computation: How neural details and dynamics are suited for implementing a variety of algorithms. *Brain Sciences*, 13(2), Article 245. <https://doi.org/10.3390/brainsci13020245>
- Dutilh, G., & Rieskamp, J. (2016). Comparing perceptual and preferential decision making. *Psychonomic Bulletin & Review*, 23, 723–737. <https://doi.org/10.3758/s13423-015-0941-1>
- Eliasmith, C. (2013). *How to build a brain: A neural architecture for biological cognition*. Oxford University Press.
- Eliasmith, C., & Anderson, C. H. (2003). *Neural engineering: Computation, representation, and dynamics in neurobiological systems*. MIT Press.
- Evans, N. J. (2020). *Same model, different conclusions: An identifiability issue in the linear ballistic accumulator model of decision-making* [Preprint]. <https://doi.org/10.31234/osf.io/2xu7f>
- Evans, N. J., Hawkins, G. E., Boehm, U., Wagenmakers, E.-J., & Brown, S. (2017). The computations that support simple decision-making: A comparison between the diffusion and urgency-gating models. *Scientific Reports*, 7(1), Article 16433. <https://doi.org/10.1038/s41598-017-16694-7>
- Faubert, J. (2002). Visual perception and aging. *Canadian Journal of Experimental Psychology/Revue canadienne de psychologie expérimentale*, 56(3), 164–176. <https://doi.org/10.1037/h0087394>
- Fiedler, K., McCaughey, L., Prager, J., Eichberger, J., & Schnell, K. (2021). Speed–accuracy trade-offs in sample-based decisions. *Journal of Experimental Psychology: General*, 150(6), 1203–1224. <https://doi.org/10.1037/xge0000986>
- Fish, S., Toumaian, M., Pappa, E., Davies, T. J., Tanti, R., Saville, C. W., Theleritis, C., Economou, M., Klein, C., & Smyrnis, N. (2018). Modelling reaction time distribution of fast decision tasks in schizophrenia: Evidence for novel candidate endophenotypes. *Psychiatry Research*, 269, 212–220. <https://doi.org/10.1016/j.psychres.2018.08.067>
- Fisher, C. R., Walsh, M. M., Blaha, L. M., & Gunzelmann, G. (2015). Act-r and lba model mimicry reveals similarity across modeling formalisms. *Proceedings of the annual meeting of the cognitive science society*. <https://escholarship.org/uc/item/9qr5t4f5>
- Forstmann, B. U., Dutilh, G., Brown, S., Neumann, J., Von Cramon, D. Y., Ridderinkhof, K. R., & Wagenmakers, E.-J. (2008). Striatum and pre-sma facilitate decision-making under time pressure. *Proceedings of the National Academy of Sciences*, 105(45), 17538–17542. <https://doi.org/10.1073/pnas.0805903105>
- Forstmann, B. U., Ratcliff, R., & Wagenmakers, E.-J. (2016). Sequential sampling models in cognitive neuroscience: Advantages, applications, and extensions. *Annual Review of Psychology*, 67, 641–666. <https://doi.org/10.1146/annurev-psych-122414-033645>
- Forstmann, B. U., Tittgemeyer, M., Wagenmakers, E.-J., Derrfuss, J., Imperati, D., & Brown, S. (2011). The speed–accuracy tradeoff in the elderly brain: A structural model-based approach. *Journal of Neuroscience*, 31(47), 17242–17249. <https://doi.org/10.1523/JNEUROSCI.0309-11.2011>
- Frank, M. J. (2006). Hold your horses: A dynamic computational role for the subthalamic nucleus in decision making. *Neural Networks*, 19(8), 1120–1136. <https://doi.org/10.1016/j.neunet.2006.03.006>
- Furlong, M., & Eliasmith, C. (2022). Fractional binding in vector symbolic architectures as quasi-probability statements. *Proceedings of the annual meeting of the cognitive science society*. <https://escholarship.org/uc/item/33j7c63v>
- Furlong, M., Stöckel, A., Stewart, T. C., & Eliasmith, C. (2022). *Learned legendre predictor: Learning with compressed representations for efficient online multistep*. Centre for Theoretical Neuroscience. <https://compneuro.uwaterloo.ca/files/publications/furlong.2022a.pdf>
- Gallego, J. A., Perich, M. G., Miller, L., & Solla, S. A. (2017). Neural manifolds for the control of movement. *Neuron*, 94(5), 978–984. <https://doi.org/10.1016/j.neuron.2017.05.025>
- Gešiarz, F., & Crockett, M. J. (2015). Goal-directed, habitual and pavlovian prosocial behavior. *Frontiers in Behavioral Neuroscience*, 9, Article 135. <https://doi.org/10.3389/fnbeh.2015.00135>
- Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. *Annual Review of Psychology*, 62, 451–482. <https://doi.org/10.1146/annurev-psych-120709-145346>
- Glimcher, P. W. (2011). Understanding dopamine and reinforcement learning: The dopamine reward prediction error hypothesis. *Proceedings of the National Academy of Sciences*, 108(Suppl. 3), 15647–15654. <https://doi.org/10.1073/pnas.1014269108>
- Gluth, S., Rieskamp, J., & Büchel, C. (2014). Neural evidence for adaptive strategy selection in value-based decision-making. *Cerebral Cortex*, 24(8), 2009–2021. <https://doi.org/10.1093/cercor/bht049>
- Grennan, G., & Stocco, A. (n.d.). *Estimating act-r declarative memory parameters using a drift diffusion model* [Conference session]. 20th International Conference on Cognitive Modeling, Toronto, CA.
- Gupta, A., Bansal, R., Alashwal, H., Kacar, A. S., Balci, F., & Moustafa, A. A. (2021). Neural substrates of the drift-diffusion model in brain disorders. *Frontiers in Computational Neuroscience*, 15, Article 678232. <https://doi.org/10.3389/fncom.2021.678232>
- Guthrie, M., Leblais, A., Garenne, A., & Boraud, T. (2013). Interaction between cognitive and motor cortico-basal ganglia loops during decision making: A computational study. *Journal of Neurophysiology*, 109(12), 3025–3040. <https://doi.org/10.1152/jn.00026.2013>

- Hanks, T., Kiani, R., & Shadlen, M. N. (2014). A neural mechanism of speed-accuracy tradeoff in macaque area LIP. *Elife*, 3, Article e02260. <https://doi.org/10.7554/eLife.02260>
- Hanks, T., Mazurek, M., Kiani, R., Hopp, E., & Shadlen, M. (2011). Elapsed decision time affects the weighting of prior probability in a perceptual decision task. *Journal of Neuroscience*, 31(17), 6339–6352. <https://doi.org/10.1523/JNEUROSCI.5613-10.2011>
- Heekeren, H. R., Marrett, S., & Ungerleider, L. G. (2008). The neural systems that mediate human perceptual decision making. *Nature Reviews Neuroscience*, 9(6), 467–479. <https://doi.org/10.1038/nrn2374>
- Heitz, R. P. (2014). The speed-accuracy tradeoff: History, physiology, methodology, and behavior. *Frontiers in Neuroscience*, 8, Article 150. <https://doi.org/10.3389/fnins.2014.00150>
- Huang, Y., & Rao, R. P. (2013). Reward optimization in the primate brain: A probabilistic model of decision making under uncertainty. *PLOS ONE*, 8(1), Article e53344. <https://doi.org/10.1371/journal.pone.0053344>
- Hurzook, A., Trujillo, O., & Eliasmith, C. (2013). Visual motion processing and perceptual decision making. *Proceedings of the annual meeting of the cognitive science society*. <https://escholarship.org/uc/item/41t5n218>
- Hutchinson, C. V., Arena, A., Allen, H. A., & Ledgeway, T. (2012). Psychophysical correlates of global motion processing in the aging visual system: A critical review. *Neuroscience & Biobehavioral Reviews*, 36(4), 1266–1272. <https://doi.org/10.1016/j.neubiorev.2012.02.009>
- Jones, M., & Dhazharov, E. N. (2014). Unfalsifiability and mutual translatability of major modeling schemes for choice reaction time. *Psychological Review*, 121(1), 1–32. <https://doi.org/10.1037/a0034190>
- Kanai, R., & Rees, G. (2011). The structural basis of inter-individual differences in human behaviour and cognition. *Nature Reviews Neuroscience*, 12(4), 231–242. <https://doi.org/10.1038/nrn3000>
- Karalunas, S. L., Huang-Pollock, C. L., & Nigg, J. T. (2012). Decomposing attention-deficit/hyperactivity disorder (ADHD)-related effects in response speed and variability. *Neuropsychology*, 26(6), 684–694. <https://doi.org/10.1037/a0029936>
- Kaysers, A. S., Buchsbaum, B. R., Erickson, D. T., & D’Esposito, M. (2010). The functional anatomy of a perceptual decision in the human brain. *Journal of Neurophysiology*, 103(3), 1179–1194. <https://doi.org/10.1152/jn.00364.2009>
- Khalvati, K., Mirbagheri, S., Park, S. A., Dreher, J.-C., & Rao, R. P. (2019). A bayesian theory of conformity in collective decision making. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, & R. Garnett (Eds.), *Advances in neural information processing systems* (p. 32). Curran Associates.
- Khalvati, K., Park, S. A., Mirbagheri, S., Philippe, R., Sestito, M., Dreher, J.-C., & Rao, R. P. (2019). Modeling other minds: Bayesian inference explains human choices in group decision-making. *Science Advances*, 5(11), Article eaax8783. <https://doi.org/10.1126/sciadv.aax8783>
- Kiani, R., Corthell, L., & Shadlen, M. N. (2014). Choice certainty is informed by both evidence and decision time. *Neuron*, 84(6), 1329–1342. <https://doi.org/10.1016/j.neuron.2014.12.015>
- Kira, S., Yang, T., & Shadlen, M. N. (2015). A neural implementation of wald’s sequential probability ratio test. *Neuron*, 85(4), 861–873. <https://doi.org/10.1016/j.neuron.2015.01.007>
- Kira, S., Zylberberg, A., & Shadlen, M. N. (2024). *Incorporation of a cost of deliberation time in perceptual decision making*. bioRxiv.
- Kofler, M. J., Rapport, M. D., Sarver, D. E., Raiker, J. S., Orban, S. A., Friedman, L. M., & Kolomeyer, E. G. (2013). Reaction time variability in ADHD: A meta-analytic review of 319 studies. *Clinical Psychology Review*, 33(6), 795–811. <https://doi.org/10.1016/j.cpr.2013.06.001>
- Krajbich, I., & Rangel, A. (2011). Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. *Proceedings of the National Academy of Sciences*, 108(33), 13852–13857. <https://doi.org/10.1073/pnas.1101328108>
- Lee, D., Seo, H., & Jung, M. W. (2012). Neural basis of reinforcement learning and decision making. *Annual Review of Neuroscience*, 35, 287–308. <https://doi.org/10.1146/annurev-neuro-062111-150512>
- Lee, D., & Usher, M. (2021). *Value certainty in drift-diffusion models of preferential choice*. Psychological Review.
- Liu, Z., Liu, S., Li, S., Li, L., Zheng, L., Weng, X., Guo, X., Lu, Y., Men, W., Gao, J., & You, X. (2022). Dissociating value-based neurocomputation from subsequent selection-related activations in human decision-making. *Cerebral Cortex*, 32(19), 4141–4155. <https://doi.org/10.1093/cercor/bhab471>
- Lo, C.-C., Wang, C.-T., & Wang, X.-J. (2015). Speed-accuracy tradeoff by a control signal with balanced excitation and inhibition. *Journal of Neurophysiology*, 114(1), 650–661. <https://doi.org/10.1152/jn.00845.2013>
- Lu, Z.-L., & Doshier, B. A. (1998). External noise distinguishes attention mechanisms. *Vision Research*, 38(9), 1183–1198. [https://doi.org/10.1016/S0042-6989\(97\)00273-3](https://doi.org/10.1016/S0042-6989(97)00273-3)
- Luo, J. (2018). The neural basis of and a common neural circuitry in different types of pro-social behavior. *Frontiers in Psychology*, 9, Article 859. <https://doi.org/10.3389/fpsyg.2018.00859>
- Machamer, P., Darden, L., & Craver, C. F. (2000). Thinking about mechanisms. *Philosophy of Science*, 67(1), 1–25. <https://doi.org/10.1086/392759>
- Moustafa, A. A., Kéri, S., Somlai, Z., Balsdon, T., Frydecka, D., Misiak, B., & White, C. (2015). Drift diffusion model of reward and punishment learning in schizophrenia: Modeling and experimental data. *Behavioural Brain Research*, 291, 147–154. <https://doi.org/10.1016/j.bbr.2015.05.024>
- Mulder, M. J., Boekel, W., Ratcliff, R., & Forstmann, B. U. (2014). Cortico-subthalamic connection predicts individual differences in value-driven choice bias. *Brain Structure and Function*, 219(4), 1239–1249. <https://doi.org/10.1007/s00429-013-0561-3>
- Mulder, M. J., Bos, D., Weusten, J. M., van Belle, J., van Dijk, S. C., Simen, P., van Engeland, H., & Durston, S. (2010). Basic impairments in regulating the speed-accuracy tradeoff predict symptoms of attention-deficit/hyperactivity disorder. *Biological Psychiatry*, 68(12), 1114–1119. <https://doi.org/10.1016/j.biopsych.2010.07.031>
- Murphy, P. R., Boonstra, E., & Nieuwenhuis, S. (2016). Global gain modulation generates time-dependent urgency during perceptual choice in humans. *Nature Communications*, 7(1), Article 13526. <https://doi.org/10.1038/ncomms13526>
- Nissim, N. R., O’Shea, A. M., Bryant, V., Porges, E. C., Cohen, R., & Woods, A. J. (2017). Frontal structural neural correlates of working memory performance in older adults. *Frontiers in Aging Neuroscience*, 8, Article 328. <https://doi.org/10.3389/fnagi.2016.00328>
- Palmer, J., Huk, A. C., & Shadlen, M. N. (2005). The effect of stimulus strength on the speed and accuracy of a perceptual decision. *Journal of Vision*, 5(5), Article 1. <https://doi.org/10.1167/5.5.1>
- Park, S., & Holzman, P. S. (1992). Schizophrenics show spatial working memory deficits. *Archives of General Psychiatry*, 49(12), 975–982. <https://doi.org/10.1001/archpsyc.1992.01820120063009>
- Pärnamets, P., Shuster, A., Reinero, D. A., & Van Bavel, J. J. (2020). A value-based framework for understanding cooperation. *Current Directions in Psychological Science*, 29(3), 227–234. <https://doi.org/10.1177/0963721420906200>
- Pietsch, A., & Vickers, D. (1997). Memory capacity and intelligence: Novel techniques for evaluating rival models of a fundamental information-processing mechanism. *The Journal of General Psychology*, 124(3), 229–339. <https://doi.org/10.1080/00221309709595520>
- Pulvermüller, F., Tomasello, R., Henningsen-Schomers, M. R., & Wennekers, T. (2021). Biological constraints on neural network models of cognitive function. *Nature Reviews Neuroscience*, 22(8), 488–502. <https://doi.org/10.1038/s41583-021-00473-5>
- Purcell, B. A., & Palmeri, T. J. (2017). Relating accumulator model parameters and neural dynamics. *Journal of Mathematical Psychology*, 76, 156–171. <https://doi.org/10.1016/j.jmp.2016.07.001>

- Rae, B., Heathcote, A., Donkin, C., Averell, L., & Brown, S. (2014). The hare and the tortoise: Emphasizing speed can change the evidence used to make decisions. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *40*(5), 1226–1243. <https://doi.org/10.1037/a0036801>
- Rao, R. P. (2010). Decision making under uncertainty: A neural model based on partially observable Markov decision processes. *Frontiers in Computational Neuroscience*, *4*, Article 146. <https://doi.org/10.3389/fncom.2010.00146>
- Rasmussen, D., Voelker, A., & Eliasmith, C. (2017). A neural model of hierarchical reinforcement learning. *PLOS ONE*, *12*(7), Article e0180234. <https://doi.org/10.1371/journal.pone.0180234>
- Ratcliff, R., & Childers, R. (2015). Individual differences and fitting methods for the two-choice diffusion model of decision making. *Decision*, *2*(4), 237–279. <https://doi.org/10.1037/dec0000030>
- 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., Brown, S., & 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>
- Recanatesi, S., Bradde, S., Balasubramanian, V., Steinmetz, N. A., & Shear-Brown, E. (2022). A scale-dependent measure of system dimensionality. *Patterns*, *3*(8), Article 100555. <https://doi.org/10.1016/j.patter.2022.100555>
- Redgrave, P., Prescott, T. J., & Gurney, K. (1999). The basal ganglia: A vertebrate solution to the selection problem? *Neuroscience*, *89*(4), 1009–1023. [https://doi.org/10.1016/S0306-4522\(98\)00319-4](https://doi.org/10.1016/S0306-4522(98)00319-4)
- Reinhardt, R. M., & Nguyen, J. A. (2019). Working memory revived in older adults by synchronizing rhythmic brain circuits. *Nature Neuroscience*, *22*(5), 820–827. <https://doi.org/10.1038/s41593-019-0371-x>
- Roxin, A. (2019). Drift-diffusion models for multiple-alternative forced-choice decision making. *The Journal of Mathematical Neuroscience*, *9*(1), Article 5. <https://doi.org/10.1186/s13408-019-0073-4>
- Rusch, T., Steininger-Kumar, S., Doshi, P., Spezio, M., & Gläscher, J. (2020). Theory of mind and decision science: Towards a typology of tasks and computational models. *Neuropsychologia*, *146*, Article 107488. <https://doi.org/10.1016/j.neuropsychologia.2020.107488>
- Schweitzer, L. R., & Lee, M. (1992). Use of the speed accuracy trade-off to characterize information processing in schizophrenics and normals. *Psychopathology*, *25*(1), 29–40. <https://doi.org/10.1159/000284751>
- Sharma, S. (2018). *Neural plausibility of Bayesian inference* [Masters thesis]. University of Waterloo.
- Shen, B., Louie, K., & Glimcher, P. (2023). Flexible control of representational dynamics in a disinhibition-based model of decision-making. *Elife*, *12*, Article e82426. <https://doi.org/10.7554/eLife.82426>
- Shooshtari, S. V., Sadrabadi, J. E., Azizi, Z., & Ebrahimpour, R. (2019). Confidence representation of perceptual decision by EEG and eye data in a random dot motion task. *Neuroscience*, *406*, 510–527. <https://doi.org/10.1016/j.neuroscience.2019.03.031>
- Sih, A., & Del Giudice, M. (2012). Linking behavioural syndromes and cognition: A behavioural ecology perspective. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *367*(1603), 2762–2772. <https://doi.org/10.1098/rstb.2012.0216>
- Simen, P. (2012). Evidence accumulator or decision threshold-which cortical mechanism are we observing? *Frontiers in Psychology*, *3*, Article 183. <https://doi.org/10.3389/fpsyg.2012.00183>
- Singh, R., & Eliasmith, C. (2006). Higher-dimensional neurons explain the tuning and dynamics of working memory cells. *Journal of Neuroscience*, *26*(14), 3667–3678. <https://doi.org/10.1523/JNEUROSCI.4864-05.2006>
- Smith, P. L., & Vickers, D. (1989). Modeling evidence accumulation with partial loss in expanded judgment. *Journal of Experimental Psychology: Human Perception and Performance*, *15*(4), 797–815. <https://doi.org/10.1037/0096-1523.15.4.797>
- Sreenivasan, K. K., Curtis, C. E., & D'Esposito, M. (2014). Revisiting the role of persistent neural activity during working memory. *Trends in Cognitive Sciences*, *18*(2), 82–89. <https://doi.org/10.1016/j.tics.2013.12.001>
- Standage, D., You, H., Wang, D., & Dorris, M. C. (2011). Gain modulation by an urgency signal controls the speed-accuracy trade-off in a network model of a cortical decision circuit. *Frontiers in Computational Neuroscience*, *5*, Article 7. <https://doi.org/10.3389/fncom.2011.00007>
- Steinemann, N. A., Stine, G. M., Trautmann, E. M., Zylberberg, A., Wolpert, D. M., & Shadlen, M. N. (2022). *Direct observation of the neural computations underlying a single decision*. bioRxiv. <https://doi.org/10.1101/2022.05.02.490321>
- Stewart, T. C., Choo, X., & Eliasmith, C. (2010). Dynamic behaviour of a spiking model of action selection in the basal ganglia. *Proceedings of the 10th international conference on cognitive modeling* (pp. 235–240).
- Stine, G. M., Trautmann, E. M., Jeurissen, D., & Shadlen, M. N. (2023). A neural mechanism for terminating decisions. *Neuron*, *111*(16), 2601–2613. <https://doi.org/10.1016/j.neuron.2023.05.028>
- Stine, G. M., Zylberberg, A., Ditterich, J., & Shadlen, M. N. (2020). Differentiating between integration and non-integration strategies in perceptual decision making. *Elife*, *9*, Article e55365. <https://doi.org/10.7554/eLife.55365>
- Stöckel, A., Stewart, T. C., & Eliasmith, C. (2021). Connecting biological detail with neural computation: Application to the cerebellar granule-golgi microcircuit. *Topics in Cognitive Science*, *13*(3), 515–533. <https://doi.org/10.1111/tops.12536>
- Suzuki, S., & O'Doherty, J. P. (2020). Breaking human social decision making into multiple components and then putting them together again. *Cortex*, *127*, 221–230. <https://doi.org/10.1016/j.cortex.2020.02.014>
- Teodorescu, A. R., & Usher, M. (2013). Disentangling decision models: From independence to competition. *Psychological Review*, *120*(1), 1–38. <https://doi.org/10.1037/a0030776>
- Thura, D., & Cisek, P. (2017). The basal ganglia do not select reach targets but control the urgency of commitment. *Neuron*, *95*(5), 1160–1170. <https://doi.org/10.1016/j.neuron.2017.07.039>
- Trueblood, J., Brown, S., & Heathcote, A. (2014). The multiattribute linear ballistic accumulator model of context effects in multialternative choice. *Psychological Review*, *121*(2), 179–205. <https://doi.org/10.1037/a0036137>
- Turner, B. M., Van Maanen, L., & Forstmann, B. U. (2015). Informing cognitive abstractions through neuroimaging: The neural drift diffusion model. *Psychological Review*, *122*(2), 312–336. <https://doi.org/10.1037/a0038894>
- Urai, A. E., & Donner, T. H. (2022). Persistent activity in human parietal cortex mediates perceptual choice repetition bias. *Nature Communications*, *13*(1), Article 6015. <https://doi.org/10.1038/s41467-022-33237-5>
- Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review*, *108*(3), 550–592. <https://doi.org/10.1037/0033-295X.108.3.550>
- Vabalas, A., Gowen, E., Poliakoff, E., & Casson, A. J. (2020). Applying machine learning to kinematic and eye movement features of a movement imitation task to predict autism diagnosis. *Scientific Reports*, *10*(1), Article 8346. <https://doi.org/10.1038/s41598-020-65384-4>
- Vallesi, A., Canalaz, F., Balestrieri, M., & Brambilla, P. (2015). Modulating speed-accuracy strategies in major depression. *Journal of Psychiatric Research*, *60*, 103–108. <https://doi.org/10.1016/j.jpsychires.2014.09.017>
- van Ravenzwaaij, D., Brown, S. D., Marley, A., & Heathcote, A. (2020). Accumulating advantages: A new conceptualization of rapid multiple choice. *Psychological Review*, *127*(2), 186–215. <https://doi.org/10.1037/rev0000166>
- Vandekerckhove, J., Tuerlinckx, F., & Lee, M. D. (2011). Hierarchical diffusion models for two-choice response times. *Psychological Methods*, *16*(1), 44–62. <https://doi.org/10.1037/a0021765>
- Vickers, D., Caudrey, D., & Willson, R. (1971). Discriminating between the frequency of occurrence of two alternative events. *Acta Psychologica*, *35*(2), 151–172. [https://doi.org/10.1016/0001-6918\(71\)90018-7](https://doi.org/10.1016/0001-6918(71)90018-7)

Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., ... SciPy 1.0 Contributors. (2020). SciPy 1.0: Fundamental algorithms for scientific computing in python. *Nature Methods*, 17, 261–272. <https://doi.org/10.1038/s41592-019-0686-2>

Voss, A., Rothermund, K., & Voss, J. (2004). Interpreting the parameters of the diffusion model: An empirical validation. *Memory & Cognition*, 32, 1206–1220. <https://doi.org/10.3758/BF03196893>

Wagenmakers, E.-J. (2009). Methodological and empirical developments for the ratcliff diffusion model of response times and accuracy. *European Journal of Cognitive Psychology*, 21(5), 641–671. <https://doi.org/10.1080/09541440802205067>

Wang, M., Gamo, N. J., Yang, Y., Jin, L. E., Wang, X.-J., Laubach, M., Mazer, J. A., Lee, D., & Arnsten, A. F. (2011). Neuronal basis of age-related working memory decline. *Nature*, 476(7359), 210–213. <https://doi.org/10.1038/nature10243>

Wang, X.-J. (2008). Decision making in recurrent neuronal circuits. *Neuron*, 60(2), 215–234. <https://doi.org/10.1016/j.neuron.2008.09.034>

Wiecki, T. V., Sofer, I., & Frank, M. J. (2013). HDDM: Hierarchical Bayesian estimation of the drift-diffusion model in python. *Frontiers in Neuroinformatics*, 7, Article 14. <https://doi.org/10.3389/fninf.2013.00014>

Yau, Y., Dadar, M., Taylor, M., Zeighami, Y., Fellows, L., Cisek, P., & Dagher, A. (2020). Neural correlates of evidence and urgency during human perceptual decision-making in dynamically changing conditions. *Cerebral Cortex*, 30(10), 5471–5483. <https://doi.org/10.1093/cercor/bhaa129>

Zhang, J., & Rowe, J. B. (2014). Dissociable mechanisms of speed-accuracy tradeoff during visual perceptual learning are revealed by a hierarchical drift-diffusion model. *Frontiers in Neuroscience*, 8, Article 69. <https://doi.org/10.3389/fnins.2014.00069>

## Appendix

### Fitted Parameter Values

**Table A1**

*Parameter Values Used in the Spiking Neural Network Model for Each Simulated Experiment*

Experiment	Figure	Details	<i>R</i>	<i>T</i>	<i>L</i>	$\sigma$	<i>t<sub>s</sub></i>
Preliminary	2	Seed 0	1.5	0.25	0.0	0.3	0.02
Preliminary	3	Seed 7	1.5	0.25	0.0	0.3	0.02
Preliminary	4	Seed 0	1.5	0.1	1.0	0.3	0.02
Preliminary	5	Seed 0	1.5	0.25	0.7	0.3	0.02
Preliminary	5	Seed 7	1.5	0.25	0.7	0.3	0.02
Preliminary	6	Seeds 0–500	1.5	0.25	0.7	0.3	0.02
1	7–10	Speed	1.18	0.22	0.13	0.62	0.04
1	7–10	Accuracy	1.18	0.33	0.13	0.62	0.04
2	11–12	Two-choice	1.52	0.19	0.7	0.4	0.045
2	11–12	Four-choice	1.96	0.26	0.81	0.32	0.08
3	13	as1t, speed	1.9	0.39	0.035	0.37	0.068
3	13	as1t, accuracy	1.9	0.38	0.035	0.37	0.068
3	13	bd6t, speed	1.9	0.22	0.47	0.067	0.013
3	13	bd6t, accuracy	1.9	0.35	0.47	0.067	0.013
3	13	bl1t, speed	1.6	0.2	0.36	0.055	0.071
3	13	bl1t, accuracy	1.6	0.3	0.36	0.055	0.071
3	13	hsft, speed	2.0	0.32	0.01	0.13	0.029
3	13	hsft, accuracy	2.0	0.4	0.01	0.13	0.029
3	13	hsgt, speed	1.9	0.28	0.18	0.052	0.082
3	13	hsgt, accuracy	1.9	0.31	0.18	0.052	0.082
3	13	kd6t, speed	1.6	0.27	0.24	0.01	0.072
3	13	kd6t, accuracy	1.6	0.31	0.24	0.01	0.072
3	13	kd9t, speed	1.9	0.32	0.094	0.3	0.03
3	13	kd9t, accuracy	1.9	0.41	0.094	0.3	0.03
3	13	kmat, speed	1.8	0.31	0.068	0.6	0.048
3	13	kmat, accuracy	1.8	0.39	0.068	0.6	0.048
3	13	ku4t, speed	1.8	0.23	0.2	0.022	0.068
3	13	ku4t, accuracy	1.8	0.27	0.2	0.022	0.068
3	13	na1t, speed	2.0	0.28	0.03	0.092	0.035
3	13	na1t, accuracy	2.0	0.32	0.03	0.092	0.035
3	13	rmbt, speed	1.2	0.2	0.18	0.33	0.082
3	13	rmbt, accuracy	1.2	0.32	0.18	0.33	0.082
3	13	rt2t, speed	1.9	0.24	0.046	0.1	0.014
3	13	rt2t, accuracy	1.9	0.32	0.046	0.1	0.014
3	13	rt3t, speed	1.9	0.3	0.49	0.041	0.072
3	13	rt3t, accuracy	1.9	0.33	0.49	0.041	0.072
3	13	rt5t, speed	1.8	0.3	0.5	0.5	0.075
3	13	rt5t, accuracy	1.8	0.35	0.5	0.5	0.075

(table continues)

(Appendix continues)

**Table A1** (continued)

Experiment	Figure	Details	$R$	$T$	$L$	$\sigma$	$t_s$
3	13	scat, speed	1.9	0.34	0.095	0.38	0.016
3	13	scat, accuracy	1.9	0.35	0.095	0.38	0.016
3	13	ta5t, speed	1.8	0.28	0.083	0.34	0.039
3	13	ta5t, accuracy	1.8	0.41	0.083	0.34	0.039
3	13	vf1t, speed	1.9	0.28	0.026	0.32	0.053
3	13	vf1t, accuracy	1.9	0.44	0.026	0.32	0.053
3	13	zk1t, speed	1.9	0.28	0.18	0.36	0.07
3	13	zk1t, accuracy	1.9	0.43	0.18	0.36	0.07
4	15-16	1	2.0	0.51	0.8		
4	15-16	2	2.4	0.35	0.47		
4	15-16	3	1.5	0.14	0.92		
4	15-16	4	2.5	1.0	0.44		
4	15-16	5	1.5	0.37	0.49		
4	15-16	6	2.6	0.23	0.89		
4	15-16	7	2.9	0.65	0.4		
4	15-16	9	2.3	0.65	0.07		
4	15-16	10	1.8	0.58	0.64		
4	15-16	11	2.0	0.26	0.92		
4	15-16	13	2.1	0.36	0.56		
4	15-16	14	2.4	0.8	0.16		
4	15-16	20	1.3	0.44	0.54		
4	15-16	21	1.2	0.59	0.14		
4	15-16	22	2.3	0.75	0.49		
4	15-16	23	2.2	0.27	0.49		
4	15-16	24	2.5	0.47	0.05		
4	15-16	25	1.8	0.51	0.05		
4	15-16	26	1.9	0.68	0.13		
4	15-16	27	2.7	0.46	0.0		
4	15-16	28	2.9	0.65	0.2		
4	15-16	29	2.2	0.48	0.57		
4	15-16	30	1.4	0.64	0.12		
4	15-16	31	1.1	0.01	0.54		
4	15-16	32	1.0	0.07	0.45		
4	15-16	33	1.5	0.64	0.24		
4	15-16	34	1.3	0.49	0.38		
4	15-16	35	2.7	0.82	0.75		
4	15-16	36	2.4	0.71	0.42		
4	15-16	37	1.4	0.72	0.05		
4	15-16	38	2.3	0.32	0.76		
4	15-16	39	2.5	0.48	0.08		
4	15-16	40	2.5	0.8	0.13		
4	15-16	41	2.4	0.81	0.38		
4	15-16	42	1.4	0.4	0.46		
4	15-16	43	1.4	0.66	0.17		
4	15-16	44	2.2	0.33	0.95		
4	15-16	45	2.5	0.37	0.78		
4	15-16	46	1.1	0.31	0.48		
4	15-16	47	1.6	0.23	0.71		
4	15-16	48	1.7	0.7	0.24		
4	15-16	49	1.2	0.45	0.33		
4	15-16	50	2.1	0.76	0.25		
4	15-16	51	2.8	0.67	0.19		
4	15-16	52	1.9	0.3	0.74		
4	15-16	53	2.1	0.75	0.33		
4	15-16	54	1.3	0.56	0.19		
4	15-16	55	2.3	0.61	0.55		
4	15-16	56	1.5	0.34	0.82		
4	15-16	57	2.8	0.44	0.76		
4	15-16	58	2.4	0.24	0.72		
4	15-16	59	1.5	0.19	0.93		
4	15-16	60	1.7	0.58	0.55		
4	15-16	61	1.4	0.32	0.26		
4	15-16	62	2.8	0.56	0.57		
Valuation		as1t	1.3	0.15	0.0	0.6	0.03
Valuation		bd6t	1.9	0.17	0.0	0.43	0.035

(table continues)

(Appendix continues)

**Table A1** (continued)

Experiment	Figure	Details	$R$	$T$	$L$	$\sigma$	$t_s$
Valuation		bl1t	1.9	0.16	0.0	0.63	0.018
Valuation		hsft	1.4	0.066	0.0	0.19	0.023
Valuation		hsgt	1.5	0.049	0.0	0.27	0.014
Valuation		kd6t	1.7	0.22	0.0	0.73	0.016
Valuation		kd9t	1.3	0.079	0.0	0.1	0.09
Valuation		kmat	1.4	0.039	0.0	0.11	0.09
Valuation		ku4t	1.8	0.15	0.0	0.51	0.04
Valuation		na1t	1.6	0.058	0.0	0.11	0.028
Valuation		rmbt	1.8	0.2	0.0	0.62	0.035
Valuation		rt2t	1.7	0.073	0.0	0.2	0.033
Valuation		rt3t	1.5	0.23	0.0	0.63	0.043
Valuation		rt5t	1.4	0.24	0.0	0.78	0.096
Valuation		scat	1.7	0.19	0.0	0.44	0.034
Valuation		ta5t	1.6	0.12	0.0	0.38	0.034
Valuation		vf1t	1.6	0.1	0.0	0.42	0.033
Valuation		zk1t	1.4	0.039	0.0	0.34	0.09
Valuation		1	1.8	0.85	0.0		
Valuation		2	2.9	0.54	0.0		
Valuation		3	2.9	0.54	0.0		
Valuation		4	1.9	1.0	0.0		
Valuation		5	2.9	0.86	0.0		
Valuation		6	2.7	0.46	0.0		
Valuation		7	2.9	0.84	0.0		
Valuation		9	2.5	0.72	0.0		
Valuation		10	1.3	0.7	0.0		
Valuation		11	2.9	0.78	0.0		
Valuation		13	1.8	0.39	0.0		
Valuation		14	2.9	0.97	0.0		
Valuation		20	1.6	0.77	0.0		
Valuation		21	1.6	0.79	0.0		
Valuation		22	1.2	0.65	0.0		
Valuation		23	2.6	0.44	0.0		
Valuation		24	2.6	0.51	0.0		
Valuation		25	2.5	0.74	0.0		
Valuation		26	2.2	0.82	0.0		
Valuation		27	2.2	0.37	0.0		
Valuation		28	1.9	0.5	0.0		
Valuation		29	2.4	0.73	0.0		
Valuation		30	1.4	0.66	0.0		
Valuation		31	2.0	0.1	0.0		
Valuation		32	2.7	0.2	0.0		
Valuation		33	1.0	0.5	0.0		
Valuation		34	1.4	0.67	0.0		
Valuation		35	1.8	0.98	0.0		
Valuation		36	2.7	0.98	0.0		
Valuation		37	1.4	0.75	0.0		
Valuation		38	2.7	0.57	0.0		
Valuation		39	2.8	0.54	0.0		
Valuation		40	2.7	0.91	0.0		
Valuation		41	2.1	0.92	0.0		
Valuation		42	2.8	0.94	0.0		
Valuation		43	1.7	0.84	0.0		
Valuation		44	2.9	0.87	0.0		
Valuation		45	1.1	0.29	0.0		
Valuation		46	2.5	0.87	0.0		
Valuation		47	2.5	0.5	0.0		
Valuation		48	1.0	0.5	0.0		
Valuation		49	1.8	0.79	0.0		
Valuation		50	1.7	0.77	0.0		
Valuation		51	2.7	0.73	0.0		
Valuation		52	2.9	0.74	0.0		
Valuation		53	1.9	0.86	0.0		
Valuation		54	2.2	0.9	0.0		
Valuation		55	2.9	0.98	0.0		
Valuation		56	2.9	1.0	0.0		

(table continues)

(Appendix continues)

**Table A1** (*continued*)

Experiment	Figure	Details	$R$	$T$	$L$	$\sigma$	$t_s$
Valuation		57	3.0	0.81	0.0		
Valuation		58	1.8	0.3	0.0		
Valuation		59	2.8	0.72	0.0		
Valuation		60	1.7	0.82	0.0		
Valuation		61	2.5	0.68	0.0		
Valuation		62	3.0	0.85	0.0		

*Note.* In Experiments 1–4, these parameters are fitted to the empirical data using the procedure described in the Parameter Fitting Procedure section. The details column describes additional conditions of the experiment, including the simulation seed used to initialize neuron parameters, speed- or accuracy-emphasis, the number of choice alternatives, or the ID of each participant. Fitted parameter values for the DD and extrema model are available on GitHub. DD = drift diffusion.

Received September 3, 2023

Revision received September 16, 2024

Accepted September 21, 2024 ■