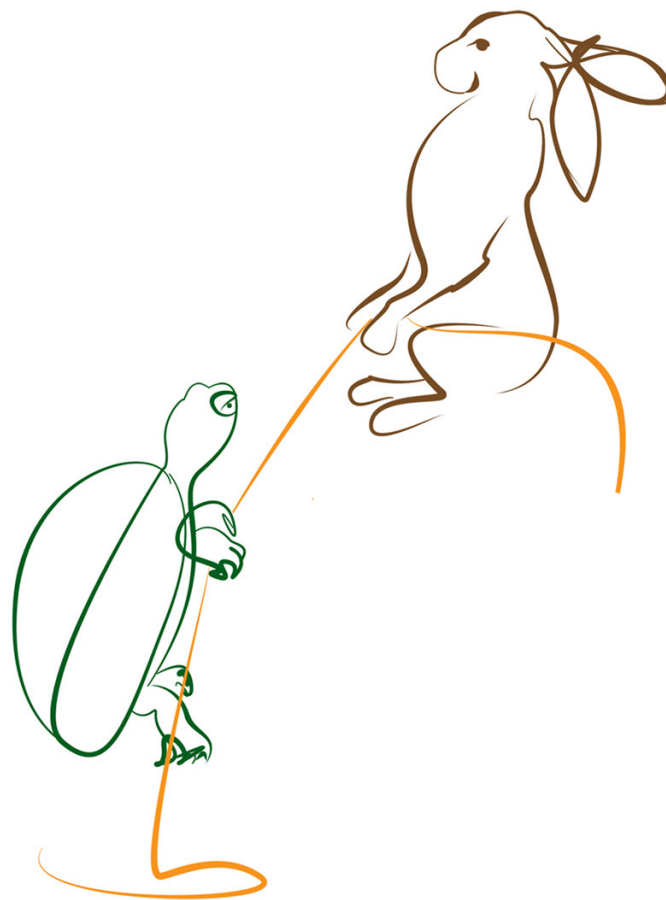# TOWARD A UNIFIED VIEW OF THE SPEED-ACCURACY TRADE-OFF: BEHAVIOUR, NEUROPHYSIOLOGY AND MODELLING

EDITED BY: Dominic Standage, Da-Hui Wang, Richard P. Heitz and Patrick Simen
PUBLISHED IN: Frontiers in Neuroscience

frontiers Research Topics

About Frontiers

Frontiers is more than just an open-access publisher of scholarly articles: it is a pioneering approach to the world of academia, radically improving the way scholarly research is managed. The grand vision of Frontiers is a world where all people have an equal opportunity to seek, share and generate knowledge. Frontiers provides immediate and permanent online open access to all its publications, but this alone is not enough to realize our grand goals.

Frontiers Journal Series

The Frontiers Journal Series is a multi-tier and interdisciplinary set of open-access, online journals, promising a paradigm shift from the current review, selection and dissemination processes in academic publishing. All Frontiers journals are driven by researchers for researchers; therefore, they constitute a service to the scholarly community. At the same time, the Frontiers Journal Series operates on a revolutionary invention, the tiered publishing system, initially addressing specific communities of scholars, and gradually climbing up to broader public understanding, thus serving the interests of the lay society, too.

Dedication to Quality

Each Frontiers article is a landmark of the highest quality, thanks to genuinely collaborative interactions between authors and review editors, who include some of the world's best academicians. Research must be certified by peers before entering a stream of knowledge that may eventually reach the public - and shape society; therefore, Frontiers only applies the most rigorous and unbiased reviews.

Frontiers revolutionizes research publishing by freely delivering the most outstanding research, evaluated with no bias from both the academic and social point of view.

By applying the most advanced information technologies, Frontiers is catapulting scholarly publishing into a new generation.

What are Frontiers Research Topics?

Frontiers Research Topics are very popular trademarks of the Frontiers Journals Series: they are collections of at least ten articles, all centered on a particular subject. With their unique mix of varied contributions from Original Research to Review Articles, Frontiers Research Topics unify the most influential researchers, the latest key findings and historical advances in a hot research area! Find out more on how to host your own Frontiers Research Topic or contribute to one as an author by contacting the Frontiers Editorial Office: **researchtopics@frontiersin.org**

TOWARD A UNIFIED VIEW OF THE SPEED-ACCURACY TRADE-OFF: BEHAVIOUR, NEUROPHYSIOLOGY AND MODELLING
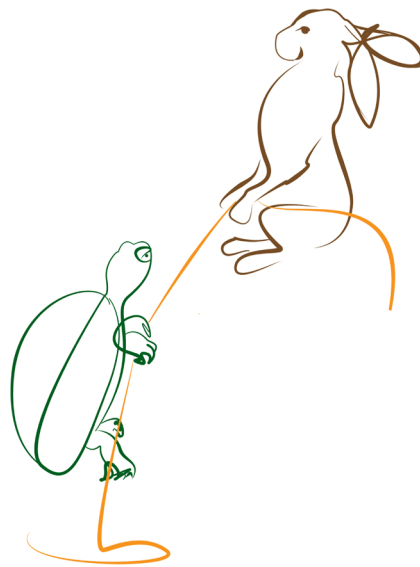
Topic Editors:
Dominic Standage, Queen's University, Canada
Da-Hui Wang, Beijing Normal University, China
Richard P. Heitz, Vanderbilt University, USA
Patrick Simen, Oberlin College, USA

Everyone is familiar with the speed-accuracy trade-off (SAT). To make good choices, we need to balance the conflicting demands of fast and accurate decision making. After all, hasty decisions often lead to poor choices, but accurate decisions may be useless if they take too long. This notion is intuitive because it reflects a fundamental aspect of cognition: not only do we deliberate over the evidence for decisions, but we can control that deliberative process. This control raises many questions for the study of choice behaviour and executive function. For example, how do we figure out the appropriate balance between speed and accuracy on a given task? How do we impose that balance on our decisions, and what is its neural basis?

Researchers have addressed these and related questions for decades, using a variety of methods and offering answers at different levels of abstraction. Given this diverse methodology, our aim is to provide a unified view of the SAT. Extensive analysis of choice behaviour suggests that we make decisions by accumulating evidence until some criterion is reached. Thus, adjusting the criterion controls how long we accumulate evidence and therefore the speed and accuracy of decisions. This simple framework provides the platform for our unified view. In the pages that follow, leading experts in decision neuroscience consider the history of SAT research, strategies for determining the optimal balance between speed and accuracy, conditions under which this seemingly ubiquitous phenomenon breaks down, and the neural mechanisms that may implement the computations of our unifying framework.

Cover image (with permission): Xu, Sisi *The Tortoise and the Hare* (2015) Queen's University, Ontario, Canada

Citation: Standage, D., Wang, D-H., Heitz, R. P., Simen, P., eds. (2016). Toward a Unified View of the Speed-Accuracy Trade-Off: Behaviour, Neurophysiology and Modelling. Lausanne: Frontiers Media. doi: 10.3389/978-2-88919-756-9

Table of Contents

frontiers
in Neuroscience

Toward a unified view of the speed-accuracy trade-off

Dominic Standage[1]*, Da-Hui Wang[2], Richard P. Heitz[3] and Patrick Simen[4]

[1] Department of Biomedical and Molecular Sciences, Queen's University, Kingston, ON, Canada, [2] Department of Systems Science, Beijing Normal University, Beijing, China, [3] Department of Psychology, Vanderbilt University, Nashville, TN, USA, [4] Department of Neuroscience, Oberlin College, Oberlin, OH, USA

Hasty decisions often lead to poor choices, whereas accurate decisions are ineffective if they take too long. Thus, good choices require cognitive mechanisms to determine the appropriate balance between speed and accuracy, and to control decision processing accordingly. This balance is referred to as the speed-accuracy trade-off (SAT) and the mechanisms by which it is determined and imposed are the subject of this Frontiers Research Topic. Given the near-ubiquity of the SAT across species and experimental tasks, it is not surprising that a wide range of methods have been used to investigate it. Our aim is to provide a unified view of the SAT in light of this diverse methodology. Computationally, decision making and the SAT are well characterized by the framework of *bounded integration*, providing a solid foundation for this view. Under this framework, noisy evidence for the available choices is added up (integrated) until the running total for one of them reaches a criterion (the bound). The SAT is readily controlled by the bound, where a higher bound favors accuracy at the expense of speed and *vice versa*. In this collection, we use bounded integration as a reference point for considering the factors that determine the optimal balance between speed and accuracy, the interpretation of behavior by different models from this general class, and the neural implementation of the computations captured by these models. Articles herein further consider conditions under which the above descriptions of the SAT and bounded integration do not explain behavior, and the utility of the SAT for manipulating the context of decisions.

The review by Heitz (2014) describes the history of the SAT as a quantifiable behavioral phenomenon and provides a critical appraisal of methodologies for its study. His historical account describes the shaping of decision theory by the SAT, a perspective that nicely sets up the original research article by Ivanoff et al. (2014), who used SAT methodology to investigate *spatial compatibility* effects, that is, how the respective locations of stimuli and responses can influence behavior. They found that SAT manipulations can systematically promote or impede the efficacy of stimulus-response mappings.

Stone (2014) investigated the relationship between speed and accuracy in his original research article, reasoning that the information gained by the observation of evidence should be reflected in both the speed and accuracy of decisions. By fitting a bounded integration model to experimental data, he used model parameters to estimate the mutual information between perceptual evidence and speed, and between perceptual evidence and accuracy. These measures provide bounds on the information gained by the observation of evidence and were used to calculate the smallest detectable change in the strength of evidence.

Salinas et al. (2014) reviewed recent studies of perceptual decisions under extreme time pressure. In this context, the respective contributions of perception and motor planning to choice behavior can be distinguished from one another, quantifying how the former guides the latter. These experiments showed that perceptual information can accelerate or decelerate the competition between ongoing motor plans, revealing the SAT as the combined effect of multiple adjustments to decision processing, not a monolithic phenomenon.

The isolation of perception from motor planning under extreme time pressure (Salinas et al., 2014) is manifest in the independence of accuracy from decision time, which constitutes a violation of the SAT. Another well-known violation is the improvement in speed *and* accuracy while learning a task. This improvement is readily captured by increasing a parameter that loosely corresponds to the difference in strength between sources of evidence, often referred to as "drift." In effect, learning mimics a decrease in task difficulty. In their original research article, Zhang and Rowe (2014) used a bounded integration model to investigate the effects of speed and accuracy emphasis during and after learning. Under accuracy emphasis, increasing the bound and the drift captured subjects' behavior at the beginning of learning, whereas only an increase in the bound captured behavior after learning. Their results suggest that learning and speed-accuracy emphasis differentially influence decision processing on different timescales.

It is widely accepted that the objective of the SAT is to optimize decisions in terms of reward rate, that is, decision makers aim to maximize the pay-off of the task at hand. Three original research articles in the collection investigated optimal decision making, each considering a different set of conditions and corresponding strategies. Khodadadi et al. (2014) considered the case of a limited time interval, during which decision makers can make as many (or as few) decisions as they wish. This task can be formulated as a search for the reward-maximizing bound in a given condition. Khodadadi et al. (2014) took a *reinforcement learning* approach to this problem, specifying a set of conditions, each corresponding to a configuration of task constraints, e.g., the difficulty of the task, the magnitude of reward and so forth. In the terminology of reinforcement learning, each condition is a "state" and the bound that maximizes the reward rate in that condition is its "action" under the optimal "policy." Their model took a conservative strategy, choosing a high, sub-optimal bound in the early stages of learning, before lowering it with experience to achieve optimality. This result is a testable prediction for behavioral experiments.

Karsilar et al. (2014) investigated decisions with deadlines, in which the optimal strategy is to reduce the bound during each decision. This strategy ensures that decisions are always made by the deadline, at a cost of lower accuracy. As such, decision makers have to estimate the upcoming deadline and have to account for the variability in these estimates. Crucially, models that implement this strategy predict that accuracy will decline to near-chance levels as the deadline approaches. Karsilar et al. (2014) tested this prediction with a perceptual choice task, finding that subjects' performance did not decline to chance levels near the deadline, and that a slight decline did not relate to timing variability. Furthermore, subjects' behavior was captured by a standard bounded integration model. These results suggest that perceptual decisions are too short for within-trial adaptation of the neural mechanisms captured by the bound.

As described above, the fundamental principle of bounded integration is that the effect of *within-trial* noise can be limited by integrating evidence. Goldfarb et al. (2014) compared several bounded integration models with a popular model that does not include within-trial noise, in which decision-time variability and error rates are determined only by *between-trial* noise,

i.e., parameter values that vary from trial to trial. Their study focused on reward-maximization tasks, in which task difficulty is held constant for a block of trials and subjects try to earn as much reward as possible, i.e., they try to optimize the trade-off between speed and accuracy for a given task difficulty. Significant differences were found between the classes of model, especially on difficult tasks. As such, the models provide different interpretations of behavior as task difficulty increases.

The issue of optimality is further considered by Pirrone et al. (2014), who took an evolutionary perspective in their opinion article. They argued that in most real-world decisions, each of the alternatives offers some quantity of reward (e.g., deciding between food items), whereas the dominant experimental approach to date has been to reward a single alternative only. Therefore, they suggest that most natural decisions are value-based, necessitating a *speed-value* trade-off, optimized by natural selection. They formalized this optimization problem and argued that bounded integration models that optimize the SAT can only account for value-based decisions if their parameter values are assigned on a case-by-case basis, limiting their generality.

The hypothesis and theory article by Standage et al. (2014b) questioned the commonly held view that the bound is implemented by the rate of decision-correlated neural activity at the time of commitment to a choice, as well as the view that the difference between this rate and a "baseline" rate controls the SAT. Using a model derived from biophysical considerations, they showed that these views may be inconsistent with widely-held principles of cortical computation, which account for the SAT. According to their hypothesis, the behavior of the bound is well-approximated by an emergent property of cortical dynamics, but not by the aforementioned difference in firing rates.

The SAT has long been investigated as a behavioral phenomenon, but studies addressing its neural basis are a recent development. Standage et al. (2014a) reviewed hypotheses on the neural basis of the SAT, considering three general mechanistic categories: modulation of the encoding of evidence under speed and accuracy emphasis, modulation of the integration of encoded evidence, and modulation of the amount of integrated evidence sufficient to make a choice. Thus, their review is structured by the principles of bounded integration, but they focused on models addressing the neural implementation of these principles, and on the explanations offered by these models for a growing body of neuroimaging and electrophysiological data. This convergence of neural and behavioral data with models at different levels of abstraction is exemplary of interdisciplinary neuroscience, and suggests a productive future for the mechanistic study of decision making, the SAT and cognition.

We believe this collection of articles provides a useful reference for future SAT research, with review articles to direct readers to relevant work in the literature, opinions and hypotheses on the interpretation of topical methodologies and data, and original research articles that make important advances in the field. Moreover, we believe that bounded integration has successfully provided a unifying framework for the collection, supporting the systematic consideration

of the SAT under different methodologies, at different levels of abstraction, and from different perspectives. A complete theory of decision making must explain the SAT. We hope this Research Topic makes a valued contribution toward this fundamental goal of cognitive neuroscience.

References

Goldfarb, S., Leonard, N. E., Simen, P., Caicedo-Nez, C. H., and Holmes, P. (2014). A comparative study of drift diffusion and linear ballistic accumulator models in a reward maximization perceptual choice task. *Front. Neurosci.* 8:148. doi: 10.3389/fnins.2014.00148

Heitz, R. P. (2014). The speed-accuracy tradeoff: history, physiology, methodology, and behavior. *Front. Neurosci.* 8:150. doi: 10.3389/fnins.2014.00150

Ivanoff, J., Blagdon, R., Feener, S., McNeil, M., and Muir, P. H. (2014). On the temporal dynamics of spatial stimulus-response transfer between spatial incompatibility and simon tasks. *Front. Neurosci.* 8:243. doi: 10.3389/fnins.2014.00243

Karsilar, H., Simen, P., Papadakis, S., and Balci, F. (2014). Speed accuracy trade-off under response deadlines. *Front. Neurosci.* 8:248. doi: 10.3389/fnins.2014.00248

Khodadadi, A., Fakhari, P., and Busemeyer, J. R. (2014). Learning to maximize reward rate: a model based on semi-markov decision processes. *Front. Neurosci.* 8:101. doi: 10.3389/fnins.2014.00101

Pirrone, A., Stafford, T., and Marshall, J. A. R. (2014). When natural selection should optimize speed-accuracy trade-offs. *Front. Neurosci.* 8:73. doi: 10.3389/fnins.2014.00073

Salinas, E., Scerra, V. E., Hauser, C. K., Costello, M. G., and Stanford, T. R. (2014). Decoupling speed and accuracy in an urgent decision-making task reveals multiple contributions to their trade-off. *Front. Neurosci.* 8:85. doi: 10.3389/fnins.2014.00085

Standage, D., Blohm, G., and Dorris, M. C. (2014a). On the neural implementation of the speed-accuracy trade-off. *Front. Neurosci.* 8:236. doi: 10.3389/fnins.2014.00236

Standage, D., Wang, D.-H., and Blohm, G. (2014b). Neural dynamics implement a flexible decision bound with a fixed firing rate for choice: a model-based hypothesis. *Front. Neurosci.* 8:318. doi: 10.3389/fnins.2014.00318

Stone, J. V. (2014). Using reaction times and binary responses to estimate psychophysical performance: an information theoretic analysis. *Front. Neurosci.* 8:35. doi: 10.3389/fnins.2014.00035

Zhang, J., and Rowe, J. B. (2014). Dissociable mechanisms of speed-accuracy tradeoff during visual perceptual learning are revealed by a hierarchical drift-diffusion model. *Front. Neurosci.* 8:69. doi: 10.3389/fnins.2014.00069

The speed-accuracy tradeoff: history, physiology, methodology, and behavior

Richard P. Heitz *

Department of Psychology, Center for Integrative and Cognitive Neuroscience, Vanderbilt Vision Research Center, Vanderbilt University, Nashville, TN, USA

There are few behavioral effects as ubiquitous as the speed-accuracy tradeoff (SAT). From insects to rodents to primates, the tendency for decision speed to covary with decision accuracy seems an inescapable property of choice behavior. Recently, the SAT has received renewed interest, as neuroscience approaches begin to uncover its neural underpinnings and computational models are compelled to incorporate it as a necessary benchmark. The present work provides a comprehensive overview of SAT. First, I trace its history as a tractable behavioral phenomenon and the role it has played in shaping mathematical descriptions of the decision process. Second, I present a "users guide" of SAT methodology, including a critical review of common experimental manipulations and analysis techniques and a treatment of the typical behavioral patterns that emerge when SAT is manipulated directly. Finally, I review applications of this methodology in several domains.

Keywords: speed-accuracy tradeoff, decision-making

"…we face a very common problem in psychology: the existence of a tradeoff between dependent variables, in this case false alarms and reaction time. The only sensible long-range strategy is, in my opinion, to study the tradeoff… and to devise some summary statistic to describe it." - Luce, 1986, p. 56.

INTRODUCTION

Prima facie, the notion of speed-accuracy tradeoff (SAT) is pedestrian. Who has not encountered that a decision, made in haste, often leads to err? Who has not felt the deleterious effects of time pressure on ultimate outcomes? The concept seems so commonsensical as to deserve little interest—an obvious product of nothing more than human limitations. Ironically, it is just this pervasiveness that demands the SAT be considered—not only as a phenomenon in and of itself—but also as a benchmark for models of the decision process. Common across task domains and in creatures ranging from house-hunting ants (Franks et al., 2003) and bumblebees (Chittka et al., 2003; for a review, see Marshall et al., 2009) to humans (Wickelgren, 1977) and monkeys (Heitz and Schall, 2012, 2013), the SAT is thus a topic of great concern. Fortunately, there has been a renewed interest in SAT, particularly in the neuroscience community. Using fMRI, EEG, and single-unit recordings, never have we been closer to understanding, at a fundamental level, how the brain takes in sensory information and transforms it into a decision variable guiding choice. As a ubiquitous phenomenon intimately tied to the decision process, the SAT is integral.

HISTORICAL OVERVIEW

The idea that response time[1] (RT) can be used to study the inner workings of the mind is as old as psychology itself. In the mid 1800's, Hermann von Helmholtz demonstrated that peripheral nerve conduction velocity was finite and measureable—a revolutionary conception for his time. The logic was simple, yet elegant. Helmholtz created a preparation of frog legs with a portion of nerve still attached; applying current to the nerve elicited muscle contraction. He then noted the difference in the latency to contraction when either a proximal or distal portion of the nerve was stimulated. Since the distance between the stimulation points was known, Helmholtz easily worked out the conduction velocity (see Foster, 1870). Helmholtz' logic was perhaps just as important as his discovery: one can use the time of an overt movement as a dependent measure, and by altering the antecedent conditions, estimate the duration of intermediary components. Perhaps one could use similar methodology to objectively measure the component processes of the mind. This philosophy guided several researchers in their exploration of the "velocity of thought," including Helmholtz' colleague Wilhelm Wundt, in what would be known as the first true psychology laboratory. Similar logic was employed by Merkel (1885), and very notably, by Donders (1868) in his study of processing stages using task comparisons. The use of RT—one of the only non-introspective measures available, became central.

That the accuracy of a response varies with the time taken to produce it was probably already known, if implicitly. However, such variation was of little interest, the field being dominated at either extreme by psychophysics experiments—which emphasize high accuracy without concern for RT—and reaction time experiments, which examine one's ability to produce predefined responses to simple visual or auditory stimuli. Outside of this asymptotic performance lay a nether-region of neither wholly

---

[1] The term "response time" and "reaction time" are typically used interchangeably, and I will make no distinction here, but there is a slight semantic difference. "Reaction time" is often associated with the limits of ability, as in making a fast, predetermined response to the onset of a visual stimulus, whereas "response time" more generally describes "time to overt action." See (Luce, 1986).

accurate nor wholly fast responding. Still, the fact that such variability exists led some early researchers to address the speed-accuracy relation empirically. The first demonstration that the accuracy of an action varies with its speed was provided in 1899, both in a dissertation by Woodworth (1899) and a contemporaneous work by Martin and Müeller (1899), though these studies focused on the speed of obligatory movements rather than choice behavior[2]. The first demonstration of a relationship between *choice* accuracy and decision time can be traced to 1911, when Henmon (1911) presented subjects with a simple discrimination task. Two lines were presented, each differing slightly in length, and subjects were to determine which line was longer (or shorter) and press the appropriate left or right button. In the first analysis of its kind, Henmon "binned" the data by RT to examine the effect of latency on accuracy. His data revealed an orderly relation, suggesting they were not independent. A short time later, Henmon's observations were replicated and the relationship dubbed the "speed-accuracy relation" for the first time in oft-neglected dissertation by Garrett (1922). The phenomenon received only sporadic attention thereafter, for nearly three decades.

In the intervening years, work conducted on statistical decision-making would ultimately provide a framework for understanding the SAT, and also bring the phenomenon to center stage. This work, carried out independently by Alan Turing[3], Abraham Wald, and others, demonstrated that decision-making under uncertainty can be bolstered through sequential sampling of information—a suggestion not previously considered by the extant literature in economics (Edwards, 1954). Consider a choice between two competing hypotheses—say, whether or not a batch of product contains sufficient defects to warrant rejection. At the outset, one may already have some prior expectation regarding which hypothesis is more likely. An updated posterior probability can be computed by simply sampling information (e.g., units of product) sequentially. The problem is that information is costly—each sample takes some quanta of time and effort (Drugowitsch et al., 2012). Therefore, it is in one's best interest to sample as little as possible to reach some specified compromise between confidence and time spent sampling. Wald's procedure, which became known as the *sequential probability ratio test* (Wald, 1947), allows one to approach a known (acceptable) error rate with a potentially enormous savings in time and resources.

Turing and Wald's application was a utilitarian approach to economical decision-making, but it did not take long for others to realize that the process may apply more generally to human choice behavior. The first instance of this was provided in 1958 by Becker (1958). Participants viewed successive presentations of cards, upon each of which was an imprinted letter. Cards were drawn from one of two or more competing distributions, described to subjects prior to each run. Viewers were asked to sample as many cards as needed to determine which distribution the cards were drawn from. Becker manipulated the difficulty of the discrimination by altering the form of the parent distributions. For instance, subjects might need to determine if a sequence of "P" and "Q" letters were sampled from a distribution with a P:Q ratio of 2:1 or 1:1. Becker found that even in this abstract situation, humans produce data conforming to Wald's predictions, at least to a first approximation.

THE INTRODUCTION OF MATHEMATICAL DECISION MODELS

Meanwhile, others were working on formulating a mathematical relationship between decision time and accuracy. The first attempts, provided by Audley (Audley and Jonckheere, 1956; Audley, 1957, 1958), demonstrated that two-choice decisions could be modeled as a stochastic process. Audley had been working with albino rats trained to push one of two buttons to earn reward. At that time, stochastic models had seen success in predicting the form of the learning curve in terms of a gain in accuracy over successive trials, but they did not accommodate decision times. Nonetheless, decision times, and the RT distributions they form, were thought to reflect the structure of the choice process (Christie and Luce, 1956), and so were likely an important component of a complete choice model. Audley demonstrated that with some simple assumptions regarding the form of the underlying RT distribution (in this case, exponential), one could simultaneously predict both choice accuracy and decision time. However, the individual quanta in this situation were single, punctate choices made by rats; the model was opaque to the cognitive events carried out *within* any given trial. Audely soon remedied this, in a model that would become known as the *Runs* model (Audley, 1960); see also (LaBerge, 1962; Audley and Pike, 1965). In a guarded conceptual leap, Audley assumed that the choice process involves a series of "implicit responses" arising from the presentation of a sensory signal. Though the definition of "implicit responses" was left open to interpretation, it seems closely related to what we might now call "perceptual accumulation." During a choice trial, observers obtain successive samples of implicit responses, and some counting mechanism keeps track of the number of consecutive runs favoring either of two potential actions. Formulated mathematically, Audley demonstrated that the model could account for choice behavior; notably, he fit the model to Henmon's data (Henmon, 1911) described earlier.

The above efforts came to a head in 1960, when Stone (1960) produced a formal mathematical model of the decision process. The model combined (1) the relation between RT and accuracy rates as a stochastic process; (2) the mathematics and optimality of the sequential probability ratio test; and (3) the presumption of information accumulation over the course of perceptual decision-making. The model, known as the *random walk*[4], made very specific, empirically testable predictions about the means and shapes of reaction time distributions, and how those distributions change with SAT. **Figure 1** presents two depictions of the random walk, adapted, respectively, from Fitts (1966) and

---

[2]As the present work is focused on choice behavior, the movement speed-accuracy tradeoff will not be considered. The reader is referred to (Hancock and Newell, 1985; Meyer et al., 1990; Plamondon and Alimi, 1997).

[3]Turing's effort was directed at breaking the Nazi enigma machine. For a fascinating review, see (Gold and Shadlen, 2002).

[4]The random walk process is by no means limited to psychology, but has seen application in physics, chemistry, and economics. It was first proposed by Pearson (1905), the same year that Albert Einsten published work on the closely related, continuous-time stochastic process describing Brownian motion, later to become known as the diffusion process.
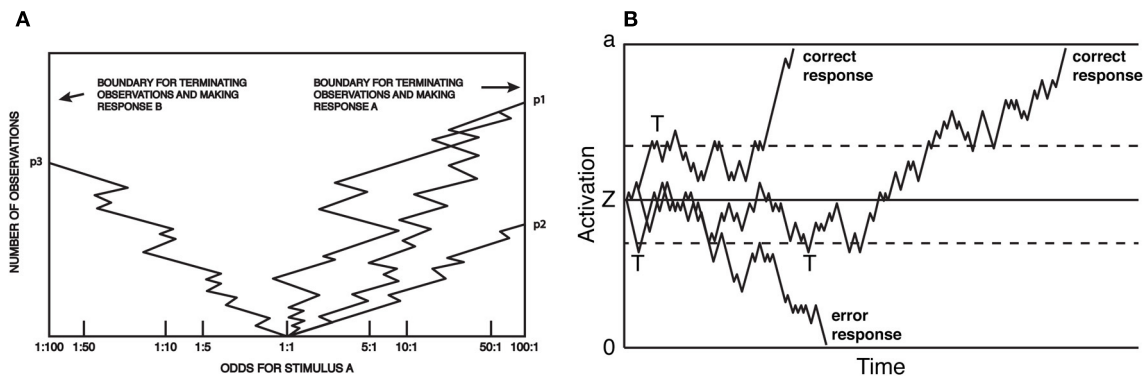
FIGURE 1 | Random-walk model of choice reaction time. (A) Each sample can be considered evidence favoring one of two options, and at each step, the observer updates an estimate of the posterior probability (here, presented as an odds ratio) based on that evidence. A response is produced at a threshold odds ratio. Reaction time is not explicit, but proportional to the total number of samples. Adapted from Fitts (1966). **(B)** The closely related diffusion model. Here, boundaries are associated with the correct or errant response and the X-axis is real-time. As in **(A)** responses are produced when activation reaches threshold, and the SAT is a function of the placement of the threshold. Adapted from Ratcliff and Rouder (1998).

Ratcliff and Rouder (1998). During a trial, subjects sample perceptual information, at each step computing a revised estimate of the likelihood of either hypothesis being true. Responses are produced when the observers' posterior probability exceeds some *threshold* odds ratio (**Figure 1A**). The same model is presented in **Figure 1B**, except that the process carries out more clearly in real time, and the response threshold is defined in an equivalent, yet more abstract dimension. **Figure 1B** illustrates how sequential sampling models implement SAT: when the decision threshold is high (solid upper and lower lines), RT tends to be longer and more likely correct, as noise in the process is allowed to average out over time. When lowered (dashed lines), the process terminates early (marked by a "T" in **Figure 1B**). This speeds RT, but also increases the probability that an error will result due to noise in the sampling process: note that the longest-latency correct response would result in an error under low but not high threshold. Moreover, the model makes very specific, empirically testable predictions about the form of the resulting RT distributions, and how they change with various manipulations. The random walk model received immediate acclaim, and was extended and revised almost immediately (Edwards, 1965; Laming, 1968).

The random walk model provided a rigorous and principled treatment of SAT, but was not favored by all. In Ollman (1966) proposed the first of what would become known as *mixture models*. Whereas sequential sampling models assume incremental evidence accumulation, Ollman suggested a mixture of dichotomous states: fast guesses and slow controlled decisions. The latency of the guess process and controlled process was assumed constant; SAT was achieved by simply changing the mixture. Note that this predicts a linear accuracy-RT relationship anchored by a theoretical *true guess RT* (corresponding to chance level accuracy) and a *true controlled RT* (corresponding to perfect accuracy). Intermediate values are simply weighted averages of the two component latencies. This *fast guess model* was tested by Yellott (1971). Subjects performed a simple color discrimination task while SAT was induced through *response deadlines*: arbitrary time limits subjects must beat in order to produce a fully correct response (see

section SAT Manipulations). The fast guess model predicts that both unknown quantities—the true guess and true controlled RT—should be invariant over deadline conditions. Yellot devised a method for estimating these latencies, and found remarkable invariance. The guess and controlled RT was constant not only across deadline conditions, but over subjects.

The idea that SAT results from a mixture of random guesses is certainly attractive from a standpoint of simplicity. It should not be controversial that subjects can, if they wish, produce a pre-selected random guess in nearly any choice task. But, there are problems with this proposal. The most obvious is the prediction that mean error RT is faster than mean correct RT. This must occur if errors are produced by guesses, which in turn are always fast. While this is a common observation (Ollman, 1966; Schouten and Bekker, 1967; Hale, 1969; Grice and Spiker, 1979), it is certainly not the rule. Further, it is likely that Yellot's color choice task may have been so simple that subjects had to begin guessing to meet the demands of the deadline manipulation. This was in fact found to be the case. One year later, Swensson (1972a) had subjects determine which of two rectangles, oriented at 45°, was longer. SAT was induced using a payoff matrix that favored accurate or fast responding. Swensson conducted a regimented trial-by-trial analysis, categorizing each as a likely guess or non-guess response. When the discrimination was simple, Swensson found data consistent with the fast guess model: subjects either used a guessing strategy or a highly accurate controlled strategy. A mixture also obtained when the discrimination was made more difficult, except for one critical detail. When the analysis was limited to non-guess trials, accuracy rate continued to vary with RT. Swensson proposed an alternative, known as the *deadline model*[5]. Like the fast guess, subjects are assumed to mix pure guesses with correct responses, but whether or not a guess is to be made is not decided prior to the trial. Instead, subjects maintain an internal

---

[5]The deadline model is usually attributed to Swensson (1972a), but it was in fact proposed earlier, both by Nickerson (1969) as well as an alternative to his own fast guess model by Yellott (1971).

timer; SAT manipulations set a limit on this timer. A response is produced either when sufficient information has been gleaned as to make a correct response, or when the deadline is reached. While intuitively appealing, the deadline model has seen little success. For one thing, one might better term the model the *slow guess*, as it predicts error RTs that are later than correct RTs—a prediction not borne out by Swensson's own data and numerous other studies (but see Estes and Wessel, 1966; Pike, 1968; Link and Tindall, 1971; Audley, 1973; Pfefferbaum et al., 1983; Ditterich, 2006a; Heitz et al., 2010). Perhaps more problematic for the deadline model—indeed all mixture models—is the observation that error RT is sometimes faster and sometimes slower than correct RT (Link and Tindall, 1971; Swensson, 1972a; Luce, 1986). Mixtures models are not flexible enough to predict both. Other efforts have rendered mixture models untenable as a sole explanation for SAT (Reed, 1973; Ruthruff, 1996; Wagenmakers et al., 2008; but see Dutilh et al., 2011; Schneider and Anderson, 2012; Donkin et al., 2013).

For several reasons, sequential sampling has emerged as the dominant decision model framework. For one, they naturally account for choice behavior under SAT without appeal to a mixture of two states, and with some assumptions, can predict either fast or slow error RT (Laming, 1968; Ratcliff and Rouder, 1998). Another is precision: they provide a quantitative account of mean correct and error RT, accuracy rate, the shapes of correct and error RT distributions, and how each of these change with experimental manipulations such as SAT, response bias, and the strength of sensory evidence. Third, they make testable predictions. For instance, when sensory evidence remains constant, there exists a unique, optimal decision threshold that maximizes reward rate (RR) (Gold and Shadlen, 2002; Bogacz et al., 2006), and humans closely match this threshold even when optimality changes between blocks of trials (Simen et al., 2009; Bogacz et al., 2010a; Balci et al., 2011). Likewise, these models can be shown to account for high-level behaviors such as visual fixations and purchasing decisions (Krajbich et al., 2010, 2012; Milosavljevic et al., 2010; Towal et al., 2013). Fourth, there is mounting evidence that something akin to sequential sampling occurs in the brain, as I will discuss later.

There exist several sequential-sampling models that embrace these strengths, notably, the Drift-Diffusion (Ratcliff, 1978; Busemeyer and Townsend, 1993; Ratcliff and Smith, 2004), Race/Accumulator (Pike, 1968; Vickers and Smith, 1985; Smith and Vickers, 1988; Logan, 2002), Leaky-Competing Accumulator (Usher and McClelland, 2001), LATER (Carpenter and Williams, 1995; Reddi and Carpenter, 2000), and Linear Ballistic Accumulator (Brown and Heathcote, 2005, 2008) among others (cf. Cisek et al., 2009; Drugowitsch and Pouget, 2012; Thura et al., 2012; Thura and Cisek, 2014). Though a full discussion is beyond the scope of this article [the reader may refer to Bogacz et al. (2006) and Ratcliff and Smith (2004)], it should be noted that nearly all assume SAT is a function of the distance (or "excursion," Churchland et al., 2008) a decision variable must travel from a start point to a threshold, sometimes called *response caution* (Forstmann et al., 2008). In many, SAT is implemented by a change in decision threshold alone (**Figure 1**). This idea has been challenged, and several efforts

now consider SAT to be a multifaceted phenomenon including changes in, for example, sensory gain (Ditterich, 2006b; Standage et al., 2011, 2013; Heitz and Schall, 2013) along with decision threshold.

SUMMARY

The SAT has long been a phenomenon of interest in behavioral science. From early on, the covariation between response speed and accuracy was seen not as a nuisance, but a signature of the decision process itself. Consequently, experimental investigations of SAT progressed largely in parallel with mathematical models of the decision process. This work is ongoing, but a consensus has emerged: agents make choices based on a sequential analysis of sensory evidence. As decades of research make clear, this decision process is adaptable: actions are dictated not only by the nature of perceptual input but also environmental constraints, internal goals, and biases. An embodiment of this flexibility, the SAT arises due to the inherent contradiction between response speed and decision accuracy. Faster responses entail less accumulated evidence, and hence less informed decisions. Sequential sampling models provide an intuitive framework for understanding SAT. Observers set a decision criterion—an amount of evidence required to commit to a choice—based on current task demands and internal goals. This begs the question: how can we know *what* decision criteria subjects employ? It would seem that without this knowledge, mean RTs and accuracy rates conflate experimental factors with strategic effects employed by the observer. The solution to this problem is to bring decision criterion[6] under experimenter control. As explained below, this not only avoids ambiguity, but also quantifies *precisely* how accuracy trades off with latency.

SAT METHODOLOGY: EXPERIMENTAL MANIPULATIONS AND ANALYSIS TECHNIQUES

A common theme in the above is the manipulation of subjects' decision criteria through experimenter influence. These *SAT experiments* quantify how accuracy covaries with RT over the range of decision criteria subjects might use. In contrast, group means obtained at a single criterion provide only a snapshot of performance that conflates decision strategy with the nature of the task (e.g., its difficulty). In other words, with decision criteria free to vary, many different group means could obtain, from very fast RT and chance accuracy to very slow RT and asymptotic accuracy. The problem is further exacerbated if the experimental conditions under comparison also encourage different SAT settings, making group means difficult to interpret and conclusions ambiguous (Wickelgren, 1977; Lohman, 1989). In this way, SAT manipulations avoid problems shared by *non-SAT experiments*, echoed in the quote that opened this work. Furthermore, deriving the pattern of performance over a variety of decision criteria, SAT experiments offer a window into the

---

[6]I use the terms SAT setting, SAT criterion, and decision criterion equivalently to refer to one's momentary willingness to trade response speed for accuracy. It is a single point along an accuracy-latency performance function (Wickelgren, 1977; Lohman, 1989). In the context of sequential sampling models, it is often referred to as decision threshold.

decision process itself. An empirical example will drive home the point.

Heitz and Engle (2007) addressed the possibility that individuals rated high or low on a measure of working memory capacity exhibit differences in processing efficiency during low-level visual (non-memory) tasks. Specifically, they proposed that those with low working memory process sensory evidence more slowly than those with high working memory capacity. To test this, high and low working memory subjects performed the Eriksen flanker task (Eriksen and Eriksen, 1974; Gratton et al., 1988). Subjects reported the identity of a central letter (H or S, mapped to key presses on different hands), each flanked on either side with response-congruent or response-incongruent stimuli. Subjects typically respond more quickly and with higher accuracy to congruent (e.g., HHHHH) than incongruent (e.g., HHSHH) strings. Heitz and Engle manipulated SAT through the use of response deadlines ranging from 200 to 700 ms. By implicating rate of perceptual accumulation, they predicted that asymptotic performance would be equivalent. That is, if given sufficient time, both groups should perform equally. This is particularly suited for SAT methodology, as obtaining group means at a single criterion would not address the question.

The data in **Figure 2A** depict accuracy rate conditionalized on RT[7] (known as a *conditional accuracy function*—a topic I will return to). The data are fit by a function known as an *exponential approach to a limit*[8], as is common (Wickelgren, 1977; McElree and Dosher, 1989; Öztekin and McElree, 2010), to obtain numerical estimates of intercept (the processing time needed to make above-chance, informed decisions), rate (gain in accuracy

with RT), and asymptote (peak accuracy). The critical pattern concerns the difference between high and low working memory groups on incompatible trials (dashed lines). It is observed that at very fast RT, both groups are equally fast and at respond at about chance level. Asymptotic accuracy also appears equivalent, suggesting that the two groups perform equally when given sufficient time. What distinguishes the groups is the rate of gain in accuracy with RT, which the authors interpreted as evidence that the groups did in fact differ in processing efficiency. The relationship is perhaps more straightforward when the negatively accelerated function is linearized using a log-odds transformation, also a common practice (**Figure 2B**). It is clear that the slope of the function relating accuracy and RT is greater for the high than low working memory group. This conclusion—quite different than the authors had expected—was made possible though SAT manipulations[9]. In sum, bringing decision criteria under experimenter control provides a detailed picture of the decision process, avoids ambiguity that may arise when SAT is not controlled, and facilitates more specific hypotheses. Numerous experimental methods accomplish this, each with strengths and weaknesses.

SAT MANIPULATIONS[10]
Verbal instructions
In the vast majority of behavioral studies, subjects are directed to maintain both high accuracy and fast RTs. This is problematic, as the two constraints are contradictory. As pointed out humorously by Edwards (1961): "These instructions are internally inconsistent. A computing machine would reject as insoluble a problem presented with such instructions" (p. 276). It is with this in mind that Howell and Kreidler carried out the first true SAT experiment (Howell and Kreidler, 1963). In a task similar to the venerable Hick paradigm (Hick, 1952), different groups of participants were asked to favor fast, accurate, or fast and accurate

---

[7]Data are collapsed over Experiments 1 and 2 of (Heitz and Engle, 2007). For fitting, initial RT-accuracy bins with chance-level responding were eliminated. The conclusions remain unaltered. See original publication for details.

[8]The exponential approach to a limit takes the form: $Acc = \lambda \left[1 - e^{-\gamma(T-\delta)} \right]$ where Acc is some measure of accuracy rate (proportion correct or d-prime), λ is asymptotic performance, γ the rate, δ the x-axis intercept, and T is RT. The use of an exponential approach to a limit has been criticized (Ratcliff, 2006) on the grounds that it is atheoretical and not necessitated by process models such as the drift-diffusion. Others might consider this a strength.

[9]For a similar application of SAT methodology to memory phenomena, see (McElree and Dosher, 1989; Kumar et al., 2008; Öztekin and McElree, 2010).
[10]Several of the below SAT methodologies were previously reviewed by Wickelgren (1977).

FIGURE 2 | Data from Heitz and Engle (2007) **Experiments 1 and 2.** **(A)** Data were fit by an exponential approach to a limit. The critical pattern concerns the incompatible condition (dashed lines). The groups do not differ in intercept or in asymptote, but do differ in rate. **(B)** The same data in **(A)** linearized using a log-odds transformation and fit with a log-linear regression.
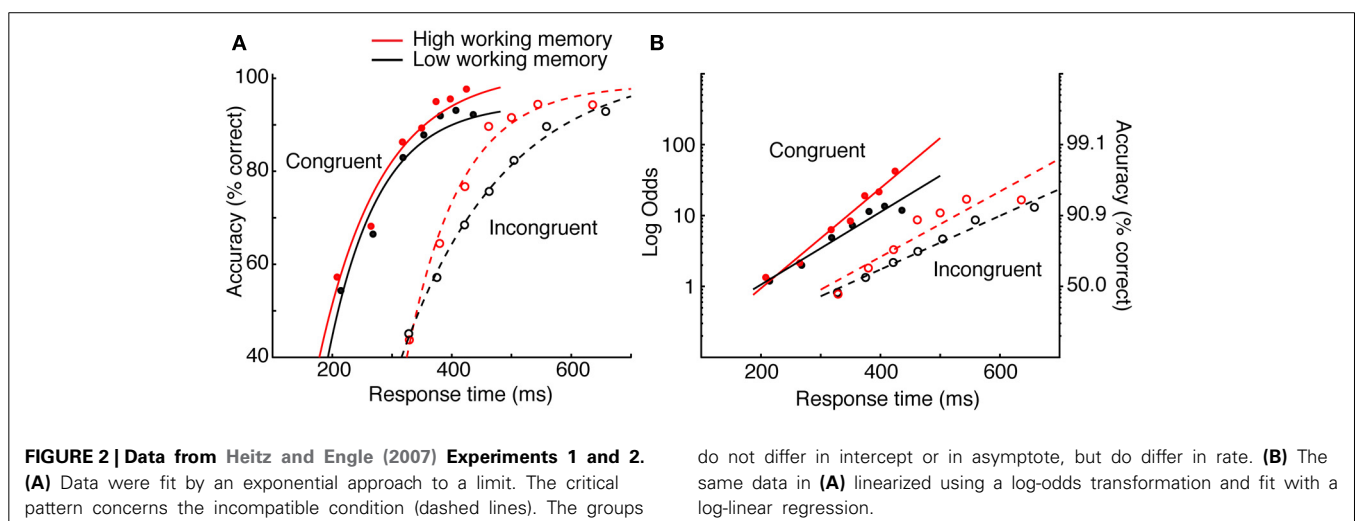
responding. In their own words, this required "...that S establish a "trade-off" between two dimensions" (p. 41). For obvious reasons, instructions remain the most common SAT manipulation: they are simple to implement, require little training, and yield large effect sizes.

Though popular, verbal instructions are not ideal in several respects. First, instructions are qualitative. It is unlikely that individuals adopt similar response criteria both within and between emphasis conditions (Lohman, 1989), which serves to both diminish effect sizes and increase experimental error (Edwards, 1961). Moreover, without a quantitative method, the potential for regression to the mean is high. Subjects may modify behavior initially, but over the course of trials in a block, settle into some less distinct mode. In fact, there is a tendency for controlled RT distributions to skew toward an individual's natural mean RT (Schouten and Bekker, 1967). Second, the number of qualitatively different emphasis conditions subjects can achieve is limited; any more than three seems difficult. This is certainly adequate for gross comparisons (e.g., Hale, 1969; Osman et al., 2000; Forstmann et al., 2008; Ivanoff et al., 2008), but may be inadequate for describing the accuracy-latency function mathematically, particularly if decision criteria are not homogenous over subjects (McClelland, 1979). Finally, and particularly important for future work, instructions are decidedly not available in non-human subject populations.

Payoffs. To combat the ambiguity of instructions, Fitts (1966) designed a payoff matrix to differentially reward correct decisions and penalize errors. Fitts defined four response categories, based on whether the response was correct and whether the RT met an arbitrary "criterion time." As shown in **Table 1**, subjects were awarded +1.0 point for fast and correct responses, and penalized −1.0 point for slow and inaccurate responses. The SAT emphasis conditions were distinguished by the penalty incurred for correct but slow or incorrect but fast responding. Under accuracy emphasis, there was a higher penalty associated with errors, whereas under speed emphasis, the penalty was greater for slow responding. This scheme worked quite well; payoff matrices induced significant covariation in RT and accuracy rate even in the absence of verbal instructions. Others have since used similar methods to manipulate SAT (Pachella and Pew, 1968; Swanson and Briggs, 1969; Lyons and Briggs, 1971; Swensson and Edwards, 1971; Gehring et al., 1993).

Payoffs have at least two advantages over verbal instructions. First, the quantitative nature of the rewards and penalties allow for a larger number of emphasis conditions. Secondly, verbal instructions become unnecessary; observers learn contingencies over the course of the experiment or in practice blocks, making this method viable for use with non-human populations. On the other hand, the payoff scheme requires one to define a "criterion time" that defines whether or not a particular response is considered fast or slow. Ideally, the criterion time is determined subject-by-subject using a data-driven method, such as some percentile of a subjects' RT distribution during the same task without time constraints. Whether arbitrary or subject-specific, the choice of the criterion time separating "fast" and "slow" RT is an important consideration, as improper values render the method ineffective. That said, some early studies have seen success using a constant, arbitrary criterion time for all subjects (Fitts, 1966; Ollman, 1966; Pachella and Pew, 1968). It is also worth noting that without additional instructions or cuing events, switching between emphasis conditions will not be immediate.

Pure payoffs. Avoiding the problem of arbitrary criterion times, Swensson designed a method making rewards and penalties linearly related to RT (Swensson and Edwards, 1971). Correct responses are rewarded $[D − k(RT)]$ and errors penalized $[−k(RT)]$. Parameter k specifies the relative gain or loss with changes in RT, while D defines the relative gain due to correct responding. When D is small, rewards and penalties are based entirely on RT; when large, the reward associated with correct responding outweighs loss due to long latency. This regime, known as "pure payoffs," has seen little use (Swensson and Edwards, 1971; Swensson, 1972a,b), but is in principle superior to a standard payoff structure. Unfortunately, it shares one weakness with the payoff matrix: learning the reward contingencies takes time, and subjects will be unable to switch between conditions immediately without ancillary cuing signals.

Deadlines. Pachella introduced a simplification of the payoff procedure described above. He demonstrated that SAT can be induced using only the criterion times that define "fast" and "slow" responses without any associated payoff matrix (Pachella and Pew, 1968; Pachella and Fisher, 1969). As is typical, a single deadline is in effect throughout a block of trials; choice latencies that do not beat the deadline are met with some tone or visual feedback to indicate the response was not fast enough[11]. Practice trials preceding each block provide an acclimation period. Numerous classic and contemporary works use this simple, highly effective manipulation (Pachella and Pew, 1968; Pachella and Fisher, 1969; Link and Tindall, 1971; Yellott, 1971; Green and Luce, 1973; Pike et al., 1974; Jennings et al., 1976; Ratcliff and Rouder, 2000; Diederich and Busemeyer, 2006; Heitz and Engle, 2007; Yamaguchi et al., 2013).

There are several considerations that warrant discussion. The first is the number of deadline conditions, which depends on both

Table 1 | Payoff matrices used by Fitts (1966) to induce SAT.

Payoff condition	Correct and fast	Correct and slow	Wrong and fast	Wrong and slow
Pretest	+1	−0.2	−0.2	−1
Speed	+1	−0.5	−0.1	−1
Accuracy	+1	−0.1	−0.5	−1

[11]How to treat "missed deadline" trials is an important issue. On one hand, it can be argued that missed deadline trials are qualitatively different from made deadline trials (e.g., subjects failed to adopt the appropriate decision criterion), and so might be eliminated. On the other hand, this leads to artificially truncated RT distributions and artifactual effects on mean RT and accuracy rate. The most conservative approach is to compute mean RT and accuracy rate for each condition as if deadlines did not exist (i.e., categorically accurate responses count as correct even when deadlines were not met). In practice, overall conclusions are robust to this choice.

the desired resolution as well as willingness to obtain increasingly more observations per subject. While as few as three are sufficient to mathematically describe the tradeoff function (McClelland, 1979), as many as 5–8 are not uncommon (Schouten and Bekker, 1967; Yellott, 1971; Jennings et al., 1976; Heitz and Engle, 2007). In regards to selecting particular deadline values, it is important to have an idea of both the mean and variance of subjects' RT during an unconstrained version of the same task. One then selects N deadlines that more than span this range. Note that spanning too large a range increases experimental complexity with diminishing returns. Deadlines that are too fast will encourage guessing, and deadlines that are too long will have little to no effect. Another concern is the order of the deadline blocks. If all subjects are presented with the same order, practice effects become confounded with SAT effects. It is desirable to present deadlines in random or pseudo-random order, ideally with multiple repetitions to account for gains in performance over time.

Deadline tracking. An even more principled method for manipulating SAT uses an adaptive tracking method coupled with a deadline procedure. Rinkenauer et al. (2004) targeted particular accuracy rates (97.5, 82.0, 66.0%) instead of RT *per se*. Accuracy rate was computed in successive blocks of trials, and deadline values increased or decreased (in 30 ms steps) accordingly. This data-driven method has the advantage of naturally accounting for practice effects, attentiveness, fatigue, etc. that may alter behavior throughout an experiment. However, because accuracy rates must be computed over sets of trials, there is considerable overhead in converging to a desired performance level. Furthermore, if practice effects are large, substantial changes in the underlying RT distributions may occur despite holding average accuracy rate constant.

Response-to-stimulus interval (RSI). In the absence of explicit SAT manipulations, subjects are thought to choose decision criteria that maximize potential reward, whether that be monetary or otherwise (Edwards, 1965; Gold and Shadlen, 2002). One's RR is simply the proportion of correct responses divided by the average length of a trial. Several factors contribute to the average length of a trial (and hence RR), including decision time, non-decision related (e.g., sensory) delays, and the interval between one's response and the beginning of the following trial (the *response-to-stimulus interval,* RSI). Recent theoretical work suggests that altering RSI should provide a means to implicitly alter one's SAT criteria (Bogacz et al., 2006). This makes intuitive sense: when RSI is long and the pace of the task is slow, the available number of decision opportunities is likely to be fewer than when RSI is short and the pace is fast. In this case, the optimal RR is attained through slow, highly accurate decision-making. Conversely, when RSI is short, the optimal RR is achieved by emitting decisions more quickly, even if many of those decisions are incorrect. This has firm empirical support: RSI manipulations lead to SAT in much the same way as conventional time limitations (Simen et al., 2009), and mathematical decision models localize the effect to decision

threshold (Simen et al., 2006; Bogacz et al., 2010a; Balci et al., 2011)[12].

The use of RSI to manipulate SAT has several advantages. First, it is divorced from any explicit time limitations and is clearly a voluntary, strategic adaptation. Second, RSI is formalized mathematically in decision models and makes contact with a theoretical literature on RR optimization. Third, RSI may be ideal for use with non-human populations. On the other hand, RSI manipulations do not take effect immediately, as observers cannot optimize decision criteria instantaneously (Simen et al., 2009; Balci et al., 2011). Even the most sensitive subjects may require as many as 20 trials before performance stabilizes, and not all subjects produce an effect (Bogacz et al., 2010a). Furthermore, the assumption that RSI operates on subjects' inherent motivation to maximize RR seems to require experimental designs that are time-limited rather than trial-limited. In practice, this point may be moot as subjects appear to remain sensitive to RSI even in fixed trial length blocks (Simen, personal communication, 4/3/2014).

Response signals. The last two methods, *response signals* and *RT Titration* are motivated by different goals. Whereas the methods above attempt to alter subjects' cognitive state, the following attempt to bring RT under experimental control while keeping SAT criteria constant. The response signal method[13] was first developed in 1973, as a direct test of the *fast guess* model (Reed, 1973). The procedure effectively prevents fast guesses by allowing subjects to respond only when cued; in this case, the disappearance of visual stimuli served as the signal. Even with fast guesses eliminated, Reed observed that accuracy rate covaried with RT, rendering the fast guess model untenable.

The strength of this method lies in the unpredictable nature of the upcoming trial. The duration of the stimulus-to-cue duration cannot be anticipated, ensuring that each trial is approached with equivalent cognitive states—exactly the opposite intention as instructions, deadlines, etc. In this case, the accuracy-latency relationship is less likely to involve strategic changes in decision criteria but rather results from the quantity of information accumulated before encountering the cue to respond. Early cues truncate processing and force a response based on partial information.

There are two weaknesses to this approach. First, for long cue delays, subjects may withhold their response when they would otherwise have emitted a choice. In sequential sampling terms, responses are obligated not by threshold crossing but by external influence, questioning its relevance to the normal choice process. (Even the deadline method allows the choice process to terminate

---

[12]Interestingly, human subjects seem to perform sub-optimally, with accuracy rates slightly too high and and mean latencies slightly too long to maximize RR (Simen et al., 2009; Bogacz et al., 2010a). Why this is so is not fully understood, but it is worth noting that humans can learn to become optimal with sufficient practice (Balci et al., 2011).

[13]There is actually an earlier example. In 1967, Schouten and Bekker presented subjects with a simple choice task and cued them to respond on the last of three acoustic "pips" (but not earlier). Critically, the duration of the stimulus-to-cue interval was blocked, such that subjects would adopt different SAT settings. In this sense it is similar to the deadline manipulation, except that early responses are not allowed.

normally on most trials.) Related to this point, the choice process has been altered such that one cannot be sure exactly *what* SAT criterion observers are using. The method simply ensures that, on average, observers use the same criterion at the beginning of each trial, or alternatively, that the criterion does not vary in any controlled way. The last method obviates this concern.

RT titration. RT Titration (Meyer et al., 1988) seeks to hold constant observers' SAT criteria trial-to-trial while ensuring subjects begin each trial as if it were a normal, no-signal, free RT task. The procedure is straightforward: subjects make choices whenever they wish, unless a response signal is encountered, at which time a response is obligated. Because many trials include no response signal, behavior on each trial is governed by the same sequential sampling process in operation during non-SAT tasks. Meanwhile, the influence of processing time on accuracy and the contribution of partial information can be gauged by those trials including a response signal. In many ways, RT Titration is superior to the response signal method, except that subjects require training in order to produce responses swiftly after encountering the relatively more rare response signal.

Methods that hold decision criteria constant (response signals and RT Titration) are fundamentally different from those that force criteria to change (instructions, deadlines, etc.). Must the form of the accuracy-latency relationship also be different? One study to test this compared the deadline and response signal methods in the same subjects during the same task (Dambacher and Hübner, 2013). Interestingly, there was surprising agreement between the two, despite a tendency for lower overall performance in the response signal method. How can there be so much agreement between such disparate methodologies? This can be explained by the constancy of the perceptual input. Whether perceptual accumulation terminates naturally due to threshold crossing or is truncated artificially by experimenter influence, the stimulus information driving the process remained constant. What *does* differ—and this may partially explain the discrepancy between the methods—is that the predictable deadline procedure allows for proactive adjustments, such as the type observed in the baseline neural firing rates in monkeys (Heitz and Schall, 2012). Additionally, the response signal method likely involves extra cognitive demand as observers must also perform signal detection.

Selecting the best SAT manipulation. All of the above methodologies are effective, but which is most appropriate? The answer is guided by at least three considerations: (1) should RT be explicitly controlled; (2) should decision criteria vary between conditions; and if so (3) must adjustments be immediate? A guide is presented in **Table 2**, but is non-exhaustive. For instance, verbal instructions might be combined with deadlines to ensure at least minimal control of mean RT (e.g., Forstmann et al., 2008), making it an instance of "explicit" RT control. Likewise, the response signal method will allow decision criteria to vary if presented in a blocked format (Schouten and Bekker, 1967).

ANALYSIS OF SPEED-ACCURACY TRADEOFF DATA

There are several methods for depicting the SAT; here I deal with the three most popular: the *speed-accuracy tradeoff function*

Table 2 | Summary of SAT methodologies.

RT control	Decision criteria	Adjustment time	Method
Indirect	Altered	Fast	Verbal instructions
Indirect	Altered	Slow	RSI
Explicit	Altered	Fast	Deadlines
Explicit	Altered	Slow	Payoffs, Pure payoffs, Deadline tracking
Explicit	Invariant	–	Response signals, RT Titration

(SATF), the *conditional accuracy function* (CAF), and the *quantile-probability plot* (QPP). To facilitate the discussion, I created a simulated SAT experiment employing three response deadlines at 225, 325, and 425 ms. The manipulation was assumed effective, with mean accuracy rates increasing linearly at 50, 70, and 90%, respectively. RT distributions for each condition were generated by drawing 10,000 observations randomly from an ex-Gaussian (van Zandt, 2000) distribution ($\sigma = 20$ ms, $\tau = 30$ ms) such that approximately 25% of all RTs fell later than the RT deadline in each condition, but these "missed deadlines" were not removed. The mean RT for error trials was set to be slightly (5 ms) faster than correct trials.

SATF

The SATF plots mean RT and accuracy rate for each SAT condition separately (**Figure 3A**). It reflects the efficacy of the experimental manipulation and quantifies how accuracy trades off with RT, on average. The SATF is robust to the variability of the component distributions: the extent to which conditions overlap has no effect, nor is it influenced by the direction of mean error RT. However, it is clear that there is considerable variation within each condition not captured by the SATF. For instance, observed RTs of ~250 ms obtain in both the shortest and middle deadlines. Are these responses qualitatively different? Restated, the question is whether or not the large-scale difference between SAT conditions (the *macro-SAT*) is due to the same factor as smaller-scale, within-condition variability (the *micro-SAT*; Pachella, 1974; Thomas, 1974; Wood and Jennings, 1976; Wickelgren, 1977; Grice and Spiker, 1979; Vickers et al., 1985). Perhaps the difference in between- and within-condition variability is only one of magnitude; the macro-SAT due to large changes in decision criteria and micro-SAT due to intrinsic variability and trial-to-trial adjustments (Ridderinkhof, 2002; Jentzsch and Leuthold, 2006).

CAF

If this were the case, it makes more sense to plot accuracy rate conditionalized on observed RT disregarding deadline condition altogether. All RT data are categorized into equal-observation quantiles, and accuracy rate is computed separately for each bin (**Figure 3B**). Though this provides a more detailed description of how accuracy trades off with RT, this *overall* CAF does not address whether similar latencies collected under different deadline conditions are psychologically equivalent. This may be accomplished by computing CAFs *individually* for each deadline condition

FIGURE 3 | Comparison of the SATF, overall CAF, and individual CAFs in the same simulated dataset. (A) The SATF is simply the mean RT and accuracy rate for each SAT condition Here, they were 225, 325, and 425 ms response deadlines. The manipulation was effective by definition, yielding accuracy rates of 50, 70, and 90%, respectively. Each condition was constructed by drawing $N = 10,000$ observations randomly from an ex-Gaussian distribution with parameters indicated in figure inset. Solid histograms depict correct trials, open histograms error trials. **(B)** The same data as **(A)** but aggregated disregarding SAT condition and plotted as a CAF. **(C)** Same data as **(A,B)** but CAFs computed separately for each deadline condition.

(**Figure 3C**). If the micro- and macro-SAT have the same source, the SATF, CAF, and individual CAFs should be overlapping (but see Grice and Spiker, 1979).

This can and does occur—two examples are presented in **Figure 4**—but it is perhaps more common to find that they disagree. The reason for this becomes apparent when two parameters are varied—the extent of overlap between the RT distributions and the direction of mean error RT. To demonstrate, I repeated the simulation described above while manipulating both the variability (and tail) of the RT distributions and the direction of mean error RT, being faster, equal, or slower than mean correct RT (Wood and Jennings, 1976). The results are presented in **Figure 5**. In the top row (**Figures 5A–C**), the standard deviation of the distributions is kept small, so as to include little overlap between the SAT conditions. In this unrealistic situation, the overall CAF is a fair representation of the SATF, but the individual CAFs may be decreasing (**A**), flat (**B**), or increasing (**C**) depending on the direction of mean error RT. It is straightforward to understand why: when error RTs are slower than correct RTs, early quantile bins necessarily contain more correct than error responses (**A**). If mean RTs are equal (**B**), each bin will on average contain the same number of errant and error-free trials. Finally, when mean error RT is faster than correct (**C**), early bins will tend to be less accurate, and later bins more accurate. The pattern is exaggerated in the more realistic situation of extensive overlap between RT distributions (**Figures 5D–F**). In this case, neither the overall CAF nor individual CAFs approximate the SATF. It would seem that the CAFs are unpredictable and dominated by the simple direction of mean error RT. This is true, but beside the point. While all sequential sampling models predict an increasing SATF, the form of the micro-SAT differs. For instance, the original random walk model (Stone, 1960) predicts flat CAFs, since correct and error RT are equivalent (Pachella, 1974). In contrast, some accumulator models (Vickers et al., 1985) and the random walk with collapsing bounds (where threshold decreases over time) predict decreasing or inverted "U" shape CAF (Pike, 1968). Increasing

CAFs are consistent with several models, including the fast guess (Pachella, 1974), variable criterion model (Grice et al., 1977), some versions of the random walk (Laming, 1968; Vickers et al., 1985), and others.

Quantile probability plots

Combining aspects of both the SATF and CAF is the *quantile probability plot* (Audley and Pike, 1965; see also Ratcliff and Tuerlinckx, 2002). The SATF and CAF describe changes in accuracy rate with RT, but do not depict distributional characteristics, aspects that are particularly important in evaluating the fit of mathematical decision models (Audley and Pike, 1965; Pike, 1968). The drift-diffusion model, for instance, makes quantitative predictions regarding the shape of correct and error RT distributions; the QPP describes this information succinctly. For each condition, RT quantiles are calculated separately for correct and error trials, commonly at the 10, 30, 50, 70, and 90th percentiles. The RT corresponding to these quantiles are then plotted against response probability for each condition. For instance, if the accuracy rate for a particular condition was 80%, the RT quantiles for correct trials would be plotted at 0.8, and corresponding error trials at $1.0 - 0.8 = 0.2$. Under most circumstances, points to the left of 0.5 represent error trials and those to the right of 0.5, correct trials (but see Simen et al., 2009). A typical QPP computed on SAT data from a single (non-human primate) subject (Heitz and Schall, 2012) is shown in **Figure 6A**. Several characteristics are apparent. First, both accuracy rate and RT tend to increase from a Speed emphasis condition to an Accuracy emphasis condition, giving the QPP a "U" shape. This convexity is diagnostic: in sequential sampling models such as the drift-diffusion, increasing decision bounds lead to a slowing of RT with an increase in accuracy rate. In contrast, a concave QPP indicates that accuracy rate is improving while RT becomes faster, a common occurrence when signal quality is manipulated (Ratcliff and Smith, 2010). Second, error RT tends to be longer than correct RT. The difference is small in the Speed and Neutral conditions (note the one
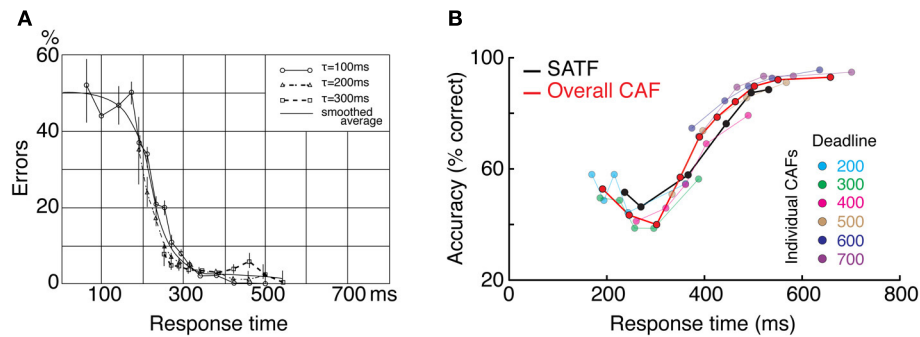
FIGURE 4 | Two empirical examples when the CAF—both the overall CAF and individual CAFs overlap with the SATF. (A) (Schouten and Bekker, 1967) forced subjects to respond to respond at target RTs during a simple two-choice response time experiment. They found that the individual CAFs overlapped significantly; the accuracy rate associated with a given RT was invariant with respect to the forced response time condition. The overall CAF and SATF are approximated by the black ogive running through individual points. Data were traced using graphics software from the original work. Note that error rate (rather than accuracy rate) is plotted on the y-axis. **(B)** (Heitz and Engle, 2007) presented subjects with a two choice response compatibility experiment under 6 response deadlines. These data, replotted from their incompatible condition, clearly indicate gross agreement between the SATF (black), overall CAF (red), and individual CAFs (colored lines). Based on this agreement, these authors used the overall CAF as their primary measure to retain time resolution.
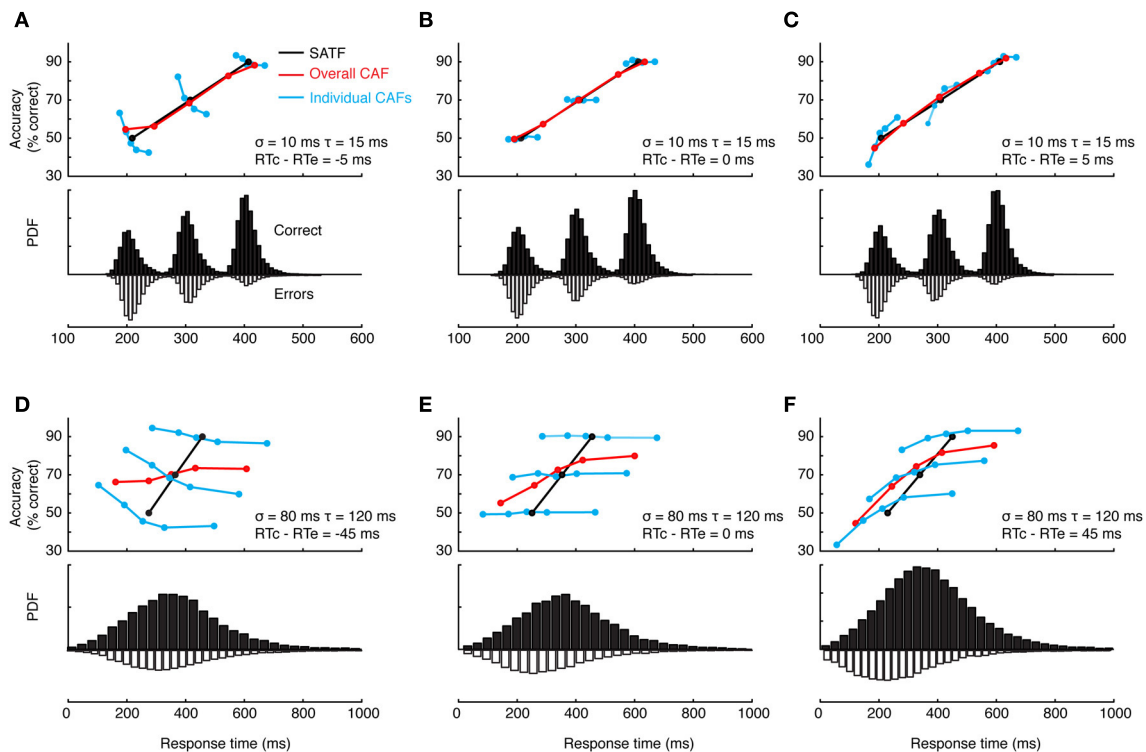
FIGURE 5 | Dependence of the CAF on component RT variance and direction of mean vs. correct RT. (A–C) With small RT variability, distributions exhibit little overlap, leading the overall CAF (red lines) to be a fair representation of the SATF with better resolution in time. The form of the individual CAFs (blue) are dictated by the direction of correct and error RT, exhibiting a downward trend for slow errors **(A)** a flat line with equal mean correct and error RT **(B)** and an upward trend for fast errors **(C)**. **(D–F)** The mismatch between SATF, overall CAF, and individual CAFs is exaggerated with more reasonable parameters. When RT distributions significantly overlap, the overall CAF no longer reflects the SATF.

point in the Neutral condition not following this trend), but quite large in the Accuracy condition. Third, the spread of the RT distributions increase with SAT stress, as might be expected given the large changes in mean RT. Fourth, the distribution of error

RTs appears roughly equivalent to correct RTs in the Speed and Neutral conditions, but noticeably larger for error trials in the Accuracy condition, particularly in the tail. The QPP provides a wealth of information absent in the SATF and CAF, yet they are

FIGURE 6 | Quantile-probability plots. (A) The QPP calculated from a single non-human primate during an SAT task. Open points to the left of 0.5 correspond to errors, closed points to the right of 0.5 are correct trials. Each vertically oriented set of 5 points mark the RT quantiles described in the text.

Lines connect quantiles between SAT conditions (red = Accuracy stress, black = Neutral, and green = Speed stress). **(B)** The QPP calculated from the same simulated dataset presented in **Figure 3**. The individual-condition RT distributions (**Figure 3C**) are reflected in the quantiles of the QPP.

related. **Figure 6B** displays this relation using the same simulated data as that of **Figure 3**.

Selecting the best analysis technique

There is no one best depiction of SAT, as each of the methods described above present different information, but there are guidelines. The SATF is the most common and straightforward approach, assuming only that the experimental design included some type of SAT manipulation. The QPP provides further detail, but requires a more sizeable dataset: estimation of RT quantiles becomes unreliable when trial counts are low, and this can be particularly problematic when errors are rare. The QPP has the additional benefit of being closely related to mathematical decision models, but less clearly depicts the rate of gain in accuracy with RT.

Overall CAFs, computed across an entire dataset, are only appropriate in specific situations. First, in the context of non-SAT experiments, the CAF might be computed to evaluate subjects' natural tendency to trade speed for accuracy (Lappin and Disch, 1972a,b) and is indeed the only available option. Second, when the CAF and SATF are overlapping, the former leads to the same conclusion as the latter while providing slightly more resolution on the RT axis (**Figure 4**). Individual-condition CAFs are useful in assessing the direction of error RT on a fine scale, but are rarely used as a sole dependent measure.

Summary

The use of SAT methodology continues to offer insight into the decision process, and how that process is altered strategically. The above provide numerous routes for obtaining and depicting the SAT. Unfortunately, SAT experiments are costly relative to non-SAT experiments, most requiring several times the number of observations. Is this gain in precision really worth the investment? In what follows, I briefly review domains outside of cognitive psychology where this has proven true.

APPLICATIONS OF SAT METHODOLOGY

NEURAL ACTIVITY UNDER SAT

A fundamental question in cognitive neuroscience concerns how the brain adapts to bring about strategic changes in decision criteria. The SAT is pervasive, and behavioral changes often large; certainly brain activity must manifest a signature of SAT. The answer to this question offers insight into the neural basis of an elementary cognitive operation, and also bears on the viability of mathematical decision models.

The sequential sampling framework described earlier has recently graduated from an abstract cognitive model to an assumed neural reality—a viable method the brain may use to carry out perceptual decision-making. Evidence supporting this claim derives from several sources, including human fMRI (Heekeren et al., 2004) and EEG (Ratcliff et al., 2009; O'Connell et al., 2012; van Vugt et al., 2012; Kelly and O'Connell, 2013), but by far the most convincing stems from single-neuron recordings in non-human primates. In the typical paradigm, monkeys view a display of static or dynamic stimuli that requires a perceptual discrimination and subsequent choice between alternative actions. Their decision is communicated through an eye movement or button press, and juice reward is delivered when the response is correct. Strikingly, activity in frontal eye field (Hanes and Schall, 1996; Kim and Shadlen, 1999; Woodman et al., 2008; Ding and Gold, 2012), lateral intraparietal area (Roitman and Shadlen, 2002; Gold and Shadlen, 2007), and superior colliculus (Horwitz and Newsome, 1999; Ratcliff et al., 2007) exhibits patterns closely resembling the sequential sampling process. Most germane is the fact that neural activity grows over time during the deliberation period and terminates at a fixed threshold at the moment an overt decision is produced. In accordance with the model, much of the variability in RT can be accounted for by the duration of the firing rate excursion—the time taken to ramp from a baseline to a fixed threshold. Further lending credence, computational (Ditterich, 2006a; Purcell et al., 2010, 2012; Zandbelt et al., 2014) and neural network models (Lo and Wang, 2006; Wong et al., 2007; Beck et al., 2008; Wang, 2008;

Zhang and Bogacz, 2010; Drugowitsch et al., 2012) inspired by the sequential sampling process capture both behavior and neurophysiology while respecting biological constraints. The neural activity associated with SAT is thus a topic of great concern, and has been examined using several techniques.

fMRI

A number of studies have used fMRI to examine neural activity during SAT manipulations. Though an fMRI approach to SAT suffers in several respects (Stark and Squire, 2001; Logothetis, 2008; Bogacz et al., 2010b), it is notable that all agree on at least one conclusion: SAT manipulations affect more than decision threshold. In fact, the most consistent finding is that relative to accuracy emphasis, placing subjects under speed stress leads to an increase in the BOLD response during *baseline* intervals (Forstmann et al., 2008; Ivanoff et al., 2008; van Veen et al., 2008; Bogacz et al., 2010b). This would seem to be interpretable within the sequential sampling framework by positing that baseline shifts are functionally identical to threshold shifts—either ultimately affects the amount of information accumulated prior to decision[14]. More interesting is the observation that more than one factor changes under SAT; at least one fMRI study has implicated changes in sensory processing with SAT (Ho et al., 2012). Further complicating the story, SAT manipulations appear to affect BOLD in region-specific ways (Vallesi et al., 2012), sometimes in opposing directions (Blumen et al., 2011). This calls into question the generality of the process: does sensory integration occur simultaneously and interactively amongst brain regions, or is there independence among sites of integration (Zhang, 2012)?

EEG

Unlike fMRI, EEG does not suffer from temporal blurring, but does not offer opportunity to definitively localize brain regions. Despite this, EEG components accurately track attention and error monitoring (Woodman and Luck, 1999; Heitz et al., 2010; Godlove et al., 2011), the chronometry of action preparation (Gratton et al., 1988), and the temporal evolution of the decision process (O'Connell et al., 2012; Kelly and O'Connell, 2013; van Vugt et al., 2014). In one early study, Gehring et al. (1993) examined the error-related negativity (ERN) under SAT using a deadline procedure. The ERN is a fronto-central negativity that appears in the moments surrounding error commission (Nieuwenhuis et al., 2001) and is though to reflect the error monitoring process. When accuracy was emphasized, the magnitude of the ERN was greater than under speed stress, when errors mattered less. This finding suggests that in addition to altering the decision process, SAT affects post-decision processing as well.

Several other studies sought to identify the processing stage locus of SAT: does speed stress affect early sensory processing or later decision and motor processing? Unfortunately, this issue remains unresolved. The first attempt to address this—in fact the first study to record neural activity under SAT—used the P3 component during a line length discrimination task under speed or accuracy emphasis (Pfefferbaum et al., 1983). The latency of the P3, thought to mark the completion of stimulus processing, was

earlier under speed than accuracy stress, suggesting that early perceptual processing was indeed facilitated. The next attempts used the *lateralized readiness potential* (LRP), a component that tracks the evolution of motor preparation. Two studies using the LRP have concluded that SAT manipulations do not affect sensory processing (Osman et al., 2000; van der Lubbe et al., 2001; see also Wenzlaff et al., 2011), while a third demonstrated that it affects both early and late processing stages (Rinkenauer et al., 2004).

Each of the above studies examined the average EEG component time-locked to some event of interest, but there is much more information in the raw signal than is immediately apparent. Understanding this, at least one study has examined the effect of SAT on the EEG frequency spectra (Pastötter et al., 2012). Using a two-choice discrimination task, subjects were cued trial-by-trial to emphasize speed or accuracy. They found that, during the baseline interval in which SAT emphasis was cued, the EEG tended to oscillate more in the lower frequency bands (4–25 Hz) under accuracy emphasis than speed emphasis (see also van Vugt et al., 2012; Heitz and Schall, 2013).

Single-unit neurophysiology

To date, there has been only one single-unit recording study employing SAT manipulations (Heitz and Schall, 2012). Monkeys were trained to perform saccade visual search under Accuracy, Neutral, or Speed emphasis, cued by the color of a fixation point. Meanwhile, neural activity was recorded from the frontal eye field, a key region in the planning and execution of eye movements. The results were diverse but can be described succinctly: SAT cues affected several stages of information processing, and speed stress generally amplifies neural activity rather than attenuate it. This was most evident for baseline neural activity (increasing under Speed stress during the pre-trial interval) and in the sensitivity of neurons to visual stimulation (responding more vigorously under Speed stress). This indicates that SAT emphasis affects perceptual processing, a suggestion that has recently gained support (Standage et al., 2011; Ho et al., 2012; Thura et al., 2012; Dambacher and Hübner, 2014; Rae et al., 2014). Surprisingly, neural threshold—the level of activity reached at saccade decision—was greater for speed than accuracy emphasis, opposite the assumption of sequential sampling models. In further analyses, it was shown that SAT affects much more than the firing rates of neurons, including the extent to which single neurons were coupled with their surrounding neural network (spike-field coherence), as well as the sensitivity of that network (Heitz and Schall, 2013).

Summary

The coupling of SAT methodology and neuroscience techniques has the potential to offer real insight into the neural mechanisms supporting decision. The consensus emerging suggests that SAT is a multifaceted phenomenon, influencing several components of the decision process and accompanied by distinct changes in brain activity. It is interesting to suppose that external changes in brain function—due to drugs, pathology, and senescence— might lead to distinct declines in cognitive performance. SAT methodology will be particularly useful in pinpointing the locus of the deficit. The next section reviews this modest, but promising literature.

---

[14]It is worth mentioning that the brain entails no such equivalence.

SAT WITH DRUGS AND PATHOLOGY

Cognitive impairments often accompany drug use, disease, injury, and pathology. For instance, individuals with schizophrenia and certain types of brain injuries exhibit impulsive, perseverative behavior on measures such as the Wisconsin Card Sort and anti-saccade tasks (Guitton et al., 1985; Fukushima et al., 1988; Kane and Engle, 2002; Thakkar et al., 2011; Cutsuridis et al., 2014). Likewise, monkeys permitted to self administer cocaine over long periods of time demonstrate increased impulsivity and reduced ability to switch between task sets (Liu et al., 2008, 2009). In contrast, aging is associated with lower performance and longer latencies (Salthouse, 2012), some of which is thought to be a "general slowing" of cognition (Kail, 1991). Do these populations simply differ in decision criteria, or has the information processing system been affected, and if so, how? A handful of studies have employed SAT methodology to address these questions.

Drugs

There have been few studies of SAT under the influence of controlled substances. The most extensively tested is the effect of alcohol. SAT was manipulated using instructions (Tiplady et al., 2001) or response deadlines (Jennings et al., 1976; Rundell and Williams, 1979) while subjects were given graded doses of alcohol and asked to perform auditory or visual discrimination tasks. In each case, alcohol reduced the slope of the SATF in a dose-dependent manner. As was the case of high and low working memory capacity described earlier (**Figure 2**), this suggests a reduction in the rate of information processing. In a more recent study, subjects performed dot motion discrimination under placebo, moderate dose, or high dose of alcohol. No SAT manipulation was included. Application of the drift-diffusion model localized the effects of alcohol to two components: drift rate and non-decision time, suggesting that perceptual accumulation was both degraded and delayed with increased intoxication (van Ravenzwaaij et al., 2012).

In other work, monkeys administered graded doses of the NMDA antagonist ketamine demonstrated both slower and more accurate performance during visual search, indicating that decision criteria may have been altered (Shen et al., 2010). Finally, a few studies have assessed the effects of stimulants on information processing, but results are inconclusive. In one, low doses of nicotine administered to non-smokers was found to benefit information processing in the absence of any SAT (Le Houezec et al., 1994). In another, the dopamine agonist bromocriptine was found to have no effect (Winkel et al., 2012) while other work suggests the dopamine reuptake inhibitor methylphenidate alters decision criteria but does not benefit information processing (Carlson et al., 1991).

Pathology and age

Research dealing with patient populations suggests a deficit in the information processing system itself rather than non-optimal decision criteria. In schizophrenics for instance, at least one modeling study suggests that relative to controls, patients suffer from increased sensory noise (Cutsuridis et al., 2014) and one explicit SAT study provides anecdotal support (Schweitzer and Lee, 1992). Similar conclusions are reached for Parkinson's Disease patients

(Wylie et al., 2009). Interestingly, the situation is quite different for one patient group of particular interest: attention-deficit hyperactivity-disorder (ADHD). Relative to controls, ADHD subjects exhibit SATFs that are shifted, but not different in slope (Sergeant and Scholten, 1985a,b) suggesting that the rate of information processing is equivalent. Recent work suggests that ADHD patients instead have a relative inflexibility in optimizing decision criteria (Mulder et al., 2010; but see Metin et al., 2013).

There is a well-characterized decline in cognitive functioning with age, but exactly what component of the decision process is altered remains unclear. On one hand, older adults tend to be more considered and cautious in their responding (Forstmann et al., 2011), suggesting a tendency to use higher decision criteria than their younger counterparts. Indeed, modeling studies suggest that older adults fail to set decision criteria optimally, often preferring overall accurate performance at the cost of speed (Phillips and Rabbit, 1995; Ratcliff et al., 2004; Starns and Ratcliff, 2010, 2012). Empirical studies using SAT methodology corroborate this, but also provide compelling evidence for an impairment in information processing (Salthouse, 1979; Madden and Allen, 1991; Hertzog et al., 1993; Kumar et al., 2008) see also (Myerson et al., 2007).

Summary

Though the cognitive impairments accompanying drug use, pathology, and age are well characterized, the underlying basis remains elusive. Traditional experimental approaches cannot dissociate performance changes due to strategic effects (e.g., preference for fast than accurate decisions) from those due to information processing *per se* (e.g., compromised perceptual sampling). By placing SAT criteria under experimental control, the true nature of the deficit becomes clear. Further research will be enlightening, and may be the key to developing targeted interventions.

SAT IN NON-HUMAN ORGANISMS

The present work has primarily dealt with human behavior; in stark contrast, this final section reviews a handful of studies assessing SAT in non-human populations (monkeys, rodents, and insects). This short discussion has two motivations. First, I wish to promote the use of SAT methodology in populations amenable to single-unit recordings. Neuroscience approaches continue to elucidate the decision process with unparalleled detail, and single-unit recordings are arguably the most definitive. This effort has been limited by the absence of methods for controlling decision criteria in non-human populations; here I show it is possible. Second, I wish to illustrate that the SAT is truly universal. Unlike humans and monkeys (and probably rodents), social insects also exhibit SAT, but in a very different way. Specifically, the decision to be made is one involving a colony, rather than a single member. Likewise, whereas many individual neurons contribute to a single decision in higher species, many individual entities contribute to a group decision in insect colonies. Whether or not these phenomena are comparable remains to be seen, but important parallels exist.

MONKEYS

There has been only one study using experimenter-induced SAT in monkeys (Heitz and Schall, 2012). Monkeys performed saccade visual search and were induced to respond at three levels of SAT emphasis: speed, neutral, and accuracy. Conditions were signaled by the color of a fixation point and presented in blocks of 10–20 trials. Emphasis conditions were defined by differential reward and punishment (time-out) contingencies, and monkeys were trained until they adapted behavior immediately upon presentation of a new emphasis condition. In several ways, the SAT in monkeys is identical to that in humans: the SATF is increasing, and the behavior is well fit by sequential sampling models with changes in decision threshold between emphasis conditions. There are slight differences, however. Whereas humans most commonly exhibit fast errors during visual search, monkeys tend to commit slow errors, leading to a decreasing (rather than increasing) CAF. Interestingly, this occurs even in tasks that do not include SAT manipulations, such as standard form visual search (Heitz et al., 2010) and the venerable random dot motion paradigm (Roitman and Shadlen, 2002; Ditterich, 2006a; Churchland et al., 2008). The origin of this disparity is not understood, but has not been systematically studied.

RODENTS

Evidence for SAT in rodent models is mixed. Using olfactory discrimination, one study has shown a lack of any relationship between accuracy rate and decision time, even when odor mixtures are highly similar (Uchida and Mainen, 2003; see also Zariwala et al., 2013; but see Abraham et al., 2004). However, a different conclusion emerges when the stimulus-sampling period is placed under experimenter control. One such study used an analog of the response signal method. During an olfactory discrimination task, mice were required to continue sniffing until an auditory buzzer signaled the availability of reward (Rinberg et al., 2006). The resulting SATFs were undeniably similar to that of humans. Moreover, the slope of the accuracy-latency relationship was altered by task difficulty: when odors were highly similar, the rate of gain of accuracy with RT was much lower than for highly dissimilar, and therefore easier, discriminations (see also Brunton et al., 2013).

INSECTS

There is some evidence for SAT in bumblebees trained to perform a type of visual search task: bees are rewarded with sucrose for choosing to land on a target "flower" presented amongst distractors. Commonly, the flowers are distinguishable through color, but other times through scent. Like humans, bees produce linear speed-accuracy relationships (Chittka et al., 2003; Kulahci et al., 2008; Riveros and Gronenberg, 2012). Those that decide more slowly tend to be more accurate than those that respond quickly. Also like humans, changing task parameters can lead to shifts of the accuracy-latency function. For instance, when errant choices are met with punishment (quinine solution), individual bees slow down and increase accuracy relative to conditions without penalty (Chittka et al., 2003). Other manipulations that lead to SAT in bees include difficulty of discrimination (Dyer and Chittka, 2004; Skorupski et al., 2006; Riveros and Gronenberg, 2012) and the

introduction of environmental stressors such as predation risk (Ings and Chittka, 2008).

Like many social insect colonies, bees choose nesting cites based on quorum sensing (Seeley and Visscher, 2004). Briefly, scout bees examine potential locations for hives and recruit others; the colony as a whole "decides" to migrate to the nest when a quorum threshold (QT) has been reached (Passino et al., 2008). It is interesting to note the parallel between the QT and the decision threshold described by sequential sampling models. Under a lower QT, fewer bees contribute to the choice of nesting cite increasing the potential for err. A computational model of bee quorum sensing confirms that changing the QT (the number of bees needed to commit to the new hive) implements SAT in an ecologically valid way (Passino and Seeley, 2006).

I am not aware of any empirical study testing this assertion in bees, but it is certainly true for ants. Like bees, ants that have found a potential nesting cite recruit others until a QT is reached. At threshold, the colony switches from individual exploration into a mode of "social carrying" in which ants pick up and carry other ants to the new cite. The SAT becomes evident when the QT is examined under different conditions. For instance, ant colonies lower their QT when placed in a harsh environment necessitating migration (Franks et al., 2003, 2009), relative to a calm environment. Similarly, QT is lowered when nests are destroyed, leading to emergency migration (Dornhaus et al., 2004); see also (Marshall et al., 2006). Interestingly, this reduction in QT has the consequences expected with SAT—faster, but less discerning migration decisions.

SUMMARY

The SAT is a truly universal phenomenon. Monkeys and rodents can be trained to vary decision criteria on cue, and exhibit behavior similar to humans. Future studies employing SAT methodology with these populations will provide critical insight into the decision process. There are parallels, too, with social insect colonies, and this has not gone unnoticed. These ecologically valid studies speak to the mechanisms of emergent behavior through the interaction of individual entities.

CONCLUSION

The SAT has been a topic of great concern for over a century. Throughout its history and still today, the SAT remains an integral component of empirical, theoretical, and mathematical explorations of the decision process. The growing popularity of SAT in the neuroscience community is particularly exciting. The last decade has witnessed incredible advances in our understanding of the neural basis of choice, and neural investigations of SAT are now gaining momentum. This work promises to detail the choice process—not just in humans but non-humans as well—and will find utility in understanding and treating common cognitive deficits. Clearly, there is much work to be done. To facilitate this, and to bring together disparate literatures and disciplines, the present work reviewed the history, methodology, physiology, and behavior associated with SAT.

REFERENCES

Abraham, N. M., Spors, H., Carleton, A., Margrie, T. W., Kuner, T., and Schaefer, A. T. (2004). Maintaining accuracy at the expense of speed: stimulus similarity defines odor discrimination time in mice. *Neuron* 44, 865–876. doi: 10.1016/j.neuron.2004.11.017

Audley, R. J. (1957). A stochastic description of the learning behaviour of an individual subject. *Q. J. Exp. Psychol.* 9, 12–20. doi: 10.1080/17470215708416215

Audley, R. J. (1958). The inclusion of response times within a stochastic description of the learning behavior of individual subjects. *Psychometrika* 23, 25–31. doi: 10.1007/BF02288976

Audley, R. J. (1960). A stochastic model for individual choice behavior. *Psychol. Rev.* 67, 1–15. doi: 10.1037/h0046438

Audley, R. J. (1973). "Some observations on theories of choice reaction time: tutorial review," in *Attention and Performance IV*, ed S. Kornblum (New York, NY: Academic Press), 509–545.

Audley, R. J., and Jonckheere, A. R. (1956). The statistical analysis of the learning process. *Br. J. Math. Stat. Psychol.* 9, 87–94. doi: 10.1111/j.2044-8317.1956.tb00176.x

Audley, R. J., and Pike, A. R. (1965). Some alternative stochastic models of choice. *Br. J. Math. Stat. Psychol.* 18, 207–225. doi: 10.1111/j.2044-8317.1965.tb00342.x

Balci, F., Simen, P., Niyogi, R., Saxe, A., Hughes, J. A., Holmes, P., et al. (2011). Acquisition of decision making criteria: reward rate ultimately beats accuracy. *Atten. Percept. Psychophys.* 73, 640–657. doi: 10.3758/s13414-010-0049-7

Beck, J. M., Ma, W. J., Kiani, R., Hanks, T., Churchland, A. K., Roitman, J., et al. (2008). Probabilistic population codes for Bayesian decision making. *Neuron* 60, 1142–1152. doi: 10.1016/j.neuron.2008.09.021

Becker, G. M. (1958). Sequential decision making: wald's model and estimates of parameters. *J. Exp. Psychol.* 55, 628–636. doi: 10.1037/h0046519

Blumen, H. M., Gazes, Y., Habeck, C., Kumar, A., Steffener, J., Rakitin, B. C., et al. (2011). Neural networks associated with the speed-accuracy tradeoff: evidence from the response signal method. *Behav. Brain Res.* 224, 397–402. doi: 10.1016/j.bbr.2011.06.004

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., and Cohen, J. D. (2006). The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol. Rev.* 113, 700–765. doi: 10.1037/0033-295X.113.4.700

Bogacz, R., Hu, P. T., Holmes, P. J., and Cohen, J. D. (2010a). Do humans produce the speed-accuracy trade-off that maximizes reward rate? *Q. J. Exp. Psychol.* 63, 863–891. doi: 10.1080/17470210903091643

Bogacz, R., Wagenmakers, E.-J., Forstmann, B. U., and Nieuwenhuis, S. (2010b). The neural basis of the speed-accuracy tradeoff. *Trends Neurosci.* 33, 10–16. doi: 10.1016/j.tins.2009.09.002

Brown, S., and Heathcote, A. (2005). A ballistic model of choice response time. *Psychol. Rev.* 112, 117–128. doi: 10.1037/0033-295X.112.1.117

Brown, S. D., and Heathcote, A. (2008). The simplest complete model of choice response time: linear ballistic accumulation. *Cogn. Psychol.* 57, 153–178. doi: 10.1016/j.cogpsych.2007.12.002

Brunton, B. W., Botvinick, M., and Brody, C. (2013). Rats and humans can optimally accumulate evidence for decision-making. *Science* 340, 95–98. doi: 10.1126/science.1233912

Busemeyer, J. R., and Townsend, J. T. (1993). Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. *Psychol. Rev.* 100, 432–459. doi: 10.1037/0033-295X.100.3.432

Carlson, C. L., Pelham, W. E., Swanson, J. M., and Wagner, J. L. (1991). A divided attention analysis of the effects of methylphenidate on the arithmetic performance of children with attention-deficit hyperactivity disorder. *J. Child Psychol. Psychiatry* 32, 463–471. doi: 10.1111/j.1469-7610.1991.tb00324.x

Carpenter, R. H., and Williams, M. L. (1995). Neural computation of log likelihood in control of saccadic eye movements. *Nature* 377, 59–62. doi: 10.1038/377059a0

Chittka, L., Dyer, A. G., Bock, F., and Dornhaus, A. (2003). Psychophysics: bees trade off foraging speed for accuracy. *Nature* 424, 388. doi: 10.1038/424388a

Christie, L. S., and Luce, R. D. (1956). Decision structure and time relations in simple choice behavior. *Bull. Math. Biophys.* 18, 89–112. doi: 10.1007/BF02477834

Churchland, A. K., Kiani, R., and Shadlen, M. N. (2008). Decision-making with multiple alternatives. *Nat. Neurosci.* 11, 693–702. doi: 10.1038/nn.2123

Cisek, P., Puskas, G. A., and El-Murr, S. (2009). Decisions in changing conditions: the urgency-gating model. *J. Neurosci.* 29, 11560–11571. doi: 10.1523/JNEUROSCI.1844-09.2009

Cutsuridis, V., Kumari, V., and Ettinger, U. (2014). Antisaccade performance in schizophrenia: a neural model of decision making in the superior colliculus. *Front. Neurosci.* 8:1–13. doi: 10.3389/fnins.2014.00013

Dambacher, M., and Hübner, R. (2013). Investigating the speed-accuracy trade-off: better use deadlines or response signals? *Behav. Res. Methods* 45, 702–717. doi: 10.3758/s13428-012-0303-0

Dambacher, M., and Hübner, R. (2014). Time pressure affects the efficiency of perceptual processing in decisions under conflict. *Psychol. Res.* doi: 10.1007/s00426-014-0542-z. [Epub ahead of print].

Diederich, A., and Busemeyer, J. R. (2006). Modeling the effects of payoff on response bias in a perceptual discrimination task: bound-change, drift-rate-change, or two-stage-processing hypothesis. *Percept. Psychophys.* 68, 194–207. doi: 10.3758/BF03193669

Ding, L., and Gold, J. I. (2012). Neural correlates of perceptual decision making before, during, and after decision commitment in monkey frontal eye field. *Cereb. Cortex* 22, 1052–1067. doi: 10.1093/cercor/bhr178

Ditterich, J. (2006a). Stochastic models of decisions about motion direction: behavior and physiology. *Neural Netw.* 19, 981–1012. doi: 10.1016/j.neunet.2006.05.042

Ditterich, J. (2006b). Evidence for time-variant decision making. *Eur. J. Neurosci.* 24, 3628–3641. doi: 10.1111/j.1460-9568.2006.05221.x

Donders, F. C. (1868). On the speed of mental processes. *Arch. Néerland.* 3, 269–317.

Donkin, C., Nosofsky, R. M., Gold, J. M., and Shiffrin, R. M. (2013). Discrete-slots models of visual working-memory response times. *Psychol. Rev.* 120, 873–902. doi: 10.1037/a0034247

Dornhaus, A., Franks, N. R., Hawkins, R. M., and Shere, H. (2004). Ants move to improve: colonies of Leptothorax albipennis emigrate whenever they find a superior nest site. *Anim. Behav.* 67, 959–963. doi: 10.1016/j.anbehav.2003.09.004

Drugowitsch, J., Moreno-Bote, R., Churchland, A. K., Shadlen, M. N., and Pouget, A. (2012). The cost of accumulating evidence in perceptual decision making. *J. Neurosci.* 32, 3612–3628. doi: 10.1523/JNEUROSCI.4010-11.2012

Drugowitsch, J., and Pouget, A. (2012). Probabilistic vs. non-probabilistic approaches to the neurobiology of perceptual decision-making. *Curr. Opin. Neurobiol.* 22, 963–969. doi: 10.1016/j.conb.2012.07.007

Dutilh, G., Wagenmakers, E.-J., Visser, I., and van der Maas, H. L. J. (2011). A phase transition model for the speed-accuracy trade-off in response time experiments. *Cogn. Sci.* 35, 211–250. doi: 10.1111/j.1551-6709.2010.01147.x

Dyer, A. G., and Chittka, L. (2004). Bumblebees (Bombus terrestris) sacrifice foraging speed to solve difficult colour discrimination tasks. *J. Comp. Physiol. A Neuroethol. Sens. Neural Behav. Physiol.* 190, 759–763. doi: 10.1007/s00359-004-0547-y

Edwards, W. (1954). The theory of decision making. *Psychol. Bull.* 51, 380–417. doi: 10.1037/h0053870

Edwards, W. (1961). Costs and payoffs are instructions. *Psychol. Rev.* 68, 275–284. doi: 10.1037/h0044699

Edwards, W. (1965). Optimal strategies for seeking information: models for statistics, choice reaction times, and human information processing. *J. Math. Psychol.* 2, 312–329. doi: 10.1016/0022-2496(65)90007-6

Eriksen, B. A., and Eriksen, C. W. (1974). Effects of noise letters upon the identification of a target letter in a nonsearch task. *Percept. Psychophys.* 16, 143–149. doi: 10.3758/BF03203267

Estes, W. K., and Wessel, D. L. (1966). Reaction time in relation to display size and correctness of response in forced-choice visual signal detection. *Percept. Psychophys.* 1, 369–373. doi: 10.3758/BF03215808

Fitts, P. M. (1966). Cognitive aspects of information processing: III. Set for speed versus accuracy. *J. Exp. Psychol.* 71, 849–857. doi: 10.1037/h0023232

Forstmann, B. U., Dutilh, G., Brown, S., Neumann, J., Cramon, Y. V., Ridderinkhof, K. R., et al. (2008). Striatum and pre-SMA facilitate decision-making

under time pressure. *Proc. Natl. Acad. Sci. U.S.A.* 105, 17538–17542. doi: 10.1073/pnas.0805903105

Forstmann, B. U., Tittgemeyer, M., Wagenmakers, E.-J., Derrfuss, J., Imperati, D., and Brown, S. (2011). The speed-accuracy tradeoff in the elderly brain: a structural model-based approach. *J. Neurosci.* 31, 17242–17249. doi: 10.1523/JNEUROSCI.0309-11.2011

Foster, M. (1870). The velocity of thought. *Nature* 2, 2–4. doi: 10.1038/002002a0

Franks, N. R., Dechaume-Moncharmont, F.-X., Hanmore, E., and Reynolds, J. K. (2009). Speed versus accuracy in decision-making ants: expediting politics and policy implementation. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 364, 845–852. doi: 10.1098/rstb.2008.0224

Franks, N. R., Dornhaus, A., Fitzsimmons, J. P., and Stevens, M. (2003). Speed versus accuracy in collective decision making. *Proc. Biol. Sci.* 270, 2457–2463. doi: 10.1098/rspb.2003.2527

Fukushima, J., Fukushima, K., Chiba, T., Tanaka, S., Yamashita, I., and Kato, M. (1988). Disturbances of voluntary control of saccadic eye movements in schizophrenic patients. *Biol. Psychiatry* 23, 670–677. doi: 10.1016/0006-3223(88)90050-9

Garrett, H. E. (1922). A study of the relation of accuracy to speed. *Arch. Psychol.* 56, 1–106.

Gehring, W., Goss, B., Coles, M., Meyer, D., and Donchin, E. (1993). A neural system for error detection and compensation. *Psychol. Sci.* 4, 385–390. doi: 10.1111/j.1467-9280.1993.tb00586.x

Godlove, D. C., Emeric, E. E., Segovis, C. M., Young, M. S., Schall, J. D., and Woodman, G. F. (2011). Event-related potentials elicited by errors during the stop-signal task. I. Macaque monkeys. *J. Neurosci.* 31, 15640–15649. doi: 10.1523/JNEUROSCI.3349-11.2011

Gold, J. I., and Shadlen, M. N. (2002). Banburismus and the brain: decoding the relationship between sensory stimuli, decisions, and reward. *Neuron* 36, 299–308. doi: 10.1016/S0896-6273(02)00971-6

Gold, J. I., and Shadlen, M. N. (2007). The neural basis of decision making. *Annu. Rev. Neurosci.* 30, 535–574. doi: 10.1146/annurev.neuro.29.051605.113038

Gratton, G., Coles, M. G., Sirevaag, E. J., Eriksen, C. W., and Donchin, E. (1988). Pre- and poststimulus activation of response channels: a psychophysiological analysis. *J. Exp. Psychol. Hum. Percept. Perform.* 14, 331–344. doi: 10.1037/0096-1523.14.3.331

Green, D. M., and Luce, D. (1973). "Speed-accuracy tradeoff in auditory detection," in *Attention and Performance IV*, ed S. Kornblum (New York, NY: Academic Press), 547–569.

Grice, G. R., Nullmeyer, R., and Spiker, A. (1977). Application of variable criterion theory to choice reaction time. *Percept. Psychophys.* 22, 431–449. doi: 10.3758/BF03199509

Grice, G. R., and Spiker, A. (1979). Speed-accuracy tradeoff in choice reaction time: within conditions, between conditions, and between subjects. *Percept. Psychophys.* 26, 118–126. doi: 10.3758/BF03208305

Guitton, D., Buchtel, H. A., and Douglas, R. M. (1985). Frontal lobe lesions in man cause difficulties in suppressing reflexive glances and in generating goal-directed saccades. *Exp. Brain Res.* 58, 455–472. doi: 10.1007/BF00235863

Hale, D. J. (1969). Speed-error tradeoff in a three-choice serial reaction task. *J. Exp. Psychol.* 81, 428–435. doi: 10.1037/h0027892

Hancock, P. A., and Newell, K. M. (1985). "The movement speed-accuracy relationship in space-time," in *Motor Behavior*, eds H. Heuer, U. Kleinbeck, and K.-H. Schmidt (Berlin; Heidelberg: Springer), 153–188. doi: 10.1007/978-3-642-69749-4_5

Hanes, D. P., and Schall, J. D. (1996). Neural control of voluntary movement initiation. *Science* 274, 427–430. doi: 10.1126/science.274.5286.427

Heekeren, H. R., Marrett, S., Bandettini, P. A., and Ungerleider, L. G. (2004). A general mechanism for perceptual decision-making in the human brain. *Nature* 431, 859–862. doi: 10.1038/nature02966

Heitz, R., Cohen, J. Y., Woodman, G. F., and Schall, J. D. (2010). Neural correlates of correct and errant attentional selection revealed through N2pc and frontal eye field activity. *J. Neurophysiol.* 104, 2433–2441. doi: 10.1152/jn.00604.2010

Heitz, R., and Engle, R. W. (2007). Focusing the spotlight: individual differences in visual attention control. *J. Exp. Psychol. Gene.* 136, 217–240. doi: 10.1037/0096-3445.136.2.217

Heitz, R., and Schall, J. D. (2012). Neural mechanisms of speed-accuracy tradeoff. *Neuron* 76, 616–628. doi: 10.1016/j.neuron.2012.08.030

Heitz, R., and Schall, J. D. (2013). Neural chronometry and coherency across speed-accuracy demands reveal lack of homomorphism between computational and

neural mechanisms of evidence accumulation. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 368, 20130071. doi: 10.1098/rstb.2013.0071

Henmon, V. A. C. (1911). The relation of the time of a judgment to its accuracy. *Psychol. Rev.* 18, 186–201. doi: 10.1037/h0074579

Hertzog, C., Vernon, M. C., and Rypma, B. (1993). Age differences in mental rotation task performance: the influence of speed/accuracy tradeoffs. *J. Gerontol.* 48, P150–P156. doi: 10.1093/geronj/48.3.P150

Hick, W. E. (1952). On the rate of gain of information. *Q. J. Exp. Psychol.* 4, 11–26. doi: 10.1080/17470215208416600

Ho, T., Brown, S., van Maanen, L., Forstmann, B. U., Wagenmakers, E. J., and Serences, J. T. (2012). The optimality of sensory processing during the speed-accuracy tradeoff. *J. Neurosci.* 32, 7992–8003. doi: 10.1523/JNEUROSCI.0340-12.2012

Horwitz, G. D., and Newsome, W. T. (1999). Separate signals for target selection and movement specification in the superior colliculus. *Science* 284, 1158–1161. doi: 10.1126/science.284.5417.1158

Howell, W. C., and Kreidler, D. L. (1963). Information processing under contradictory instructional sets. *J. Exp. Psychol.* 65, 39–46. doi: 10.1037/h0038982

Ings, T. C., and Chittka, L. (2008). Speed-accuracy tradeoffs and false alarms in bee responses to cryptic predators. *Curr. Biol.* 18, 1520–1524. doi: 10.1016/j.cub.2008.07.074

Ivanoff, J., Branning, P., and Marois, R. (2008). fMRI evidence for a dual process account of the speed-accuracy tradeoff in decision-making. *PLoS ONE* 3:e2635. doi: 10.1371/journal.pone.0002635

Jennings, J. R., Wood, C. C., and Lawrence, B. E. (1976). Effects of graded doses of alcohol on speed-accuracy tradeoff in choice reaction time. *Percept. Psychophys.* 19, 85–91. doi: 10.3758/BF03199391

Jentzsch, I., and Leuthold, H. (2006). Control over speeded actions: a common processing locus for micro- and macro-trade-offs? *Q. J. Exp. Psychol.* 59, 1329–1337. doi: 10.1080/17470210600674394

Kail, R. (1991). Developmental change in speed of processing during childhood and adolescence. *Psychol. Bull.* 109, 490–501. doi: 10.1037/0033-2909.109.3.490

Kane, M. J., and Engle, R. W. (2002). The role of prefrontal cortex in working-memory capacity, executive attention, and general fluid intelligence: an individual-differences perspective. *Psychon. Bull. Rev.* 9, 637–671. doi: 10.3758/BF03196323

Kelly, S. P., and O'Connell, R. G. (2013). Internal and external influences on the rate of sensory evidence accumulation in the human brain. *J. Neurosci.* 33, 19434–19441. doi: 10.1523/JNEUROSCI.3355-13.2013

Kim, J. N., and Shadlen, M. N. (1999). Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nat. Neurosci.* 2, 176–185. doi: 10.1038/5739

Krajbich, I., Armel, C., and Rangel, A. (2010). Visual fixations and the computation and comparison of value in simple choice. *Nat. Neurosci.* 13, 1292–1298. doi: 10.1038/nn.2635

Krajbich, I., Lu, D., Camerer, C., and Rangel, A. (2012). The attentional drift-diffusion model extends to simple purchasing decisions. *Front. Psychol.* 3:193. doi: 10.3389/fpsyg.2012.00193

Kulahci, I. G., Dornhaus, A., and Papaj, D. R. (2008). Multimodal signals enhance decision making in foraging bumble-bees. *Proc. Biol. Sci.* 275, 797–802. doi: 10.1098/rspb.2007.1176

Kumar, A., Rakitin, B. C., Nambisan, R., Habeck, C., and Stern, Y. (2008). The response-signal method reveals age-related changes in object working memory. *Psychol. Aging* 23, 315–329. doi: 10.1037/0882-7974.23.2.315

LaBerge, D. (1962). A recruitment theory of simple behavior. *Psychometrika* 27, 375–396. doi: 10.1007/BF02289645

Laming, D. (1968). *Information Theory of Choice-Reaction Times.* New York, NY: Academic Press.

Lappin, J. S., and Disch, K. (1972a). The latency operating characteristic. I. Effects of stimulus probability on choice reaction time. *J. Exp. Psychol.* 92, 419–427. doi: 10.1037/h0032360

Lappin, J. S., and Disch, K. (1972b). The latency operating characteristic. II. Effects of visual stimulus intensity on choice reaction time. *J. Exp. Psychol.* 93, 367–372. doi: 10.1037/h0032465

Le Houezec, J., Halliday, R., Benowitz, N. L., Callaway, E., Naylor, H., and Herzig, K. (1994). A low dose of subcutaneous nicotine improves information processing in non-smokers. *Psychopharmacology* 114, 628–634. doi: 10.1007/BF02244994

Link, S. W., and Tindall, A. D. (1971). Speed and accuracy in comparative judgments of line length. *Percept. Psychophys.* 9, 284–288. doi: 10.3758/BF03212649

Liu, S., Heitz, R., and Bradberry, C. W. (2009). A touch screen based stop signal response task in rhesus monkeys for studying impulsivity associated with chronic cocaine self-administration. *J. Neurosci. Methods* 177, 67–72. doi: 10.1016/j.jneumeth.2008.09.020

Liu, S., Heitz, R. P., Sampson, A. R., Zhang, W., and Bradberry, C. W. (2008). Evidence of temporal cortical dysfunction in rhesus monkeys following chronic cocaine self-administration. *Cereb. Cortex* 18, 2109–2116. doi: 10.1093/cercor/bhm236

Lo, C.-C., and Wang, X.-J. (2006). Cortico-basal ganglia circuit mechanism for a decision threshold in reaction time tasks. *Nat. Neurosci.* 9, 956–963. doi: 10.1038/nn1722

Logan, G. D. (2002). An instance theory of attention and memory. *Psychol. Rev.* 109, 376–400. doi: 10.1037/0033-295X.109.2.376

Logothetis, N. K. (2008). What we can do and what we cannot do with fMRI. *Nature* 453, 869–878. doi: 10.1038/nature06976

Lohman, D. (1989). Individual differences in errors and latencies on cognitive tasks. *Learn. Indiv. Differ.* 1, 179–202. doi: 10.1016/1041-6080(89)90002-2

Luce, D. (1986). *Response Times: Their Role in Inferring Elementary Mental Organization.* New York, NY: Oxford University Press.

Lyons, J. J., and Briggs, G. E. (1971). Speed-accuracy trade-off with different types of stimuli. *J. Exp. Psychol.* 91, 115–119. doi: 10.1037/h0031815

Madden, D. J., and Allen, P. A. (1991). Adult age differences in the rate of information extraction during visual search. *J. Gerontol.* 46, P124–P126. doi: 10.1093/geronj/46.3.P124

Marshall, J. A. R., Bogacz, R., Dornhaus, A., Planqué, R., Kovacs, T., and Franks, N. R. (2009). On optimal decision-making in brains and social insect colonies. *J. R. Soc. Interface* 6, 1065–1074. doi: 10.1098/rsif.2008.0511

Marshall, J. A. R., Dornhaus, A., Franks, N. R., and Kovacs, T. (2006). Noise, cost and speed-accuracy trade-offs: decision-making in a decentralized system. *J. R. Soc. Interface* 3, 243–254. doi: 10.1098/rsif.2005.0075

Martin, L. J., and Müeller, G. E. (1899). *Zur Analyse der Unterschiedsempfindlichkeit.* Leipzig: J. A. Barth.

McClelland, J. L. (1979). On the time relations of mental processes: an examination of systems of processes in cascade. *Psychol. Rev.* 86, 287. doi: 10.1037/0033-295X.86.4.287

McElree, B., and Dosher, B. A. (1989). Serial position and set size in short-term memory: the time course of recognition. *J. Exp. Psychol. Gene.* 118, 346–373. doi: 10.1037/0096-3445.118.4.346

Merkel, J. (1885). Die zeitlichen verhältnisse der willenstätigkeit. *Philos. Stud.* 2, 73–127.

Metin, B., Roeyers, H., Wiersema, J. R., van der Meere, J. J., Thompson, M., and Sonuga-Barke, E. (2013). ADHD performance reflects inefficient but not impulsive information processing: a diffusion model analysis. *Neuropsychology* 27, 193–200. doi: 10.1037/a0031533

Meyer, D. E., Irwin, D. E., Osman, A. M., and Kounios, J. (1988). The dynamics of cognition and action: mental processes inferred from speed-accuracy decomposition. *Psychol. Rev.* 95, 183–237. doi: 10.1037/0033-295X.95.2.183

Meyer, D. E., Keith-Smith, J. E., Kornblum, S., Abrams, R. A., and Wright, C. E. (1990). "Speed-accuracy tradeoffs in aimed movements: toward a theory of rapid voluntary action," in *Attention and Performance XIII*, ed M. Jeannerod (Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.), 173–226.

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

Mulder, M. J., Bos, D., Weusten, J. M. H., van Belle, J., van Dijk, S. C., Simen, P., et al. (2010). Basic impairments in regulating the speed-accuracy tradeoff predict symptoms of attention-deficit/hyperactivity disorder. *Biol. Psychiatry* 68, 1114–1119. doi: 10.1016/j.biopsych.2010.07.031

Myerson, J., Robertson, S., and Hale, S. (2007). Aging and intraindividual variability in performance: analyses of response time distributions. *J. Exp. Anal. Behav.* 88, 319–337. doi: 10.1901/jeab.2007.88-319

Nickerson, R. S. (1969). Same-different response times: a model and a preliminary test. *Acta Psychologica* 30, 257–275. doi: 10.1016/0001-6918(69)90054-7

Nieuwenhuis, S., Ridderinkhof, K. R., Blom, J., Band, G. P., and Kok, A. (2001). Error-related brain potentials are differentially related to awareness of response errors: evidence from an antisaccade task. *Psychophysiology* 38, 752–760. doi: 10.1111/1469-8986.3850752

O'Connell, R. G., Dockree, P. M., and Kelly, S. P. (2012). A supramodal accumulation-to-bound signal that determines perceptual decisions in humans. *Nat. Neurosci.* 15, 1729–1735. doi: 10.1038/nn.3248

Ollman, R. T. (1966). Fast guesses in choice reaction time. *Psychon. Sci.* 6, 155–156. doi: 10.3758/BF03328004

Osman, A., Lou, L., Müller-Gethmann, H., Rinkenauer, G., Mattes, S., and Ulrich, R. (2000). Mechanisms of speed-accuracy tradeoff: evidence from covert motor processes. *Biol. Psychol.* 51, 173–199. doi: 10.1016/S0301-0511(99)00045-9

Öztekin, I., and McElree, B. (2010). Relationship between measures of working memory capacity and the time course of short-term memory retrieval and interference resolution. *J. Exp. Psychol. Learn. Mem. Cogn.* 36, 383–397. doi: 10.1037/a0018029

Pachella, R. G. (1974). "The interpretation of reaction time in information processing research," in *Human Information Processing*, ed B. Kantowitz (New York, NY: Lawrence Erlbaum), 41–82.

Pachella, R. G., and Fisher, D. F. (1969). Effect of stimulus degradation and similarity on the trade-off between speed and accuracy in absolute judgments. *J. Exp. Psychol.* 81, 7–9. doi: 10.1037/h0027431

Pachella, R. G., and Pew, R. W. (1968). Speed-accuracy tradeoff in reaction time: effect of discrete criterion times. *J. Exp. Psychol.* 76, 19–24. doi: 10.1037/h0021275

Passino, K. M., and Seeley, T. D. (2006). Modeling and analysis of nest-site selection by honeybee swarms: the speed and accuracy trade-off. *Behav. Ecol. Sociobiol.* 59, 427–442. doi: 10.1007/s00265-005-0067-y

Passino, K. M., Seeley, T. D., and Visscher, P. K. (2008). Swarm cognition in honey bees. *Behav. Ecol. Sociobiol.* 62, 401–414. doi: 10.1007/s00265-007-0468-1

Pastötter, B., Berchtold, F., and Bäuml, K.-H. T. (2012). Oscillatory correlates of controlled speed-accuracy tradeoff in a response-conflict task. *Hum. Brain Mapp.* 33, 1834–1849. doi: 10.1002/hbm.21322

Pearson, K. (1905). The problem of the random walk. *Nature* 72, 294. doi: 10.1038/072294b0

Pfefferbaum, A., Ford, J., Johnson, R., Wenegrat, B., and Kopell, B. S. (1983). Manipulation of P3 latency: speed vs. accuracy instructions. *Electroencephalogr. Clin. Neurophysiol.* 55, 188–197. doi: 10.1016/0013-4694(83)90187-6

Phillips, L. H., and Rabbit, P. (1995). Impulsivity and speed-accuracy strategies in intelligence test performance. *Intelligence* 21, 13–29. doi: 10.1016/0160-2896(95)90036-5

Pike, A. R. (1968). Latency and relative frequency of response in psychophysical discrimination. *Br. J. Math. Stat. Psychol.* 21, 161–182. doi: 10.1111/j.2044-8317.1968.tb00407.x

Pike, R., McFarland, K., and Dalgleish, L. (1974). Speed-accuracy tradeoff models for auditory detection with deadlines. *Acta Psychol.* 38, 379–399. doi: 10.1016/0001-6918(74)90042-0

Plamondon, R., and Alimi, A. M. (1997). Speed/accuracy trade-offs in target-directed movements. *Behav. Brain Sci.* 20, 279–303. discussion: 303–349. doi: 10.1017/S0140525X97001441

Purcell, B. A., Heitz, R., Cohen, J. Y., Schall, J. D., Logan, G. D., and Palmeri, T. J. (2010). Neurally constrained modeling of perceptual decision making. *Psychol. Rev.* 117, 1113–1143. doi: 10.1037/a0020311

Purcell, B. A., Schall, J. D., Logan, G. D., and Palmeri, T. J. (2012). From salience to saccades: multiple-alternative gated stochastic accumulator model of visual search. *J. Neurosci.* 32, 3433–3446. doi: 10.1523/JNEUROSCI.4622-11.2012

Rae, B., Heathcote, A., Donkin, C., Averell, L., and Brown, S. (2014). The hare and the tortoise: emphasizing speed can change the evidence used to make decisions. *J. Exp. Psychol. Learn. Mem. Cogn.* doi: 10.1037/a0036801. [Epub ahead of print].

Ratcliff, R. (1978). A theory of memory retrieval. *Psychol. Rev.* 85, 59–108. doi: 10.1037/0033-295X.85.2.59

Ratcliff, R. (2006). Modeling response signal and response time data. *Cogn. Psychol.* 53, 195–237. doi: 10.1016/j.cogpsych.2005.10.002

Ratcliff, R., Philiastides, M. G., and Sajda, P. (2009). Quality of evidence for perceptual decision making is indexed by trial-to-trial variability of the EEG. *Proc. Natl. Acad. Sci. U.S.A.* 106, 6539–6544. doi: 10.1073/pnas.0812589106

Ratcliff, R., Hasegawa, Y. T., Hasegawa, R. P., Smith, P. L., and Segraves, M. A. (2007). Dual diffusion model for single-cell recording data from the superior colliculus in a brightness-discrimination task. *J. Neurophysiol.* 97, 1756–1774. doi: 10.1152/jn.00393.2006

Ratcliff, R., and Rouder, J. (1998). Modeling response times for two-choice decisions. *Psychol. Sci.* 9, 347–356. doi: 10.1111/1467-9280.00067

Ratcliff, R., and Rouder, J. N. (2000). A diffusion model account of masking in two-choice letter identification. *J. Exp. Psychol. Hum. Percept. Perform.* 26, 127–140. doi: 10.1037/0096-1523.26.1.127

Ratcliff, R., and Smith, P. L. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychol. Rev.* 111, 333–367. doi: 10.1037/0033-295X.111.2.333

Ratcliff, R., and Smith, P. L. (2010). Perceptual discrimination in static and dynamic noise: the temporal relation between perceptual encoding and decision making. *J. Exp. Psychol. Gene.* 139, 70–94. doi: 10.1037/a0018128

Ratcliff, R., Thapar, A., Gomez, P., and Mckoon, G. (2004). A diffusion model analysis of the effects of aging in the lexical-decision task. *Psychol. Aging* 19, 278–289. doi: 10.1037/0882-7974.19.2.278

Ratcliff, R., and Tuerlinckx, F. (2002). Estimating parameters of the diffusion model: approaches to dealing with contaminant reaction times and parameter variability. *Psychon. Bull. Rev.* 9, 438–481. doi: 10.3758/BF031 96302

Reddi, B. A., and Carpenter, R. H. (2000). The influence of urgency on decision time. *Nat. Neurosci.* 3, 827–830. doi: 10.1038/77739

Reed. (1973). Speed-accuracy trade-off in recognition memory. *Science* 181, 574–576. doi: 10.1126/science.181.4099.574

Ridderinkhof, K. R. (2002). Micro- and macro-adjustments of task set: activation and suppression in conflict tasks. *Psychol. Res.* 66, 312–323. doi: 10.1007/s00426-002-0104-7

Rinberg, D., Koulakov, A., and Gelperin, A. (2006). Speed-accuracy tradeoff in olfaction. *Neuron* 51, 351–358. doi: 10.1016/j.neuron.2006.07.013

Rinkenauer, G., Osman, A., Ulrich, R., Muller-Gethmann, H., and Mattes, S. (2004). On the locus of speed-accuracy trade-off in reaction time: inferences from the lateralized readiness potential. *J. Exp. Psychol. Gene.* 133, 261–282. doi: 10.1037/0096-3445.133.2.261

Riveros, A. J., and Gronenberg, W. (2012). Decision-making and associative color learning in harnessed bumblebees (Bombus impatiens). *Anim. Cogn.* 15, 1183–1193. doi: 10.1007/s10071-012-0542-6

Roitman, J. D., and Shadlen, M. N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J. Neurosci.* 22, 9475–9489.

Rundell, O. H., and Williams, H. L. (1979). Alcohol and speed-accuracy tradeoff. *Hum. Fact.* 21, 433–443.

Ruthruff, E. (1996). A test of the deadline model for speed-accuracy tradeoffs. *Percept. Psychophys.* 58, 56–64. doi: 10.3758/BF03205475

Salthouse, T. (1979). Adult age and the speed-accuracy trade-off. *Ergonomics* 22, 811–821. doi: 10.1080/00140137908924659

Salthouse, T. (2012). Consequences of age-related cognitive declines. *Ann. Rev. Psychol.* 63, 201–226. doi: 10.1146/annurev-psych-120710-100328

Schneider, D. W., and Anderson, J. R. (2012). Modeling fan effects on the time course of associative recognition. *Cogn. Psychol.* 64, 127–160. doi: 10.1016/j.cogpsych.2011.11.001

Schouten, J. F., and Bekker, J. A. (1967). Reaction time and accuracy. *Acta Psychol.* 27, 143–153. doi: 10.1016/0001-6918(67)90054-6

Schweitzer, L. R., and Lee, M. (1992). Use of the speed accuracy trade-off to characterize information processing in schizophrenics and normals. *Psychopathology* 25, 29–40. doi: 10.1159/000284751

Seeley, T. D., and Visscher, P. K. (2004). Quorum sensing during nest-site selection by honeybee swarms. *Behav. Ecol. Sociobiol.* 56, 594–601. doi: 10.1007/s00265-004-0814-5

Sergeant, J. A., and Scholten, C. A. (1985a). On data limitations in hyperactivity. *J. Child Psychol. Psychiatry* 26, 111–124. doi: 10.1111/j.1469-7610.1985.tb01632.x

Sergeant, J. A., and Scholten, C. A. (1985b). On resource strategy limitations in hyperactivity: cognitive impulsivity reconsidered. *J. Child Psychol. Psychiatry* 26, 97–109. doi: 10.1111/j.1469-7610.1985.tb01631.x

Shen, K., Kalwarowsky, S., Clarence, W., Brunamonti, E., and Paré, M. (2010). Beneficial effects of the NMDA antagonist ketamine on decision processes in visual search. *J. Neurosci.* 30, 9947–9953. doi: 10.1523/JNEUROSCI.6317-09.2010

Simen, P., Cohen, J. D., and Holmes, P. (2006). Rapid decision threshold modulation by reward rate in a neural network. *Neural Netw.* 19, 1013–1026. doi: 10.1016/j.neunet.2006.05.038

Simen, P., Contreras, D., Buck, C., Hu, P., Holmes, P., and Cohen, J. D. (2009). Reward rate optimization in two-alternative decision making: empirical tests of theoretical predictions. *J. Exp. Psychol. Hum. Percept. Perform.* 35, 1865–1897. doi: 10.1037/a0016926

Skorupski, P., Spaethe, J., and Chittka, L. (2006). Visual search and decision making in bees: time, speed, and accuracy. *Int. J. Comp. Psychol.* 19, 342–357.

Smith, P. L., and Vickers, D. (1988). The accumulator model of two-choice discrimination. *J. Math. Psychol.* 32, 135–168. doi: 10.1016/0022-2496(88)90043-0

Standage, D., You, H., Wang, D.-H., and 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. *Front. Comput. Neurosci.* 5:7. doi: 10.3389/fncom.2011.00007

Standage, D., You, H., Wang, D.-H., and Dorris, M. C. (2013). Trading speed and accuracy by coding time: a coupled-circuit cortical model. *PLoS Comput. Biol.* 9:e1003021. doi: 10.1371/journal.pcbi.1003021

Stark, C. E., and Squire, L. R. (2001). When zero is not zero: the problem of ambiguous baseline conditions in fMRI. *Proc. Natl. Acad. Sci. U.S.A.* 98, 12760–12766. doi: 10.1073/pnas.221462998

Starns, J. J., and Ratcliff, R. (2010). The effects of aging on the speed-accuracy compromise: boundary optimality in the diffusion model. *Psychol. Aging* 25, 377–390. doi: 10.1037/a0018022

Starns, J. J., and Ratcliff, R. (2012). Age-related differences in diffusion model boundary optimality with both trial-limited and time-limited tasks. *Psychon. Bull. Rev.* 19, 139–145. doi: 10.3758/s13423-011-0189-3

Stone, M. (1960). Models for choice-reaction time. *Psychometrika* 25, 251–260. doi: 10.1007/BF02289729

Swanson, J. M., and Briggs, G. E. (1969). Information processing as a function of speed versus accuracy. *J. Exp. Psychol.* 81, 223–229. doi: 10.1037/h0027774

Swensson, R. G. (1972a). The elusive tradeoff: speed vs accuracy in visual discrimination tasks. *Percept. Psychophys.* 12, 16–32. doi: 10.3758/BF03212837

Swensson, R. G. (1972b). Trade-off bias and efficiency effects in serial choice reactions. *J. Exp. Psychol.* 95, 397. doi: 10.1037/h0033674

Swensson, R. G., and Edwards, W. (1971). Response strategies in a two-choice reaction task with a continuous cost for time. *J. Exp. Psychol.* 88, 67. doi: 10.1037/h0030646

Thakkar, K. N., Schall, J. D., Boucher, L., Logan, G. D., and Park, S. (2011). Response inhibition and response monitoring in a saccadic countermanding task in schizophrenia. *Biol. Psychiatry* 69, 55–62. doi: 10.1016/j.biopsych.2010.08.016

Thomas, E. A. (1974). The selectivity of preparation. *Psychol. Rev.* 81, 442. doi: 10.1037/h0036945

Thura, D., Beauregard-Racine, J., Fradet, C.-W., and Cisek, P. (2012). Decision making by urgency gating: theory and experimental support. *J. Neurophysiol.* 108, 2912–2930. doi: 10.1152/jn.01071.2011

Thura, D., and Cisek, P. (2014). Deliberation and commitment in the premotor and primary motor cortex during dynamic decision making. *Neuron* 81, 1401–1416. doi: 10.1016/j.neuron.2014.01.031

Tiplady, B., Drummond, G. B., Cameron, E., Gray, E., Hendry, J., Sinclair, W., et al. (2001). Ethanol, errors, and the speed-accuracy trade-off. *Pharmacol. Biochem. Behav.* 69, 635–641. doi: 10.1016/S0091-3057(01)00551-2

Towal, R. B., Mormann, M., and Koch, C. (2013). Simultaneous modeling of visual saliency and value computation improves predictions of economic choice. *Proc. Natl. Acad. Sci. U.S.A.* 110, E3858–E3867. doi: 10.1073/pnas.1304429110

Uchida, N., and Mainen, Z. F. (2003). Speed and accuracy of olfactory discrimination in the rat. *Nat. Neurosci.* 6, 1224–1229. doi: 10.1038/nn1142

Usher, M., and McClelland, J. L. (2001). The time course of perceptual choice: the leaky, competing accumulator model. *Psychol. Rev.* 108, 550–592. doi: 10.1037/0033-295X.108.3.550

Vallesi, A., McIntosh, A. R., Crescentini, C., and Stuss, D. T. (2012). fMRI investigation of speed-accuracy strategy switching. *Hum. Brain Mapp.* 33, 1677–1688. doi: 10.1002/hbm.21312

van der Lubbe, R. H., Jaœkowski, P., Wauschkuhn, B., and Verleger, R. (2001). Influence of time pressure in a simple response task, a choice-by-location task, and the Simon task. *J. Psychophysiol.* 15, 241. doi: 10.1027/0269-8803.15.4.241

van Ravenzwaaij, D., Dutilh, G., and Wagenmakers, E.-J. (2012). A diffusion model decomposition of the effects of alcohol on perceptual decision making. *Psychopharmacology* 219, 1017–1025. doi: 10.1007/s00213-011-2435-9

van Veen, V., Krug, M. K., and Carter, C. S. (2008). The neural and computational basis of controlled speed-accuracy tradeoff during task performance. *J. Cogn. Neurosci.* 20, 1952–1965. doi: 10.1162/jocn.2008.20146

van Vugt, M. K., Simen, P., Nystrom, L., Holmes, P., and Cohen, J. D. (2014). Lateralized readiness potentials reveal properties of a neural mechanism for implementing a decision threshold. *PLoS ONE* 9:e90943. doi: 10.1371/journal.pone.0090943

van Vugt, M. K., Simen, P., Nystrom, L. E., Holmes, P., and Cohen, J. D. (2012). EEG oscillations reveal neural correlates of evidence accumulation. *Front. Neurosci.* 6:106. doi: 10.3389/fnins.2012.00106

van Zandt, T. (2000). How to fit a response time distribution. *Psychon. Bull. Rev.* 7, 424–465. doi: 10.3758/BF03214357

Vickers, D., Burt, J., Smith, P., and Brown, M. (1985). Experimental paradigms emphasising state or process limitations: I Effects on speed-accuracy tradeoffs. *Acta Psychol.* 59, 129–161. doi: 10.1016/0001-6918(85)90017-4

Vickers, D., and Smith, P. (1985). Accumulator and random-walk models of psychophysical discrimination: a counter-evaluation. *Perception* 14, 471–497. doi: 10.1068/p140471

Wagenmakers, E., Ratcliff, R., Gomez, P., and McKoon, G. (2008). A diffusion model account of criterion shifts in the lexical decision task. *J. Mem. Lang.* 58, 140–159. doi: 10.1016/j.jml.2007.04.006

Wald, A. (1947). *Sequential Analysis.* New York, NY: Wiley & Sons.

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

Wenzlaff, H., Bauer, M., Maess, B., and Heekeren, H. R. (2011). Neural characterization of the speed-accuracy tradeoff in a perceptual decision-making task. *J. Neurosci.* 31, 1254–1266. doi: 10.1523/JNEUROSCI.4000-10.2011

Wickelgren, W. A. (1977). Speed-accuracy tradeoff and information processing dynamics. *Acta Psychol.* 41, 67–85. doi: 10.1016/0001-6918(77)90012-9

Winkel, J., van Maanen, L., Ratcliff, R., van der Schaaf, M. E., van Schouwenburg, M. R., Cools, R., et al. (2012). Bromocriptine does not alter speed-accuracy tradeoff. *Front. Neurosci.* 6:126. doi: 10.3389/fnins.2012.00126

Wong, K.-F., Huk, A. C., Shadlen, M. N., and Wang, X.-J. (2007). Neural circuit dynamics underlying accumulation of time-varying evidence during perceptual decision making. *Front. Comput. Neurosci.* 1:6. doi: 10.3389/neuro.10.006.2007

Wood, C. C., and Jennings, J. R. (1976). Speed-accuracy tradeoff functions in choice reaction time: experimental designs and computational procedures. *Attent. Percept. Psychophys.* 19, 92–102. doi: 10.3758/BF03199392

Woodman, G. F., Kang, M.-S., Thompson, K., and Schall, J. D. (2008). The effect of visual search efficiency on response preparation: neurophysiological evidence for discrete flow. *Psychol. Sci.* 19, 128–136. doi: 10.1111/j.1467-9280.2008.02058.x

Woodman, G. F., and Luck, S. J. (1999). Electrophysiological measurement of rapid shifts of attention during visual search. *Nature* 400, 867–869. doi: 10.1038/23698

Woodworth, R. S. (1899). The accuracy of a voluntary movement. *Psychol. Rev.* 3, 1–114.

Wylie, S. A., van den Wildenberg, W. P. M., Ridderinkhof, K. R., Bashore, T. R., Powell, V. D., Manning, C. A., et al. (2009). The effect of speed-accuracy strategy on response interference control in Parkinson's disease. *Neuropsychologia* 47, 1844–1853. doi: 10.1016/j.neuropsychologia.2009.02.025

Yamaguchi, M., Crump, M. J. C., and Logan, G. D. (2013). Speed-accuracy trade-off in skilled typewriting: decomposing the contributions of hierarchical control loops. *J. Exp. Psychol. Hum. Percept. Perform.* 39, 678–699. doi: 10.1037/a0030512

Yellott, J. I. (1971). Correction for fast guessing and the speed-accuracy tradeoff in choice reaction time. *J. Math. Psychol.* 8, 159–199. doi: 10.1016/0022-2496(71)90011-3

Zandbelt, B., Purcell, B. A., Palmeri, T. J., Logan, G. D., and Schall, J. D. (2014). Response times from ensembles of accumulators. *Proc. Natl. Acad. Sci. U.S.A.* 111, 2848–2853. doi: 10.1073/pnas.1310577111

Zariwala, H. A., Kepecs, A., Uchida, N., Hirokawa, J., and Mainen, Z. F. (2013). The limits of deliberation in a perceptual decision task. *Neuron* 78, 339–351. doi: 10.1016/j.neuron.2013.02.010

Zhang, J. (2012). The effects of evidence bounds on decision-making: theoretical and empirical developments. *Front. Psychol.* 3:263. doi: 10.3389/fpsyg.2012.00263

Zhang, J., and Bogacz, R. (2010). Optimal decision making on the basis of evidence represented in spike trains. *Neural Comput.* 22, 1113–1148. doi: 10.1162/neco.2009.05-09-1025

On the temporal dynamics of spatial stimulus-response transfer between spatial incompatibility and Simon tasks

Jason Ivanoff[1]*, Ryan Blagdon[1], Stefanie Feener[1], Melanie McNeil[1] and Paul H. Muir[2]

[1] Department of Psychology, Saint Mary's University, Halifax, NS, Canada
[2] Department of Mathematics and Computing Science, Saint Mary's University, Halifax, NS, Canada

The *Simon effect* refers to the performance (response time and accuracy) advantage for responses that spatially correspond to the task-irrelevant location of a stimulus. It has been attributed to a natural tendency to respond toward the source of stimulation. When location is task-relevant, however, and responses are intentionally directed away (incompatible) or toward (compatible) the source of the stimulation, there is also an advantage for spatially compatible responses over spatially incompatible responses. Interestingly, a number of studies have demonstrated a reversed, or reduced, Simon effect following practice with a spatial incompatibility task. One interpretation of this finding is that practicing a spatial incompatibility task disables the natural tendency to respond toward stimuli. Here, the temporal dynamics of this stimulus-response (S-R) transfer were explored with speed-accuracy trade-offs (SATs). All experiments used the mixed-task paradigm in which Simon and spatial compatibility/incompatibility tasks were interleaved across blocks of trials. In general, bidirectional S-R transfer was observed: while the spatial incompatibility task had an influence on the Simon effect, the task-relevant S-R mapping of the Simon task also had a small impact on congruency effects within the spatial compatibility and incompatibility tasks. These effects were generally greater when the task contexts were similar. Moreover, the SAT analysis of performance in the Simon task demonstrated that the tendency to respond to the location of the stimulus was not eliminated because of the spatial incompatibility task. Rather, S-R transfer from the spatial incompatibility task appeared to partially mask the natural tendency to respond to the source of stimulation with a conflicting inclination to respond away from it. These findings support the use of SAT methodology to quantitatively describe rapid response tendencies.

Keywords: speed-accuracy trade-off, stimulus-response compatibility, Simon effect, spatial compatibility, S-R associations

INTRODUCTION

The spatial configuration of stimuli and responses greatly affects human performance (Fitts and Seeger, 1953; Fitts and Deininger, 1954). Studies of stimulus-response (S-R) compatibility provide an opportunity to explore which sorts of S-R associations are more natural, and perhaps more automatic, than others. *Spatial incompatibility* tasks, where the stimulus location is *task-relevant* and the goal is to respond away from a stimulus, are generally performed more slowly and with greater errors than *spatial compatibility* tasks, where responses are directed toward stimuli (Fitts and Deininger, 1954). Fitts and Deininger proposed that the number of transformations between stimulus and response was a partial determinant of speeded responding under S-R compatible/incompatible conditions. Others have taken a slightly different approach, suggesting that the number or complexity of rules in an incompatibility task is greater than it is in a compatible condition (Duncan, 1977, 1978). It is generally thought that it is easier to respond when there is some kind of conceptual match between stimulus and response features (Kornblum et al., 1990).

The location of a stimulus, even when task-irrelevant, affects spatial responding (Simon and Rudell, 1967; Simon, 1969), suggesting there is some sort of well-established or automatic pathway extending from neural regions responsible for processing stimulus location to neural regions responsible for response selection. The *Simon effect* refers to the performance advantage for spatially corresponding responses over non-corresponding responses, when the location of the stimulus is task-irrelevant. It was originally attributed to "a 'natural' tendency to react toward the source of stimulation" (Simon, 1969, p. 175). Dual-route models (de Jong et al., 1994) usually incorporate this natural tendency as a feature of the automatic, or direct, pathway that speeds (corresponding), or slows (non-corresponding), responding. Although other accounts of the Simon effect have emphasized various mechanisms (e.g., see Lu and Proctor, 1995; Proctor, 2011; Van der Lubbe and Abrahamse, 2011; Hommel, 2011 for reviews), most accounts do tend to incorporate some kind of "natural tendency" for location information to influence response selection.
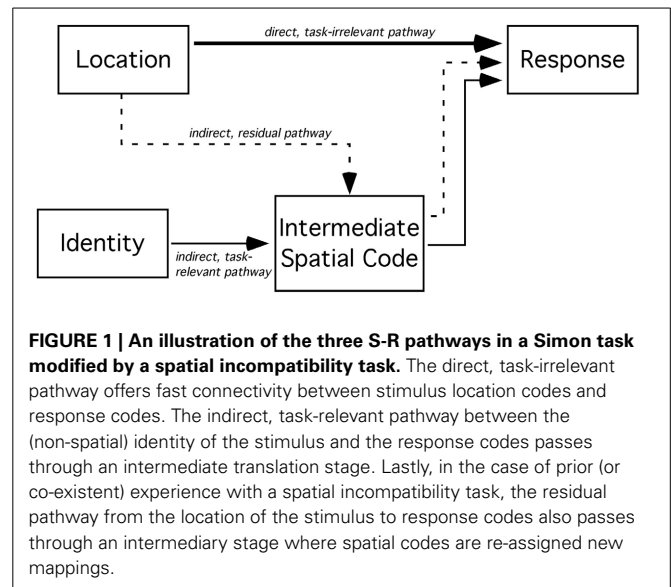
TRANSFER OF S-R PATHWAY ACTIVITY ACROSS SIMON AND SPATIAL INCOMPATIBILITY TASKS

In recent years, there has been growing interest in the transfer of S-R mappings between spatial incompatibility and Simon tasks. Proctor and Lu (1999) demonstrated that the Simon effect reversed (i.e., spatially non-corresponding responses were faster than spatially corresponding responses) when the Simon task was preceded by a spatial incompatibility task. In other studies, transfer from the spatial incompatibility task to the Simon task has eliminated, but not reversed, the Simon effect (Tagliabue et al., 2000). Tagliabue et al. (2000) attributed this discrepancy to the greater number of practice trials in the spatial incompatibility task in Proctor and Lu's (1999) study (~1800 trials) compared to that of their study (72 trials). The reverse (or absent) Simon effect following a spatial incompatibility task has been explained in one of two ways.

The first account of the reverse (or absent) Simon effect following a spatial incompatibility task is, perhaps, the most pragmatic of the two proposals. The Simon effect has routinely been attributed to "automatic" response priming from the corresponding stimulus location. This priming is thought to occur along the direct, spatial S-R pathway. Proctor and Lu (1999) suggested that activation of the direct pathway is not necessarily immutable. In their description of the connectivity between spatial features of the stimulus and the response they state, "[t]hese associations have been described as *unconditional* (de Jong et al., 1994), *permanent* (Barber and O'Leary, 1997), and as being either *hard-wired* or *learned from a lifetime's experience* (Umiltà and Zorzi, 1997). The implication of such descriptions - that the associations are essentially unmodifiable—is incorrect" (Proctor and Lu, 1999, p. 76). Thus, the learned associations from the spatial incompatibility task may simply "overwrite" the direct pathway thereby reversing, or eliminating, the Simon effect.

Tagliabue et al. proposed a different account of the effect of a spatial incompatibility task on the Simon effect. Their account includes three pathways (see **Figure 1** for a graphical representation of the three pathways). The direct spatial S-R pathway, connecting location stimulus codes directly to response codes, has a quick, yet evanescent, onset. One of the slow, indirect S-R pathways (sometimes called the *conditional route* or the *controlled pathway*) is task-relevant: it translates non-spatial, symbolic stimulus codes to intermediary spatial codes that, in turn, connect to response codes. Dual pathway models have long been presumed to encompass the cognitive architecture necessary for the Simon effect (e.g., de Jong et al., 1994). The other indirect pathway is spatial and is the result of residual activity from the spatial incompatibility task. It is likely slower than the direct spatial pathway. It connects stimulus location information to intermediary spatial codes that, in turn, recode spatial stimulus information for spatial response selection (e.g., left → right and right → left). Tagliabue et al. (2000) argued that this particular model accounts for the time course of the Simon effect, following the performance of a spatial incompatibility task, quite well.

In contrast to the S-R transfer evident from a spatial incompatibility task to a Simon task, there is currently little evidence for S-R transfer from a spatial compatibility task to a Simon task. Proctor and Lu (1999) observed a 21 ms Simon effect following practice in

FIGURE 1 | An illustration of the three S-R pathways in a Simon task modified by a spatial incompatibility task. The direct, task-irrelevant pathway offers fast connectivity between stimulus location codes and response codes. The indirect, task-relevant pathway between the (non-spatial) identity of the stimulus and the response codes passes through an intermediate translation stage. Lastly, in the case of prior (or co-existent) experience with a spatial incompatibility task, the residual pathway from the location of the stimulus to response codes also passes through an intermediary stage where spatial codes are re-assigned new mappings.

a task with central (neutral) stimuli and a 21 ms Simon effect following practice with a spatial compatibility task. Tagliabue et al. (2000) noted a baseline Simon effect of 38.5 ms (Experiment 6), and a Simon effect of 26.5, 35, and 33 ms (Experiments 3-5) when preceded by spatial compatibility task. Tagliabue et al. (2000) argue that the spatial S-R mappings from a spatial compatibility task cannot further strengthen the direct pathway. Accordingly, the (absent) effect of a spatial compatibility task on the Simon effect provides reasonable experimental control to evaluate the adverse effect of a spatial incompatibility task on the Simon effect.

Interestingly, S-R transfer does not seem to be particular to a set of stimuli as it occurs when different stimuli sets are used across tasks (Proctor and Lu, 1999). S-R transfer also occurs across different stimulus modalities (Tagliabue et al., 2002), although perhaps more weakly, given that the spatial incompatibility task did not reverse the Simon effect in this study. With a sufficient number of practice trials, there is even some evidence for S-R transfer when the spatial incompatibility task is presented along a different spatial axis from the Simon task (e.g., the practice stimuli and responses in the spatial incompatibility task are presented along the horizontal axis while the transfer stimuli and responses in the Simon task are presented along the vertical axis), suggesting that in some cases a S-R rule (e.g., a "respond opposite" procedure) may transfer across tasks (Vu, 2007). S-R transfer from a spatial incompatibility task to a Simon task may be relatively persistent. Transfer effects have been observed when the interval between spatial incompatibility and Simon tasks has ranged from 5 min to days (Tagliabue et al., 2000, 2002). S-R transfer across tasks is also evident in so-called mixing tasks, where the spatial incompatibility task alternates, or is interleaved, with a Simon task (Marble and Proctor, 2000; Proctor et al., 2000).

Despite its ubiquity, there is little known about three facets of S-R transfer between spatial incompatibility and Simon tasks. First, current research has emphasized S-R transfer in one direction (i.e., the effect of practicing a spatial incompatibility task

on the Simon effect). Very little is known about the impact of non-spatial, task-relevant S-R mappings in a Simon task on spatial incompatibility tasks. The potential existence of bidirectional S-R transfer has implications for our understanding of the limitations of S-R transfer. Secondly, there has been little work on the time course of Simon effects following transfer from a spatial incompatibility task. Speed-accuracy trade-off (SAT) approaches to the time course of the Simon effect do not possess the same disadvantages as other, more common, RT distributional analyses (e.g., Zhang and Kornblum, 1997). Lastly, context of the task is known to play a critical role in memory transfer (Smith and Vela, 2001), but its place in S-R transfer has not yet been firmly established.

BIDIRECTIONAL S-R TRANSFER

Very little research has explored the effect of task-relevant S-R mappings from a Simon task on performance in a spatial compatibility or incompatibility task. One reason for the paucity of attention on bidirectional S-R transfer is paradigmatic. Practice tasks (e.g., Proctor and Lu, 1999; Tagliabue et al., 2000) typically only include two blocks of trials—practice and test—in a fixed order (i.e., spatial compatibility or incompatibility task followed by the Simon task), thus not permitting an evaluation of the effects of S-R mapping in the Simon task on the spatial compatibility and incompatibility tasks. The other reason that bidirectional transfer is typically not explored is methodological. In mixing tasks (Marble and Proctor, 2000; Proctor et al., 2000), a non-spatial feature of the stimulus informs the participant to perform a left, right or spatial compatibility (or incompatibility) task. For example, the color of the stimulus in Marble and Proctor's (2000) task informed the participant to make a particular response (i.e., a red or green stimulus informed participants to make a left or right response), while another color (white) instructed the participant to make a spatially compatible (or incompatible) response. Accordingly, this particular methodology does not permit the researcher to explore the effect of task-relevant (non-spatial) Simon task S-R mappings on performance in the spatial compatibility or incompatibility tasks.

The exception to this lack of attention to bidirectional transfer is Proctor and Lu (1999; Experiment 2). In their task, participants made left or right responses to letters (S or H) presented to the left or right side of the screen. They practiced this Simon task in three sessions before transfer to a spatial compatibility or incompatibility task. Although the authors initially failed to record letter identity in the transfer session, a subsequent study corrected this oversight and they found no effect of letter identity on RTs within spatial compatibility/incompatibility tasks; however, there was an effect of letter identity on error rate (i.e., there were more errors when the response assigned to the letter was incongruent with the location in both compatibility and incompatibility tasks). It is not clear why the congruency effect only influenced error rates. Proctor and Lu did not discuss the implication of their finding in great detail.

There are a few theoretical implications for considering the effect of task-relevant non-spatial mappings from a Simon task on performance in a spatial compatibility or incompatibility task. Firstly, location information generally precedes selection for color

or shape (Hillyard and Munte, 1984). If prior S-R mappings between (slow) non-spatial features do not affect (fast) responses to location, then S-R transfer may hinge on temporal precedence. Second, the lack of evidence for S-R transfer from Simon to spatial compatibility/incompatibility tasks might suggest that S-R transfer is closely tied to spatial features of the stimulus. Lastly, it is possible that non-spatial S-R mappings are relatively weak and, consequently, do not transfer once the task is abandoned. In the mixed-task paradigm (Marble and Proctor, 2000), however, task-relevant S-R mappings from the Simon task should not be completely abandoned in the spatial compatibility/incompatibility tasks because they will be needed once again, once the Simon task cue is reintroduced. In the current study, the stimuli on Simon and spatial compatibility/incompatibility trials were identical in a variant of the mixed-task paradigm to provide a fertile opportunity to detect bidirectional S-R transfer. A cue precedes a block of trials informing participants to engage in a particular task (i.e., a Simon or spatial compatibility/incompatibility task). This methodology allows for an examination of the effects of non-spatial S-R mappings from the Simon task on performance in the spatial compatibility and incompatibility tasks.

TIME COURSE OF THE SIMON EFFECT FOLLOWING A SPATIAL INCOMPATIBILITY TASK: VINCENTIZING REACTION TIMES AND SPEED-ACCURACY TRADE-OFFS

The time course of the Simon effect has played a critical role in the theoretical development of purported mechanisms behind the effect (Ridderinkhof, 2002). Although a number of chronometric approaches purport to measure the unfolding of mental processes (Meyer et al., 1988), one approach in particular has been widely used in the Simon effect literature. de Jong et al. (1994) were the first to use *vincentized* RTs (Ratcliff, 1979) to study the time course of the Simon effect. According to this approach, RTs are rank ordered, divided into *bins* (quartiles, quintiles, and deciles are most commonly used), and then averaged within a bin for each condition. When the corresponding mean RT for each bin is subtracted from the non-corresponding mean RT, it is referred to as a *delta plot* (Ridderinkhof, 2002). The delta plot of the Simon effect has been interpreted as a direct measure of task-irrelevant spatial response activity. Most studies of the standard Simon effect have demonstrated negative-going slopes with the delta plot approach (Schwarz and Miller, 2012), although there are some exceptions (see Proctor et al., 2011). The interpretation of the decreasing Simon effect has been controversial, with some suggesting a passive decay of task-irrelevant activity along the direct pathway, while others suggest the direct pathway is actively suppressed (see Proctor et al., 2011, for a review of the literature).

The interpretation of delta plots is not without its challenges (Zhang and Kornblum, 1997; Schwarz and Miller, 2012). Zhang and Kornblum (1997) pointed out that the negative-going slopes of delta plots from Simon tasks simply derive from the shapes of corresponding and non-corresponding RT distributions. In particular, smaller variance in the non-corresponding condition, relative to the corresponding condition, gives rise to a negative-going slope (see also Pratte et al., 2010). Keep in mind that this description of the RT distribution does not, however, presuppose

a particular mechanism (Schwarz and Miller, 2012). Moreover, delta plots of RTs do not account for error rates.

Error rates are often considered secondary to RT in many tasks, even though they can reveal valuable information about performance. For instance, Hilchey et al. (2011) recently examined the Simon effect using two different measures of response accuracy within the context of an SAT task. A symbol (i.e., ⊗ or a ⊕) was used to instruct participants to make a left (L) or right (R) response. The location of the task-relevant symbol could be to the left, l, or right, r, of fixation. Hilchey et al. calculated the sensitivity (d') to the task-relevant (identity-based: d'_{id}), and task-irrelevant (location-based: d'_{loc}), features of the target. These calculations were possible because of the orthogonal relationship between the identity (⊗ or ⊕) and location (l or r) of the stimulus. First, ignoring the spatial correspondence between stimulus and response, d'_{id} was calculated according to the identity of the stimulus (and task instructions):

$$ d'_{id} = \frac{z\left[p(L\,|\,\otimes)\right] - z\left[p(L\,|\,\oplus)\right]}{\sqrt{2}}, \qquad (1) $$

where $z[\,]$ is the inverse of the standard normal cumulative distribution. The divisor is a standard correction when the signal detection approach is applied to alternative forced choice designs (Macmillan and Creelman, 2005). The probability of responding with a "left" response given the ⊗ stimulus, $p(L|\otimes)$, is also a *hit* within the framework of signal detection theory. On the other hand, the probability of responding L given the ⊕ stimulus [$p(L|\oplus)$] is a *false alarm* error. Hilchey et al. (2011) observed that d'_{id} increased with time, presumably reflecting evidence accrual along the task-relevant, indirect, non-spatial S-R pathway.

The second way in which Hilchey et al. (2011) assessed sensitivity was according to the location of the stimulus. Sensitivity to the location of the stimulus (d'_{loc}) was calculated with the signal detection framework, this time ignoring the non-spatial stimulus identity. It was calculated according to the following equation:

$$ d'_{loc} = \frac{z\left[p\,(L|l)\right] - z[p(L|r)]}{\sqrt{2}}. \qquad (2) $$

Here, the probability of responding with a left response, to a stimulus presented on the left side of space, $p(L|l)$, is a *hit*. The probability of responding with a left response to a stimulus on the right, $p(L|r)$, is a *false alarm* error. Hilchey et al. (2011) observed that this measure of sensitivity decreased with time. This measure is strongly related to the performance difference between corresponding and non-corresponding trials (see Hilchey et al., 2011). In signal detection theoretic terms, d'_{loc} most closely captures Simon's (1969) interpretation of the Simon effect: it reflects the sensitivity to the location of the stimulus. In other words, the d'_{loc} score presumably reflects the combined impact of the direct, and indirect, spatial S-R pathways (illustrated in **Figure 1**) on response selection. Although this measure approximated an exponential decay function with response lag (time), this kind of function has yet to be quantitatively fitted to data.

Speed-accuracy tradeoffs: methodology and functions

Although there are a number of methodological approaches for measuring SAT functions, the response-signal technique is arguably one of the most common (e.g., Schouten and Bekker, 1967; Reed, 1973; McElree and Carrasco, 1999; Carrasco and McElree, 2001). With this procedure, participants are presented with a target stimulus (usually a visual stimulus) and they withhold responding until the onset of a response signal (usually a simple auditory tone). Following the response signal, there is a short (≤ 300 ms) window in which responses are collected. Responses that precede the window, or follow it, are typically discarded. Varying the interval between the target onset and the response signal controls reaction time. SAT functions allow for a quantitative description of the time course of an effect as an alternative to the delta plot approach (Pachella, 1974; Wickelgren, 1977; Salthouse and Hedden, 2002). SAT functions typically plot d' as a function of response lag (i.e., time). Response lag is the sum of the mean reaction time to the response signal (i.e., RTs within the response window) and the stimulus onset asynchrony (SOA) between the target and the response signal. The following is a general equation for the SAT function (Wickelgren, 1977) that has been widely used to describe the trading relation between speed and accuracy:

$$ d'_{id}(t) = \lambda\left[1 - e^{-\beta(t-\delta)}\right], \; for\; t > \delta, \; else\; 0, \qquad (3) $$

where t is the mean response lag, λ is the asymptotic d' value, β is a rate parameter, and δ is the intercept. This SAT function, describing accumulation of evidence to a maximum, is used to fit many different sorts of SAT datasets and seems to fit just as well as other equations (McElree and Dosher, 1989).

Although the exponential SAT function in Equation 3 is quite common, Wickelgren once suggested that "no one knows the correct mathematical form for the speed-accuracy tradeoff function for any cognitive process, so the exponential approach to a limit… should be taken solely as an example" (Wickelgren, 1977, p. 70). One potentially serious challenge to this function is that, in practice, early data points close to the intercept sometimes rise slowly from the baseline, not as abruptly as is assumed in Equation 3. The standard SAT equation does not account for any changes in d'_{id} between $t = 0$ and δ (one of the parameters to be determined by the fitting process). Thus, Equation 3 is rather unusual as many psychometric functions generally follow an ogive, or an S-shaped, function (Gescheider, 1997) where there is gradual change in the dependent measure (plotted along the y-axis) at the extremes of the independent variable (plotted along the x-axis). Thus, as an alternative approach to the standard SAT function in Equation 3, it seems reasonable to include gradual, rather than abrupt, evidence accrual into the function. Accordingly, a hyperbolic tangent function might capture the slight accumulation of evidence from an assumed $d' = 0$ (at $t = 0$):

$$ d'_{id}(t) = \frac{\lambda}{2}\left[1 + \tanh\left(\frac{t - \omega}{\kappa}\right)\right], for\; t \geq 0, \qquad (4) $$

where λ is the asymptotic value, ω is a shift parameter (i.e., reflecting the time at which the function reaches 50% of λ) and κ reflects

the speed of the transition from the initial region where $d' = 0$ to the final region where d' takes on its asymptotic value of λ. Unlike the standard SAT equation, Equation 4 models the entire timecourse, from $t = 0$ to asymptote (λ). This hyperbolic tangent function produces an ogive-shaped curve that permeates much of psychophysics (Gescheider, 1997).

Neither of these functions, however, adequately captures the decreasing sensitivity to location information Hilchey et al. (2011) observed. However, it does appear that d'_{loc} may fit a simple exponential decay function:

$$d'_{loc}(t) = \delta e^{(-\beta t)}, \; for \; t \geq \; 0, \tag{5}$$

where δ is the peak d'_{loc} value at $t = 0$ and β is a decay rate parameter. It has yet to be determined how well d'_{loc} data fit this function.

The goodness of fit of SAT functions is typically assessed using an adjusted R^2 (Dosher et al., 2004) which includes a penalty for increasing the number of parameters:

$$R^2_{adj} = 1 - \frac{\sum_{i=1}^{n} \left(d_i - \widehat{d_i} \right)^2 / (n-k)}{\sum_{i=1}^{n} \left(d_i - \overline{d} \right)^2 / (n-1)}, \tag{6}$$

where k is the number of free parameters, n is the number of data values, d_i are the observed d_i values, $\widehat{d}_i$ are the predicted d_i from the model, and $\overline{d}$ is the mean.

Using SAT functions to explore the time course of spatial information processing in a Simon task has a possible benefit over distributional analyses (e.g., vincentizing or delta plots) in that it captures response *decisions* at a given time and is therefore practically immune to the different distributional properties of corresponding and non-corresponding RTs (Zhang and Kornblum, 1997).

The time course of the Simon effect that follows a spatial incompatibility task is unlike what one usually sees with a standard Simon task. In studies that have included a vincentized analysis of RT, the reverse Simon effect, resulting from prior or concurrent experience with a spatial incompatibility task, *increases* with increasing RT (Marble and Proctor, 2000; Proctor and Vu, 2009). This time course seems rather unnatural, as there is no *a priori* theoretical reason to suppose that a reverse Simon effect should not be actively suppressed or naturally decay with time (but see Tagliabue et al., 2000). The use of vincentized RTs as a measure of time course is convenient, but as previously discussed, it is not without its interpretational challenges. Here, we use SAT functions to explore the full temporal dynamics of the reverse Simon effect that follows from mixing a spatial incompatibility task with a Simon task.

THE ROLE OF TASK CONTEXT ON S-R TRANSFER

Surprisingly, there has been little investigation into the effect of environmental context on S-R transfer effects in Simon tasks. Recognition performance is often best when the testing conditions resemble those in training (e.g., Godden and Baddeley, 1975). Context plays an important role in memory (Smith, 1994; Murnane et al., 1999), perhaps because incidental environmental

features are usually encoded with task-relevant information, unless intentionally suppressed (Smith and Vela, 2001). One recent study (Milanese et al., 2011) explored the effect of practicing a spatial incompatibility task with a partner on a subsequent *social* Simon task, also performed with a partner. Like the standard version of this paradigm, where only one individual performs the task, the social Simon effect reverses when it follows practice with a spatial incompatibility task (Milanese et al., 2010). Milanese et al. (2011) observed that switching partners between tasks did not eliminate the reverse social Simon effect. Given that the identity of the partner was not integral to the task, it is likely that it would not be a salient feature of the task context. When the partners changed positions (i.e., from the left side to the right side), however, there was no effect of the spatial incompatibility task on the Simon effect. In this task, one's position relative to the partner is a stimulus feature that is critical to performing the task properly. Thus S-R transfer may depend on task-relevant, salient features.

Another paper (Yamaguchi and Proctor, 2009) considered response mode to be an integral part of context. Yamaguchi and Proctor's (2009) participants performed a spatial incompatibility task by responding to stimuli on a keyboard or a joystick. Participants then performed a Simon task (with a keyboard or a joystick), where the color of the stimulus was task-relevant and the location was task-irrelevant. When the response mode was consistent across tasks the reduction of the Simon effect (from the spatial incompatibility task) was generally greater than when the response mode did not match. Thus, response mode may provide a context that modulates S-R transfer.

There are two reasons to expect a contextual modulation of S-R transfer across tasks in the present study. First, the response-signal methodology (used to acquire SAT functions) is quite different from standard RT tasks in which instructions emphasize both the speed and accuracy of performance. The response-signal methodology includes auditory signals and visual feedback that are not present in the standard RT tasks. These components are necessary to control RT in SAT tasks. Second, previous work has demonstrated a switch cost when switching between tasks with different speed-accuracy instructions (Gopher et al., 2000). This switch cost suggests that SAT settings constitute part of a task-set. Thus, it is expected that when spatial compatibility/ incompatibility and Simon task contexts are similar (i.e., they are both SAT or standard RT tasks) maximal S-R transfer should occur.

THE PRESENT STUDY

The current investigation used a mixing task, where Simon and spatial compatibility (or incompatibility) tasks were signaled with a task cue and alternated predictably every eight trials. Unlike previous experiments using the mixed-task methodology (Marble and Proctor, 2000; Experiment 1), the stimuli in the present study were identical in both the Simon and spatial compatibility/incompatibility tasks. In Experiment 1, both the Simon and spatial compatibility effects were measured in standard RT tasks. In Experiment 4, they were measured in SAT tasks. To date, no study has used SAT methodology to study the temporal dynamics of S-R transfer from spatial incompatibility tasks to Simon tasks.

In Experiment 2, the spatial compatibility/incompatibility task was administered with the response-signal methodology, while the Simon task was a standard RT task. In Experiment 3, the reverse was true. Unlike Experiments 1 and 4, the spatial compatibility/incompatibility and Simon task contexts in Experiments 2 and 3 do not match.

EXPERIMENT 1

Participants were provided with a visual cue every eight trials instructing them to perform the Simon task or the spatial compatibility (or incompatibility) task. The instructions of each task equally emphasized the speed and accuracy of responding. One group of participants performed a spatial incompatibility task with the Simon task while another group performed a spatial compatibility task with the Simon task. The stimuli in all tasks are identical. A cue presented at the onset of a block of trials informed participants of the task to perform. The purpose of this experiment was to (1) replicate the reversal of the Simon effect when paired with a spatial incompatibility task, (2) identify the effect of the task-relevant S-R mappings from a Simon task on spatial compatibility and incompatibility tasks, and (3) determine whether transfer occurs in a version of mixed-task design (Marble and Proctor, 2000; Proctor et al., 2000) where the task is predictably cued and stimuli are identical across tasks.

METHODS

Participants

Sixteen undergraduate participants from Saint Mary's University took part in the spatial compatibility condition and sixteen took part in the spatial incompatibility condition. All participants were between 18 and 30 years of age. All experiments were approved by the Saint Mary's University Research Ethics Board (REB) in accordance with the Tri-council Policy Statement on Ethical Conduct for Research Involving Humans (Canadian Institutes of Health Research, Natural Sciences and Engineering Research Council of Canada, and Social Sciences and Humanities Research Council of Canada, Tri-Council Policy Statement: Ethical Conduct for Research Involving Humans, December 2010).

Apparatus and stimuli

The experiment was conducted on an Apple iMac G3/400 DV computer, powered by a 400MHz Power PC 750 (G3) processor, running OS9. Superlab (ver 1.75; Cedrus, CA) was used to present stimuli. The experiment took place in a quiet room with ambient lighting. Responses were executed by pressing, with index fingers, the "z" and "/" keys on a standard QWERTY Apple keyboard.

The viewing distance was approximately 57 cm. There were three types of cues: (1) "Sym" (symbol task) to signal the Simon task, (2) "Same" (same sided response) to signal a spatially compatible response, and (3) "Opp" (opposite sided response) to signal a spatial incompatibility task. The task cues were 0.75° vertically and 1.5° ("Sym"/"Opp") or 2.0° ("Same") wide. Three horizontally arranged square box outlines (1.2° × 1.2°) were used as placeholders for the stimuli. The peripheral placeholders were 5.3° (edge-to-edge) from the central placeholder. The fixation point, a circle with a diameter of 0.8°, was presented in the center placeholder. The target stimuli, ⊗ and ⊕, were presented within a circle of 1.2° in diameter. These targets were placed inside either the left or right placeholder. All images were black on a white background.

Procedure and design

Each participant underwent 128 trials, equally split between Simon and spatial compatibility/incompatibility tasks. All stimuli and responses were equally balanced between left and right positions. The starting task was randomly determined. A block of eight trials in a particular task alternated with roughly half the group starting with the Simon task the rest starting with the spatial compatibility/incompatibility task. Each block of eight trials was preceded by the 900 ms presentation of the task cue. Following the task cue, a trial was presented. The sequence of trial events was as follows: blank screen (300 ms), fixation display (450 ms), and target (until response).

Each group took part in two tasks: Simon and spatial compatibility tasks or Simon and spatial incompatibility tasks. No feedback was provided for these tasks, and participants were told to respond as fast and accurately as possible.

Spatial compatibility task. Participants were presented with "Same" cue (900 ms) at the beginning of the first trial for every block of spatial compatibility task trials, indicating response to the same-side as stimulus location. Therefore, stimuli presented on the right of the fixation point required "/" key responses and stimuli on the left required "z" key responses.

Spatial incompatibility task. The spatial incompatibility task was the same as the spatial compatibility task with the following exceptions. Participants were presented with "Opp" cue (900 ms) at the beginning of the first trial for every block of the spatial incompatibility task, indicating response to the opposite-side of stimulus location. Therefore, stimuli presented on the right of the fixation point required left ("z") key response and stimuli presented to the left of fixation required "/" key response.

Simon task. Participants were presented with "Sym" cue (900 ms) at the beginning of the first trial for every location-irrelevant block, indicating they were to respond to the non-spatial identity of the target (i.e., the symbol). Presentation of the ⊗ stimulus indicated a left response while the presentation of the ⊕ stimulus indicated a right response, regardless of the location of the stimulus.

RESULTS AND DISCUSSION

RTs in each condition were subject to a recursive procedure eliminating trials with RTs that were less than or greater than 3.5 SDs from the mean. This procedure generally eliminated fewer than 5% of all trials across subjects.

Simon task

Table 1 presents the mean RTs for the Simon task. A 2 (Simon correspondence: corresponding and non-corresponding) × 2 (group: spatial compatibility and spatial incompatibility) mixed ANOVA revealed the expected interaction between Simon correspondence and group [$F_{(1, 30)} = 14.35$, $MSE = 599.93$,

Table 1 | RTs in the Simon tasks in Experiments 1 and 2.

SOA	Compatibility group			Incompatibility group		
	Corr.	Non-corr.	Simon effect	Corr.	Non-corr.	Simon effect
EXP. 1						
	570	599	29**	602	565	−37**
EXP. 2						
60	491	508	17	491	506	15
120	491	509	18*	509	505	−4
240	483	514	31**	509	525	16
360	496	517	21	525	525	0
480	520	548	28*	552	568	16
960	546	551	5	539	554	15
1440	517	559	42**	547	551	4

SOA, stimulus-onset asynchrony; Corr., Corresponding; Non-corr., Non-corresponding.

*Value is significantly different from zero, p < 0.05.

**Value is significantly different from zero, Bonferroni corrected.

$p < 0.001$]. There was a standard 29 ms Simon effect in the spatial compatibility group and a -37 ms Simon effect in the spatial incompatibility group. This finding replicates a number of papers in the literature demonstrating a reverse Simon effect when it is presented following, or within the context of, a spatial incompatibility task (e.g., Marble and Proctor, 2000; Tagliabue et al., 2000).

The mean sensitivity (d'_{loc} and d'_{id}) values for the Simon task are presented in **Table 2**. The d'_{loc} and the d'_{id} were compared with an unpaired t-test across the incompatible and compatibility groups. Sensitivity to the task-relevant instructions (d'_{id}) was significantly higher in the compatibility group ($d'_{id} = 2.27$) than it was in the incompatibility group ($d'_{id} = 1.73$), $t_{(30)} = 3.96$, $p < 0.001$. Sensitivity to the location of the stimulus also differed significantly [$t_{(30)} = 9.92$, $p < 0.001$] between the two groups (compatible $d'_{loc} = 0.13$; incompatible $d'_{loc} = -0.25$). Both of these effects were significantly different from $d'_{loc} = 0$ ($ps < 0.001$), suggesting that engaging a spatial incompatibility task reverses the tendency to respond toward the source of stimulation.

The Simon effect reversed when the Simon task alternated with a spatial incompatibility task. No such reversal was evident in the control Simon task that alternated with the spatial compatibility task. This finding is consistent with other studies using a different mixed-task design (Marble and Proctor, 2000; Proctor et al., 2000). The spatial incompatibility task also affected the d' measures in the Simon task. Consistent with the pattern of RTs, sensitivity to the location of the stimulus (d'_{loc}) was *positive* in the compatibility group, indicating a tendency to respond *toward* the location of the stimulus. In contrast, the same measure was *negative* in the incompatibility group, indicative of a tendency to respond *away* from the stimulus.

Spatial compatibility and incompatibility tasks

Trials were sorted into *congruent* and *incongruent* conditions for each task, where congruency reflects the match between the *response assigned to the identity* of the stimulus (i.e., $\otimes$ or $\oplus$)

Table 2 | d′ scores in the Simon tasks in Experiments 1–4 as a function of SOA in the other (spatial compatibility and incompatibility) task (Experiment 2), and as a function of SOA in the Simon (SAT) tasks (Experiments 3 and 4).

SOA	Compatibility group		Incompatibility group	
	d'_{id}	d'_{loc}	d'_{id}	d'_{loc}
EXP. 1				
	2.27**	0.12**	1.73**	−0.25**
EXP. 2				
60	1.74**	0.21*	1.81**	−0.04
120	1.97**	0.11	1.93**	0.00
240	2.12**	0.83	2.06**	−0.07
360	2.20**	0.13**	2.24**	−0.04
480	2.36**	0.06*	2.30**	0.04
960	2.26**	0.03	2.27**	0.00
1440	2.20**	0.11*	2.27**	−0.03
EXP. 3 (SAT)				
60	0.21*	1.04**	0.13	0.15
120	0.19	0.72**	0.50**	−0.18
240	0.86**	0.39*	0.96**	−0.16
360	1.71**	0.06	1.82**	−0.19
480	2.33**	0.10	2.29**	−0.04
960	2.88**	−0.01	2.72**	−0.02
1440	2.90**	0.03	2.88**	0.00
EXP. 4 (SAT)				
60	0.02	1.14**	0.06	0.22*
120	0.02	1.21**	0.26*	−0.03
240	0.63**	0.62**	0.65**	−0.04
360	1.60**	0.41*	1.61**	−0.14
480	2.17**	0.16*	1.84**	−0.06
960	2.22**	0.26*	2.67**	−0.05
1440	2.43**	0.14*	2.57**	−0.07*

SAT, speed-accuracy trade-off; SOA, stimulus-onset asynchrony.

*Value is significantly different from zero, p < 0.05.

**Value is significantly different from zero, Bonferroni corrected.

in the Simon task and the *location of the stimulus* in the spatial compatibility (or incompatibilty) task. The RTs were entered into a 2 (congruency: congruent and incongruent) × 2 (group: compatible and incompatible) mixed ANOVA. Only the interaction between congruency and group was significant, $F_{(1, 30)} = 12.15$, $MSE = 204.97$, $p < 0.005$. The congruency effect was *negative* (incongruent − congruent = −14.3 ms) and significant in the compatibility group [$t_{(15)} = 2.55$, $p < 0.05$]. In the incompatible condition, the congruency effect was *positive* (+11 ms), $t_{(15)} = -2.39$, $p < 0.05$. The mean RTs are presented in **Table 3**.

As in the Simon task, we compared the measures of identity and location sensitivity (d'_{id} and the d'_{loc}, respectively) across the incompatible and compatibility groups. There was no difference in d'_{id} in the compatible ($d'_{id} = 0.13$) and incompatibility groups ($d'_{id} = 0.09$); however both of these effects were significantly different from 0 ($p < 0.005$), suggesting a small, but significant, sensitivity to the target's identity (an irrelevant feature within the context of the spatial compatibility

Table 3 | RTs in the spatial compatibility and incompatibility tasks in Experiments 1 and 3 as a function of congruency and SOA (Experiment 3) in the Simon task.

SOA	Compatibility task			Incompatibility task		
	Con.	Incon.	Con. Effect	Con.	Incon.	Con. Effect
EXP. 1						
	370	356	−14**	381	392	11*
EXP. 2						
60	278	282	4	319	313	−6
120	290	282	−8	331	332	1
240	284	287	3	331	330	−1
360	290	286	−4	344	343	−1
480	296	293	−3	362	357	−5
960	308	322	14	368	361	−7
1440	291	296	5	348	352	4

Con., congruent; Incon., incongruent; Con. Effect, congruency effect; SOA, stimulus-onset asynchrony.

*Value is significantly different from zero, p < 0.05.

**Value is significantly different from zero, Bonferroni corrected.

task). As expected, d'_{loc} was significantly different [$t_{(30)} = 27.42$, $p < 0.001$] between the compatible ($d'_{loc} = 2.42$) and incompatibility groups ($d'_{loc} = -2.56$), demonstrating that participants were following instructions. The mean d' values are provided in **Table 4**.

Previous studies have addressed the effect of spatial compatibility tasks on the Simon effect. However, few investigations have addressed the effect of task-relevant S-R mappings from the Simon task on performance in spatial compatibility or incompatibility tasks. Proctor and Lu (1999; Experiment 2), assessed the effects of repeated practice with a Simon task on performance of a spatial compatibility and incompatibility tasks. While there was no effect of congruency on RTs, they did observe an effect on error rates. In the present study, the congruency effect in the incompatibility group was consistent in RTs and d'_{id}, demonstrating a performance *advantage* when the identity-response mapping in the Simon effect converges with the location-response mapping in the incompatibility task. Interestingly, a different pattern emerged with the compatibility task. While responses were more accurate when the identity-response mapping in the Simon effect was congruent with the location-response mapping in the compatibility task, responses were faster when the mappings were incongruent. This trade-off between accuracy and speed suggests that S-R mappings in a Simon task have distinct effects on performance in spatial compatibility and incompatibility tasks.

EXPERIMENT 2

In Experiment 2, the spatial compatibility and incompatibility tasks were subjected to the response-signal methodology for a SAT analysis. The Simon task was the same as it was in Experiment 1 (i.e., standard RT task) and alternated with the spatial compatibility/incompatibility block across all response-signal SOAs. Thus, in Experiment 2, the Simon and spatial compatibility/incompatibility task contexts did not match.

Table 4 | d' Scores in the spatial compatibility and incompatibility tasks in Experiment 1, as a function of SOA in the Simon (SAT) task in Experiment 3, and as a function of SOA in the spatial compatibility and incompatibility tasks (SAT) in Experiments 2 and 4.

SOA	Compatibility task		Incompatibility task	
	d'_{id}	d'_{loc}	d'_{id}	d'_{loc}
EXP. 1				
	0.13**	2.42**	0.09**	−2.56**
EXP. 2 (SAT)				
60	0.00	2.48**	0.03	−1.35**
120	−0.03	2.70**	0.01	−2.10**
240	−0.02	2.83**	0.02	−2.44**
360	0.10**	2.71**	0.04	−2.75**
480	0.07	2.64**	0.09	−2.49**
960	0.03	2.84**	0.03	−2.87**
1440	−0.01	2.91**	−0.02	−2.84**
EXP. 3				
60	0.00	2.55**	0.01	−2.36**
120	0.01	2.65**	−0.01	−2.46**
240	−0.02	2.71**	0.02	−2.66**
360	0.04	2.79**	0.01	−2.67**
480	0.01	2.71**	−0.01	−2.77**
960	0.01	2.78**	−0.01	−2.75**
1440	0.01	2.78**	−0.03	−2.75**
EXP. 4 (SAT)				
60	0.02	2.46**	0.07	−0.92**
120	−0.03	2.81**	−0.02	−1.86**
240	0.00	2.76**	0.03	−2.36**
360	0.06	2.38**	0.12*	−2.36**
480	0.02	2.68**	0.06	−2.02**
960	0.10	2.46**	0.03	−2.72**
1440	0.19*	2.45**	0.14*	−2.58**

SAT, speed-accuracy tradeoff task; SOA, stimulus-onset asynchrony.

*Value is significantly different from zero, p < 0.05.

**Value is significantly different from zero, Bonferroni corrected.

METHODS

Participants

Fifteen undergraduate participants took part in each of the compatible/Simon and incompatible/Simon tasks for course credit and monetary bonuses. To encourage participants to make timely responses in the SAT task, they received a penny for every response that fell within the 240 ms response window and an additional penny for a correct response.

Procedure

The general procedure was the same as it was in Experiment 1 with the following exceptions. The Simon task was exactly like it was in Experiment 1, however it alternated with the spatial compatibility or incompatibility task in each block of trials. There were seven blocks of trials with 128 trials each. The response-signal methodology was applied to the compatibility and incompatibility tasks, but not the Simon task. In the spatial compatibility and incompatibility tasks, the ⊗ or ⊕ appeared in

one of the two peripheral placeholders, left or right of fixation, for 60ms. As before, the $\otimes$ and $\oplus$ stimuli, presented to the left or right, were presented with equal frequencies. The response signal, an auditory tone (44.1 KHz, 15 ms), was presented following the onset of the stimulus after a delay (i.e., the SOA). The SOA was fixed within a block. There were seven SOAs: 60, 120, 240, 360, 480, 960, and 1440 ms. Participants were required to respond within a 240 ms response window following the tone. They were also provided visual feedback with respect to the timing, but not the accuracy, of their response. Participants were presented feedback "HIT" when responding within 240 ms of tone, "MISS" when responding more than 240 ms after the tone, and "TOO SOON" when responding prior to the tone. Thus, responding within the response window took precedence over response accuracy. This prioritization reliably encouraged participants to trade accuracy for speed at the shorter SOAs.

RESULTS AND DISCUSSION

The RTs in the Simon task were trimmed as before. The d' scores were calculated as they were in Experiment 1. The analysis of the spatial compatibility and incompatibility task was much like other SAT analyses. First, only responses that fell within the 240 ms time window following the tone were analyzed. Second, *response lag* was measured in the SAT task, not RT. Response lag is an estimate of the average response time, relative to the response signal (the tone), within the response window *plus* the SOA (for example, if a response was made 129 ms following the tone when the SOA was 360 ms, the response lag would be 489 ms for that particular trial). Lastly, d'_{loc} and d'_{id} were estimated for each SOA. The d' estimate for each participant was the mean of a bootstrapping procedure. Ten thousand samples (with replacement) were taken from each SOA using the base number of trials found in the SOA with the most trials discarded (i.e., because the responses fell outside of the response window). This bootstrapping procedure was used to ensure that d' values were not artificially deflated across SOAs due to missing trials (i.e., when responses did not fall within the response window). Trials with perfect scores were adjusted according to the conventional $0.5f$ recommendation (Kadlec, 1999).

Simon task

RTs are presented in **Table 1** as a function of SOA in the spatial compatibility task. The Simon trials were separated according to the spatial compatibility/incompatibility group, correspondence (corresponding and non-corresponding), and the SOA from the spatial compatibility/incompatibility task (60, 120, 240, 360, 480, 960, and 1440 ms) and entered into a 2 × 2 × 7 mixed ANOVA. There was a main effect of correspondence [$F_{(1, 28)} = 12.32, MSE = 2238.14, p < 0.005$], with a 15 ms Simon effect overall. There was also a main effect of SOA [$F_{(6, 168)} = 8.58, MSE = 3294.13, p < 0.001$], where the overall RT in the Simon task increased with the SOA in the spatial compatibility/incompatibility task. Surprisingly, there was no significant interaction between correspondence and group, although there was a numerical reduction in the Simon effect in the incompatibility group (incompatibility group: 9 ms [$t_{(14)} = 1.19, p > 0.25$]; compatibility group: 23 ms [$t_{(14)} = 4.39, p < 0.001$]).

The d' values are presented in **Table 2** as a function of SOA in the spatial compatibility task. The d'_{id} values were entered into a 2 (group) × 7 (SOA) ANOVA. Only the main effect of SOA was significant, $F_{(6, 168)} = 8.69, MSE = 0.14, p < 0.001$. With increasing SOA, the d'_{id} values also increased. This suggests that the speed-accuracy setting in the spatial compatibility/incompatibility task transferred to the Simon task. The d'_{loc} values were also entered into the same 2 × 7 ANOVA. The main effect of group was significant [$F_{(1, 28)} = 14.82, p < 0.005$], demonstrating an overall greater effect of location in the Simon task within the compatibility group ($d'_{loc} = 0.15$) than with the incompatibility group ($d'_{loc} = -0.03$). The interaction between SOA and group was significant, $F_{(6, 168)} = 3.08, MSE = 0.17, p < 0.01$. We compared the d'_{loc} values between compatible and incompatibility groups at each SOA. d'_{loc} was greater for the compatibility group than the incompatibility group at 60 ms [$t_{(28)} = 3.56, p < 0.005$], 240 ms [$t_{(28)} = 2.54, p < 0.05$], 360 ms [$t_{(28)} = 4.17, p < 0.0005$], and 1440 ms [$t_{(28)} = 3.16, p < 0.005$].

Unlike Experiment 1, there was no evidence of a reversal of the Simon effect in Experiment 2. In fact, there was surprisingly weak evidence of an influence of the spatial incompatibility task on the Simon effect. Individual mean corresponding and non-corresponding RTs from the Simon tasks for the spatial incompatibility groups in Experiment 1 and 2 were entered into 2 (correspondence) × 2 (Experiment) ANOVA. The interaction between correspondence and Experiment was significant [$F_{(1, 29)} = 15.82, MSE = 501.62, p < 0.0005$], backing the claim that S-R transfer was not as strong in Experiment 2 as it was in Experiment 1. Moreover, the same analysis on the Simon effect for the spatial compatibility groups revealed no interaction between Experiment and correspondence [$F_{(1, 29)} = 0.54, MSE = 423.05, p = 0.47$]. This finding supports the idea that there is no S-R transfer from spatial compatibility tasks to Simon tasks (Tagliabue et al., 2000). Together, this supplementary analysis suggests that the context of the task has an important modulating influence on S-R transfer effects from spatial incompatibility tasks to Simon tasks.

Spatial compatibility and incompatibility tasks

The performance in the spatial compatibility task proved quite easy, as d'_{loc} was near ceiling across all SOAs (see **Table 4**). This suggests that information processing along the direct spatial pathway is very quick. For those in the incompatibility task, d'_{loc} was slightly impaired at the earliest lags, but still not enough for proper curve-fitting as values were still quite far from chance. We analyzed the spatial compatibility task by entering the d' values into a 2 (group) × 7 (SOA) ANOVA. The analysis of d'_{loc} revealed a main effect of SOA [$F_{(6, 168)} = 10.33, MSE = 0.13, p < 0.001$], an expected large main effect of group [$F_{(1, 28)} = 2671.44, MSE = 0.52, p < 0.001$], and the SOA × group interaction [$F_{(6, 168)} = 25.80, MSE = 0.13, p < 0.001$]. Although the SOA effect was significant in both groups, the difference between d'_{loc} at the longest SOA and the shortest SOA was much greater in the incompatible condition than it was in the compatible condition (**Table 4**).

The analysis of d'_{id} also included SOA and group as factors. Only the main effect of SOA was significant [$F_{(6, 168)} = 2.30$, $MSE = 0.02$, $p < 0.05$]. As seen in **Table 4**, there was a significant increase in d'_{id} values at 360 ms, but only in the compatibility group did the d'_{id} values deviate significantly from zero.

There was little evidence of S-R transfer from the Simon task to the spatial compatibility or incompatibility tasks. The only effect of S-R mappings from the Simon task on the spatial compatibility task (i.e., the congruency effect) occurred at the 360 ms SOA. However, there was no *a priori* reason to expect the effect to be restricted to a single SOA. There was also little reason to expect that S-R transfer would not occur from the Simon task to the spatial incompatibility task. Thus, Experiment 2 did not replicate the observation in Experiment 1 of S-R identity transfer from the Simon task to the spatial compatibility and incompatibility tasks. There are two reasons for this apparent discrepancy. First, the response-signal methodology was applied to the spatial compatibility and incompatibility tasks in Experiment 2, while in Experiment 1 they were standard RT tasks. It is possible, though unlikely, that the effects of S-R transfer are not measurable in SAT tasks. Second, the difference in the context of the task may have hampered S-R transfer. Task contexts were reversed in Experiment 3 to assess this latter possibility.

EXPERIMENT 3

In this experiment the Simon task was an SAT task while the spatial compatibility and incompatibility tasks were standard RT tasks. The SAT functions from the Simon task were analyzed in three ways. First, the d'_{id} values were analyzed using a hierarchical modeling approach (e.g., see McElree and Carrasco, 1999; Carrasco and McElree, 2001). Second, fits with the standard SAT equation (Equation 3) were compared to the fits achieved with the proposed hyperbolic tangent equation (Equation 4). Lastly, the d'_{loc} data were fit with an exponential decay function (Equation 5).

METHODS

Participants

Fifteen undergraduates took part in each condition (compatible and incompatible) for course credit and monetary incentives (for the SAT task) as in Experiment 2.

Procedure

The general procedure was the same as it was in Experiment 2 with the exception that the response-signal methodology was applied in the Simon task while the spatial compatibility/incompatibility tasks were "fast and accurate" standard RT tasks.

RESULTS AND DISCUSSION

Simon task

SAT analysis of task-relevant identity information. The d'_{id} vs. response lag data were fit using the standard SAT function (Equation 3) and the hyperbolic SAT function (Equation 4). Fit was quantitatively and qualitatively assessed using a hierarchical model-testing approach, commonly used in SAT studies (McElree and Dosher, 1989; Carrasco and McElree, 2001; Giordano et al., 2009). The models ranged from all factorial combinations that

ranged from single fit (1 λ, 1 β, 1 δ, and 1 λ, 1 ω, 1 κ) to both datasets to a fully saturated model (2 λ, 2 β, 2 δ, and 2 λ, 2 ω, 2 κ). Model error was assessed using a least squares approach wherein normalized residuals were scaled to the total error for the model.

The analysis of the SAT data was accomplished in two stages. In the first stage, the best fit parameters of the group mean were identified for the compatibility and incompatibility groups separately. Goodness of fit was assessed with the adjusted R^2 method (Dosher et al., 2004). These fit parameters were then used as starting points for the hierarchical modeling approach, where the mean data for both groups were concurrently fit using nonlinear data-fitting optimization routines (i.e., with the lsqnonlin function in Matlab; Mathworks Inc., Natick, MA). The second stage determined the best fit parameters for each individual participant using Equations 3 and 4. These parameter values were statistically compared across compatibility and incompatibility groups using unpaired t-tests.

The analysis of the d'_{id} data, using the standard SAT equation (Equation 3), revealed that the model with a single set of parameters (1 λ, 1 β, 1 δ) across the datasets for the spatial incompatibility and compatibility groups had the best fit overall ($R^2_{adj} = 0.98$). The group mean, and the best fit, are presented in **Figure 2**. Equation 3 was then fit to the individual data for the compatible and incompatibility groups. In general, the fits were very good (average fit for compatibility group: $R^2_{adj} = 0.86$; average fit for the incompatibility group: $R^2_{adj} = 0.86$). The parameters from the fits for each group were compared and no differences were significant.

The hyperbolic equation (Equation 4) was also fit to the group mean. Again, a model that assumes a single set of parameters (1 λ, 1 ω, 1 κ) had the best fit ($R^2_{adj} = 0.99$), slightly better than the fit of the standard SAT equation. The best hyperbolic fit and the group mean are presented in **Figure 2**. The individual fits for the spatial compatibility (mean $R^2_{adj} = 0.95$) and incompatibility (mean $R^2_{adj} = 0.94$) groups were also quite good. The parameters from the fits from each group were compared, and again, there were no significant differences.

SAT analysis of task-irrelevant location information. Unlike the effect of task-relevant information (d'_{id}) on response choice, the effect of location-based information (d'_{loc}) lessened with time (see **Figure 3**). Neither the standard SAT (Equation 3) nor the hyperbolic (Equation 4) function fit the data particularly well. While Equation 5 (i.e., single exponential decay) fit the data for the compatibility group well, it failed to fit the data for the incompatibility group. As previously discussed, d'_{loc} reflects the impact of spatial information on response selection. At any given moment, t, d'_{loc} may be jointly influenced by the direct and/or indirect spatial pathways depicted in **Figure 1**. The activity along each spatial pathway is believed to lessen with time and have a summative effect on response selection. Thus, a second exponential component was included to account for these two sources of spatial information,

$$d'_{loc}(t) = \delta_1 e^{(\beta_1 t)} + \delta_2 e^{(\beta_2 t)}, \text{ for } t > 0. \qquad (7)$$
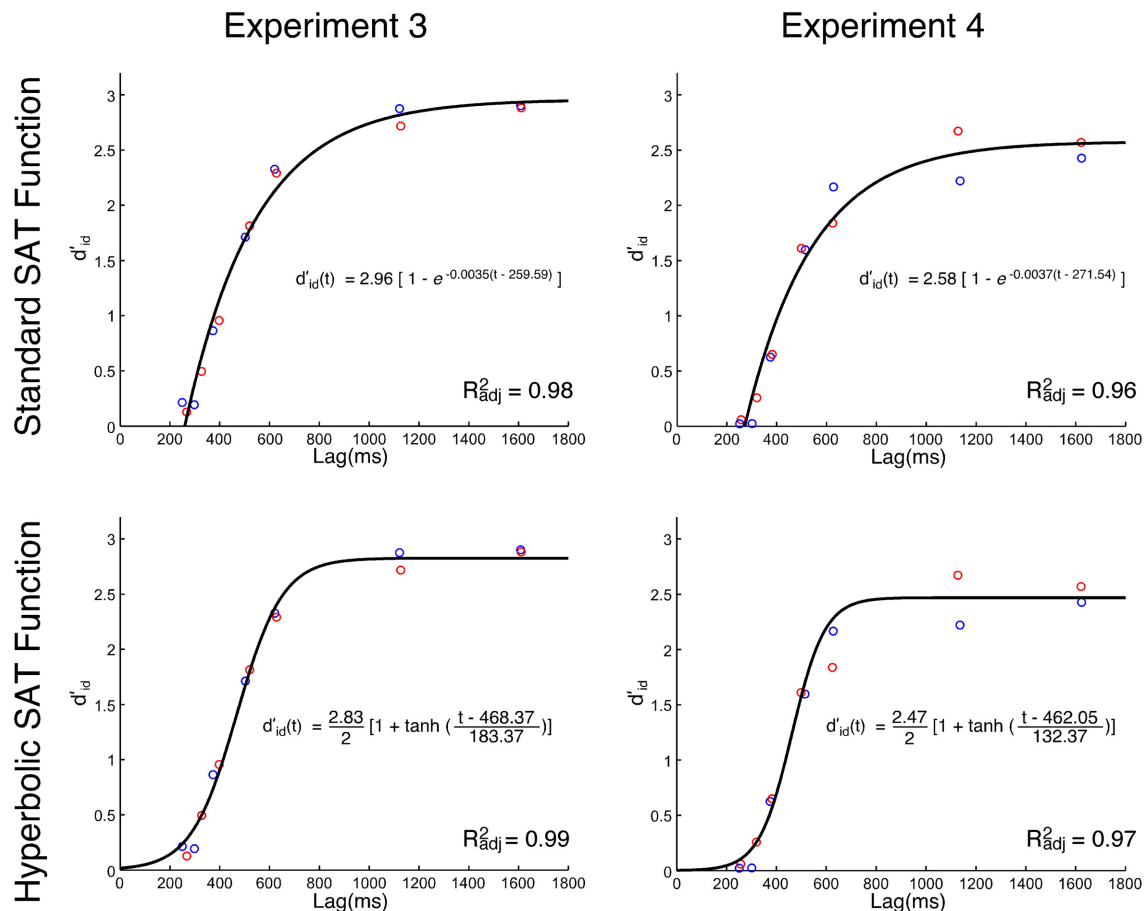
FIGURE 2 | Group mean d'_{id} values as a function of processing lag (ms) in Experiments 3 and 4. Top-row: best fit of the standard speed-accuracy tradeoff function (Equation 3). **Bottom-row:** best fit of the hyperbolic tangent speed-accuracy tradeoff function (Equation 4). The Simon SAT data from the spatial incompatibility group are red. The Simon SAT data from the spatial compatibility group are blue.
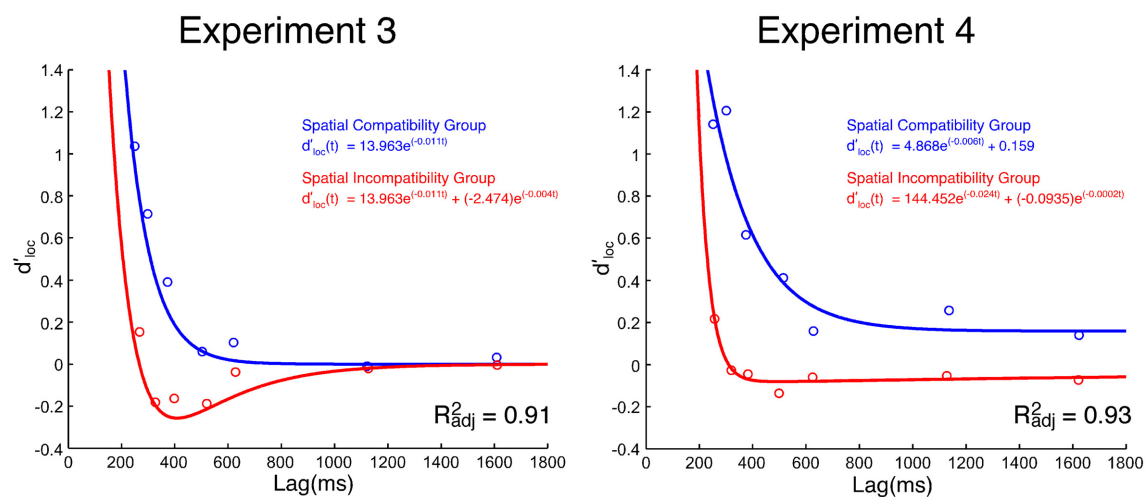
FIGURE 3 | Group mean d'_{loc} values as a function of processing lag (ms) in Experiments 3 and 4. The Simon SAT data from the spatial incompatibility group are red. The Simon SAT data from the spatial compatibility group are blue.

Because the compatible and incompatible datasets could not be fit by the same function, we abandoned the hierarchical modeling approach. To analyze the decay of the task-irrelevant location information, we developed two models. The models were derived from the architecture depicted in **Figure 1**. Both models were fit to the mean group data using nonlinear data-fitting optimization routines in Matlab (Mathworks Inc., Natick, MA).

Both models presume that the d'_{loc} values in Simon tasks, when combined with an incompatibility task, are the result of two exponential functions (Equation 7): (1) a positive component resulting from the direct pathway, and (2) a negative component resulting from the spatial incompatible mapping (i.e., the indirect spatial pathway). The models only differ in their characterization of the S-R transfer from the spatial compatibility task to the Simon task.

The first model (Model 1) specifically presumes there is S-R transfer from the spatial compatibility task to the Simon task. The model holds that the component that is transferred from the spatial compatibility task to the Simon task is similar in magnitude, but opposite in direction (i.e., toward, not away, from the location), to the negative component passed along from the spatial incompatibility task to the Simon task (i.e., $1\beta_1$, $1\delta_1$, $2\beta_2$, $1\delta_2$; with the constraint that β_2 in the spatial compatibility task is equal to $-\beta_2$ in the spatial incompatibility task). This account presumes that there are not only two exponential components (Equation 7) in the spatial incompatibility group, but also two exponential components in the spatial compatibility group. The best fit for this model (Model 1) was quite poor ($R^2_{adj} = 0.38$).

The second model (Model 2) presumes that, although the spatial incompatibility task introduces a third pathway to the Simon task, the spatial compatibility task has no effect on the Simon effect. There are only two studies (Proctor and Lu, 1999; Tagliabue et al., 2000) that have directly compared the Simon effect in a neutral condition to one that follows a spatial compatibility task. In both cases, there was no evidence of S-R transfer from a spatial compatibility task to the Simon task. Accordingly, Model 2 includes only a single exponential function for the Simon task d'_{loc} data in the spatial compatibility group (Equation 5: $1\beta_1$, $1\delta_1$) and the same exponential component and a negative-going exponential component in the spatial incompatibility group (Equation 7: $1\beta_1$, $1\delta_1$, β_2, δ_2). Thus, only one of the exponential components is shared, while the function for the incompatibility group also includes a second exponential component reflecting the third, indirect (residual) pathway. This model fit the group mean reasonably well ($R^2_{adj} = 0.91$). The group mean is plotted in **Figure 3** along with the fitted parameters from Model 2.

It was not possible to directly compare parameters from the group-derived models for the compatible and incompatible groups because the best fits were achieved with different functions. Moreover, the fits of Model 2 to the individual data were quite variable, with some being quite good (e.g., $R^2_{adj} = 0.97$) and others failing to reach a meaningful convergence. Thus, as a second step in the analysis, we performed *post-hoc*, unpaired *t*-tests on the d'_{loc} values for each SOA. This approach does not presume any particular model. The d'_{loc} value was significantly greater for the compatibility group than the incompatibility group at the 60 ms [$t_{(28)} = 3.33$, $p < 0.005$], 120 ms [$t_{(28)} = 3.68$, $p < 0.005$],

and 240 ms [$t_{(28)} = 2.49$, $p < 0.05$] SOAs. No other difference was significant. As shown in **Table 2**, none of the d'_{loc} values differed from 0 in the spatial incompatibility group while the d'_{loc} at the three earliest SOAs did differ from 0 in the spatial compatibility group.

The SAT analysis of the Simon task revealed two key findings. First, there were no effects of the spatial incompatibility task on d'_{id}. The model fits and inferential statistics suggest that the spatial compatibility and incompatibility tasks had no impact on the ability to identify the task-relevant stimulus features (i.e., shape and/or orientation) in the Simon task. This finding is in accord with the model proposed in **Figure 1** and suggests there is some independence between the indirect, residual pathway and the indirect, task-relevant pathway. Second, the spatial incompatibility task did have a noticeable effect on the sensitivity to the location of the stimulus (d'_{loc}). The spatial incompatibility task appeared to weaken, but not reverse, the Simon effect (as measured with d'_{loc}) in a SAT task. This pattern is similar to what was observed in Experiment 2 in a standard RT Simon task. Interestingly, the evidence for an early tendency to respond to the location of the stimulus, while clear in the spatial compatibility group, was mixed in the spatial incompatibility group. The inferential statistics suggest that the d'_{loc} data do not differ from zero[1]. The modeling work, however, suggests that an early exponential component is being masked by a second component. It is possible, that these effects are the result of reduced S-R transfer because of a mismatch between task contexts.

Spatial compatibility and incompatibility tasks

The RTs for the spatial compatibility and incompatibility tasks were analyzed with a 2 (group) × 2 (congruency) × 7 (SOA) ANOVA. The mean RTs of the spatial compatibility group ($M = 292$ ms) were significantly faster than those of the incompatibility group ($M = 342$ ms), $F_{(1, 28)} = 8.46$, $MSE = 31539.06$, $p < 0.01$. RTs also increased with increasing SOA, $F_{(6, 168)} = 6.63$, $MSE = 1718.17$, $p < 0.0001$. No other main effect or interaction was significant. **Table 3** provides the non-significant mean congruency effects for the spatial compatibility and incompatibility tasks.

The d'_{loc} data were entered into a 2 (group) × 7 (SOA) ANOVA. As expected, there was a large group effect [$F_{(1, 28)} = 2361.80$, $MSE = 0.93$, $p < 0.0001$] indicating that participants were following directions (i.e., responding to the target's location in the spatial compatibility task and away from the target's location in the spatial incompatibility task). There was also a group × X

---

[1]Given that there was no significant impact of task-irrelevant spatial information on responding for the spatial incompatibility group, fit was assessed with a third model (Model 3). This model assumes a single exponential function (Equation 5) for the spatial compatibility group and no impact (i.e., $d'_{loc} = 0$) of task-irrelevant spatial responding for the spatial incompatibility group. Model 3 yielded a reasonable fit ($R^2_{adj} = 0.91$), virtually indistinguishable from Model 2. However, when the spatial compatibility group was excluded from the analysis, the nil model ($d'_{loc} = 0$) fit the data of the spatial incompatibility group more poorly ($R^2_{adj} = -0.11$) than a double exponential model (Equation 7; $R^2_{adj} = 0.85$), suggesting that any description of the data presuming there is no effect of task-irrelevant spatial information on response decisions is likely false.

SOA interaction, $F_{(6, 168)} = 5.47$, $MSE = 0.10$, $p < 0.0001$. The SOA effect was only significant in the incompatibility group, $F_{(6, 84)} = 5.18$, $MSE = 0.07$, $p < 0.0005$. The same analysis was performed on the d'_{id} data. None of the effects were significant. None of the d'_{id} values differed significantly from zero (**Table 4**).

The key finding from the spatial compatibility and incompatibility tasks was the absence of a congruency effect on RTs and d'. The possibility that this was the result of the disparity in task context (SAT and standard RT) was addressed in Experiment 4 where both tasks were SAT tasks.

EXPERIMENT 4

In this experiment the response-signal methodology was applied to both Simon and spatial compatibility/incompatibility tasks. Thus, like Experiment 1, the task contexts were identical.

METHODS

Participants

There were 15 undergraduate participants in the compatibility group and 15 in the incompatibility group. Participants earned course credit and small performance bonuses, as in the previous experiments.

Procedure

In this experiment, both the Simon and spatial compatibility tasks were subject to the response-signal methodology. Thus, the Simon task was identical to the Simon task in Experiment 3 and the spatial compatibility/incompatibility tasks were identical to the spatial compatibility/incompatibility tasks in Experiment 2. The same SOA was used in each block of Simon and spatial compatibility/incompatibility trials.

RESULTS AND DISCUSSION

Simon task

SAT analysis of task-relevant identity information. The analysis of the data using the standard SAT equation (Equation 3), revealed—again—that the single fit [1 λ, 1 β, 1 δ] model had the best fit ($R^2_{adj} = 0.96$; see **Figure 2**). The model was fit to the individual data for the compatible and incompatibility groups. In general the fits were very good (average fit for compatibility group: $R^2_{adj} = 0.89$; average fit for the incompatibility group: $R^2_{adj} = 0.87$). The parameters from the fits for each group were compared and no differences were significant.

The d'_{id} analysis using the hyperbolic equation (Equation 4) was fit to the group using the hierarchical model-testing approach, as described above. Again, a model that assumes a single set of parameters (1 λ, 1 ω, 1 κ) had the best fit ($R^2_{adj} = 0.97$), slightly better than the fit of the standard SAT equation. The average of the fits to individual data was also good for the compatible (mean $R^2_{adj} = 0.90$) and incompatibility groups (mean $R^2_{adj} = 0.91$). The only comparison between individual fits that was significant was that between the asymptote, λ, [$t_{(28)} = 2.20$, $p < 0.05$]. The mean asymptote of individual fits was slightly higher for the spatial incompatibility group ($M = 2.66$) than it was for the spatial compatibility group ($M = 2.31$). This difference is apparent in the mean d_{id} values presented in the

late SOAs in **Figure 2**. A *post-hoc* analysis of the group differences in d'_{id} for each SOA only revealed a difference at the 120 ms [$t_{(28)} = 2.21$, $p < 0.05$] and the 960 ms [$t_{(28)} = 2.16$, $p < 0.05$] SOA, although these effects do not survive a Bonferroni correction for multiple comparisons. Thus, the evidence that the spatial compatibility task had an effect on the sensitivity to the non-spatial, task-relevant feature (d'_{id}) of the target in the Simon task was generally poor, and mixed, at best.

SAT analysis of task-irrelevant location information. The SAT functions of the group mean d'_{loc} values are presented in **Figure 3**. The first pass of fitting the group mean, using the same models in Experiment 3, was unsuccessful. Model 1 fit the data very poorly ($R^2_{adj} = 0.15$). Model 2 fared better, but the fit was less than spectacular ($R^2_{adj} = 0.71$). That the mean data for the spatial compatibility group did not return to the zero baseline likely explains these poor fits. Using Equation 5, but including a constant, for the spatial compatibility group did not improve Model 2 ($R^2_{adj} = 0.64$). In fact, the best model was one where the group mean for the spatial compatibility group was fit with a constant and the group mean for the spatial incompatibility group was fit with Equation 7 independently ($R^2_{adj} = 0.93$).

The d'_{loc} values were significantly different between the spatial compatibility and incompatibility groups at each SOA ($ps < 0.05$), with the exception of the 480 ms SOA. The d'_{loc} values were also compared to 0 for each group and SOA. The d'_{loc} values for the spatial compatibility group were significantly greater than 0 at all SOAs ($ps < 0.05$, uncorrected). For the spatial incompatibility group, the d'_{loc} value at the 60 ms SOA was significantly greater than 0 ($p < 0.05$, uncorrected) and at the 1440 ms SOA the d'_{loc} value was significantly less than 0 ($p < 0.05$, uncorrected; see **Table 2**)[2].

The present findings suggest that there are fundamental differences between the temporal dynamics of task-irrelevant spatial information processing in Simon tasks when mixed with spatial compatibility and incompatibility tasks. Unlike Experiment 3, the model with the best fit was one in which there were no shared parameters between spatial compatibility and incompatibility groups in the Simon task. A potential implication of this fully saturated model is that the direct spatial pathway may be compromised by the spatial compatibility/incompatibility task.

Spatial compatibility and incompatibility tasks

Accuracy was near ceiling in all conditions, as it was in Experiment 2, so the data were not subjected to a curve-fitting procedure. The d'_{loc} values were entered into a 2 (group) $\times$ 7 (SOA) ANOVA (see **Table 4** for means). The main effects

---

[2]When the mean for the spatial incompatibility group was fit to a nil model (i.e., $d'_{loc} = 0$), as in Experiment 3 (see Footnote 1), and the spatial compatibility group mean was fit with a single exponential function (Equation 5; with a constant), it produced a reasonable fit ($R^2_{adj} = 0.92$). However, excluding the spatial compatibility group from the analysis, again, tells a different story. The fit of the spatial incompatibility group mean to a nil model was poor ($R^2_{adj} = 0.09$), while the fit to a double exponential function (Equation 7) was good ($R^2_{adj} = 0.88$). Once more, any description of the data suggesting there is no effect of task-irrelevant spatial information on responding is likely false.

of group $[F_{(1, 28)} = 1524.45, MSE = 0.76, p < 0.0001]$, SOA $[F_{(6, 168)} = 9.70, MSE = 0.32, p < 0.0001]$, and the interaction $[F_{(6, 168)} = 8.70, MSE = 0.32, p < 0.0001]$ were all significant. The interaction was the result of a much larger SOA effect in the spatial incompatibility task than the spatial compatibility task.

The d'_{id} values were also entered into a 2 (group) × 7 (SOA) ANOVA. Only the SOA effect was significant, $F_{(6, 168)} = 3.43$, $MSE = 0.032$, $p < 0.005$. d'_{id} values increased with SOA. However, only the d'_{id} values at the 360 and 1440 ms SOA were significantly different from 0 (see **Table 4**).

GENERAL DISCUSSION

When a spatial incompatibility task is intermixed with a Simon task, the Simon effect is reversed (Marble and Proctor, 2000; Proctor et al., 2000, 2003; Proctor and Vu, 2002). In Experiment 1, this finding was replicated in a different paradigm where tasks predictably alternated between spatial incompatibility and Simon tasks. The most common explanation for this finding is that the spatial incompatibility task activates an additional, indirect pathway that connects nodes representing spatial features of the stimulus with response nodes (**Figure 1**). The current work addressed three features of this paradigm: bidirectional S-R transfer across Simon and spatial compatibility/incompatibility tasks, the modulating effects of task context similarity on S-R transfer, and the time course of task-irrelevant S-R location information on response selection.

BIDIRECTIONAL S-R TRANSFER BETWEEN SIMON AND SPATIAL COMPATIBILITY TASKS

This was the first study to explore bidirectional S-R transfer between Simon and spatial compatibility/incompatibility tasks in the mixed-tasks paradigm. Evidence for S-R transfer from the spatial compatibility/incompatibility task to the Simon effect was evident in all experiments. In general, performing the spatial incompatibility task with the Simon task reduced or reversed the tendency to respond to the location of the stimulus. This pattern has been observed in a number of studies in a variety of different paradigms (e.g., Tagliabue et al., 2000, 2002; Marble and Proctor, 2000; Proctor et al., 2007, 2013; Proctor and Vu, 2009; Yamaguchi and Proctor, 2009).

The evidence for S-R transfer from the Simon task to the spatial compatibility/incompatibility task was best when task contexts (SAT or standard RT) matched. Congruent responses, in the spatial compatibility and incompatibility tasks, were those in which the response associated with the non-spatial identity of the stimulus in the Simon task matched the location of the stimulus. In Experiment 1, the congruency effect for the spatial compatibility group was a speed-accuracy tradeoff: responses were faster and less accurate for incongruent trials. On the other hand, for those participants in the spatial incompatibility condition, congruent trials were faster and more accurate than incongruent trials. In Experiment 3, when the task contexts did not match, there was no effect of congruency on RTs or d'_{id}. Congruency effects were rarely seen in the d'_{id} measure in SAT tasks with response-signal methodology (Experiments 2 and 4). Thus, transfer from the Simon task to the spatial compatibility/incompatibility tasks was weak and sporadic, suggesting that S-R transfer between Simon and

spatial compatibility/incompatibility is bidirectional and asymmetric. S-R transfer from spatial incompatibility tasks to the Simon task was much more convincing than S-R transfer in the other direction. It may be that the precedence for location information (Hillyard and Munte, 1984) offers greater opportunities to influence tasks wherein the task-relevant information comes from slower (non-spatial) S-R pathways. Further research is needed to assess the precise reason for asymmetrical S-R transfer. The clearest evidence, in the current work, for a congruency effect came when (i) responding was slow (i.e., with the spatial incompatibility task in Experiment 1 and with long SOAs in Experiment 4) and (ii) the two tasks shared a task context (i.e., in Experiments 1 and 4). The context of the task, thus, appears to play a key role in S-R transfer.

TASK-CONTEXT DEPENDENT S-R TRANSFER

Environmental context plays a critical role in memory performance. When features of the encoding environment match features of the retrieval environment, memory performance is generally better than when the environmental features do not match (Godden and Baddeley, 1975). Smith and Vela (2001) noted that manipulations that draw attention to the task or away from the environmental context tend to reduce task-dependent memory effects. Thus, context plays an important role when it is attended during encoding and retrieval.

Yamaguchi and Proctor (2009) observed evidence for context-dependent S-R transfer from a spatial incompatibility task to a Simon task when the response mode (key-press vs. joystick response) was the same for both tasks. Response modality (as in Yamaguchi and Proctor, 2009) is one feature of task context; yet the context of the task may also include other features. In the current work, response-signal (i.e., SAT) methodology affected S-R transfer in a context-dependent manner. The SAT task not only included the same stimuli presented in the standard RT task, but also included other task-relevant stimuli such as an auditory response-signal tone and post-response feedback. These features likely contributed to the unique context of the task and were quite different from the context of the standard RT task. The results of the present investigation support this claim. The spatial incompatibility task reversed the Simon effect in Experiment 1 (where both tasks were standard RT tasks), but not in Experiment 2 where the Simon (standard RT methodology) and spatial incompatibility (response-signal methodology) tasks were different. In the response-signal (SAT) Simon tasks (i.e., Experiments 3 and 4), there was evidence for a late reversal of d'_{loc} in Experiment 4 (task contexts match), but not in Experiment 3 (task contexts do not match). Together, the evidence suggests that the opportunity for S-R transfer is greatest when task features are closely matched. Moreover, the current works also demonstrates that the context of the task plays an important role in the mixed-task experimental design.

THE TIME COURSE OF TASK-IRRELEVANT LOCATION INFORMATION ON RESPONSE SELECTION

A number of previous studies have used vincentizing approaches to study the time course of the Simon effect. The challenge with this approach is that it relies on differences in the shape of RT

distributions (Zhang and Kornblum, 1997; Pratte et al., 2010; Schwarz and Miller, 2012). The shape of an RT distribution can be affected by a number of factors like fast guesses, fatigue, or inattention. It can be troubling if these factors differ systematically across conditions. It is, perhaps, even more troubling that distributional approaches, like vincentizing, completely ignore error rates. Wickelgren (1977) argued that it "…may not be defensible …to attempt to test quantitative theories of information processing dynamics …by functions which use reaction time as the sole dependent variable, without simultaneously predicting accuracy." (p. 81). Thus, researchers should be cautious not to overvalue the contribution of vincentized approaches (e.g., delta plots) to the temporal dynamics of information processing.

The response-signal (SAT) approach is similar to another approach that has been commonly used to study the time course of the Simon effect (Ridderinkhof, 2002). Conditional accuracy functions (CAFs) partition RTs, and error rates, into a small number of bins (Wood and Jennings, 1976; Ridderinkhof, 2002; Band et al., 2003), not unlike the vincentization approach. This analytic approach produces the so-called *micro*-SAT (Pachella, 1974). Micro-SAT analyses have also depicted the influence of task-irrelevant spatial information on response selection as an exponential decay function (e.g., Ridderinkhof, 2002). This approach, while less cumbersome than a full SAT analysis, may be criticized on two grounds. First, it can be argued that, not unlike vincentizing, different processes (guesses, fatigue, inattention, etc.) are not equally represented along the RT distribution. The response-signal approach avoids this pitfall by capturing a point along the SAT within a single block of trials. Secondly, as Pachella (1974) pointed out, the relationship between a micro-SAT and the standard SAT (sometimes called the *macro*-SAT) is unknown, but what is known is that they do not seem to tap into the same underlying function (Luce, 1986). Given this, some caution when interpreting CAFs is warranted (Wickelgren, 1975, 1977).

The current work extended the first SAT analysis of the Simon effect presented by Hilchey et al. (2011). A dissociation between two measures of sensitivity, d'_{id} (sensitivity to the task-relevant target feature) and d'_{loc} (sensitivity to the task-irrelevant spatial feature of the target), in the context of a Simon task was revealed. While d'_{id} increased with time (a standard SAT), d'_{loc} decreased with time. The d'_{id} data were fit with the standard SAT function and a hyperbolic tangent function. Both fits were excellent, although the hyperbolic tangent function fit was slightly superior. This is not to suggest that the hyperbolic tangent function should replace the standard SAT function. Future research is needed to determine which function might best describe performance in a wider range of tasks.

The spatial incompatibility task had virtually no impact on d'_{id} in the Simon task, suggesting independence between spatial S-R transfer and the processes involved in the identification of non-spatial, task-relevant, target features. The only fly in the ointment was seen in the hyperbolic tangent fits in Experiment 4: the asymptotic parameter (λ) was significantly higher for the spatial incompatibility group than the spatial compatibility group. There are, however, a number of reasons to be skeptical about this finding. First, there was little reason to expect, from any *a priori* theoretical perspective, that the ability to identify the

non-spatial, task-relevant feature in a Simon task should be better when the alternate task is a spatial incompatibility task than a spatial compatibility task. Second, there was no significant difference between the asymptotic parameters, derived from the standard SAT function (Equation 3), for the spatial compatibility and incompatibility groups in both Experiments 3 and 4. Third, a *post-hoc* analysis suggested the d'_{id} difference between groups was only significant at one of the late SOAs (960 ms) near asymptote. Lastly, in Experiments 3 and 4 the best fits to the d'_{id} group data in the Simon task assumed only a single set of parameters, suggesting the alternate task (i.e., the spatial compatibility or incompatibility task) had no impact on the accumulation of task-relevant information. Thus, the asymptotic difference between the groups found in Experiment 4 is, at best, equivocal.

Although the spatial incompatibility task had no influence on d'_{id} in the Simon task, it had a robust effect on d'_{loc}. This effect provides another example of a single dissociation between d'_{loc} and d'_{id}. The d'_{loc} data were fit with an exponential decay function. The exponential models used in the current investigation were simply initial attempts at providing a quantitative description of the time course of task-irrelevant, spatial S-R activity. It could be argued that an exponential decay model is psychophysically implausible, as irrelevant S-R location information should follow a Gaussian, biphasic, accumulation-decay pattern (e.g., Kornblum et al., 1999). Unfortunately we did not capture an early accumulation phase. Future SAT investigations of the Simon effect may consider manipulations (e.g., Ivanoff et al., 2002) that might possibly delay the d'_{loc} function in order to capture an early accumulation phase. For now, it is worth noting that the decrease in d'_{loc} with time was fit reasonably well with an exponential decay model.

The pattern of d'_{loc} across time lag in the Simon task with the spatial compatibility group is very similar to the pattern of vincentized RTs for Simon effects when there is prior or concurrent experience with a spatial compatibility task (Tagliabue et al., 2000; Proctor and Vu, 2009; Proctor et al., 2013). Both approaches demonstrate a standard pattern of declining influence of task-irrelevant location information on response selection with time. In the current work, that a single exponential component described the time course of d'_{loc} in the Simon task with the spatial compatibility group is consistent with at least three mechanisms. First, there may be no S-R transfer from spatial compatibility tasks to Simon tasks. This proposal is consistent with Tagliabue et al's (2000) assertion that spatial compatibility tasks have no impact on the direct spatial pathway. Second, a spatial compatibility task may induce some activity along an indirect spatial pathway that is largely masked by robust activity along the spatial direct pathway. It is possible that this activity may be unmasked at later SOAs given conditions that favor S-R transfer. The evidence for this possibility comes from Experiment 4, where the best fit to the d'_{loc} Simon data for the spatial compatibility group included a constant because d'_{loc} did not decline to zero. Unfortunately, this particular finding is ambiguous and may be explained by another mechanism. It is possible that the spatial compatibility task modulates the decline (decay or suppression) of the spatial direct pathway. This account is generally

consistent with Proctor and Lu's (1999) original proposal that the direct spatial pathway is not "unmodifiable." It is not consistent, however, with some modeling approaches (e.g., Tagliabue et al., 2000). Future research is needed to disentangle and dissociate the effects of different spatial S-R pathways on response decisions.

Perhaps the most important contribution of the current work stemmed from the observation that the d'_{loc} time course in the Simon task was different across the spatial compatibility/incompatibility groups. The time course of d'_{loc} in the Simon task, performed concomitantly with a spatial incompatibility task, was unlike that observed with previous research using vincentization approaches (Marble and Proctor, 2000; Tagliabue et al., 2000; Proctor and Vu, 2009) where the reverse Simon effect generally increased with time. There was no evidence for a monotonically increasing reverse Simon effect in the current study. In Experiment 3, although none of the d'_{loc} values in the Simon task differed significantly from zero across SOAs, the spatial incompatibility group Simon data were fit well to a model that included two exponential decay components: (i) the identical exponential decay component found in the spatial compatibility group, presumably capturing activity along the spatial direct pathway, and (ii) a negative exponential decay component that captured a slight tendency to respond away from the location of the stimulus. In Experiment 4, the spatial incompatibility group mean data were also fit to a double exponential function (Equation 7) quite well. Moreover, in Experiment 4, the d'_{loc} value of the spatial incompatibility group at the earliest SOA was greater than 0, indicating a tendency to respond to the location of the stimulus). Interestingly, at the longest SOA, the opposite pattern emerged (i.e., indicating a tendency to respond away from the location of the stimulus). The data-fitting approaches espoused herein appeared to be particularly sensitive to the time course of d'_{loc} and the findings generally support the tripartite pathway model depicted in **Figure 1**. In summary, the findings from the current study suggest that the early activity along the direct, task-irrelevant, spatial S-R pathway is indeed masked by late (and relatively persistent) residual activity from the indirect spatial S-R pathway. The current findings are also consistent with modeling approaches that presume activity along the task-irrelevant direct spatial pathway is unaffected by prior, or concurrent, experience with a spatial incompatibility task (Tagliabue et al., 2000).

CONCLUSIONS

The present findings firmly establish Simon's (1969) original claim that there is a natural tendency to respond toward the source of stimulation. Performing a spatial incompatibility task can reverse or eliminate this tendency. However, the current results suggest that activity along the indirect spatial pathway may mask this natural tendency to respond to the source of stimulation. The present work also suggests that response-signal (i.e., SAT) methodology provides a task context that that may promote or impede S-R transfer. Lastly, these findings also demonstrate that transfer between spatial compatibility/incompatibility tasks and the Simon task can be bidirectional, although asymmetric.

REFERENCES

Band, G. P., Ridderinkhof, K. R., and van der Molen, M. W. (2003). Speed-accuracy modulation in case of conflict: the roles of activation and inhibition. *Psychol. Res.* 67, 266–279. doi: 10.1007/s00426-002-0127-0

Barber, P., and O'Leary, M. (1997). "The relevance of salience: towards an activational account of irrelevant stimulus-repsonse compatibility effects," in *Theoretical Issues in Stimulus-response Compatibility,* eds B. Hommel and W. Prinz (Amsterdam: North-Holland), 135–172.

Carrasco, M., and McElree, B. (2001). Covert attention accelerates the rate of visual information processing. *Proc. Natl. Acad. Sci. U.S.A.* 98, 5363–5367. doi: 10.1073/pnas.081074098

de Jong, R., Liang, C.-C., and Lauber, E. J. (1994). Conditional and unconditional automaticity: A dual-process model of the effects of spatial stimulus-response correspondence. *J. Exp. Psychol. Hum. Percept. Perform.* 21, 731–750. doi: 10.1037/0096-1523.20.4.731

Dosher, B. A., Han, S., and Lu, Z. L. (2004). Parallel processing in visual search asymmetry. *J. Exp. Psychol. Hum. Percept. Perform.* 30, 3–27. doi: 10.1037/0096-1523.30.1.3

Duncan, J. (1977). Response selection errors in spatial choice reaction tasks. *Q. J. Exp. Psychol.* 29, 415–423. doi: 10.1080/14640747708400618

Duncan, J. (1978). Response selection in spatial choice reaction: Further evidence against associative models. *Q. J. Exp. Psychol.* 30, 429–440. doi: 10.1080/00335557843000034

Fitts, P. M., and Deininger, R. L. (1954). S-R compatibility: Correspondence among paired elements within stimulus and response codes. *J. Exp. Psychol.* 48, 483–492. doi: 10.1037/h0054967

Fitts, P. M., and Seeger, C. M. (1953). S-R compatibility: Spatial characteristics of stimulus and response codes. *J. Exp. Psychol.* 46, 199–210. doi: 10.1037/h0062827

Gescheider, G. A. (1997). *Psychophysics: The Fundamentals, 3rd Edn.* Mahwah, NJ: Lawrence Erlbaum Associates.

Giordano, A. M., McElree, B., and Carrasco, M. (2009). On the automaticity and flexibility of covert attention: a speed-accuracy trade-off analysis. *J. Vis.* 9, 1–10. doi: 10.1167/9.3.30

Godden, D. R., and Baddeley, A. D. (1975). Context-dependent memory in two natural environments: on land and under water. *Br. J. Psychol.* 66, 325–331. doi: 10.1111/j.2044-8295.1975.tb01468.x

Gopher, D., Armony, L., and Greenshpan, Y. (2000). Switching tasks and attention policies. *J. Exp. Psychol. Gen.* 129, 308–339. doi: 10.1037/0096-3445.129.3.308

Hilchey, M. D., Ivanoff, J., Taylor, T. L., and Klein, R. M. (2011). Visualizing the temporal dynamics of spatial information processing responsible for the Simon effect and its amplification by inhibition of return. *Acta Psychol.* 136, 235–244. doi: 10.1016/j.actpsy.2010.09.003

Hillyard, S. A., and Munte, T. F. (1984). Selective attention to color and location: an analysis with event-related brain potentials. *Percept. Psychophys.* 36, 185–198. doi: 10.3758/BF03202679

Hommel, B. (2011). The Simon effect as tool and heuristic. *Acta Psychol.* 136, 189–202. doi: 10.1016/j.actpsy.2010.04.011

Ivanoff, J., Klein, R. M., and Lupiáñez, J. (2002). Inhibition of return interacts with the Simon effect: An omnibus analysis and its implications. *Percept. Psychophys.* 64, 318–327. doi: 10.3758/BF03195794

Kadlec, H. (1999). Statistical properties of d' and ßestimates of signal detection theory. *Psychol. Methods* 4, 22–43. doi: 10.1037/1082-989X.4.1.22

Kornblum, S., Hasbroucq, T., and Osman, A. (1990). Dimensional overlap: cognitive basis for stimulus-response compatibility–a model and taxonomy. *Psychol. Rev.* 97, 253–270. doi: 10.1037/0033-295X.97.2.253

Kornblum, S., Stevens, G. T., Whipple, A., and Requin, J. (1999). The effects of irrelevant stimuli: 1. The time course of stimulus-stimulus and stimulus-response consistency effects with Stroop-like stimuli, Simon-like tasks, and their factorial combinations. *J. Exp. Psychol. Hum. Percept. Perform.* 25, 688–714. doi: 10.1037/0096-1523.25.3.688

Lu, C. H., and Proctor, R. W. (1995). The influence of irrelevant location information on performance: a review of the Simon and spatial Stroop effects. *Psychon. Bull. Rev.* 2, 174–207. doi: 10.3758/BF03210959

Luce, R. D. (1986). *Response Times: Their Role in Inferring Elementary Mental Organization.* New York, NY: Oxford University Press.

Macmillan, N., and Creelman, C. D. (2005). *Detection Theory: A User's Guide, 2nd Edn.* New York, NY: Lawrence Erlbaum Associates, Inc.

Marble, J. G., and Proctor, R. W. (2000). Mixing location-relevant and location-irrelevant choice-reaction tasks: influences of location mapping on the Simon effect. *J. Exp. Psychol. Hum. Percept. Perform.* 26, 1515–1533. doi: 10.1037/0096-1523.26.5.1515

McElree, B., and Carrasco, M. (1999). The temporal dynamics of visual search: evidence for parallel processing in feature and conjunction searches. *J. Exp. Psychol. Hum. Percept. Perform.* 25, 1517–1539. doi: 10.1037/0096-1523.25.6.1517

McElree, B., and Dosher, B. A. (1989). Serial position and set size in short-term memory: The time course of recognition. *J. Exp. Psychol. Gen.* 118, 346–373. doi: 10.1037/0096-3445.118.4.346

Meyer, D. E., Osman, A. M., Irwin, D. E., and Yantis, S. (1988). Modern mental chronometry. *Biol. Psychol.* 26, 3–67. doi: 10.1016/0301-0511(88)90013-0

Milanese, N., Iani, C., and Rubichi, S. (2010). Shared learning shapes human performance: Transfer effects in task sharing. *Cognition* 116, 15–22. doi: 10.1016/j.cognition.2010.03.010

Milanese, N., Iani, C., Sebanz, N., and Rubichi, S. (2011). Contextual determinants of the social-transfer-of-learning effect. *Exp. Brain Res.* 211, 415–422. doi: 10.1007/s00221-011-2679-y

Murnane, K., Malmberg, K., and Phelps, M. P. (1999). Context-dependent recognition memory: the ICE theory. *J. Exp. Psychol. Gen.* 128, 403–415. doi: 10.1037/0096-3445.128.4.403

Pachella, R. G. (1974). "The interpretation of reaction time in information-processing research," in *Human Information Processing: Tutorials in Performance and Cognition,* ed B. H. Kantowitz (Hillsdale, NJ: Lawrence Erlbaum), 41–82.

Pratte, M. S., Rouder, J. N., Morey, R. D., and Feng, C. (2010). Exploring the differences in distributional properties between Stroop and Simon effects using delta plots. *Atten. Percept. Psychophys.* 72, 2013–2024. doi: 10.3758/APP.72.7.2013

Proctor, R. W. (2011). Playing the Simon game: use of the Simon task for investigating human information processing. *Acta Psychol.* 136, 182–188. doi: 10.1016/j.actpsy.2010.06.010

Proctor, R. W., and Lu, C. H. (1999). Processing irrelevant location information: practice and transfer effects in choice-reaction tasks. *Mem. Cognit.* 27, 63–77. doi: 10.3758/BF03201214

Proctor, R. W., Marble, J. G., and Vu, K. P. (2000). Mixing incompatibly mapped location-relevant trials with location-irrelevant trials: Effects of stimulus mode on the reverse Simon effect. *Psychol. Res.* 64, 11–24. doi: 10.1007/s004260000041

Proctor, R. W., Miles, J. D., and Baroni, G. (2011). Reaction time distribution analysis of spatial correspondence effects. *Psychon. Bull. Rev.* 18, 242–266. doi: 10.3758/s13423-011-0053-5

Proctor, R. W., and Vu, K. P. (2002). Mixing location-irrelevant and location-relevant trials: Influence of stimulus mode on spatial compatibility effects. *Mem. Cogn.* 30, 281–293. doi: 10.3758/BF03195289

Proctor, R. W., and Vu, K. P. (2009). Stimulus-response compatibility for mixed mappings and tasks with unique responses. *Q. J. Exp. Psychol.* 63, 320–340. doi: 10.1080/17470210902925270

Proctor, R. W., Vu, K. P., and Marble, J. G. (2003). Mixing location-relevant and irrelevant tasks: Spatial compatibility effects eliminated by stimuli that share the same spatial codes. *Vis. Cogn.* 10, 15–50. doi: 10.1080/713756673

Proctor, R. W., Yamaguchi, M., Dutt, V., and Gonzalez, C. (2013). Dissociation of S-R compatibility and Simon effects with mixed tasks and mappings. *J. Exp. Psychol. Hum. Percept. Perform.* 39, 593–609. doi: 10.1037/a0029923

Proctor, R. W., Yamaguchi, M., and Vu, K. P. (2007). Transfer of noncorresponding spatial associations to the auditory Simon task. *J. Exp. Psychol. Learn. Mem. Cogn.* 33, 245–253. doi: 10.1037/0278-7393.33.1.245

Ratcliff, R. (1979). Group reaction time distributions and an analysis of distribution statistics. *Psychol. Bull.* 86, 446–461. doi: 10.1037/0033-2909.86.3.446

Reed, A. V. (1973). Speed-accuracy trade-off in recognition memory. *Science* 181, 574–576. doi: 10.1126/science.181.4099.574

Ridderinkhof, K. R. (2002). Micro- and macro-adjustments of task set: activation and suppression in conflict tasks. *Psychol. Res.* 66, 312–323. doi: 10.1007/s00426-002-0104-7

Salthouse, T. A., and Hedden, T. (2002). Interpreting reaction time measures in between-group comparisons. *J. Clin. Exp. Neuropsychol.* 24, 858–872. doi: 10.1076/jcen.24.7.858.8392

Schouten, J. F., and Bekker, J. A. M. (1967). Reaction time and accuracy. *Acta Psychol.* 27, 143–153. doi: 10.1016/0001-6918(67)90054-6

Schwarz, W., and Miller, J. (2012). Response time models of delta plots with negative-going slopes. *Psychon. Bull. Rev.* 19, 555–574. doi: 10.3758/s13423-012-0254-6

Simon, J. R. (1969). Reactions toward the source of stimulation. *J. Exp. Psychol.* 81, 174–176. doi: 10.1037/h0027448

Simon, J. R., and Rudell, A. P. (1967). The effect of an irrelevant cue on information processing. *J. Appl. Psychol.* 51, 300–304. doi: 10.1037/h0020586

Smith, S. M. (1994). "Theoretical principles of context-dependent memory," in *Theoretical Aspects of Memory, 2nd Edn.,* eds P. Morris and M. Gruneberg (New York, NY: Routledge), 168–195.

Smith, S. M., and Vela, E. (2001). Environmental context-dependent memory: a review and meta-analysis. *Psychon. Bull. Rev.* 8, 203–220. doi: 10.3758/BF03196157

Tagliabue, M., Zorzi, M., and Umilta, C. (2002). Cross-modal re-mapping influences the Simon effect. [Empirical Study]. *Mem. Cogn.* 30, 18–23. doi: 10.3758/bf03195261

Tagliabue, M., Zorzi, M., Umilta, C., and Bassignani, F. (2000). The role of long-term-memory and short-term-memory links in the Simon effect. *J. Exp. Psychol. Hum. Percept. Perform.* 26, 648–670. doi: 10.1037/0096-1523.26.2.648

Umiltà, C., and Zorzi, M. (1997). "Commentary on Barber and O'Leary: Learning and attention in S-R compatibility," in *Theoretical Issues in Stimulus-response Compatibility,* eds B. Hommel and W. Prinz (Amsterdam: North-Holland), 173–178.

Van der Lubbe, R. H. J., and Abrahamse, E. L. (2011). The premotor theory of attention and the Simon effect. *Acta Psychol.* 136, 259–264. doi: 10.1016/j.actpsy.2010.09.007

Vu, K.-P. L. (2007). Influences on the Simon effect of prior practice with spatially incompatible mappings: transfer within and between horizontal and vertical dimensions. *Mem. Cogn.* 35, 1463–1471. doi: 10.3758/bf03193616

Wickelgren, W. A. (1975). "Dynamics of retrieval," in *Short-term Memory,* eds D. Deutsch and J. A. Deutsch (New York, NY: Academic Press), 233–355.

Wickelgren, W. A. (1977). Speed-accuracy tradeoff and information processing dynamics. *Acta Psychol.* 41, 67–85. doi: 10.1016/0001-6918(77)90012-9

Wood, C. C., and Jennings, J. R. (1976). Speed-accuracy tradeoff functions in choice reaction time: Experimental designs and computational procedures. *Percept. Psychophys.* 19, 92–101. doi: 10.3758/BF03199392

Yamaguchi, M., and Proctor, R. W. (2009). Transfer of learning in choice reactions: contributions of specific and general components of manual responses. *Acta Psychol.* 130, 1–10. doi: 10.1016/j.actpsy.2008.09.008

Zhang, J., and Kornblum, S. (1997). Distributional analysis and De Jong, Liang, and Lauber's (1994) dual-process model of the Simon effect. *J. Exp. Psychol. Hum. Percept. Perform.* 23, 1543–1551. doi: 10.1037/0096-1523.23.5.1543

Using reaction times and binary responses to estimate psychophysical performance: an information theoretic analysis

James V. Stone *

Psychology Department, Sheffield University, Sheffield, UK

As the strength of a stimulus increases, the proportions of correct binary responses increases, which define the psychometric function. Simultaneously, mean reaction times (RT) decrease, which collectively define the chronometric function. However, RTs are traditionally ignored when estimating psychophysical parameters, even though they may provide additional Shannon information. Here, we extend Palmer et al's (2005) proportional-rate diffusion model (PRD) by: (a) fitting individual RTs to an inverse Gaussian distribution, (b) including lapse rate, (c) point-of-subjective-equality (PSE) parameters, and, (d) using a two-alternative forced choice (2AFC) design based on the proportion of times a variable comparison stimulus is chosen. Maximum likelihood estimates of mean RT values (from fitted inverse Gaussians) and binary responses were fitted both separately and in combination to this extended PRD (EPRD) model, to obtain psychophysical parameter values. Values estimated from binary responses alone (i.e., the psychometric function) were found to be similar to those estimated from RTs alone (i.e., the chronometric function), which provides support for the underlying diffusion model. The EPRD model was then used to estimate the mutual information between binary responses and stimulus strength, and between RTs and stimulus strength. These provide conservative bounds for the average amount of Shannon information the observer gains about stimulus strength on each trial. For the human experiment reported here, the observer gains between 2.68 and 3.55 bits/trial. These bounds are monotonically related to a new measure, the *Shannon increment*, which is the expected value of the smallest change in stimulus strength detectable by an observer.

Keywords: psychometric function, chronometric function, point of subjective equality, diffusion model, reaction time, threshold, Shannon information, mutual information

1. INTRODUCTION

For over a 100 years, it has been known that the ability to discriminate between two stimuli increases as a sigmoidal function of the difference between those stimuli, where this is traditionally measured using binary observer responses. However, when an observer makes a response, there is a trade-off between speed, or reaction time (RT), and accuracy of responses. This speed-accuracy trade-off has been the subject of numerous papers, notably (Ratcliff, 1978; Harvey, 1986; Swanson and Birch, 1992; Wichmann and Hill, 2001; Palmer et al., 2005), and more recently in Bonnet et al. (2008).

Here, we propose four extensions to the proportional-rate diffusion model (PRD) proposed in Palmer et al. (2005). First, we introduce a new parameter, the point-of-subjective-equality (PSE), which takes account of systematic shifts or bias in observer perception. This parameter is incorporated into the chronometric and psychometric functions. Second, we use a maximum likelihood estimate (MLE) of the RT mean based on a physically motivated diffusion model of RTs which involves fitting individual RTs to an inverse Gaussian distribution. Third, we take account of lapses in observer concentration by introducing a lapse rate parameter, which is estimated simultaneously with

other psychophysical parameters. Fourth, we use a two-alternative forced choice (2AFC) design where the psychometric function is defined, not by the proportion of correct responses (range 50–100%), but by the proportion of times a variable comparison stimulus is chosen in preference to a fixed reference stimulus (range 0–100%). Note that the 2AFC experimental procedure is the same whether one chooses to measure the proportion of correct responses or the proportion of times a variable comparison stimulus is chosen.

Once the model has been fitted to these data, it can be used to estimate the mutual information (Shannon and Weaver, 1949; MacKay, 2003; Stone, 2014) between binary responses and stimulus strength, and between RT and stimulus strength. Finally, the mutual information provides a value for the Shannon increment, which is the expected value of the smallest change in stimulus strength detectable by an observer.

2. THE PROPORTIONAL-RATE DIFFUSION MODEL

We provide a brief summary of Palmer et al's PRD model (Palmer et al., 2005) here, and describe extensions below. In the experiment described in Palmer et al. (2005), an observer is presented with an array of moving dots. Stimulus strength x is defined by

coherence (i.e., the percentage of dots moving in the same direction), and the observer is required to indicate which one of two directions the dots are moving. Note that coherence, and therefore stimulus strength x, varies between zero and some upper bound.

The PRD model is based on a diffusion model of RT, where the *mean* RT $\bar{\tau}_{PRD}$ varies as a sigmoidal function of x

$$\bar{\tau}_{PRD} = \frac{A}{Kx} \tanh(KAx) + \bar{\tau}_{res}, \tag{1}$$

where K is a measure of observer sensitivity, and A represents a decision boundary associated with RT. The first term on the right hand side represents the time to make a decision, and $\bar{\tau}_{res}$ is a fixed residual RT (e.g., time to respond after a decision is made). Notice that this model requires that the mean RT $\bar{\tau}_{PRD}$ decreases monotonically as the motion signal increases above zero, a requirement which will be relaxed in the model proposed below.

Within the PRD model, the probability P_{PRD} of making a correct response is defined by the logistic psychometric function

$$P_{PRD} = \frac{1}{1 + e^{-2AK|x|}}, \tag{2}$$

where $|x|$ indicates the absolute value of x. In Equation (2), the product AK acts as a single parameter which modulates the steepness of the sigmoidal function, and therefore acts as a measure of sensitivity to changes in stimulus strength. Note that the stimulus strength cannot fall below zero in Palmer et al's moving dot experiment, and that, when the stimulus motion strength is $x = 0\%$, the observer has to guess, so that $P_{PRD} = 0.5$, whereas if $x = 100\%$ then $P_{PRD} = 1.0$.

3. THE EXTENDED PROPORTIONAL-RATE DIFFUSION (EPRD) MODEL

The model proposed here is based on the assumption that responses arise from a two-alternative forced choice (2AFC) procedure. On each trial, the observer is presented with two stimuli, and the task is to choose the stronger stimulus, where strength can be defined in terms of differences in any physical quantity, such as speed, luminance, or contrast. The two stimuli are a *reference stimulus* with a stimulus value s_R that remains constant within a specific subset of trials, and a *comparison stimulus* with a value s_C that varies between trials. A *comparison* response is obtained if the observer chooses the comparison stimulus. The stimulus strength x within one trial is defined as the difference between the reference value s_R and the comparison value s_C, specifically $x = s_C - s_R$.

We measure performance in terms of the proportion P of times that a variable comparison stimulus is chosen in preference to the fixed reference stimulus, which we define as a comparison stimulus response, so P varies between zero and one. A direct translation from P_{PRD} to P would guarantee that a stimulus strength of zero corresponds to $P = 0.5$. However, if observer perception is biased, such that a stimulus difference of $x = 0$ is not perceived as zero, then a stimulus strength of zero would not coincide with $P = 0.5$. This perceptual bias can be accommodated with a second modification, a new parameter s_{PSE}, which is the point-of-subjective-equality (PSE) between the comparison and reference stimuli. Specifically, s_{PSE} is the value s_C of the comparison stimulus which is perceived to be the same as the value s_R of reference stimulus.

Given that the stimulus strength is $x = s_C - s_R$, the *perceived stimulus strength* x' is

$$x' = s_C - s_{PSE} \tag{3}$$

$$= x - \Delta x, \tag{4}$$

where Δx is the error in the perceived value of s_C. The probability of choosing the comparison stimulus is defined as

$$P = \frac{1}{1 + e^{-2AKx'}}. \tag{5}$$

Note that the product AK effectively acts as a single parameter, and will be treated as such for binary response data (but not for RT data, see below).

In order to take account of observer lapses in concentration, which result in a pure guess, we introduce a lapse rate parameter γ. Evidence presented in Wichmann and Hill (2001) suggests that failure to take account of the lapse rate can lead to substantial errors in estimated psychophysical parameter values. If the lapse rate were zero then we would expect that $P = 0$ for highly negative stimulus strengths, and that $P = 1$ for highly positive stimulus strengths, so that observed deviations from $P = 0$ and $P = 1$ at extreme stimulus strengths can be used to provide an estimate of the lapse rate. Thus, the lapse rate parameter limits the lower and upper bounds of the psychometric function to $P_{min} = \gamma/2$ and $P_{max} = 1 - \gamma/2$, respectively, such that[1]

$$P = \left[\frac{1}{1 + e^{-2AKx'}} - 0.5 \right] (1 - \gamma) + 0.5. \tag{6}$$

Thus, the three parameters to be estimated for Equation (6) define the vector variable

$$\theta_P = (s_{PSE}, AK, \gamma). \tag{7}$$

Similarly, we model the observer's mean RT for a perceived stimulus strength x' as

$$\bar{\tau} = \frac{A}{kx'} \tanh(KAx') + \bar{\tau}_{res}. \tag{8}$$

Here, the effects of A and K are separable, and so the four parameters to be estimated for Equation (8) define the vector variable

$$\theta_\tau = (s_{PSE}, A, K, \bar{\tau}_{res}). \tag{9}$$

The lapse rate parameter is not included here because lapses have no predictable effect on RT.

Finally, we can adapt results from Luce (1986) and Palmer et al. (2005) to relate RT to response probability. The mean decision

---

[1]Notice that, if the lapse rate is $\gamma = 0.01$ then the upper and lower bounds are 0.995 and 0.005, respectively, because half of the observer's guesses will be correct, on average.

time is defined as $\bar{\tau}_{\mathrm{dec}} = \bar{\tau}_i - \bar{\tau}_{\mathrm{res}}$, so that Equations (5, 8) can be combined to provide a mapping between mean decision time $\bar{\tau}_{\mathrm{dec}}$ and the probability P of choosing the comparison stimulus

$$\bar{\tau}_{\mathrm{dec}} = (A/K)\frac{(2P-1)}{x'}. \qquad (10)$$

Thus, if the perceived stimulus strength x' has a large positive or negative value then $P = 0$ or $P = 1$ (respectively), and so $\bar{\tau}_{\mathrm{dec}} = A/(K|x'|)$ in both cases. This predicts that, for a given perceived stimulus strength, the probability of choosing the comparison stimulus is proportional to the mean decision time.

4. USING OBSERVER RESPONSES

For each trial, we obtain a RT and a binary response from the observer, which indicates whether the observer has chosen the comparison stimulus or the reference stimulus. At each stimulus strength x_i, the comparison and reference stimuli are presented to the observer on N_i trials, and the number of times the observer chooses the comparison and reference stimulus is recorded as n_i and $N_i - n_i$, respectively. For a given putative value of P_i, a standard binomial model gives the probability of the observed binary responses as

$$p(n_i|N_i, P_i) = C_{N_i}^{n_i} \times P_i^{n_i} \times (1 - P_i)^{N_i - n_i}, \qquad (11)$$

where P_i is a function of the parameters Ak, γ and PSE as defined in Equation (6). The maximum likelihood estimate of P_i is the proportion of comparison stimulus responses $P_i' = n_i/N_i$.

When considered over all N_x values of x, the probability of observing the set of all binary responses is defined by the log likelihood function

$$L_P = \log \prod_{i=1}^{N_x} C_{N_i}^{n_i} P_i^{n_i} (1 - P_i)^{N_i - n_i} \qquad (12)$$

$$= \sum_{i=1}^{N_x} n_i \log P_i + \sum_{i=1}^{N_x} (N_i - n_i) \log(1 - P_i) + \sum_{i=1}^{N_x} \log C_{N_i}^{n_i}, \quad (13)$$

where the final term does not depend on parameter values, and can be discarded unless the exact value of the likelihood is required. Recall that each P_i is determined by Equation (6), which is a function of the EPRD parameter values $\theta_P = (A, K, \gamma, \mathrm{PSE})$. The maximum likelihood estimate (MLE) of θ_P is obtained by finding EPRD parameter values θ_P that maximize L_P.

If the number of trials at each stimulus strength is large then Equation (13) can be approximated by a Gaussian function. At a given stimulus strength x_i, the observed proportion of binary responses is P_i', which is assumed to be the probability P_i plus a noise term η_P, so that $P_i' = P_i + \eta_P$. If the noise η_P has a Gaussian distribution with variance $v_{P,i}$ then

$$p(P_i'|A, k, x_i') = \frac{1}{\sqrt{2\pi v_{P,i}}} \exp \frac{-\left(P_i' - P_i\right)^2}{2 v_{P,i}}, \qquad (14)$$

where P_i is defined as a function of A, k, x' in Equation (6), and the variances $v_{P,i}$ can be estimated from the data as

$v_{P,i} = N_i P_i'(1 - P_i')$. Results for the Gaussian approximation in Equation (14) were found to be very similar to those for Equation (13). Results reported here are based on Equation (13).

5. USING REACTION TIMES

RTs tend to be short if the comparison stimulus value is very different from the reference stimulus, but as the comparison and reference stimuli become more similar, so the RT increases, as shown in **Figure 4B**. Here, we use RTs in a two stage process. First, a mean RT value is estimated at each stimulus strength. These mean RT values are then used as data for the $RT_{\bar{\tau}}$ model, which is used to estimate EPRD model parameters.

5.1. INVERSE GAUSSIAN MODEL OF INDIVIDUAL RTs

It is commonly assumed that the RT is the time required for the cumulative amount of perceptual evidence to reach some criterion value (Ratcliff, 1978; Smith, 1990). Specifically, this evidence accumulation is assumed to consist of a Brownian diffusion process with positive drift, which can be likened to a the total distance traveled in a one-dimensional biased random walk. If a Brownian process is allowed to run for a *fixed time* then it is well known that the final distribution of values (e.g., evidence) has a Gaussian distribution. However, it is less well known that if a Brownian diffusion process is allowed to run until it reaches a *fixed criterion value* then the time taken to reach that value has an *inverse Gaussian* or *Wald* distribution (see **Figure 3**). Therefore, if the amount of evidence required to make a response is stable for a given observer then RTs are appropriately modeled using an inverse Gaussian distribution[2].

If RTs have an inverse Gaussian distribution with mean $\bar{\tau}_i'$ then the probability of a single observed RT τ_{ij} associated with the jth presentation of the stimulus value x_i is

$$p(\tau_{ij}|\bar{\tau}_i', \lambda_i) = \left(\frac{\lambda_i}{2\pi \tau_{ij}^3}\right)^{1/2} \times \exp\left[\frac{-\lambda_i(\tau_{ij} - \bar{\tau}_i')^2}{2 \bar{\tau}_i'^2 \tau_{ij}}\right], \quad (15)$$

where the variance of this distribution is

$$v_{\tau_i} = \bar{\tau}_i'^3/\lambda_i. \qquad (16)$$

Each of the N_x stimulus strengths is presented N_i times. For one model RT mean, the probability of the observed N_i RTs (one RT per trial) defines the log likelihood function

$$L_{\tau,i} = \log \prod_{j=1}^{N_i} p(\tau_{ij}|\bar{\tau}_i', \lambda_i). \qquad (17)$$

Maximizing Equation (17) with respect to the parameters $\bar{\tau}_i'$ and λ_i yields a maximum likelihood estimate (MLE) of both parameters at one stimulus strength x_i. Even though the algebraic mean and the MLE mean are identical (Tweedie, 1957) for the inverse

---

[2]For reference, the Wald distribution is the distribution of first passage times of a biased Brownian process, and is qualitatively similar to the log-normal distribution, which is often used to model RT.

Gaussian, the fitting process provides the parameter estimate λ_i, which is vital for subsequent calculations.

5.2. MODEL $RT_{\bar{\tau}}$: USING MEAN REACTION TIMES

For a given stimulus strength x_i, the predicted mean RT $\bar{\tau}_i$ varies as a tanh function of x_i, as defined in Equation (8). The central limit theorem allows us to assume that the distribution of mean RTs of the inverse Gaussian pdf at a given stimulus strength x_i is Gaussian with mean $\bar{\tau}'_i$ and variance $v_{\bar{\tau},i}$. Therefore, the likelihood of the EPRD mean $\bar{\tau}_i$ from Equation (8) is

$$p(\bar{\tau}'_i|\bar{\tau}_i(\theta_\tau)) = \frac{1}{\sqrt{2\pi v_{\bar{\tau},i}}} e^{-(\bar{\tau}'_i - \bar{\tau}_i)^2/(2v_{\bar{\tau},i})}. \quad (18)$$

The variance of an inverse Gaussian distribution of RT values with mean $\bar{\tau}'_i$ is $v_{\tau i}$ (Equation 16), so the variance $v_{\bar{\tau} i}$ of a distribution of means (where each mean is based on N_i samples) is

$$v_{\bar{\tau} i} = \frac{\bar{\tau}'^3_i}{\lambda_i N_i}. \quad (19)$$

Thus, we can assess the fit of the inverse Gaussian mean RTs $\bar{\tau}'_i$ to the EPRD mean RTs $\bar{\tau}_i$ of Equation (8) as follows. The probability of the N_x mean RTs $\bar{\tau}'_i$ (one mean RT per stimulus strength) defines the log likelihood function

$$L_{\bar{\tau}} = \log \prod_{i=1}^{N_x} p(\bar{\tau}'_i|\bar{\tau}_i) \quad (20)$$

$$= -1/2 \sum_{i=1}^{N_x} \frac{(\bar{\tau}'_i - \bar{\tau}_i)^2}{v_{\bar{\tau},i}} - 1/2 \sum_{i=1}^{N_x} \log 2\pi v_{\bar{\tau},i}, \quad (21)$$

where $\bar{\tau}_i$ is defined in Equation (8), so that the parameters to be estimated for model $RT_{\bar{\tau}}$ are $\theta_\tau = (A, k, \gamma, PSE, \bar{\tau}_{res})$ to fit the overall variation in mean RT with stimulus strength x.

In summary, we have three estimates of the mean RT at each stimulus strength: the algebraic mean $\bar{\tau}'_{obsi}$, the MLE mean of the inverse Gaussian or Wald pdf $\bar{\tau}'_i$ (from Equation 17), which collectively are used as data to estimate the means $\bar{\tau}_i$ (one per stimulus strength) obtained from the fitted EPRD model (from Equation 21). The MLE means $\bar{\tau}'_i$ are shown as crosses in **Figure 4B**, and the means $\bar{\tau}_i$ are corresponding points on the fitted curve, respectively.

We also have two estimates of the probability of a comparison stimulus response at each stimulus strength: the observed proportion of comparison stimulus responses (which is the MLE $P'_i = n_i/N_i$), and the mean P_i (one per stimulus strength) obtained from fitting the EPRD model (Equation 13) to the MLE means P'_i. These are shown as dots in **Figure 4A**, and as corresponding points on the fitted curve, respectively.

6. USING BINARY RESPONSES AND RTs

In the absence of knowledge regarding the covariance between the noise in mean RT and binary response probability, we are forced to assume this covariance is zero. In other words, we assume that L_P and $L_{\bar{\tau}}$ provide independent estimates of the EPRD model

parameters. In this case, estimates based on combined RT and binary response probability are obtained by maximizing the sum of likelihoods

$$L_C = L_P + L_{\bar{\tau}}. \quad (22)$$

However, the implausibility of this independence assumption means that we will not take seriously any results based on Equation (22).

7. INFORMATION THEORY

The amount of Shannon information (Shannon and Weaver, 1949; MacKay, 2003; Stone, 2014) that the observer gains about the stimulus is reflected in both the binary responses and RTs. Specifically, the average Shannon information that each mean RT provides about the stimulus strength x is the mutual information $I(x, \bar{\tau})$ between x and the mean RT. Similarly, the average Shannon information that binary responses provide about the stimulus strength x is the mutual information $I(x, P)$ between x and the probability of a comparison stimulus binary response.

More importantly, the total amount of Shannon information that the observer has about the stimulus cannot be less than the amount of Shannon information implicit in the observer's combined binary and RT responses. In other words, the total mutual information, as measured by an experimenter, between observer responses and stimulus strength provides a lower bound for the amount of Shannon information that the observer has about the stimulus strength. Thus, each the mutual information value provided in this paper constitutes a conservative estimate of the amount of information that the observer gains about the stimulus.

7.1. EVALUATING $I(x, P)$

The mutual information $I(x, P)$ between stimulus strength s and the probability P that the observer chooses the comparison stimulus (i.e., $r = 1$) is

$$I(x, P) = \int_x \int_P p(x, P) \log \frac{p(x, P)}{p(x)p(P)} \, dP \, dx \quad (23)$$

$$= H(x) + H(P) - H(x, P) \text{ bits}, \quad (24)$$

where $H(x)$ and $H(P)$ are the differential entropies of $p(x)$ and $p(P)$, respectively, and $H(x, P)$ is the differential entropy of the joint distribution $p(x, P)$. All logarithms in this paper use base 2, so information is measured in bits. Substituting $p(x, P) = p(P|x)p(x)$, yields

$$I(x, P) = \int_x p(x) \int_P p(P|x) \log \frac{p(P|x)}{p(P)} \, dP \, dx \quad (25)$$

$$= H(P) - H(P|x) \text{ bits}, \quad (26)$$

where $H(P|x)$ is the differential entropy of the noise in the measurements P. Given Bayes' rule, $p(P|x) = p(x|P)p(P)/p(x)$, we can recognize the mutual information as the differential entropy $H(P)$ of the prior distribution minus the differential entropy $H(P|x)$ of the posterior distribution.

We can evaluate Equation (25) by summing over discrete versions of the variables x and P. Recall that the observed proportion of responses $r = 1$ at a given stimulus strength x_i is $P'_i = n_i/N_i$, so that

$$I(x, P) = \sum_{k=1}^{N_x} p(x_k) \left[\sum_{i=1}^{N_x} p(P'_i|x_k) \log \frac{p(P'_i|x_k)}{p(P'_i)} \right] \text{ bits.} \quad (27)$$

We assume that the probability of stimulus values is locally uniform, so that $p(x_k) = 1/N_k$. In order to evaluate Equation (27), we require expressions for $p(P'_i|x_k)$ and $p(P'_i)$.

7.1.1. Evaluating the posterior $p(P'_i|x_k)$
Using Equation (5) across a range of x values, the fitted value of P at x_k is P_k. Assuming a binomial distribution, the probability of the observed proportion P'_i given a fitted value P_k at x_k is

$$p(P'_i|x_k) = C_{N_i}^{n_i} P_k^{n_i} (1 - P_k)^{N_i - n_i}, \quad (28)$$

where $p(P'_i|x_k) = p(P'_i|P_k)$, and $p(P'_i|x_k)$ values are normalized to ensure that $\sum_i p(P'_i|x_k) = 1$.

7.1.2. Evaluating the prior $p(P'_i)$
The distribution of binary responses is binomial with a mean equal to the grand mean P_G of all N_G binary responses of an observer

$$P_G = \frac{1}{N_G} \sum_{i=1}^{N_G} r_i, \quad (29)$$

where $r_i = 1$ if and only if a response corresponds to the observer choosing the comparison stimulus. The observer's prior probability of the binary responses for the ith stimulus strength is therefore

$$p(P'_i) = C_{N_i}^{n_i} P_G^{n_i} (1 - P_G)^{N_i - n_i}, \quad (30)$$

where $p(P'_i)$ values are normalized to ensure that $\sum_i p(P'_i) = 1$.

7.2. EVALUATING $I(x, \bar{\tau})$
Following the same line of reasoning as above, the mutual information $I(x, \bar{\tau})$ between stimulus strength and mean RT is

$$I(x, \bar{\tau}) = \int_x p(x) \int_{\bar{\tau}} p(\bar{\tau}|x) \log \frac{p(\bar{\tau}|x)}{p(\bar{\tau})} \, d\bar{\tau} \, dx \quad (31)$$

$$= H(\bar{\tau}) - H(\bar{\tau}|x) \text{ bits,} \quad (32)$$

where $H(\bar{\tau}|x)$ is the differential entropy of the noise in the measurements $\bar{\tau}$.

We can evaluate Equation (31) by summing over discrete versions of the variables x and $\bar{\tau}$

$$I(x, \bar{\tau}) = \sum_{k=1}^{N_x} p(x_k) \left[\sum_{i=1}^{N_i} p(\bar{\tau}'_i|x_k) \log \frac{p(\bar{\tau}'_i|x_k)}{p(\bar{\tau}'_i)} \right] \text{ bits,} \quad (33)$$

where $p(\bar{\tau}'_i|x_k)$ is defined by the EPRD model (Equation 8) with a fitted value $\bar{\tau}_k$, so that

$$p(\bar{\tau}'_i|x_k) = p(\bar{\tau}'_i|\bar{\tau}_k(\theta_\tau)), \quad (34)$$

as in Equation (18). As before, we assume that the probability of stimulus values is uniform, so that $p(x_k) = 1/N_i$.

7.2.1. Evaluating the posterior $p(\bar{\tau}'_i|x_k)$
The posterior is defined in Equation (18), but is repeated here with changed subscripts for clarity

$$p(\bar{\tau}'_i|x_k) = \frac{1}{\sqrt{2\pi v_{\bar{\tau}k}}} \exp \left[\frac{-(\bar{\tau}'_i - \bar{\tau}_k)^2}{2 v_{\bar{\tau}k}} \right], \quad (35)$$

where $v_{\bar{\tau}k}$ is defined in Equation (19), and $p(\bar{\tau}'_i|x_k)$ values are normalized to ensure that $\sum_i p(\bar{\tau}'_i|x_k) = 1$.

7.2.2. Evaluating the prior $p(\bar{\tau}'_i)$
A parametric form for the observer's prior probability distribution $p(\tau)$ of individual RTs was estimated from the entire set of that observer's grand total of N_G RTs. These were fitted to an inverse Gaussian distribution to obtain a grand mean $\bar{\tau}_G$ and a parameter λ_G. This pdf has a variance

$$v_G = \bar{\tau}_G^3 / \lambda_G. \quad (36)$$

At each stimulus strength x_i, the RT mean is based on a sample of N_i RTs, and the central limit theorem suggests that the distribution of means is approximately Gaussian with a variance

$$v_g = v_G / N_i. \quad (37)$$

Therefore, the prior probability density of each inverse Gaussian mean $\bar{\tau}'_i$ is

$$p(\bar{\tau}'_i) = \frac{1}{\sqrt{2\pi v_g}} \exp \left[\frac{-(\bar{\tau}'_i - \bar{\tau}_G)^2}{2 v_g} \right], \quad (38)$$

where $p(\bar{\tau}'_i)$ values are normalized to ensure that $\sum_i p(\bar{\tau}'_i) = 1$.

7.3. THE SHANNON INFORMATION OF A SINGLE RESPONSE
So far we have derived expressions for the Shannon information implicit in the average RT $\bar{\tau}_i$ and also in the average binary response, which is summarized as the proportion P_i of comparison responses, for a stimulus strength x_i. Here, we derive an expression for the Shannon information associated with a single trial; first for RTs, and then for binary responses.

As the number of trials at each stimulus strength is increased, so the variance in each mean RT decreases, and the central limit theorem ensures that the distribution of means becomes increasingly Gaussian. The mutual information between two variables (e.g., mean RT and stimulus strength) depends on the signal to noise ratio SNR

$$I \leq 1/2 \log_2(1 + \text{SNR}), \quad (39)$$

where SNR is the signal variance expressed as a fraction of the noise variance in the measurement (Shannon and Weaver, 1949).

If the distribution of mean RTs is Gaussian then the distribution of differences $\Delta\bar{\tau}$ between mean RT $\bar{\tau}$ and the grand mean RT (at one stimulus strength) must also be Gaussian. Because the mutual information is defined in Equation (32) to be the differential entropy of $\bar{\tau}$ minus the differential entropy of the noise $\Delta\bar{\tau}$ in $\bar{\tau}$, we can assume equality in Equation (39) (Rieke et al., 1997). In fact, we do not need to rely on the central limit theorem here, because even if the perturbing noise $\Delta\bar{\tau}$ is not Gaussian, Shannon's Theorem 18 (Shannon and Weaver, 1949) implies equality in Equation (39), so that

$$I = 1/2 \, \log_2(1 + \text{SNR}) \text{ bits.} \tag{40}$$

We already have a value for the mutual information $I(x, \bar{\tau})$ from Equation (27), so we can re-arrange Equation (40) to find the SNR associated with $\bar{\tau}$

$$\text{SNR}_{\bar{\tau}} = 2^{2I(x,\bar{\tau})} - 1 \text{ bits.} \tag{41}$$

However, the mutual information $I(x, \bar{\tau})$ obtained from Equation (27) tells us how much average Shannon information each *mean* RT provides about stimulus strength, whereas we want to know how much average information each *individual* RT provides about stimulus strength. Because the value of SNR in Equation (41) is based on mean RTs, each of which involves N_i trials, the variance of the measurement noise has been reduced by a factor of N_i relative to the noise in the RT of a single trial (provided this noise is iid). This implies that the value of SNR for a single trial is

$$\text{SNR}_{\tau} = \text{SNR}_{\bar{\tau}}/N_i \tag{42}$$

$$= (2^{2I(x,\bar{\tau})} - 1)/N_i \text{ bits.} \tag{43}$$

If we substitute SNR_{τ} into Equation (40) then we obtain an estimate of the average Shannon information $I(x, \tau)$ implicit in the observer's RT in a single trial

$$I(x, \tau) = \frac{1}{2} \log_2 \left[1 + \frac{(2^{2I(x,\bar{\tau})} - 1)}{N_i} \right] \text{ bits.} \tag{44}$$

A similar line of reasoning implies that the average Shannon information $I(x, r)$ implicit in the observer's binary response r in a single trial is

$$I(x, r) = \frac{1}{2} \log_2 \left[1 + \frac{(2^{2I(x,P)} - 1)}{N_i} \right] \text{ bits.} \tag{45}$$

In order to compare mutual information estimates for the different variables τ and r, the calculations for $I(x, \tau)$ and $I(x, r)$ should be based on the same range of stimulus strengths x.

7.4. DEFINING THE SHANNON INCREMENT
The mutual information between stimulus strength and (binary or RT) responses can be used to define the smallest average detectable difference in stimulus strength, which we call the *Shannon increment* (SI). We first define the effective stimulus range x_{range} as the range of stimulus strengths x associated with

response probabilities between $P = \epsilon$ and $P = 1 - \epsilon$, for some small value ϵ. Then the SI is related to the mutual information I by

$$SI = \frac{x_{\text{range}}}{2^I}, \tag{46}$$

where the value 2 is based on the assumption that information is measured in bits (i.e., using log to the base 2), and SI has the same units as stimulus strength. Because SI decreases monotonically with mutual information, it should become asymptotically closer to the true value of SI as the number of trials or stimulus strengths is increased.

A brief explanation for this definition is as follows. Consider a range of stimulus strengths x_{range} which give rise to "noisy" observer responses $y = f(x)$, where these responses are samples from a probability density function $p(y(x))$, and where the mutual information between x and y is I bits. One way to interpret SI involves assuming that $p(y(x))$ is uniform. In this case, on average, knowing the value of y reduces the possible range of x values to an interval $\Delta x = x_{\text{range}}/2^I$, which we can recognize as being equal to the SI.

8. FAT-FACE THIN: A DEMONSTRATION EXPERIMENT
We used the EPRD models described above to estimate the PSE and other key parameters for a simple demonstration experiment using a human observer. On each trial, the observer was presented with a colored picture of an upright face and an inverted face (see **Figure 2**) on a computer screen, and was required to indicate which one appeared to be wider by pressing a left/right computer key. For half of the trials, the reference stimulus was an upright face, and the comparison stimulus was an inverted version of the same face, and these were swapped for the other half of the trials. The width of the comparison image was determined by 1 of 21 stretch factors $s = 0.90, 0.91, \ldots, 1.10$, but the height of both stimuli was kept constant. The stimulus strength was defined to be $x = s - 1$, so that x varied between -0.1 and 0.1. For a given value of s_i, the observer was presented with the same stimulus pair for a total of $N_i = 20$ trials. Stimuli were shown in random order, and the left/right position of reference/comparison stimuli was counterbalanced across trials.

8.1. RESULTS
Each of three models defined by L_P, $L_{\bar{\tau}}$, and L_C was used to fit a psychometric and/or a chronometric function to the data from one subject, as shown in **Figure 4**. Maximum likelihood parameter estimation was implemented in MatLab using the Nelder–Mead simplex method. The parameter estimates for each model are summarized in **Table 1**.

8.2. USING BINARY RESPONSES: MODEL L_P
Based on 420 binary responses, maximizing L_P (Equation 12) yields a psychometric function similar to that in **Figure 4A**, and a PSE of $s_{\text{PSE}} = 1.031$. This maximum likelihood estimate implies that an inverted face must be 3.1% wider than an upright face in order for the two faces to be perceived as the same width. Numerical estimation of the Hessian matrix

FIGURE 1 | How the entropy $H(x)$ in stimulus strength x is accounted for by the entropy $H(\tau)$ in RT (τ) and entropy $H(P)$ in the probability P of a particular binary response r. The entropies of x, P, and τ are represented by the discs X, Y, and Z, respectively. The mutual information between x and P is $I(x, P) = (a + b)$, and the mutual information between x and τ is $I(x, \tau) = (a + c)$.

FIGURE 2 | Schematic illustration of typical stimulus shown to observer on a single trial. The observer has to choose the face that looks wider. The stimulus in the experiment used was a picture of the actor James Corden's face, with all background details removed (see http://illusionoftheyear.com/2010/the-fat-face-thin-fft-illusion).

FIGURE 3 | Reaction times fitted with an inverse Gaussian (Equation 15). Each dot represents 1 of 20 RTs for a stimulus value (width scaling) of $s = 1.05$.

of second derivatives of Equation (12) at s_{PSE} yields a standard error (se) of 0.003, which implies that s_{PSE} is significantly different from $s = 1$ ($p < 0.001$). The values of three parameters were estimated for this model, the PSE, Ak, and γ, and the product Ak is quoted in **Table 1** for comparison with other works.

8.3. USING MEAN REACTION TIMES: MODEL $L_{\bar{\tau}}$

Each of 21 mean RTs (one per stimulus strength) was first estimated by maximizing Equation (17), based on 20 RTs per stimulus strength. Using these 21 mean RTs, $L_{\bar{\tau}}$ (Equation 21), was maximized with respect to four parameters (PSE, A, k, and $\bar{\tau}_{res}$) to yield a chronometric function similar to that in **Figure 4B**. The estimated PSE is $s_{PSE} = 1.034$ (se = 0.004, $p < 0.001$).

8.4. USING MEAN RTs AND OBSERVER RESPONSES: MODEL L_C

Based on 42 data points (the 21 estimated mean RTs used for $L_{\bar{\tau}}$ plus 21 corresponding binary response probabilities used for L_P),
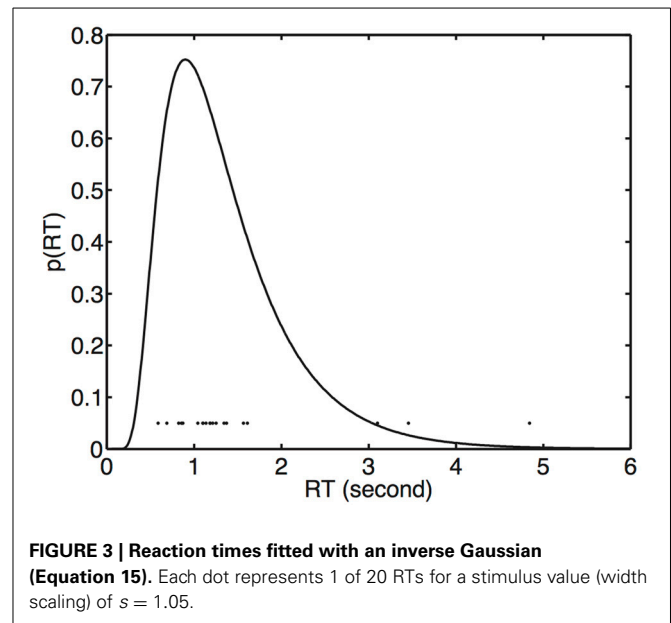
maximizing L_C (Equation 22) yields the psychometric function and the chronometric function in **Figures 4A,B**, respectively, and a PSE of 1.032 (se = 0.003, $p < 0.001$). There are five parameters to be estimated for this model, the PSE, A, k, $\bar{\tau}_{res}$, and γ.

8.5. SHANNON INFORMATION

The mutual information $I(x, \bar{\tau})$ between x and $\bar{\tau}$ is the entropy in $p(\bar{\tau})$ and $p(x)$ shared by the joint distribution $p(x, \bar{\tau})$. Using Equation (33), this evaluates to $I(x, \bar{\tau}) = 2.79$ bits. Using Equation (44) with $N_i = 20$, this implies that the mutual information $I(x, \tau)$ for a single RT is $I(x, \tau) = 0.87$ bits, and is represented by the intersection of regions X and Z.

Similarly, Equation (27) can be used to estimate the mutual information between x and P, which comes to $I(x, P) = 4.82$ bits. Using Equation (45) with $N_i = 20$, this implies that the mutual information $I(x, r)$ for a single binary response r is $I(x, r) = 2.68$ bits, and is represented by the intersection of regions X and Y.

We can use $I(x, \tau)$ and $I(x, r)$ to provide lower and upper bounds on the total amount of mutual information I_{tot} between x and the combined variables (r, τ), which can be considered to be a vector variable. If τ and r provide independent information about x (i.e., if $a = 0$ in **Figure 1**) then the maximum value of I_{tot} is

$$\max(I_{tot}) = I(x, \tau) + I(x, r) \qquad (47)$$

$$= 0.87 + 2.68 \qquad (48)$$

$$= 3.55 \text{ bits.} \qquad (49)$$

However, if all of the information $I(x, \tau)$ provided by τ about x is the same as part of the information provided by r about x (i.e., if $c = 0$ in **Figure 1**) then I_{tot} cannot be less than $I(x, r)$. To take account of the possibility that all of the information $I(x, r)$ provided by r about x is the same as part of the information provided
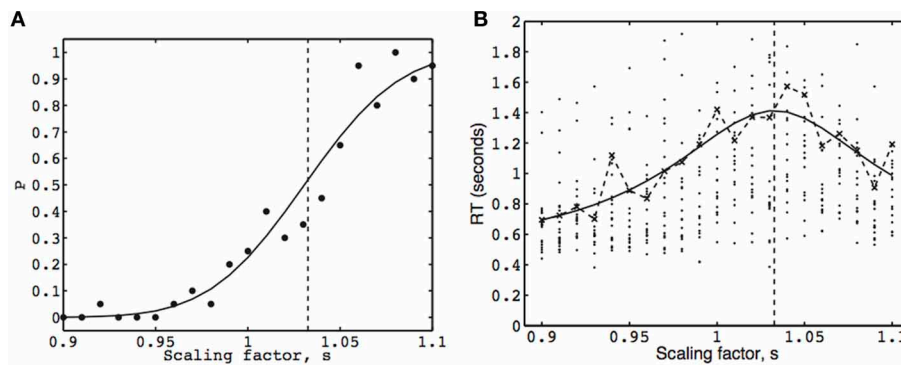
FIGURE 4 | The psychometric function (A) and chronometric function (B), from the face inversion experiment for one observer. The width scaling factor s applied to the comparison image is indicated on the abscissa. The vertical dashed line marks the point-of-subjective-equality (PSE) at $s = 1.032$. **(A)** Each dot represents the observed proportion of trials for which the observer chose the comparison stimulus, and the fitted psychometric function is defined in Equation 6. **(B)** Each dot represents the RT of a single trial for the same responses as in

Figure 4A (RTs greater than 2 s are not shown). The fitted chronometric function is defined in Equation 8. The dashed curve joins the fitted (inverse Gaussian) mean RTs, each of which was obtained by maximizing Equation 17. The solid curves in **(A, B)** (Equations 6, 8, respectively) were fitted using combined binary and mean RT data by maximizing Equation 22. A graph similar to **(A)** was obtained for model L_P (i.e., using only binary responses), and a graph similar to **(B)** was obtained for model $L_{\bar{\tau}}$ (i.e., using only mean RTs).

Table 1 | Results for three models.

Model	PSE	A	k	A × K	$\bar{\tau}_{res}$ (s)	γ	LLik	MI (bits)
Binary L_P	1.031 ± 0.003	NA	NA	22.32	NA	0.005	-31.13	2.68
RT $L_{\bar{\tau}}$	1.034 ± 0.004	0.998	28.37	28.32	0.437	NA	18.7	0.87
Comb L_C	1.032 ± 0.003	1.016	23.12	23.50	0.354	0.011	-13.10	3.18

Binary model: based only on binary response probability (Equation 12).

RT model: based only on mean RT (Equation 17).

Comb (combined model): based on binary response probability and mean RT (Equation 22).

PSE, point of subjective equality (± indicates standard error); A and k are EPRD parameters, τ_{res} is the fixed part of RT; γ, lapse rate; LLik, log likelihood; and MI, mutual information between stimulus strength and RT or binary responses or both (see text). The final number (3.18 bits) represents $I(x, r) = 2.68$ plus $I(x, \tau) = 0.497$, computed using parameter values obtained from Equation 22.

by τ about x, we can write

$$\min(I_{\text{tot}}) = \max(I(x, \tau), I(x, r)) \qquad (50)$$

$$= \max(0.87, 2.68) \qquad (51)$$

$$= 2.68 \text{ bits.} \qquad (52)$$

Thus, on average, each trial provides the observer with between 2.68 and 3.55 bits.

8.6. SHANNON INCREMENT

Using a conservative estimate of mutual information of $I = 2.68$ bits suggests that the observer can discriminate differences between the reference and comparison stimulus with an average resolution of about one part in 6.39 ($= 2^{2.68}$) of the effective range x_{range} of stimulus strengths. Note that the range of scaling values used $s_{\text{range}} = 0.2$ (i.e., $0.9 \ldots 1.1$) equals the range of stimulus strengths $x_{\text{range}} = 0.2$ (i.e., $-0.1 \ldots 0.1$). Therefore, the SI for the width scaling factor is

$$SI = x_{\text{range}}/2^I \qquad (53)$$

$$= 0.2/6.39 \qquad (54)$$

$$= 0.031, \qquad (55)$$

where we have assumed $\epsilon = 0$ here. Thus, on average, the smallest change in scaling factor (between reference and comparison stimulus) detectable by the observer is SI = 0.031.

9. DISCUSSION

We have shown how the PRD model from Palmer et al. (2005) can be extended to make use of individual RTs, which can be combined with binary observer responses to estimate key psychophysical parameters in a 2AFC design.

A key feature of diffusion-based models is that they treat each RT as the end-point of an accumulation of evidence. If we take this type of evidence-accumulation process seriously then it makes sense to model the distribution of RT values as an inverse Gaussian distribution (for reasons described in section 5).

A striking result is the difference between the log likelihoods associated with the binary response model and the RT model, despite the fact that the binary response model has fewer free parameters than the RT model, and that both models provide similar PSE estimates which (based on their sems, not shown) are not significantly different. These log likelihood values suggest that the EPRD model provides a better fit to the RT data than it does to the binary response data. This difference in likelihoods suggests that the parameter estimates obtained using the combined RT and response data is dominated by the binary data likelihood term.

Self-evidently, both the RT and binary responses of an observer depend on the stimulus strength x. However, in general, it is not known if RT or binary response data provide more Shannon information about the value of x. More importantly, and more subtly, it is not known if they provide the same information about x, or if they merely provide the same *amount* of information about x (see **Figure 1**).

We can gain some insight into the nature of this problem by considering the proportion of the differential entropy in stimulus values accounted for by the corresponding differential entropy in observer responses. At one extreme, if an observer is told to respond as quickly as possible then the RTs should provide relatively large amounts of mutual information regarding stimulus strength, whereas the binary responses carry relatively little mutual information (because speeded responses tend to be inaccurate Hanks et al., 2011). In this case, the RT entropy at a given stimulus strength will be relatively small, because RTs will be tightly coupled to the stimulus strength, whereas the binary response entropy at a given stimulus strength will be relatively large (because these responses are inaccurate, and therefore not tightly coupled to the stimulus strength). However, when considered across different stimulus strengths, the tight coupling between RT and stimulus strength will give rise to a relatively large RT entropy, and most of this entropy will be shared with stimulus strength entropy (which defines a large mutual information between RT and stimulus strength). In contrast, these fast, inaccurate responses across stimulus strengths will be associated with a relatively small range of response probability values (e.g., $P \approx 0.5$), which will therefore have a relatively small entropy, most of which is not shared with the stimulus strength entropy (which defines a small mutual information between binary responses and stimulus strength). In summary, fast responses should yield high entropy RT values, which share a large proportion of their entropy with the stimulus strength, combined with low entropy P values which share a small proportion of their entropy with the stimulus strength. At the other extreme, if an observer is told to be as accurate as possible then this should yield high entropy P values which share a large proportion of their entropy with the stimulus strength, combined with low entropy RT values which share a small proportion of their entropy with the stimulus strength. In summary, the entropy in stimulus strength can be shared with entropy in both accuracy (P) and speed (RT). However, as there is probably only a finite amount of such shared entropy (mutual information) available, we predict that it can be realized experimentally as maximum speed or maximum accuracy, but not both.

The scenario considered above can be represented geometrically, as in **Figure 1**. If we compare the mutual information between τ and x with the mutual information between r and x then it is possible that they have the same magnitude [e.g., $(a + c) = (a + b)$, as in **Figure 1**]. However, the fact that both τ and x have the same *amount* of mutual information (i.e., they account for the same amount of entropy in x) does not imply that they account for the same *entropy* in x. Formally, the fact that $(a + c) = (a + b)$ does not imply that $(a + c) \equiv (a + b)$. This matters because, even if $I(x, \tau) = I(x, r)$, we could not conclude that $I(x, \tau) \equiv I(x, r)$, and so we could not conclude that τ and r provide mutually redundant information. Thus, we cannot dismiss τ simply because r accounts for more entropy in x than τ does (or vice versa). Indeed, this is precisely the situation that we have in the results reported here, and provides reasonable grounds for making use of both RT and binary response data in general.

Unfortunately, we have been unable to derive an expression for the total mutual information between the joint variables (RT and binary responses) and stimulus strength $I(\bar{\tau}, P; x')$ (i.e., the area $[a + b + c]$ in **Figure 1**), although it may be possible to do so using Equation (10) [where the entropy of the difference between P and $\bar{\tau}$ is $H(\bar{\tau}, P | x')$]. The precise effect of the instructions given to observers on mutual information, and the proposed invariance of the total mutual information with respect to instructions, clearly require further research (Soukoreff and MacKenzie, 2009).

The Shannon increment (SI) is similar in spirit to the more conventional just noticeable difference (JND). However, the JND has an arbitrary value, and (despite its name) there is no reason to suppose that a JND is indeed just noticeable. The SI is monotonically related to the average amount of Shannon information an observer gains regarding a single presentation of a stimulus, and is a measure of the perceptual resolution with which a parameter is represented by the observer.

10. CONCLUSION

We have presented an extended proportional-rate diffusion model, which takes account of both individual RTs and binary responses for maximum likelihood estimation of key psychophysical parameters (e.g., PSE, slope) of the psychometric and chronometric functions. The fact that these psychophysical parameters have similar estimated values when computed independently for two models based on RTs alone or on binary responses alone provides support for the underlying physical basis of this class of diffusion models.

An information-theoretic analysis was used to estimate the average amount of Shannon information that each RT provided about the stimulus value, and also the average amount of Shannon information that each binary response provided about the stimulus value. This analysis provides bounds for the average amount of Shannon information that the observer gains about the stimulus value from one presentation, which was found to be between 2.68 and 3.55 bits/trial for the experiment used here.

REFERENCES

Bonnet, C., Ars, J., and Ferrer, S. (2008). Reaction times as a measure of uncertainty. *Psicothema* 20, 43–48.

Hanks, T., Mazurek, M., Kiani, R., Hopp, E., and Shadlen, M. (2011). Elapsed decision time affects the weighting of prior probability in a perceptual decision task. *J. Neurosci.* 31, 6339–6352. doi: 10.1523/JNEUROSCI.5613-10.2011

Harvey, L. (1986). Efficient estimation of sensory thresholds. *Behav. Res. Methods Instrum. Comput.* 18, 623–632. doi: 10.3758/BF03201438

Luce, R. D. (1986). *Response Times: Their Role in Inferring Elementary Mental Organization*. New York, NY: Oxford University Press.

MacKay, D. (2003). *Information Theory, Inference, and Learning Algorithms*. Cambridge: Cambridge University Press.

Palmer, J., Huk, A., and Shadlen, M. (2005). The effect of stimulus strength on the speed and accuracy of a perceptual decision. *J. Vis.* 5, 376–404. doi: 10.1167/5.5.1

Ratcliff, R. (1978). A theory of memory retrieval. *Psychol. Rev.* 85, 59–108. doi: 10.1037/0033-295X.85.2.59

Rieke, F., Warland, D., van Steveninck, R., and Bialek, W. (1997). *Spikes: Exploring the Neural Code*. Cambridge, MA: MIT Press.

Shannon, C., and Weaver, W. (1949). *The Mathematical Theory of Communication*. Urbana, IL: University of Illinois Press.

Smith, P. (1990). A note on the distribution of response times for a random walk with Gaussian increments. *J. Math. Psychol.* 34, 445–459. doi: 10.1016/0022-2496(90)90023-3

Soukoreff, R. W., and MacKenzie, I. S. (2009). "An informatic rationale for the speed-accuracy tradeoff," in *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics, 2009* (San Antonio, TX), 2890–2896. doi: 10.1109/ICSMC.2009.5346580

Stone, J. (2014). *Information Theory: A Tutorial Introduction*. Sheffield: Sebtel Press.

Swanson, W., and Birch, E. (1992). Extracting thresholds from noisy psychophysical data. *Percept. Psychophys.* 51, 409–422. doi: 10.3758/BF032 11637

Tweedie, I. M. C. K. (1957). Statistical properties of inverse Gaussian distributions. *Ann. Math. Stat.* 28, 362–377. doi: 10.1214/aoms/1177706881

Wichmann, F. A., and Hill, N. (2001). The psychometric function: I. Fitting, sampling, and goodness of fit. *Percept. Psychophys.* 63, 1293–1313. doi: 10.3758/BF03194544

APPENDIX

MATHEMATICAL SYMBOLS AND ABBREVIATIONS

A an EPRD model parameter which is the amount of evidence required to trigger a response.

comparison stimulus response: a response indicating the comparison stimulus was chosen.

EPRD: extended proportional-rate diffusion model.

SI: Shannon increment, the smallest detectable change in a stimulus.

γ EPRD lapse rate parameter.

i index over stimulus strength x, with range $i = 1, \ldots, N_x$.

j index over trials at one stimulus strength x_i, with range $j = 1, \ldots, N_i$.

k index over stimulus strength, with range $k = 1, \ldots, N_x$.

K is a measure of sensitivity to changes in x in the EPRD model.

N_i number of trials at stimulus strength x_i.

N_x number of different stimulus strengths.

PSE: point of subjective equality.

P_i proportion of comparison stimulus responses at stimulus strength x_i, predicted by EPRD model.

P'_i MLE mean, equal to observed proportion of comparison responses at stimulus strength x_i.

r binary observer response (e.g., observer chooses comparison or reference stimulus).

s_C variable stimulus value of the comparison stimulus.

s_R fixed stimulus value of the reference stimulus.

s_{PSE} value of the comparison stimulus which the observer perceives as being the same as the reference stimulus.

$\bar{\tau}'_i$ MLE mean of inverse Gaussian RT at stimulus strength x_i.

$\bar{\tau}_i$ mean RT at stimulus strength x_i, as predicted by EPRD model.

$\bar{\tau}_{dec,i}$ mean decision RT at stimulus strength x_i, as predicted by EPRD model.

$\bar{\tau}_{res}$ mean residual RT (assumed the same at all stimulus strengths), as predicted by EPRD model, where $\bar{\tau}_{res} = \bar{\tau}_{dec,i} - \bar{\tau}_i$.

$\theta_\tau = (s_{PSE}, A, K, \gamma, \bar{\tau}_{res})$, five parameters for the RT component of the EPRD model.

$\theta_P = (s_{PSE}, AK, \gamma)$, three parameters for the binary response component of the EPRD model.

$v_{\bar{\tau},i}$ variance in mean RT.

x_i stimulus strength.

x'_i perceived strength of stimulus with strength x_i.

Decoupling speed and accuracy in an urgent decision-making task reveals multiple contributions to their trade-off

Emilio Salinas*, Veronica E. Scerra , Christopher K. Hauser , M. Gabriela Costello and Terrence R. Stanford

Department of Neurobiology and Anatomy, Wake Forest School of Medicine, Winston-Salem, NC, USA

A key goal in the study of decision making is determining how neural networks involved in perception and motor planning interact to generate a given choice, but this is complicated due to the internal trade-off between speed and accuracy, which confounds their individual contributions. Urgent decisions, however, are special: they may range between random and fully informed, depending on the amount of processing time (or stimulus viewing time) available in each trial, but regardless, movement preparation always starts early on. As a consequence, under time pressure it is possible to produce a psychophysical curve that characterizes perceptual performance independently of reaction time, and this, in turn, makes it possible to pinpoint how perceptual information (which requires sensory input) modulates motor planning (which does not) to guide a choice. Here we review experiments in which, on the basis of this approach, the origin of the speed-accuracy trade-off becomes particularly transparent. Psychophysical, neurophysiological, and modeling results in the "compelled-saccade" task indicate that, during urgent decision making, perceptual information—if and whenever it becomes available—accelerates or decelerates competing motor plans that are already ongoing. This interaction affects both the reaction time and the probability of success in any given trial. In two experiments with reward asymmetries, we find that speed and accuracy can be traded in different amounts and for different reasons, depending on how the particular task contingencies affect specific neural mechanisms related to perception and motor planning. Therefore, from the vantage point of urgent decisions, the speed-accuracy trade-off is not a unique phenomenon tied to a single underlying mechanism, but rather a typical outcome of many possible combinations of internal adjustments within sensory-motor neural circuits.

Keywords: choice, computational model, decision making, discrimination, mental chronometry, race to threshold, saccade, subtraction method

1. THE PROBLEM OF PARSING THE REACTION TIME

In daily life, some decisions are rather abstract (should I trust the financial adviser?) whereas others require a specific action (should I press the brake or the accelerator?). Within the latter category, speed and accuracy are inversely related in virtually every task (Woodworth, 1899; Hick, 1952; Wickelgren, 1977; Chittka et al., 2009); the faster the decision, the less accurate the outcome. This means that the traditional, key quantities that are used to measure cognitive performance, the reaction time (**RT**) and the percentage of correct responses, are fundamentally intertwined. What is the underlying cause of this interdependence? How does it emerge from the structure and dynamics of neural circuits? We consider these questions in the context of decisions that are coupled to immediate actions.

Intuitive models of the speed-accuracy trade-off have been formulated (Reddi and Carpenter, 2000; Bogacz et al., 2010; Standage et al., 2013), but the empirical investigation of these questions reveals further complexity (Cook and Maunsell, 2002;

DiCarlo and Maunsell, 2005; Battaglia and Schrater, 2007; Cohen et al., 2009; Heitz and Schall, 2012). Part of the problem is that the RT reflects the total amount of time consumed by all the subsystems that contribute to a choice or decision process. Thus, when a subject executes an action in response to a sensory scene, the RT must comprise, at the very least, the time necessary for analyzing the sensory information plus the amount of time required to plan the motor action that is congruent with that information. Discerning just these two components has been challenging because the underlying neural networks are themselves strongly interrelated: neurons that encode a subject's perceptual decision, that participate in motor planning, or that do both, are typically found within the same, local microcircuits (Horwitz and Newsome, 1999; Shadlen and Newsome, 2001; Hernández et al., 2010; Costello et al., 2013; Mante et al., 2013). Furthermore, other distinct cognitive processes may contribute to the RT too; for instance, deploying visuospatial attention or accessing information stored in memory could represent separate processing

steps requiring a certain amount of time to unfold independently of the perceptual and motor-planning stages (Sternberg, 1966; Monsell, 2003; Horowitz et al., 2004; Busse et al., 2008). As such, the RT must reflect a total sum over the times consumed by multiple covert processes (Sternberg, 1969), each of which could conceivably constrain or be traded against the others.

Broadly speaking, three techniques have been used to distinguish the two major components of the RT during relatively fast perceptual decision-making tasks, i.e., the contributions of perceptual and motor-planning processes. (1) A common approach is to introduce a delay between the perceptual evaluation and the motor report required in each trial. This strategy is widely used to characterize neuronal activity as sensory-, memory-, or movement-related neurons (Shadlen and Newsome, 2001; Sommer and Wurtz, 2001; Lemus et al., 2007). (2) Another possibility is to limit the amount of cue viewing time (Bergen and Julesz, 1983; Ratcliff and Rouder, 2000; Bodelón et al., 2007; Kiani et al., 2008). The idea is that neurons, or any processing component in general, whose responses vary systematically as functions of cue viewing time may be strongly involved in the analysis of perceptual information. This manipulation is not as straightforward as it may seem, though, because controlling very short stimulus durations is difficult and typically requires additional masking stimuli to prevent stimulus persistence, and such masking introduces other potential problems (Breitmeyer and Ogmen, 2000, 2006). (3) An alternative that is not quite as intuitive, is to do the reverse of 1: inform choices on the basis of urgent perceptual decisions. That is, start preparing a motor response first, before the relevant cue information becomes available (Ghez et al., 1989; Chapman et al., 2010). That way, the initial motor planning stage stays relatively constant.

That is the approach we have taken (Stanford et al., 2010; Shankar et al., 2011; Costello et al., 2013). It provides a simple and highly effective way to dissociate motor and perceptual performance, and thus a different set of tools with which to characterize and quantify perceptual decision-making mechanisms. Here we review previously published results of experiments in which urgent decisions inform rapid choices (Stanford et al., 2010; Shankar et al., 2011; Costello et al., 2013), but focus specifically on their implications for understanding the origin of the speed-accuracy trade-off. As discussed below, under this light it is possible to see not only how perceptual capacity and motor execution interact to determine the response speed and success rate of a subject, but also how additional factors such as motivation or internal preference may alter that interaction. In this way, it becomes quite clear that the speed-accuracy trade-off is not a unitary phenomenon derived from a unique, underlying mechanism, but is instead the result of multiple, semi-independent moving parts that interact with each other within sensory-motor neural circuits.

2. PERCEPTUAL DECISIONS UNDER TIME PRESSURE

As a means to disambiguate perceptual and motor processes, we designed a compelled-response task wherein participants are given the instruction to respond *before* the relevant perceptual information appears (Stanford et al., 2010). In the oculomotor version, the compelled-saccade task (**Figure 1A**), the response is

FIGURE 1 | Dissociating perceptual and motor performance in the compelled-saccade task. (A) Sequence of events in the task. The subject is required to make a saccade when the fixation point disappears (go). If the chosen target matches the color of the fixation point (red, in this example), the choice is correct and a reward is obtained. The go instruction is given first, before the relevant sensory information is revealed (cue). The gap (50—250 ms) is the time interval between the go and the cue. The rPT is the amount of time during which the color information can potentially inform the saccadic choice. **(B)** Percentage of correct responses as a function of gap, or psychometric curve. **(C)** Mean RT (±1 SD) as a function of gap, or chronometric curve. Both correct and incorrect trials are included. **(D)** Percentage of correct responses as a function of rPT (equal to RT − gap), or tachometric curve. In **(B–D)**, blue and black lines/symbols correspond to psychophysical and simulation results, respectively. See Shankar et al. (2011) for details about the experimental data and modeling methods.

an eye movement. First, the observer fixates on a central spot, the color of which indicates the color of the eventual target. Then two yellow (neutral) dots appear in the periphery; these are simply placeholders indicating the possible response locations. Next, the central fixation point disappears, and this is the "go" signal

that tells the observer "respond now!" Note that, when the go is given, the identities of the target and distracter are still unknown, but the observer must begin planning a movement to one of the two potential targets nonetheless. Then, after a variable time gap (from 50 to 250 ms) the peripheral dots change color, revealing one to be the target and the other the distracter. The onset of the subject's response occurs when the eyes just start moving, and marks the end of the RT period that started at the go signal (**Figure 1A**).

The logic behind this design is that, by telling the subject when to respond, the motor choice process is initiated early, and so perceptual information, once presented, influences a motor plan that is already developing. By unpredictably varying the time delay between the go signal and the appearance of the color cue (i.e., the gap), the subject generates responses that range between fully informed choices (for gaps that are much shorter than the typical saccadic RT) and fully uninformed choices, or guesses (for gaps that are comparable to the typical saccadic RT), all with the same underlying distribution of motor plans. So, it becomes possible to dissociate the effect of motor preparation from the perceptual decision-making process in an otherwise standard saccadic choice task.

The crucial event in the task is the go instruction, which compels the subject to respond before the target and distracter are revealed. But, why is it that subjects do not simply wait for the color cue to appear before making a choice? In essence, there are three reasons. First, because responding is natural; with the fixation point gone and two salient objects present, it takes effort not to look at one of them. Second, because throughout training, the subjects learn two separate rules, (1) that the offset of the fixation point means "respond now!" and (2) that the correct choice is the one matching the color of the fixation point. Rule 1 is learned first, and if necessary, which is not always the case, it is practiced independently of rule 2. And third, during the compelled-saccade task subjects have a limited time window for making a valid response, so a trial is scored as incorrect—and no reward is given—if the RT is too long, regardless of the choice. It should be noted, however, that consistent with the first two points, such trials in which the RT limit is exceeded are extremely rare (<2%). For a detailed analysis of possible waiting strategies see Salinas et al. (2010).

Performance in the task is expected to decline toward chance as a function of the gap, and indeed this is what happens, as illustrated with representative data from two monkeys (**Figure 1B**). In contrast, RTs are expected to remain approximately—but not exactly—constant, and this is also the case: mean RTs change by less than 30 ms, or approximately 10%, while performance varies between chance and near 100% correct (**Figure 1C**). In other perceptual decision-making tasks, RTs often show comparable variations of a few tens of milliseconds, although the difference over the full performance range sometimes reaches hundreds of milliseconds, or several fold (Wolfe, 1998; Ratcliff and Smith, 2004; Palmer et al., 2005; Reinagel, 2013a,b). The variation of ~30 ms in mean RT seen in the compelled-saccade task is modest, but more importantly, it and the systematic increase in the spread of the RTs with gap (error bars in **Figure 1C**) can be fully accounted for by a simple model; the key notion is that even though the initial motor planning process is statistically the same

for all the gaps, the motor conflict is resolved sooner or later depending on when the perceptual information arrives (more on this below).

2.1. THE TACHOMETRIC CURVE

Although the gap is the main control parameter in the task, the variable that fundamentally determines the probability of success in each trial is the raw processing time, or **rPT** (**Figure 1A**), which is the amount of time before the onset of the motor response during which the color information is available to inform the choice. It is important to stress that this theoretical limit to the maximum amount of cue viewing time is a trial-specific quantity, which can be easily computed via

$$rPT = RT - gap \qquad (1)$$

based on the gap and the RT recorded in each trial. Using this equation we can determine how long the perceptual information was available for guiding each saccadic choice. Furthermore, by plotting the percentage of correct responses versus rPT we obtain a "tachometric curve," a curve that characterizes the perceptual performance of a subject (**Figure 1D**).

This curve has a sigmoidal shape with parameters that are readily interpretable in terms of psychophysical capacity. For the data in **Figure 1D**, the saturation values are very near 100% correct, so the color information is fully exploited by the subjects when they have enough time to view the cue. The center point of the curve is the rPT at which the percent correct is halfway between chance and the saturation value. It is an indication of how much viewing time is necessary for perceptual information to have a significant impact on performance. For the curves in **Figure 1D**, the center points are 134 ± 2 ms for monkey S and 157 ± 2 ms for monkey G. Trials to the left of the center point are mostly near chance performance and correspond predominantly to uninformed choices, or guesses, whereas trials to the right of the center point correspond mostly to fully informed choices. Finally, the steepness of the curve near the center point provides a measure of the speed with which perceptual information influences the choice once this information has begun having an impact. For monkeys S and G, half of the performance range (from 62.5% to 87.5% correct) is covered within 24 ± 2 ms and 40 ± 2 ms, respectively, so the color discrimination unfolds extremely rapidly once it gets going.

2.2. MULTIPLE MECHANISMS FOR GENERATING TRADE-OFFS

In principle, simultaneous variations in RT and percent correct may result from changes in motor planning alone, in perception alone, or in both, and to distinguish these options it is essential to have independent, quantifiable measures of their impact on choice behavior. That is the key advantage of the compelled-response approach, it provides independent assessments of perceptual and motor performance in the tachometric and chronometric (RT versus gap) curves, respectively. This is illustrated in detail further below with data from two experiments, but before discussing those, it is useful, first, to consider some simplified examples, and second, to gain some mechanistic intuition, via a heuristic model, about the ways in which

perceptual and motor-planning processes interact when decisions are made under time pressure. The three scenarios that follow are meant simply to illustrate, based on the model, how variations in the three psychophysical curves obtained in the compelled-saccade task (**Figures 1B–D**) may relate to each other.

The speed-accuracy trade-off is often explained in terms of a change in threshold (Reddi and Carpenter, 2000; Bogacz et al., 2010; Hanks et al., 2011). That is, a motor response is triggered after a "decision variable" reaches a particular value (**Figure 2A**), and increasing that value produces both higher RTs and a higher proportion of correct responses. This is because it takes longer for the variable to go from baseline to threshold in each trial, and a longer RT means more time during which the perceptual information can advance the decision variable in the correct direction. Although our theoretical framework is somewhat different (and considers the threshold to be fixed; see below), in the compelled-saccade task a change in threshold would have precisely the expected effects associated with a standard trade-off

(**Figures 2B,C**), but notably, it would have absolutely no impact on the tachometric curve (**Figure 2D**). This is because, in contrast to the psychometric and chronometric curves, the tachometric curve is highly insensitive to the dynamics of the motor planning process. In essence, it reflects how soon (after the cue is revealed) and how strongly the perceptual information modulates motor activity that is already rising. In the context of this urgent decision-making task, both the rising activity—which specifically represents a motor plan—and the threshold are properties intrinsic to the motor circuitry, in agreement with neurophysiological evidence (Costello et al., 2013 see also Hanes and Schall, 1996; Heitz and Schall, 2012). So, variations in threshold could produce a standard trade-off between speed and accuracy in the compelled-saccade task, but this mechanism would leave the tachometric curve intact.

The response threshold is not the only quantity that may be altered to produce a trade-off. In theory, varying the baseline level of activity would be essentially equivalent (see Bogacz

FIGURE 2 | Changes in speed and accuracy generated by three distinct mechanisms. All results are expectations based on model simulations of the compelled-saccade task. **(A)** Schematic of a developing oculomotor plan. A saccade is triggered shortly after motor-related activity (blue trace) reaches a threshold (dotted lines). The threshold-crossing time, and thus the RT, varies with the threshold value. **(B–D)** Model results. Variations in threshold produce a standard speed-accuracy trade-off. As the threshold increases, both performance **(B)** and mean RT **(C)** increase, but the tachometric curve **(D)** does not change. **(E)** Schematic illustrating how threshold-crossing time varies with the build-up rate of the motor plan. **(F–H)** Model results. Variations in mean build-up rate also produce a trade-off. As the mean build-up rate decreases,

both performance **(F)** and mean RT **(G)** increase, whereas the tachometric curve **(H)** changes minimally. **(I)** Schematic illustrating how threshold-crossing time varies with the latency of the go signal. **(J–L)** Model results. Variations in visual latency alone do not produce a trade-off. When the latency of the visual stimuli (go signal and color cue) decreases, performance does not change **(E)**, but mean RT **(F)** decreases and the tachometric curve **(G)** shifts to the left, indicating that perception informs the subject's choices systematically sooner. Results are from model simulations either with identical parameters as in **Figure 1** (for monkey S), or with a 25% increase or a 25% decrease in the value of one parameter, either the threshold **(B–D)**, the mean build-up rate of the motor plans **(F–H)**, or the latency of the visual information **(J–L)**.

et al., 2010). But beyond that, in the context of urgent-decision tasks in which motor planning starts before perceptual analysis, changes in the mean build-up rate of the developing motor activity (**Figure 2E**) would produce qualitatively similar effects (**Figures 2F–H**). The intuition is simple: when motor plans rise more quickly, the excursion from baseline to threshold takes less time and there is, consequently, less opportunity for the perceptual information to influence those ongoing motor plans. This case would again correspond to behavioral changes driven exclusively by modulations in the dynamics of the motor circuitry.

Finally, consider another hypothetical scenario in which the only difference between three conditions is in the latency with which the visual stimuli, i.e., the go signal (**Figure 2I**) and the cue, may start informing the motor plans. This latency could depend on multiple factors, such as contrast or alertness, for instance, but regardless of the cause, everything else being equal, a decrease in visual latency would manifest in a very specific way: it would decrease the mean RTs (**Figure 2K**), because effectively all afferent delays would be shorter; it would produce a leftward shift of the tachometric curve (**Figure 2L**), indicating that perception starts guiding performance sooner relative to cue onset; and it would have no effect on the observed percentage of correct responses (**Figure 2J**), because the motor plans would still have the same amount of time to rise before the arrival of the cue information (stated differently, from the point of view of the motor planning circuit, the time elapsed between the arrival of the go signal and the arrival of the color cue would not change). So, in this case the RT would drop without a trade-off, and the underlying mechanism would be purely sensory/perceptual.

Now, if the mechanisms illustrated in **Figure 2** could be combined arbitrarily, it would be possible to produce a range of trade-offs with widely different magnitudes in terms of the ratio of change in percent correct to change in RT. Also keep in mind that, for simplicity, these effects were illustrated based on just three parameters, saccade threshold, mean build-up rate, and visual latency, but various other parameters of the motor and perceptual circuits could serve to modulate performance. So, more generally, different combinations of alterations in the intrinsic dynamics of the motor plans and in the perceptual discrimination process could lead to large or small changes in RT coupled to large or small changes in accuracy.

In conclusion, while overall it may still be true that increases in performance are accompanied by increases in RT, and vice versa, a trade-off may occur for very different reasons, and its magnitude may vary enormously. As will be shown below, this is precisely what appears to be happening under realistic experimental conditions.

3. A HEURISTIC MODELING FRAMEWORK FOR DESCRIBING URGENT DECISIONS

This section presents a model that replicates the performance of subjects in the compelled-saccade task and is consistent with the neurophysiology of the underlying neural circuits. Although the model and associated theoretical framework have been described before (for details and parameter values see Salinas et al., 2010; Shankar et al., 2011), they are important for interpreting the experimental data relevant to the speed-accuracy trade-off that

are discussed below. We review key findings that establish the model's credibility.

The preparation for action in the context of the compelled-saccade task can be viewed as a competition between two opposing motor plans that develop concurrently, racing to a threshold for triggering a movement to one of the two potential target locations. The direction in which the eyes move is determined by whichever plan reaches the threshold first, but crucially, when perceptual information is available, it modulates the ongoing plans to favor the correct choice. That is the essence of the "accelerated race-to-threshold" model (Salinas et al., 2010; Stanford et al., 2010; Shankar et al., 2011). Such a model is useful because it provides a quantitative, yet intuitive, link between the measured psychophysical behavior and its neural basis.

The dynamics of the model are determined by two key assumptions. (1) That the cue information accelerates the motor plan developing toward the target and decelerates the plan developing toward the distracter. And (2), that in each trial, the competing motor plans begin rising toward threshold shortly after the go signal, with initial build-up rates drawn randomly from a distribution. In this way, the outcome of any given trial depends on when the cue information becomes available relative to how advanced each of the developing oculomotor plans is at that time, and notably, this interaction can take just five distinct forms (**Figures 3A–E**).

In these examples (**Figures 3A–E**), a correct choice is produced when the cyan motor plan wins the race. In all trials, the initial build-up rates randomly favor one of the two potential targets, and reflect the subject's initial predisposition. So, when a saccade is triggered before the cue information arrives (**Figures 3D,E**), the result is an uninformed choice, i.e., a guess. Note that the probability of such a random outcome increases both for longer gap durations and for higher initial build-up rates. In contrast, when the cue information arrives early enough to guide the ongoing motor plans, the result is an informed choice (**Figures 3A–C**). However, the initial build-up rates still play an important role: when the motor plan that is congruent with the target starts as the leader, it curves upward slightly and triggers a correct saccade with a short RT (**Figure 3A**), but when this target-related plan lags behind, it starts out slowly and has more ground to cover once the acceleration kicks in, so it takes longer to reach threshold (**Figure 3B**). On such trials success also requires that the distracter-related plan be decelerated, but if this leading plan is sufficiently advanced, the influence of the cue information may be insufficient to prevent it from reaching threshold and producing an incorrect saccade (**Figure 3C**).

Thus, the mechanistic signature of the model is this: when the motor activity evoked shortly after the go signal is intense and strongly committed to one of the potential targets, the result is typically an uninformed choice (with a short or negative rPT), whereas when the initial motor activity is more moderate and less biased, the later arriving perceptual information is more likely to resolve the motor conflict in favor of the correct choice (with a long rPT). We think that similar dynamics are, in general, the underlying substrate of rapid perceptual decisions lasting just a fraction of the RT (see Discussion; Cisek and Kalaska, 2010). The accelerated race-to-threshold model, which instantiates these

FIGURE 3 | The accelerated race-to-threshold model closely reproduces psychophysical data in the compelled-saccade task.
(A–E) Simulated trials illustrating the five essential types of interaction between competing motor plans. In each panel, two competing variables represent oculomotor activity that triggers an eye movement either to the right (cyan) or to the left (magenta). In these examples the target is on the right, so races in which the cyan trace reaches threshold first correspond to correct choices. The two variables start racing 75 ms (afferent delay) after the go signal and a saccade is triggered 15 ms (efferent delay) after threshold is crossed. Initial build-up rates are drawn randomly in each trial. Gray shades indicate the time during which the cue information is available to modulate the motor activity. During informed choices (**A–C**, gap = 100 ms), the motor plan toward the target accelerates (its build-up rate increases) and the plan toward the distracter decelerates (its build-up rate decreases), whereas during guesses (**D,E**, gap = 250 ms) the build-up rates do not change. **(F)** Reaction time distributions in correct (cyan) and incorrect (magenta) trials at specific gaps (indicated on upper left corners). Results are shown for two monkeys, S (left) and G (right). In each plot, the black curves correspond to model simulations. Vertical lines indicate the center point of the tachometric curve of the corresponding monkey. Results are based on the same experimental and simulated data as in **Figure 1**.

interactions quantitatively, is consistent with both psychophysical and neurophysiological data, as discussed next.

3.1. ACCOUNTING FOR THE MICROSTRUCTURE OF BEHAVIOR

With the correct parameter values, the model can replicate a monkey's psychometric, chronometric and tachometric curves very accurately (**Figures 1B–D**, compare Data versus Model), but each point in these curves aggregates many trials with motor competitions (races) of different types (**Figures 3A–E**), so the three curves provide a relatively coarse summary of the subject's behavior. Matching the full RT distributions for correct and error trials at each individual gap (**Figure 3F**) is a much more stringent benchmark for any model (Salinas et al., 2010), because the shapes of these distributions are directly related to the more limited mixtures of race trajectories that occur at each gap.

For example, the distribution of RTs for correct responses, or hits (**Figure 3F**, cyan histograms), clearly transitions from unimodal to bimodal. According to the accelerated race-to-threshold model, this is because short gaps contain a large proportion of fast informed decisions (**Figure 3A**), whereas long gaps contain a mixture of correct guesses (**Figure 3E**) and slow informed decisions (**Figure 3B**). Similarly, the distribution of RTs for incorrect responses (**Figure 3F**, magenta histograms) contains mostly wrong guesses (**Figure 3D**) and a small proportion of informed choices that were nonetheless incorrect (**Figure 3C**), which occupy the rightward tail of the histograms.

These combinations are easy to distinguish by noticing that the rPT that corresponds to the center point of the tachometric curve can be marked as a line in each plot (vertical lines in **Figure 3F**), and that this line divides each RT distribution into two parts: the

trials to the right are all informed choices, whereas the trials to the left are, except for those very near the line, uninformed choices. With this in mind, it becomes immediately obvious that correct trials at short gaps are almost always informed (cyan histograms in **Figure 3F**, top; RTs are predominantly to the right of the line), whereas correct trials at long gaps are almost always lucky guesses (cyan histograms in **Figure 3F**, bottom; RTs are predominantly to the left of the line). This also explains why, when looking at the correct responses going from long to short gaps, the peak to the right of the line moves progressively to the left: as the perceptual information arrives earlier and earlier, more and more trials that would have otherwise ended up on the rightward tail of the distribution are accelerated, resulting instead in short RTs. The position of the line itself shifts to the left as the gap decreases because rPT and RT differ precisely by the gap value (Equation 1), but the center point of the tachometric curve remains a fixed number for each monkey—a number that, as mentioned earlier and illustrated in **Figure 3**, is crucial for assessing the degree to which perceptual information determines the outcome of each trial.

3.2. MODEL PARAMETERS AND THEIR INTERPRETATION

As implemented here, the accelerated race-to-threshold model has 11 parameters that can be adjusted to fit the psychophysical data of individual monkeys, as in **Figure 3F** (Salinas et al., 2010; Shankar et al., 2011; Costello et al., 2013). Although this number may seem large, the effect of any given parameter is quite limited; each one affects the dynamics of the two competing motor plans in a very specific way and has a well-defined neurophysiological interpretation.

Three parameters describe the distribution from which the initial build-up rates of the two motor plans are drawn in each trial. A description based on three numbers, corresponding to the mean, variance, and correlation of the build-up rates, is quite minimal for a two-dimensional (joint) distribution.

Two parameters, one for the mean and another for the variance, determine the visual latency in each trial. This latency is agnostic about the underlying causes (afferent delay, additional visual processing stages, etc.); it simply describes when the relevant visual information (go and cue) reaches the model circuit. For the results presented here, we assume that the mean latencies of the go signal and the color cue are the same, but this is not necessarily the case in general.

Three parameters describe how perception alters the trajectories of the ongoing motor plans (as in **Figures 3A,B**); they specify the magnitude of the acceleration and deceleration and how long they last. Using fixed acceleration and deceleration coefficients is the simplest possible way to describe motor plans that are not perfectly straight, i.e., for which the build-up rates are not constant.

One other parameter, the probability of confusion or lapse rate, accounts for incorrect responses that occur at long processing times and cannot be attributed to insufficient cue viewing time. There are many possible reasons for such lapses; here they are simply considered random events.

Finally, two additional parameters are included to replicate a subtle but systematic feature seen in distributions of RTs that are bimodal (as in **Figure 3F**, for 175–225 ms gap), a dip that is slightly more pronounced than expected. This corresponds closely to a phenomenon known as "saccadic inhibition" that occurs when a distracting stimulus appears while a saccade is already being programmed (Reingold and Stampe, 2002; Buonocore and McIntosh, 2008, 2012; Bompas and Sumner, 2011). The race model accounts for this deviation via a brief interruption in the rise of the motor activity linked to the detection of the cue. The two corresponding parameters determine the onset and offset of the brief pause, and have a relatively minor impact on other aspects of the data.

Thus, the model starts with a simple description of the motor choice process and is augmented with a mechanism whereby perception can guide it. So, is the model activity comparable to saccade-related neural responses evoked during perceptually driven choices?

3.3. LINKING BEHAVIOR AND NEUROPHYSIOLOGY

The accelerated race-to-threshold model provides excellent fits to the RT distributions at fixed gaps for all the monkeys we have trained in the compelled-saccade task (Shankar et al., 2011). Although this is certainly reassuring, psychophysical data alone cannot fully constrain or validate such a mechanistic model, even if the fits were perfect; this is true not only for our model (Salinas et al., 2010) but also in general (e.g., Ratcliff and Smith, 2004; Brunton et al., 2013; Miller and Katz, 2013). However, the activity of single neurons recorded in the frontal eye field (**FEF**) of behaving monkeys is consistent with key, non-trivial predictions of the model (Salinas et al., 2010; Stanford et al., 2010; Costello et al., 2013), suggesting that, indeed, its basic layout is correct.

To generate specific predictions directly comparable to neurophysiological data, the model was run with parameter values that fitted the behavioral data of monkey S, and expected neural responses (**Figures 4B,C**) were computed by averaging separately the simulated motor plans obtained in short- and long-rPT trials. The short- and long-rPT intervals were defined according to the tachometric curve so that they would include chiefly guesses and informed choices, respectively (**Figure 4A**, shaded areas). In this way we could ask: how should the mean neural responses differ between correct, uninformed guesses and correct, informed discriminations?

The answer to this question comprises essentially two predictions about the relative amounts of activity for saccades in the preferred (red) versus the antipreferred (green) direction of oculomotor neurons. First, during uninformed choices (short rPTs), the motor plan into the movement field should demonstrate a strong advantage shortly after the go signal (**Figure 4B**; arrows on left column). This preference should be evident before the cue is even presented (**Figure 4B**; middle column), and corresponds to a heavily biased motor competition that is decided well in advance of saccade onset (**Figure 4B**; right column). Second, during informed choices (long rPTs), the two motor plans should start building up more slowly and without a strong bias (**Figure 4C**; arrow on left column). In fact, in this case the expectation is somewhat counterintuitive: during the prolonged period of motor ambivalence, the motor plan in the direction of the target should, on average, lag behind the plan favoring the distracter (red traces below green), but ultimately the conflict must

FIGURE 4 | Comparison between model and FEF neuronal responses.
The accelerated race-to-threshold model was simulated with parameters that fitted the behavioral data of Money S, as in **Figure 3**, and the simulated activity averaged across trials was compared to that recorded from motor-related neurons in FEF. **(A)** Simulated tachometric curve. Gray shades indicate ranges used to sort the simulated trials into short- (left shade, guesses) and long-rPT (right shade, informed choices) groups. **(B,C)** Average model responses for short- **(B)** and long-rPT **(C)** trials aligned either on the go signal (left column), the cue (middle column), or saccade onset (right column).

All data are from correct responses. Separate averages were calculated for choices in the preferred (red traces) and antipreferred (green traces) direction of the model neurons. **(D–F)** As in **A–C**, but for 45 FEF neurons (motor and visuomotor) that differentiated significantly between movements into and away from the movement field before the saccade. Shaded areas indicate ±1 SE across neurons. In all plots, the y axis corresponds to normalized firing rate. Blue arrows mark key differences in evoked activity during guesses versus informed choices. See Costello et al. (2013) for details about the experimental data and modeling methods.

be resolved in favor of the correct choice. The reason for this effect is that, as discussed earlier, correct choices with long rPTs often correspond to trials in which the target-related motor plan is initially weak (**Figure 3B**), so a similar pattern emerges when averaging over multiple trials (**Figure 4C**).

The mean evoked responses of FEF neurons (motor and visuomotor) with significant movement-related activity were generally in excellent agreement with the expectations based on the model (Stanford et al., 2010; Costello et al., 2013). In particular, during informed choices, there was, indeed, a prolonged period of motor conflict during which the plan in favor of the distracter showed a slight initial advantage (**Figure 4F**), whereas no ambiguity was seen during correct guesses (**Figure 4E**). Observed differences between correct and incorrect responses were also in agreement with the model (Costello et al., 2013). Finally, to compare the model and recorded responses quantitatively, mean traces were calculated and analyzed as continuous functions of rPT via a sliding window, and the ensuing results led to two additional conclusions: (1) that the motor plans favoring the target and the distracter do accelerate and decelerate, respectively, and (2) that the acceleration and deceleration vary as functions of cue viewing time (rPT) as expected given the center point of the tachometric curve (Stanford et al., 2010; Costello et al., 2013).

These results are extremely important because they support the two fundamental elements of the accelerated race-to-threshold model. First, that in the compelled-saccade task, ongoing motor plans are modulated by perceptual information if and when that information becomes available

to the motor circuitry, but a motor choice is made either with (informed) or without it (uninformed). And second, that in spite of a profound impact on behavioral performance, the effect of perception on neural activity is rather subtle, particularly for eye movements into the receptive field of the neurons, because acceleration manifests as a slight difference in the curvature of the motor plan as it rises to threshold (**Figures 4E,F**, right column; compare red traces). Note that it did not have to be this way, as the psychophysical data alone can be replicated very accurately by a model based on completely different assumptions and dynamics (Salinas et al., 2010).

4. A TRADE-OFF DRIVEN BY MOTIVATIONAL BIAS

The tachometric curve is highly sensitive to task manipulations (Shankar et al., 2011; Hauser et al., 2013). Thus, many effects—for instance, subtle changes in performance due to perceptual learning (Shankar et al., 2011)—are clearly seen that would normally be impossible to resolve from the raw chronometric and psychometric data. In this section and the next we exploit this to discern, from the results of two experiments, the possible underlying mechanisms whereby speed and accuracy may be traded.

The first experiment consisted of a variant of the compelled-saccade task in which the monkey knew at the beginning of each trial whether a large or a small reward was at stake (all details are described by Shankar et al., 2011). The color of the target conveyed this information, and the association between color and reward amount was kept constant for blocks of 150–250 trials. So, during a block, correct movements to the red target would yield

a higher reward than correct movements to the green target, but the high- and low-reward colors were reversed in the next block. Here, because the color of the fixation spot indicates the color of the target, in each trial the subject knows how much reward can be gained, but that is all: given that target color and target location vary randomly and independently across trials, this knowledge provides no objective advantage, although it should affect the subject's motivation to perform the task correctly.

Comparison of responses in the high- and low-reward conditions revealed what appeared to be a classic trade-off between speed and accuracy. When working for a large reward (high incentive), on average the monkeys performed better (**Figures 5A,F**) and responded more slowly (**Figures 5B,G**) than when a small reward was at stake (low incentive). Both effects were relatively moderate in absolute terms, but the gain of the trade-off was high: an increase in performance of roughly 10% was accompanied by an increase in RT on the order of 10 ms—a change in RT that is quite small as a fraction of its mean value (~4%). So, based on these data alone, it would seem that the increase in performance incurred a very small cost in RT, and that the system is such that a small flexibility in RT affords a large benefit in success rate. Interpreted in terms of the two motor mechanisms discussed earlier, this would mean that a tiny increase in threshold (**Figures 2A–D**) or a tiny decrease in the mean build-up rates

(**Figures 2E–H**) would allow the sensory information to have a considerably stronger influence on the outcome of each trial.

However, analysis of the data in terms of processing time paints a much more nuanced picture in which both motor and perceptual mechanisms vary across conditions. In trials in which a high reward was at stake, the tachometric curves of monkeys G and R (**Figures 5D,I**) shifted to the left by about 30 and 20 ms, respectively, relative to when a low reward was at stake. This suggests that the decision-making process itself starts sooner or advances more rapidly when the incentive to perform accurately is high. By fitting the empirical tachometric curves to continuous functions (**Figures 5D,I**, thin black lines) and applying resampling techniques to estimate the likely error in these fits (**Figures 5E,J**), we found that the shifts were very highly significant (Shankar et al., 2011). A leftward shift, however, does not necessarily imply a higher percentage of correct responses, as illustrated earlier (**Figure 2J**), and would likely be accompanied by *lower* RTs too (**Figure 2K**), the opposite of the observed effect. So why the discrepancy?

Intuitively, the answer is that at least two mechanisms must be at work across conditions, given that the chronometric and tachometric curves are highly independent. A faster onset of the perceptual process could account for the leftward shift of the tachometric curve, a slow-down in motor activity could account
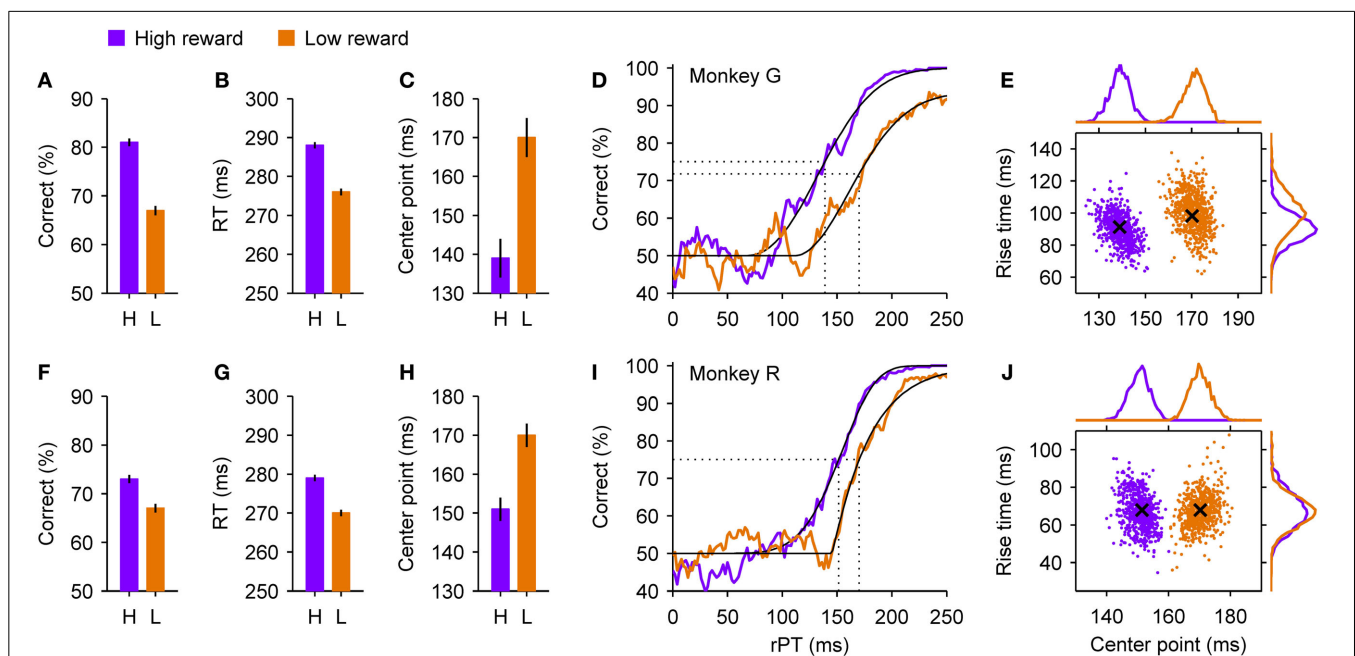
FIGURE 5 | Psychophysical performance of two monkeys in a motivational bias experiment. At the beginning of each trial of the compelled-saccade task, the monkey knew whether a correct response would result in a small or a large reward. The shown data were sorted *post hoc* according to the reward that was at stake in each trial, as indicated. **(A–C)** Summary statistics for monkey G. When a high reward (purple) rather than a low reward (orange) was at stake, the overall success rate **(A)** and mean RT **(B)** increased, and the tachometric curve shifted to the left **(C)**, indicating an earlier onset of the perceptual discrimination. Error bars indicate ±1 SE. **(D)** Tachometric curves from monkey G. Fitted Weibull functions (black curves) are shown together with the experimental data

(colored traces). A vertical dotted line marks the center point of each curve (indicated in **C**) derived from the fit, i.e., the time point at which the percent correct is halfway between chance and the maximum value. **(E)** Joint distributions of center points and rise times obtained from bootstrapping and re-fitting of monkey G's data, based on 2000 resamplings. The rise time is the time that it would take for the curve to go from 50% to 100% correct if its slope were always equal to the slope at the center point. Crosses mark the values of the original fits shown in **(D)**. Histograms at the top and on the right show the corresponding marginal distributions. **(F–J)** As in **(A–E)** but for monkey R. See Shankar et al. (2011) for details about experimental and statistical methods.

for the increase in performance, and the net effect on RT could be a combination of both.

The accelerated race-to-threshold model confirmed this intuition quantitatively. The model reproduced all the observed effects very accurately, and although this required modifying all of its parameters to various degrees across the two conditions, notably, these parameter differences were qualitatively the same for three monkeys. As discussed earlier, some of the parameters in the model relate fundamentally to perceptual processing and the tachometric curve (e.g., visual latency; magnitude of acceleration/deceleration), whereas others impact the initial motor plans only (e.g., mean and variance of the initial build-up rates). To tease apart their individual contributions to the observed experimental results (**Figure 5**), we first ran the model that fitted the low-reward condition and then compared the results to those of additional runs in which only selected parameters were modified as required to fit the high-reward condition.

The results were clear: although all parameters changed across conditions and had some impact, the experimental data could be largely explained by the two mechanisms illustrated in **Figures 2E–L** acting simultaneously. Specifically, according to the model, motor activity developed considerably more slowly during high- than low-reward trials. This slow-down accounted for virtually the full increase in the percentage of correct responses, and in the case of monkey G, if acting alone it would have yielded an increase in mean RT of $\sim$35 ms. This tendency, however, was largely offset by a smaller value of the visual latency parameter that determines when the go signal and the color cue start informing the motor circuit. This change explained most of the shift of the tachometric curve and, by itself, would have produced a drop in mean RT of $\sim$30 ms. So, motor and perceptual mechanisms exerted independent effects on accuracy but opposing effects on speed. As a consequence, the net change in RT produced by the model, with the contributions of all parameters taken into consideration, was relatively small, $\sim$10 ms, the same as found experimentally.

This simple computational dissection indicates (1) that multiple, distinct neural mechanisms are required to simultaneously explain all the experimental findings in the motivation experiment, and (2) that the coincident changes measured in speed (RT) and accuracy (percent correct) do not reflect a single, fundamental trade-off, but rather the combined action of cognitive factors on separate motor and perceptual processes.

Additional experimental observations supported these conclusions. For instance, note that the maximum percent correct reached by the tachometric curves (**Figures 5D,I**) was not the same in the two conditions. This means that, during trials in which a large reward was on offer, the monkeys rarely made a mistake when provided ample time to discriminate target from distracter, whereas in trials in which the potential reward was small, the monkeys made many more "careless" mistakes, errors that could not be attributed to insufficient viewing time. The frequency with which such errors occur is captured by one model parameter, the lapse rate, and when target and distracter are easily discriminable, as in the experiment, its effect is rather unique—it cannot be reproduced or even approximated by other

combinations of parameters—which suggests that it involves yet another mechanism that is distinct from those discussed above.

Therefore, to restate the main conclusion of this experiment, motivation affects choice behavior by simultaneously altering speed and accuracy, and there is good reason to believe that the cognitive signals that mediate these effects are diverse and exert at least partially independent control over motor and perceptual processes (see Discussion). This suggests that, in general, one-parameter descriptions of the speed-accuracy trade-off are likely to be oversimplifications, and should be interpreted with great caution.
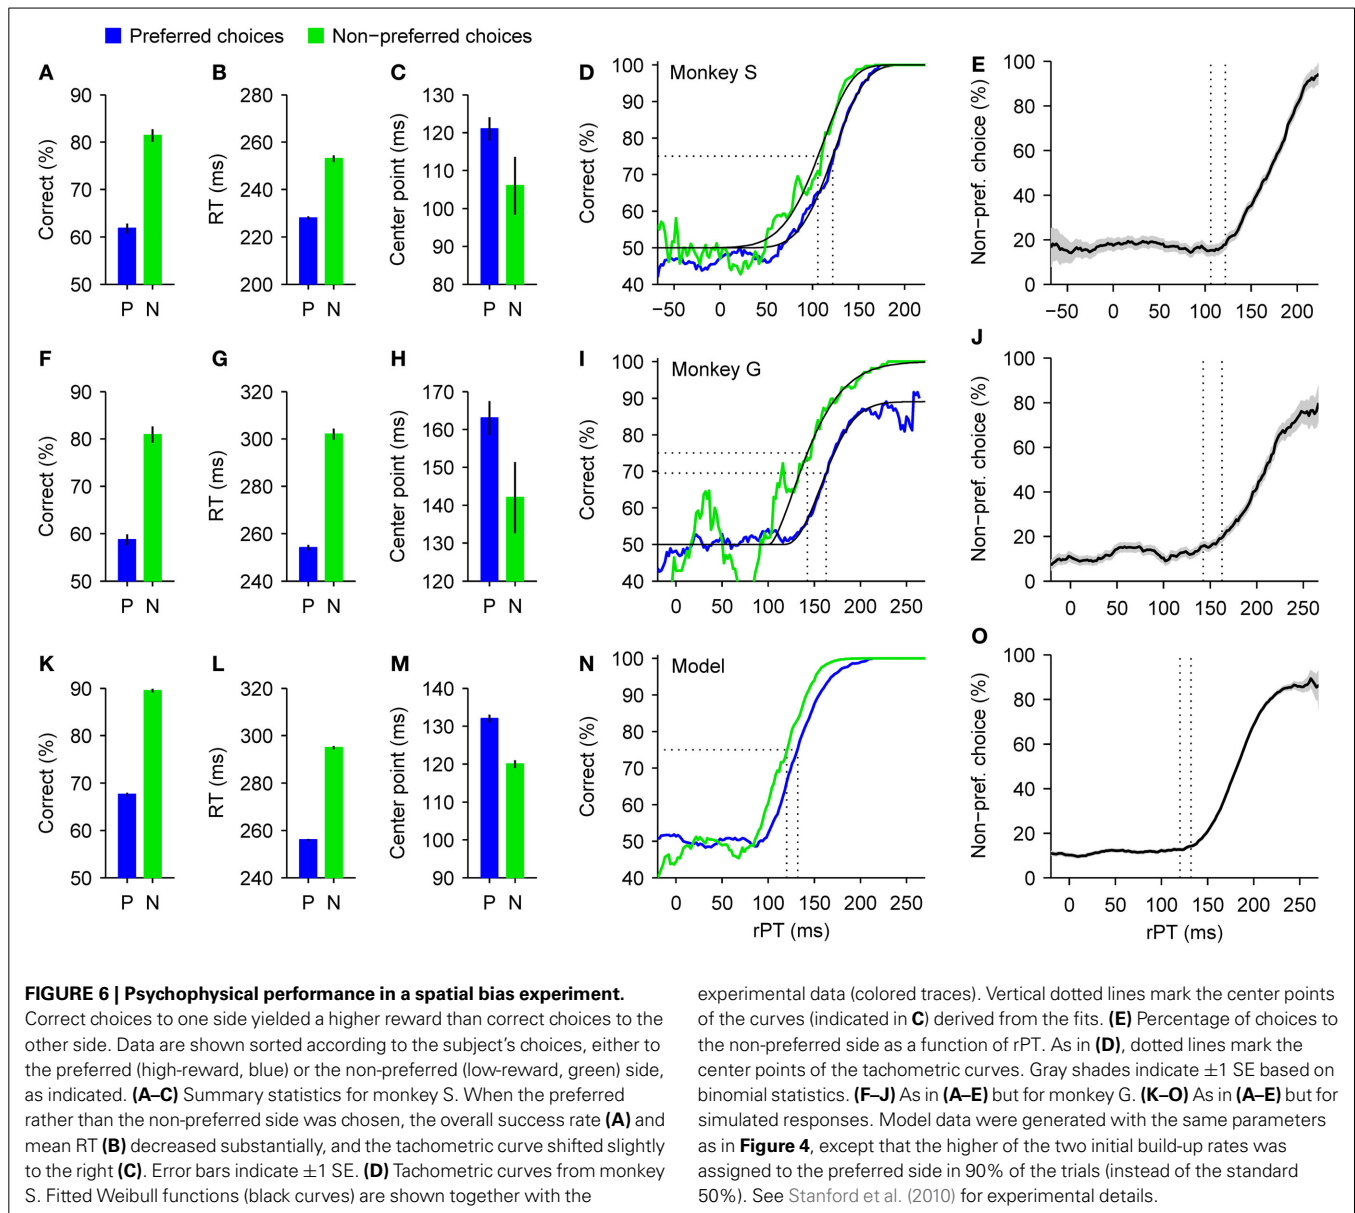
5. A TRADE-OFF DRIVEN BY SPATIAL BIAS

Next, we consider a second experiment with asymmetric rewards in the compelled-saccade task. It provides an interesting counterpoint to that in the previous section because it shows that the same motor and perceptual mechanisms may be engaged quite differently across tasks, giving rise to stronger or weaker trade-offs.

In this case, the monkeys received a large reward following correct saccades to one side and a small reward following correct saccades to the other (all details are described by Stanford et al., 2010). As a consequence, they developed a spatial bias, a strong tendency to respond more often to one side than the other. On average, the two animals that participated in this experiment chose the high-reward side about 76% of the time (but this number understates the strength of the preference; see below). The high-reward side, left or right, was kept constant during a block of 150–250 trials and was then switched. As always, target colors and locations were randomly interleaved. The collected data were then sorted according to the subject's choices; that is, trials were partitioned into two groups, those that resulted in movements in the preferred (high-reward) direction, and those that resulted in movements in the non-preferred (low-reward) direction. These two data subsets were then analyzed separately.

The behavior of the animals was strikingly different for the two types of choice. Responses in the preferred direction were much more prone to errors than those in the non-preferred direction (**Figures 6A,F**), and were also initiated much sooner (**Figures 6B,G**). In other words, the spatial bias induced a trade-off between speed and accuracy across conditions whereby an increase in performance of approximately 20% was accompanied by an increase in RT of 25 or 48 ms, depending on the subject. This behavior can be intuitively understood as follows: the high-reward side is chosen by default, so many choices toward that side are incorrect; in contrast, the low-reward side is chosen only if there is little uncertainty that the target is actually there, but this happens only when the red and green spots are discriminated accurately, i.e., when the rPT (and thus the RT) is long. This can be seen quantitatively by plotting the fraction of choices made to the low-reward side as a function of rPT (**Figures 6E,J**). The resulting choice curve rises quite sharply, so the monkey's preference is indeed dictated by the amount of cue viewing time. This curve also shows that, in the absence of sensory evidence (rPT $\lesssim$100 ms), the monkey's guess is to the high-reward side between 80% and 90% of the time.

FIGURE 6 | Psychophysical performance in a spatial bias experiment.
Correct choices to one side yielded a higher reward than correct choices to the other side. Data are shown sorted according to the subject's choices, either to the preferred (high-reward, blue) or the non-preferred (low-reward, green) side, as indicated. **(A–C)** Summary statistics for monkey S. When the preferred rather than the non-preferred side was chosen, the overall success rate **(A)** and mean RT **(B)** decreased substantially, and the tachometric curve shifted slightly to the right **(C)**. Error bars indicate ±1 SE. **(D)** Tachometric curves from monkey S. Fitted Weibull functions (black curves) are shown together with the

experimental data (colored traces). Vertical dotted lines mark the center points of the curves (indicated in **C**) derived from the fits. **(E)** Percentage of choices to the non-preferred side as a function of rPT. As in **(D)**, dotted lines mark the center points of the tachometric curves. Gray shades indicate ±1 SE based on binomial statistics. **(F–J)** As in **(A–E)** but for monkey G. **(K–O)** As in **(A–E)** but for simulated responses. Model data were generated with the same parameters as in **Figure 4**, except that the higher of the two initial build-up rates was assigned to the preferred side in 90% of the trials (instead of the standard 50%). See Stanford et al. (2010) for experimental details.

In general, the effects on speed and accuracy (**Figures 6A,B,F,G**) were considerably larger than in the motivational bias experiment (**Figures 5A,B,F,G**). Interestingly, however, although the main effect on the tachometric curve in this case was again a leftward shift congruent with the condition with higher overall performance (**Figures 6C,D,H,I**), the magnitude of the shift was smaller than in the motivational bias experiment (**Figures 5C,D,H,I**). This suggests that the perceptual process itself was affected less by the spatial bias than by the motivational bias, and therefore, that the observed trade-off in the former case may be accounted for almost entirely by an internal adjustment in motor planning alone. Indeed, that is precisely what a more thorough analysis of the data showed.

Again we used the accelerated race-to-threshold model to estimate the contributions of different mechanisms to the biases

found empirically. However, instead of discussing the full model fits to the psychophysical data, which involve numerous parameter differences across conditions, in this case we discuss a much simpler manipulation that illustrates the main result more plainly. It goes as follows. First we simulated N trials of the model with a fixed set of parameters. This set was exactly the same one used earlier to reproduce the behavior of monkey S (**Figure 3F**); everything was balanced, unbiased. Then we divided the simulated trials into two groups with approximately $N/2$ trials each: one group included all the trials in which the motor plan to the right had led initially, before the cue information was presented, and the other group included all other trials, in which the plan to the left had drawn a higher initial build-up rate. For this, trial outcome was irrelevant; only the initial build-up rates were considered. Next, we designated the right side as the preferred,

highly-rewarded side, and threw away 89% of the trials in the second group, in which the non-preferred (left) plan had led initially. Finally, we merged the remaining simulated trials back into a single data set, erased the information about which plan led initially, and analyzed them exactly as if they had been collected in the experiment. With this method, we produced a biased data set without changing the influence of the perceptual information or the dynamics of the motor plans at all; all we did was create a hypothetical subject, just like monkey S, that made 90% of its initial guesses toward a preferred side (combining $N/2$ preferred guesses with $0.11 \times N/2$ non-preferred guesses makes the former 90% of the total).

When the synthetic trials thus generated were sorted according to choice, as was done with the monkey data, the results qualitatively mimicked all the effects found experimentally: choices in the preferred direction were less accurate (**Figure 6K**) and faster (**Figure 6L**), the probability of making a non-preferred choice varied sharply as a function of rPT (**Figure 6O**), and the tachometric curves derived from preferred and non-preferred choices were slightly shifted relative to each other (**Figures 6M,N**). The underlying reason why such large differences emerged is that, by selecting trials based on the direction of the leading motor plan, the proportions of the five basic types of motor competition (**Figures 3A–E**) became drastically different for the two possible choices. Such proportions alone have an enormous impact on the average RT and success rate, even when the dynamics remain identical within each type of race. So, all the relevant differences between preferred and non-preferred choices—and in particular the bulk of the speed-accuracy trade-off—can be explained by a simple asymmetry in the way the motor plans are initially deployed.

This is not to say that other properties of the motor plans or of the perceptual process that informs them remained perfectly intact. In fact, there are hints that they did not. One is that the maximum percent correct was significantly different for the two tachometric curves of monkey G (**Figure 6I**), and another is that the shifts seen in the real data were larger than that in the simulation (**Figures 6D,I,N**). Additional adjustments to the parameters of the model would be required to account for these effects. However, these discrepancies are relatively small and do not affect the main conclusion, which is that in the spatial bias experiment the trade-off is larger than in the motivational bias experiment and depends predominantly on the way the motor plans for the two choices are deployed at the beginning of each trial.

Perhaps somewhat counterintuitively, these results also imply that average RTs may decrease in one condition relative to another without any explicit slow-down of the motor circuitry. If this circuitry naturally produces a wide distribution of RTs, then the apparent difference in response speed may result simply because one condition samples more fast and fewer slow trials than the other. In this sense, a change in response speed may not necessarily reflect a change in dynamics.

Taken together, the results reviewed in this and the previous section indicate that the individual contributions of motor and perceptual mechanisms to a given, experimentally observed trade-off may vary widely depending on the particular circumstances of an experiment.

6. BROADER IMPLICATIONS

Here we have reviewed behavioral, neurophysiological and modeling results in an urgent decision-making task in which independent, quantitative measures of motor and perceptual capacity (chronometric and tachometric curves) can be obtained. Based on this unique dissociation, we investigated how motor and perceptual mechanisms interact to determine a subject's response speed (RT) and accuracy (percentage of correct choices). In other words, we were able to decouple these quantities and investigate the potential sources of their trade-off.

Based on a combination of behavioral and neurophysiological constraints, the accelerated race-to-threshold model provides a parsimonious description of how perceptual information may resolve an ongoing motor selection process during relatively rapid choices. This heuristic model is key because it lets us evaluate the functional roles that meaningful neural elements or features play in the choice process. It shows, for instance, that the build-up rates with which competing motor plans are deployed initially, before perceptual information arrives, are absolutely critical in determining the fate of any given task trial (**Figure 4**, see arrows Salinas et al., 2010; Shankar et al., 2011). Likewise, the tachometric curve demonstrates that the response latencies—neuronal, not behavioral—to the go signal and the cue are much more flexible than one might have expected (**Figures 5D,I**), and the model serves to evaluate quantitatively the consequences of this (**Figures 2I–L**; see also Salinas and Stanford, 2013). Of course, other neural parameters may be important too; the point is simply that many specific properties of perceptual and motor-planning circuits may be quantitatively related to simultaneous changes in speed and accuracy.

When seen under the light of this framework, the experimental results obtained in the two biased versions of the compelled-saccade task lead to three conclusions: (1) that both motor and perceptual mechanisms may contribute to an observed trade-off, (2) that each of these mechanisms may weigh in more or less heavily, depending on the particulars of the task, and (3) that, as a consequence, small or large trade-offs may result from various combinations of motor and perceptual contributions.

This would also explain why, under certain circumstances, it is possible to observe a decrease in RT and/or an increase in accuracy with no apparent trade-off (Bendiksby and Platt, 2006; Takikawa et al., 2002). Other studies are also consistent with an intricate, fluid link between perception and action (Battaglia and Schrater, 2007; Cardoso-Leite and Gorea, 2010; Simoncini et al., 2012; see below).

6.1. LIFE WITHOUT THE TACHOMETRIC CURVE

It is interesting to ponder how the two bias experiments would be interpreted without the tachometric curve. In the case of the motivational bias, the trade-off would seem small (**Figures 5A,B,F,G**), and there would be no reason to think that the perceptual evaluation itself would or should change from one condition to another. The results could be explained as a small increase in a response criterion leading to slightly better

performance and slightly higher RTs. Instead, the tachometric curve reveals significant changes in perceptual performance (**Figures 5D,I**), and it is only because of the model that those changes can be reconciled with the relatively small observed trade-off, and that a rather substantial adjustment in motor planning can be inferred.

In the spatial bias experiment the speed-accuracy trade-off is large and evident (**Figures 6A,B,F,G**), but without the tachometric curve it again would be virtually impossible to ascertain whether or not changes in perception are involved—such changes are there (**Figures 6D,I**), but are noticeably smaller and less important in proportion to the magnitude of the trade-off in this case. Furthermore, the choice curve (**Figures 6E,J**) and the model (**Figures 6K–O**) provide a clear and parsimonious account of the results: the subjects' strategy is to almost always make an initial guess toward the preferred side, and override that initial plan only when the perceptual evidence against it arrives early enough and is strong enough. Without this insight, which depends critically on the distinction between RT and rPT, it would be very difficult to understand why, at a given gap, the subjects choose the low-reward side on some trials but not on others.

Interestingly, if the goal of the internal circuitry is to implement said strategy, then the observed trade-off may be a plain byproduct of the implementation, because the results can be largely accounted for simply by appropriately redistributing simulated trials across conditions, without altering any parameters or interactions in the model. In other words, the internal circuitry may not be directly attempting to find an optimal compromise in the exchange of RT for percent correct; rather, the observed exchange may be an inevitable consequence of a different trade-off, that between the possibility of a large reward versus the certainty of a small one.

6.2. UBIQUITY OF FAST DECISIONS

A few other tasks used in past studies compel participants to make a response before the correct answer is fully specified (Schouten and Bekker, 1967; Becker and Jürgens, 1979; Ghez et al., 1989; Hening et al., 1998; Chapman et al., 2010; Wood et al., 2011). In particular, the countermanding or stop-signal task is very similar to our compelled-saccade task, except for two main differences: it is a go/no-go task, and the relevant sensory evaluation is a detection rather than a discrimination—but notably, a tachometric curve can be constructed in this case too (Salinas and Stanford, 2013). Numerous experimental manipulations of the countermanding task have led to simultaneous changes in RT and percent correct (Cabel et al., 2000; Cavina-Pratesi et al., 2001; Ramautar et al., 2004; Emeric et al., 2007; Stevenson et al., 2009; Leotti and Wager, 2010), and modeling work indicates that, in different experiments, the observed trade-off may result either from adjustments in motor planning alone, in the perceptual detection process alone, or in both (Salinas and Stanford, 2013). The parallels with the experiments reviewed here are striking. For instance, variations in response latency associated with the detection of the saccadic target and the stop signal seem to be major determinants of perceptual performance. Overall, the spectrum of potential speed-accuracy trade-offs in the countermanding task is just as wide as illustrated here, if not wider, in terms of their

magnitude and variety of underlying neural mechanisms (Salinas and Stanford, 2013).

These results notwithstanding, how general are the conclusions presented here? Perhaps compelled-response tasks put subjects in an unnatural setting in which the mechanisms that control speed and accuracy are engaged in rather anomalous ways. To the contrary, we think that compelled tasks are good models for many real-life situations in which choices are made quickly (see Uchida et al., 2006).

For instance, eye movements (2–3/s) show similar distributions of fixation times and intersaccadic intervals under a wide variety of viewing conditions (Berg et al., 2009; Castelhano et al., 2009), which suggests that they are normally programmed continuously, without waiting for particular perceptual events to happen (McPeek et al., 2000; Hafed and Ignashchenkova, 2013). Furthermore, the ability to quickly modify ongoing motor plans is essential in situations that demand extreme performance, such as high-speed chases (Ghose et al., 2006, 2009). Competitive sports provide many familiar examples too. To return a tennis serve, hit a curveball, or stop a penalty, movements must be prepared early and the corresponding motor plans must take into account relevant perceptual information as soon as it becomes available (Abernethy, 1990; Land and McLeod, 2000; Yarrow et al., 2009). Interestingly, athletic skill may be thought of as an unusually weak speed-accuracy trade-off, in that a professional squash player can strike the ball both faster and more accurately than a beginner, and there is evidence that when the skill level achieved is exceptional, it is so in both perceptual and motor domains (Yarrow et al., 2009).

In this respect, note that the "urgency" of the compelled-saccade task refers to the perceptual analysis process rather than to motor execution. The saccadic RTs obtained in the task (**Figure 1C**) are well within the normal range for choice behaviors (e.g., DiCarlo and Maunsell, 2005; Berg et al., 2009); it is the color discrimination that is time-limited. For a participant, the decision is urgent in the same way as for a batter trying to hit a baseball: there is ample time to perform a required movement (a saccade or a swing), but very little time to make the relevant judgment (red/green or curveball/fastball) *and* inform the ongoing motor plan so that the movement is correct. In contrast, by specifically requiring that subjects remain still while the critical sensory information is displayed, the majority of laboratory tasks used to study perceptual decision making abolish this temporal conflict, both in fixed-duration and RT paradigms. This, however, makes it extremely difficult to determine when the perceptual discrimination finishes and when the motor plans start (e.g., Kiani et al., 2008; Port and Wurtz, 2009; Zariwala et al., 2013)—and thus to attribute a given change in mean RT to either of these events.

6.3. URGENT VERSUS NON-URGENT DECISIONS

The distinction between urgent and non-urgent tasks parallels a broad conceptual division in the ways in which sensory, cognitive and motor circuits may interact to carry out goal-directed actions or choices. In one scenario, they operate in a strictly serial fashion whereby perceptual analysis needs to reach a conclusion first, before the motor selection process can begin. In the alternative

scenario, the simultaneous activation of multiple uninformed motor plans marks the start of the choice process, and the competition is subsequently guided by perceptual information on the fly, if and whenever it becomes available. Each of these possibilities is likely to apply under certain circumstances. Cisek and Kalaska (2010), Cisek (2012) and Padoa-Schioppa (2011) discuss this issue at length. Here, we make two observations about this distinction in regard to our results.

First, we note that the former, serial account is incompatible with the compelled-saccade task (Salinas et al., 2010), but beyond that, one could argue that, for time scales below roughly 1000 ms, the idea of sequentially ordered perceptual and motor stages is inconsistent even with results from ostensibly serial decision tasks. This can be appreciated in two limit cases in which the trade-off between speed and accuracy essentially disappears. At one extreme, performance in many tasks does not benefit from prolonged deliberation times beyond 250–300 ms (Uchida et al., 2006), so that the optimal behavior is to respond rapidly (within <300 ms) regardless of difficulty. This is precisely what Mainen and colleagues found in an odor categorization task in rats (Zariwala et al., 2013). At the other extreme, note that the rise in choice-related firing activity is often interpreted as a pure accumulation of sensory evidence (Gold and Shadlen, 2001), but the notion that sensory evidence must achieve a critical threshold before the effector system is engaged is difficult to reconcile with choices made on the basis of little or no sensory evidence. Consider, for instance, the zero-coherence trials in the random-dot motion discrimination task (Shadlen and Newsome, 2001); what drives choice commitment when the sensory input to be integrated consists exclusively of noise? A choice under such condition is typically framed and modeled as the result of a lower threshold or collapsed decision bound wherein the evidence criterion is relaxed so that less (or no) sensory evidence is required to engage the motor circuitry (Ditterich, 2006; Beck et al., 2008; Hanks et al., 2011). But this is essentially a matter of interpretation: a collapsing decision bound is functionally equivalent to an increasing motor plan or urgency signal (Cisek et al., 2009; Thura et al., 2012). So, viewed from a different perspective, the "perceptual threshold" can be interpreted as the point in time at which a commitment to a motor choice curtails the evidence accumulation phase that had been informing the emerging motor plan to that point. Importantly, current neurophysiological evidence (Hanes and Schall, 1996; Heitz and Schall, 2012; see also Hayden et al., 2011) indicates that there is indeed such thing as threshold crossing, at least for saccadic choices, but it is a decidedly motor event. Furthermore, as the choice-related activity rises, its level relative to threshold is directly related to the degree of motor commitment (Gold and Shadlen, 2000).

Second, several studies within the latter camp, which considers the scheduling of motor actions to be independent of perceptual events, resonate particularly well with our approach. In particular, Goodale and colleagues studied the hand trajectories that result when humans perform a compelled-reaching task (Chapman et al., 2010; Wood et al., 2011). Participants were obliged to begin execution of a pointing movement toward one of various stimuli, but information identifying the true target was released only after the onset of the reach. The characteristic spatial patterns that

resulted indicated that, initially, multiple reaching plans toward various potential targets develop in parallel, with the initial movement direction reflecting an underlying vector-averaging operation; the final movement direction is disambiguated later, when the true target is revealed. Interestingly, they also found that stimuli of greater salience (through greater contrast or pixel density) confers greater initial weight to their corresponding motor plans, even when such saliency is unlikely to signal the true target location (Wood et al., 2011; see also Schütz et al., 2012). Notably, this pop-out effect went away when participants were allowed to briefly view the stimulus cue before initiating the reach. This means that motor plans associated with salient stimuli are activated more strongly, but unless the observer has reason to believe that a stimulus is important beyond its perceptual salience, this increased weight dwindles rapidly. So, perceptual information continuously modulates ongoing motor plans, likely via multiple pathways (e.g., bottom-up versus top-down).
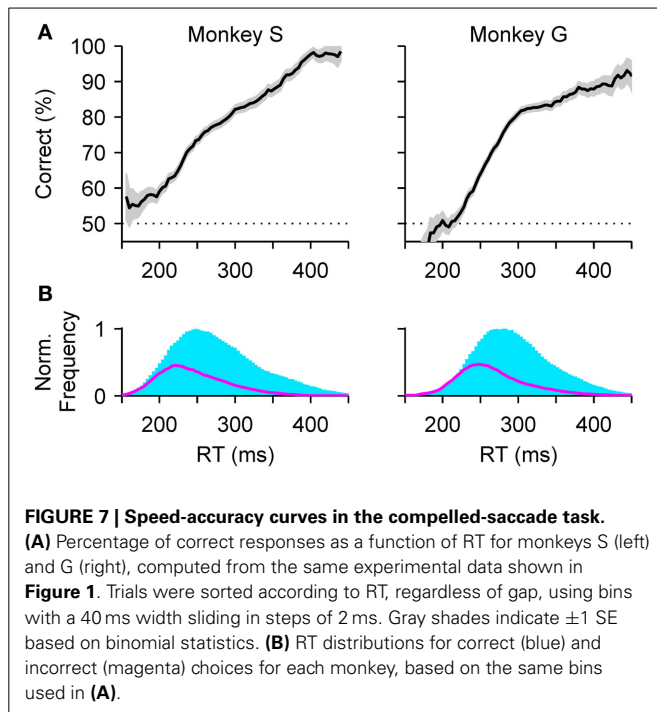
In agreement with the aforementioned findings in FEF (Stanford et al., 2010; Costello et al., 2013), this conclusion is highly consistent with analyses of single-neuron activity recorded in the parietal reach region of monkeys, which show (1) that competing motor plans are initially activated when multiple reach targets are presented and a choice needs to be made (Scherberger and Andersen, 2007), and (2) that the motor conflict is resolved either spontaneously or once the relevant cue information is provided (Klaes et al., 2011). Similar ideas have also been advocated by Cisek and colleagues based on recordings from premotor areas (Cisek and Kalaska, 2005; Pastor-Bernier and Cisek, 2011), giving rise to a powerful modeling framework, the "urgency-gating model" (Cisek, 2006; Cisek et al., 2009; Thura et al., 2012), that is similar in spirit to our accelerated race-to-threshold model (see Costello et al., 2013).

These findings demonstrate that, during rapid choices, perceptual and motor planning processes overlap extensively in time and are likely to contribute jointly to RT and accuracy under many circumstances. Their interaction is evident even in the absence of motor competition, when the upcoming movement is certain (Buonocore and McIntosh, 2008, 2012; Welchman et al., 2010; Bompas and Sumner, 2011). As a consequence, pinpointing the mechanisms that give rise to an observed trade-off is likely to be exceedingly difficult in general—unless additional experimental constraints independent of RT and percent correct are available.

6.4. BACK TO THE FUTURE: A HISTORICAL NOTE

The existence of a speed-accuracy trade-off has been acknowledged for many years (Woodworth, 1899; Hick, 1952), and it once was considered to have "great potential to advance all areas of cognitive psychology" (Wickelgren, 1977).

In 1977, Wickelgren passionately argued that generating speed-accuracy functions—the curves obtained by plotting the percentage of correct responses versus RT—would be vastly superior to simply evaluating RT and performance in single, independent experiments. He reasoned that a prototypical speed-accuracy curve would have three essential components: (1) an initial delay period during which performance would be at chance, (2) a ceiling value reached at long RTs beyond which performance could not increase further, and (3) a steep rise in performance around

FIGURE 7 | Speed-accuracy curves in the compelled-saccade task.
(A) Percentage of correct responses as a function of RT for monkeys S (left)
and G (right), computed from the same experimental data shown in
Figure 1. Trials were sorted according to RT, regardless of gap, using bins
with a 40 ms width sliding in steps of 2 ms. Gray shades indicate ±1 SE
based on binomial statistics. **(B)** RT distributions for correct (blue) and
incorrect (magenta) choices for each monkey, based on the same bins
used in **(A)**.

a short window of RTs. All three features would be informative and potentially interpretable in terms of separate cognitive mechanisms. Wickelgren (1977) further distinguished two ways to create such a curve, both potentially useful. One version used the "macro-trade-off," which is what commonly results when experimental manipulations are introduced (i.e., via deadlines, asymmetric payoffs, instructions emphasizing speed or accuracy, etc.); the other version used the "micro-trade-off," which is seen by the *post hoc* partitioning of RTs from a single experiment into small bands for analysis. Building on the work of Pachella (1974), Wickelgren suggested that internal variations in response criteria due to arousal, attention, and other covert factors creates variability within the RT distribution that macro-plots might not account for, but that would manifest in the micro-curves.

These ideas faded somewhat (but see, e.g., Giordano et al., 2009), most likely, we suspect, because the shapes of the speed-accuracy curves obtained experimentally were not stereotypical, as was hoped initially, nor consistent across experiments. For example, when the curves are generated from the data in the standard compelled-saccade task (**Figure 7A**), the resulting shapes are essentially meaningless. The framework presented here makes it easy, in retrospect, to see the reason for such failure: RT is not the same thing as processing time, and it is the relationship between performance and *processing time* that is stereotypical. That relationship—which is none other than the tachometric curve—describes precisely how much accuracy is gained for a given amount of time. It does this within a given experiment, as the micro-curve was supposed to do, and also decouples any true changes in perception from purely motor variations in RT, as may occur during a macro-trade-off. For the speed-accuracy curve to work as envisioned, the RT would need to correlate very tightly with rPT, but in general it does not, because it additionally

depends on many cognitive processes such as attention, memory, or motor planning, that contribute to its variance (**Figure 7B**).

Wickelgren (1977) recognized the enormous utility of a curve that would accurately reveal the dependence of performance on time. It could serve as a powerful tool for studying the dynamics of information processing across subjects, modalities, and task conditions, and by extension, for studying the neural mechanisms underlying fundamental cognitive functions. We submit that it is the tachometric curve, not the speed-accuracy curve, that fulfills this promise.

REFERENCES

Abernethy, B. (1990). Anticipation in squash: differences in advance cue utilization between expert and novice players. *J. Sports Sci.* 8, 17–34. doi: 10.1080/02640419008732128

Battaglia, P. W., and Schrater, P. R. (2007). Humans trade off viewing time and movement duration to improve visuomotor accuracy in a fast reaching task. *J. Neurosci.* 27, 6984–6994. doi: 10.1523/JNEUROSCI.1309-07.2007

Beck, J. M., Ma, W. J., Kiani, R., Hanks, T., Churchland, A. K., Roitman, J., et al. (2008). Probabilistic population codes for Bayesian decision making. *Neuron* 60, 1142–1152. doi: 10.1016/j.neuron.2008.09.021

Becker, W., and Jürgens, R. (1979). An analysis of the saccadic system by means of double step stimuli. *Vis. Res.* 19, 967–983. doi: 10.1016/0042-6989(79)90222-0

Bendiksby, M. S., and Platt, M. L. (2006). Neural correlates of reward and attention in macaque area LIP. *Neuropsychologia* 44, 2411–2420. doi: 10.1016/j.neuropsychologia.2006.04.011

Berg, D. J., Boehnke, S. E., Marino, R. A., Munoz, D. P., and Itti, L. (2009). Free viewing of dynamic stimuli by humans and monkeys. *J. Vis.* 9, 19.1–19.15. doi: 10.1167/9.5.19

Bergen, J. R., and Julesz, B. (1983). Parallel versus serial processing in rapid pattern discrimination. *Nature* 303, 696–698. doi: 10.1038/303696a0

Bodelón, C., Fallah, M., and Reynolds, J. H. (2007). Temporal resolution for the perception of features and conjunctions. *J. Neurosci.* 27, 725–730. doi: 10.1523/JNEUROSCI.3860-06.2007

Bogacz, R., Wagenmakers, E. J., Forstmann, B. U., and Nieuwenhuis, S. (2010). The neural basis of the speed-accuracy tradeoff. *Trends Neurosci.* 33, 10–16. doi: 10.1016/j.tins.2009.09.002

Bompas, A., and Sumner, P. (2011). Saccadic inhibition reveals the timing of automatic and voluntary signals in the human brain. *J. Neurosci.* 31, 12501–12512. doi: 10.1523/JNEUROSCI.2234-11.2011

Breitmeyer, B. G., and Ogmen, H. (2000). Recent models and findings in visual backward masking: a comparison, review, and update. *Percept. Psychophys.* 62, 1572–1595. doi: 10.3758/BF03212157

Breitmeyer, B. G., and Ogmen, H. (2006). *Visual Masking: Time Slices Through Conscious and Unconscious Vision.* Oxford: Oxford University Press. doi: 10.1093/acprof:oso/9780198530671.001.0001

Brunton, B. W., Botvinick, M. M., and Brody, C. D. (2013). Rats and humans can optimally accumulate evidence for decision-making. *Science* 340, 95–98. doi: 10.1126/science.1233912

Buonocore, A., and McIntosh, R. D. (2008). Saccadic inhibition underlies the remote distractor effect. *Exp. Brain Res.* 191, 117–122. doi: 10.1007/s00221-008-1558-7

Buonocore, A., and McIntosh, R. D. (2012). Modulation of saccadic inhibition by distractor size and location. *Vis. Res.* 69, 32–41. doi: 10.1016/j.visres.2012.07.010

Busse, L., Katzner, S., and Treue, S. (2008). Temporal dynamics of neuronal modulation during exogenous and endogenous shifts of visual attention in macaque area MT. *Proc. Natl. Acad. Sci. U.S.A.* 105, 16380–16385. doi: 10.1073/pnas.0707369105

Cabel, D. W., Armstrong, I. T., Reingold, E., and Munoz, D. P. (2000). Control of saccade initiation in a countermanding task using visual and auditory stop signals. *Exp. Brain Res.* 133, 431–441. doi: 10.1007/s002210000440

Cardoso-Leite, P., and Gorea, A. (2010). On the perceptual/motor dissociation: a review of concepts, theory, experimental paradigms and data interpretations. *Seeing Perceiving* 23, 89–151. doi: 10.1163/187847510X503588

Castelhano, M. S., Mack, M. L., and Henderson, J. M. (2009). Viewing task influences eye movement control during active scene perception. *J. Vis.* 9, 1–15. doi: 10.1167/9.3.6

Cavina-Pratesi, C., Bricolo, E., Prior, M., and Marzi, C. A. (2001). Redundancy gain in the stop-signal paradigm: implications for the locus of coactivation in simple reaction time. *J. Exp. Psychol. Hum. Percept. Perform.* 27, 932–941. doi: 10.1037/0096-1523.27.4.932

Chapman, C. S., Gallivan, J. P., Wood, D. K., Milne, J. L., Culham, J. C., and Goodale, M. A. (2010). Reaching for the unknown: multiple target encoding and real-time decision-making in a rapid reach task. *Cognition* 116, 168–176. doi: 10.1016/j.cognition.2010.04.008

Chittka, L., Skorupski, P., and Raine, N. E. (2009). Speed-accuracy trade-offs in animal decision making. *Trends Ecol. Evol.* 24, 400–407. doi: 10.1016/j.tree.2009.02.010

Cisek, P. (2006). Integrated neural processes for defining potential actions and deciding between them: a computational model. *J. Neurosci.* 26, 9761–9770. doi: 10.1523/JNEUROSCI.5605-05.2006

Cisek, P. (2012). Making decisions through a distributed consensus. *Curr. Opin. Neurobiol.* 22, 927–936. doi: 10.1016/j.conb.2012.05.007

Cisek, P., and Kalaska, J. F. (2005). Neural correlates of reaching decisions in dorsal premotor cortex: specification of multiple direction choices and final selection of action. *Neuron* 45, 801–814. doi: 10.1016/j.neuron.2005.01.027

Cisek, P., and Kalaska, J. F. (2010). Neural mechanisms for interacting with a world full of action choices. *Annu. Rev. Neurosci.* 33, 269–298. doi: 10.1146/annurev.neuro.051508.135409

Cisek, P., Puskas, G. A., and El-Murr, S. (2009). Decisions in changing conditions: the urgency-gating model. *J. Neurosci.* 29, 11560–11571. doi: 10.1523/JNEUROSCI.1844-09.2009

Cohen, J. Y., Heitz, R. P., Woodman, G. F., and Schall, J. D. (2009). Neural basis of the set-size effect in frontal eye field: timing of attention during visual search. *J. Neurophysiol.* 101, 1699–1704. doi: 10.1152/jn.00035.2009

Cook, E. P., and Maunsell, J. H. R. (2002). Dynamics of neuronal responses in macaque MT and VIP during motion detection. *Nat. Neurosci.* 10, 985–994. doi: 10.1038/nn924

Costello, M. G., Zhu, D., Salinas, E., and Stanford, T. R. (2013). Perceptual modulation of motor—but not visual—responses in the frontal eye field during an urgent-decision task. *J. Neurosci.* 33, 16394–16408. doi: 10.1523/JNEUROSCI.1899-13.2013

DiCarlo, J. J., and Maunsell, J. H. (2005). Using neuronal latency to determine sensory-motor processing pathways in reaction time tasks. *J. Neurophysiol.* 93, 2974–2986. doi: 10.1152/jn.00508.2004

Ditterich, J. (2006). Evidence for time-variant decision making. *Eur. J. Neurosci.* 24, 3628–3641. doi: 10.1111/j.1460-9568.2006.05221.x

Emeric, E. E., Brown, J. W., Boucher, L., Carpenter, R. H., Hanes, D. P., Harris, R., et al. (2007). Influence of history on saccade countermanding performance in humans and macaque monkeys. *Vis. Res.* 47, 35–49. doi: 10.1016/j.visres.2006.08.032

Ghez, C., Hening, W., and Favilla, M. (1989). Gradual specification of response amplitude in human tracking performance. *Brain Behav. Evol.* 33, 69–74. doi: 10.1159/000115902

Ghose, K., Horiuchi, T. K., Krishnaprasad, P. S., and Moss, C. F. (2006). Echolocating bats use a nearly time-optimal strategy to intercept prey. *PLoS Biol.* 4:e108. doi: 10.1371/journal.pbio.0040108

Ghose, K., Triblehorn, J. D., Bohn, K., Yager, D. D., and Moss, C. F. (2009). Behavioral responses of big brown bats to dives by praying mantises. *J. Exp. Biol.* 212, 693–703. doi: 10.1242/jeb.019380

Giordano, A. M., McElree, B., and Carrasco, M. (2009). On the automaticity and flexibility of covert attention: a speed-accuracy trade-off analysis. *J. Vis.* 9, 30.1–30.10. doi: 10.1167/9.3.30

Gold, J. I., and Shadlen, M. N. (2000). Representation of a perceptual decision in developing oculomotor commands. *Nature* 404, 390–394. doi: 10.1038/35006062

Gold, J. I., and Shadlen, M. N. (2001). Neural computations that underlie decisions about sensory stimuli. *Trends Cogn. Sci.* 5, 10–16. doi: 10.1016/S1364-6613(00)01567-9

Hafed, Z. M., and Ignashchenkova, A. (2013). On the dissociation between microsaccade rate and direction after peripheral cues: microsaccadic inhibition revisited. *J. Neurosci.* 33, 16220–16235. doi: 10.1523/JNEUROSCI.2240-13.2013

Hanes, D. P., and Schall, J. D. (1996). Neural control of voluntary movement inititation. *Science* 274, 427–430. doi: 10.1126/science.274.5286.427

Hanks, T. D., Mazurek, M. E., Kiani, R., Hopp, E., and Shadlen, M. N. (2011). Elapsed decision time affects the weighting of prior probability in a perceptual decision task. *J. Neurosci.* 31, 6339–6352. doi: 10.1523/JNEUROSCI.5613-10.2011

Hauser, C. K., Shankar, S., Salinas, E., and Stanford, T. R. (2013). Impact of multifeature integration on performance dissociated from reaction time effects via an urgent decision-making task. *Soc. Neurosci. Abstr.* Program No 826.24.

Hayden, B. Y., Pearson, J. M., and Platt, M. L. (2011). Neuronal basis of sequential foraging decisions in a patchy environment. *Nat. Neurosci.* 14, 933–939. doi: 10.1038/nn.2856

Heitz, R. P., and Schall, J. D. (2012). Neural mechanisms of speed-accuracy tradeoff. *Neuron* 76, 616–628. doi: 10.1016/j.neuron.2012.08.030

Hening, W., Favilla, M., and Ghez, C. (1998). Trajectory control in targeted force impulses. V. Gradual specification of response amplitude. *Exp. Brain Res.* 71, 116–128. doi: 10.1007/BF00247527

Hernández, A., Nácher, V., Luna, R., Zainos, A., Lemus, L., Alvarez, M., et al. (2010). Decoding a perceptual decision process across cortex. *Neuron* 66, 300–314. doi: 10.1016/j.neuron.2010.03.031

Hick, W. E. (1952). On the rate of gain of information. *Q. J. Exp. Psychol.* 4, 11–26. doi: 10.1080/17470215208416600

Horowitz, T. S., Holcombe, A. O., Wolfe, J. M., Arsenio, H. C., and DiMase, J. S. (2004). Attentional pursuit is faster than attentional saccade. *J. Vis.* 4, 585–603. doi: 10.1167/4.7.6

Horwitz, G. D., and Newsome, W. T. (1999). Separate signals for target selection and movement specification in the superior colliculus. *Science* 284, 1158–1161. doi: 10.1126/science.284.5417.1158

Kiani, R., Hanks, T. D., and Shadlen, M. N. (2008). Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. *J. Neurosci.* 28, 3017–3029. doi: 10.1523/JNEUROSCI.4761-07.2008

Klaes, C., Westendorff, S., Chakrabarti, S., and Gail, A. (2011). Choosing goals, not rules: deciding among rule-based action plans. *Neuron* 70, 536–548. doi: 10.1016/j.neuron.2011.02.053

Land, M. F., and McLeod, P. (2000). From eye movements to actions: how batsmen hit the ball. *Nat. Neurosci.* 3, 1340–1345. doi: 10.1038/81887

Lemus, L., Hernández, A., Luna, R., Zainos, A., Nácher, V., and Romo, R. (2007). Neural correlates of a postponed decision report. *Proc. Natl. Acad. Sci. U.S.A.* 104, 17174–17179. doi: 10.1073/pnas.0707961104

Leotti, L. A., and Wager, T. D. (2010). Motivational influences on response inhibition measures. *J. Exp. Psychol. Hum. Percept. Perform.* 36, 430–447. doi: 10.1037/a0016802

Mante, V., Sussillo, D., Shenoy, K. V., and Newsome, W. T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* 503, 78–84. doi: 10.1038/nature12742

McPeek, R. M., Skavenski, A. A., and Nakayama, K. (2000). Concurrent processing of saccades in visual search. *Vis. Res.* 40, 2499–2516. doi: 10.1016/S0042-6989(00)00102-4

Miller, P., and Katz, D. B. (2013). Accuracy and response-time distributions for decision-making: linear perfect integrators versus nonlinear attractor-based neural circuits. *J. Comput. Neurosci.* 35, 261–294. doi: 10.1007/s10827-013-0452-x

Monsell, S. (2003). Task switching. *Trends. Cogn. Sci.* 7, 134–140. doi: 10.1016/S1364-6613(03)00028-7

Pachella, R. G. (1974). "The interpretation of reaction time in information processing research," in *Human Processing Information*, ed B. Kantowitz (Potomac, MD: Lawrence Erlbaum), 41–82.

Padoa-Schioppa, C. (2011). Neurobiology of economic choice: a good-based model. *Annu. Rev. Neurosci.* 34, 333–359. doi: 10.1146/annurev-neuro-061010-113648

Palmer, J., Huk, A. C., and Shadlen, M. N. (2005). The effect of stimulus strength on the speed and accuracy of a perceptual decision. *J. Vis.* 5, 376–404. doi: 10.1167/5.5.1

Pastor-Bernier, A., and Cisek, P. (2011). Neural correlates of biased competition in premotor cortex. *J. Neurosci.* 31, 7083–7088. doi: 10.1523/JNEUROSCI.5681-10.2011

Port, N. L., and Wurtz, R. H. (2009). Target selection and saccade generation in monkey superior colliculus. *Exp. Brain Res.* 192, 465–477. doi: 10.1007/s00221-008-1609-0

Ramautar, J. R., Kok, A., and Ridderinkhof, K. R. (2004). Effects of stop-signal probability in the stop-signal paradigm: the N2/P3 complex further validated. *Brain Cogn.* 56, 234–252. doi: 10.1016/j.bandc.2004.07.002

Ratcliff, R., and Rouder, J. N. (2000). A diffusion model account of masking in two-choice letter identification. *J. Exp. Psychol. Hum. Percept. Perform.* 26, 127–140. doi: 10.1037/0096-1523.26.1.127

Ratcliff, R., and Smith, P. L. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychol. Rev.* 111, 333–367. doi: 10.1037/0033-295X.111.2.333

Reddi, B. A., and Carpenter, R. H. (2000). The influence of urgency on decision time. *Nat. Neurosci.* 3, 827–830. doi: 10.1038/77739

Reinagel, P. (2013a). Speed and accuracy of visual motion discrimination by rats. *PLoS ONE* 8:e68505. doi: 10.1371/journal.pone.0068505

Reinagel, P. (2013b). Speed and accuracy of visual image discrimination by rats. *Front. Neural Circ.* 7:200. doi: 10.3389/fncir.2013.00200

Reingold, E. M., and Stampe, D. M. (2002). Saccadic inhibition in voluntary and reflexive saccades. *J. Cogn. Neurosci.* 14, 371–388. doi: 10.1162/089892902317361903

Salinas, E., Shankar, S., Costello, M. G., Zhu, D., and Stanford, T. R. (2010). Waiting is the hardest part: comparison of two computational strategies for performing a compelled-response task. *Front. Comput. Neurosci.* 4:1–17. doi: 10.3389/fncom.2010.00153

Salinas, E., and Stanford, T. R. (2013). The countermanding task revisited: fast stimulus detection is a key determinant of psychophysical performance. *J. Neurosci.* 33, 5668–5685. doi: 10.1523/JNEUROSCI.3977-12.2013

Scherberger, H., and Andersen, R. A. (2007). Target selection signals for arm reaching in the posterior parietal cortex. *J. Neurosci.* 27, 2001–2012. doi: 10.1523/JNEUROSCI.4274-06.2007

Schouten, J. F., and Bekker, J. A. M. (1967). Reaction time and accuracy. *Acta Psychol.* 27, 143–153. doi: 10.1016/0001-6918(67)90054-6

Schütz, A. C., Trommershäuser, J., and Gegenfurtner, K. R. (2012). Dynamic integration of information about salience and value for saccadic eye movements. *Proc. Natl. Acad. Sci. U.S.A.* 109, 7547–7552. doi: 10.1073/pnas.1115638109

Shadlen, M. N., and Newsome, W. T. (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J. Neurophysiol.* 86, 1916–1936.

Shankar, S., Massoglia, D. P., Zhu, D., Costello, M. G., Stanford, T. R., and Salinas, E. (2011). Tracking the temporal evolution of a perceptual judgment using a compelled-response task. *J. Neurosci.* 31, 8406–8421. doi: 10.1523/JNEUROSCI.1419-11.2011

Simoncini, C., Perrinet, L. U., Montagnini, A., Mamassian, P., and Masson, G. S. (2012). More is not always better: adaptive gain control explains dissociation between perception and action. *Nat. Neurosci.* 15, 1596–1603. doi: 10.1038/nn.3229

Sommer, M. A., and Wurtz, R. H. (2001). Frontal eye field sends delay activity related to movement, memory, and vision to the superior colliculus. *J. Neurophysiol.* 85, 1673–1685.

Standage, D., You, H., Wang, D.-H., and Dorris, M. C. (2013). Trading speed and accuracy by coding time: a coupled-circuit cortical model. *PLoS Comput. Biol.* 9:e1003021. doi: 10.1371/journal.pcbi.1003021

Stanford, T. R., Shankar, S., Massoglia, D. P., Costello, M. G., and Salinas, E. (2010). Perceptual decision making in less than 30 milliseconds. *Nat. Neurosci.* 13, 379–385. doi: 10.1038/nn.2485

Sternberg, S. (1966). High speed scanning in human memory. *Science* 153, 652–654. doi: 10.1126/science.153.3736.652

Sternberg, S. (1969). The discovery of processing stages: extensions of Donders' method. *Acta Psychol.* 30, 276–315. doi: 10.1016/0001-6918(69)90055-9

Stevenson, S. A., Elsley, J. K., and Corneil, B. D. (2009). A "gap effect" on stop signal reaction times in a human saccadic countermanding task. *J. Neurophysiol.* 101, 580–590. doi: 10.1152/jn.90891.2008

Takikawa, Y., Kawagoe, R., Itoh, H., Nakahara, H., and Hikosaka, O. (2002). Modulation of saccadic eye movements by predicted reward outcome. *Exp. Brain Res.* 142, 284–291. doi: 10.1007/s00221-001-0928-1

Thura, D., Beauregard-Racine, J., Fradet, C. W., and Cisek, P. (2012). Decision making by urgency gating: theory and experimental support. *J. Neurophysiol.* 108, 2912–2930. doi: 10.1152/jn.01071.2011

Uchida, N., Kepecs, A., and Mainen, Z. F. (2006). Seeing at a glance, smelling in a whiff: rapid forms of perceptual decision making. *Nat. Rev. Neurosci.* 7, 485–491. doi: 10.1038/nrn1933

Welchman, A. E., Stanley, J., Schomers, M. R., Miall, R. C., and Bülthoff, H. H. (2010). The quick and the dead: when reaction beats intention. *Proc. R. Soc. B.* 277, 1667–1674. doi: 10.1098/rspb.2009.2123

Wickelgren, W. A. (1977). Speed-accuracy tradeoff and information processing dynamics. *Acta Psychol.* 41, 67–85. doi: 10.1016/0001-6918(77)90012-9

Wolfe, J. M. (1998). "Visual search," in *Attention*, ed H. Pashler (London, UK: University College London Press), 13–74.

Wood, D. K., Gallivan, J. P., Chapman, C. S., Milne, J. L., Culham, J. C., and Goodale, M. A. (2011). Visual salience dominates early visuomotor competition in reaching behavior. *J. Vis.* 11, 10. doi: 10.1167/11.10.16

Woodworth, R. S. (1899). The accuracy of voluntary movement. *Psychol. Rev. Monogr. Suppl.* 3, 1–119.

Yarrow, K., Brown, P., and Krakauer, J. W. (2009). Inside the brain of an elite athlete: the neural processes that support high achievement in sports. *Nat. Rev. Neurosci.* 10, 585–596. doi: 10.1038/nrn2672

Zariwala, H. A., Kepecs, A., Uchida, N., Hirokawa, J., and Mainen, Z. F. (2013). The limits of deliberation in a perceptual decision task. *Neuron* 78, 339–351. doi: 10.1016/j.neuron.2013.02.010

Dissociable mechanisms of speed-accuracy tradeoff during visual perceptual learning are revealed by a hierarchical drift-diffusion model

Jiaxiang Zhang[1]* and James B. Rowe[1,2,3]

[1] Cognition and Brain Sciences Unit, Medical Research Council, Cambridge, UK
[2] Department of Clinical Neurosciences, University of Cambridge, Cambridge, UK
[3] Behavioural and Clinical Neuroscience Institute, University of Cambridge, Cambridge, UK

Two phenomena are commonly observed in decision-making. First, there is a speed-accuracy tradeoff (SAT) such that decisions are slower and more accurate when instructions emphasize accuracy over speed, and *vice versa*. Second, decision performance improves with practice, as a task is learnt. The SAT and learning effects have been explained under a well-established evidence-accumulation framework for decision-making, which suggests that evidence supporting each choice is accumulated over time, and a decision is committed to when the accumulated evidence reaches a decision boundary. This framework suggests that changing the decision boundary creates the tradeoff between decision speed and accuracy, while increasing the rate of accumulation leads to more accurate and faster decisions after learning. However, recent studies challenged the view that SAT and learning are associated with changes in distinct, single decision parameters. Further, the influence of speed-accuracy instructions over the course of learning remains largely unknown. Here, we used a hierarchical drift-diffusion model to examine the SAT during learning of a coherent motion discrimination task across multiple training sessions, and a transfer test session. The influence of speed-accuracy instructions was robust over training and generalized across untrained stimulus features. Emphasizing decision accuracy rather than speed was associated with increased boundary separation, drift rate and non-decision time at the beginning of training. However, after training, an emphasis on decision accuracy was only associated with increased boundary separation. In addition, faster and more accurate decisions after learning were due to a gradual decrease in boundary separation and an increase in drift rate. The results suggest that speed-accuracy instructions and learning differentially shape decision-making processes at different time scales.

Keywords: speed-accuracy tradeoff, perceptual learning, drift-diffusion model, Bayesian parameter estimation, motion discrimination

INTRODUCTION

When making choices under time and resources constraints, more accurate decisions are often achievable at a cost of longer time, while faster responses are more error-prone. This phenomenon of *speed-accuracy tradeoff* (SAT) is ubiquitous across species and tasks (Schouten and Bekker, 1967; Wickelgren, 1977; Chittka et al., 2009), from collective foraging behavior in insects (Chittka et al., 2003; Franks et al., 2003; Marshall et al., 2006) to simple perceptual decisions in mammals (Uchida and Mainen, 2003; Heitz and Schall, 2012), and to complex strategic judgments in human (Beersma et al., 2003).

Most studies on the SAT compare behavioral performance under instructions of speed or accuracy emphasis. Humans can effectively trade accuracy for speed when instructed to respond as fast as possible, or *vice versa* when instructed to respond accurately. A change between speed and accuracy instructions can rapidly switch one's behavior between short blocks of trials (Ratcliff and Rouder, 1998; Mulder et al., 2013) or even between

two single trials (Forstmann et al., 2008; Ivanoff et al., 2008), suggesting that such instruction-induced SAT is embodied in the decision-making process. This is consistent with recent findings that the SAT in sensory-motor tasks is associated with neural activities in areas involved in perceptual decisions and cognitive control, such as (pre-) supplementary motor area, the frontal eye field, the anterior cingulate cortex, the striatum, and the dorsolateral prefrontal cortex (Forstmann et al., 2008; Ivanoff et al., 2008; Van Veen et al., 2008; Wylie et al., 2009; Blumen et al., 2011; Heitz and Schall, 2012).

While decisions can be rapidly adjusted in response to speed-accuracy instructions, they are also largely influenced by training and practice over a much longer time frame. It is well-established that prolonged practice gradually improves task performance, resulting in higher accuracy and faster responses (Logan, 1992; Heathcote et al., 2000). Similar to the SAT, the effect of *perceptual learning* is observed across species (Trobalon et al., 1992; Li et al., 2004) and sensory modalities (Fahle and Poggio, 2002), but

there are clear distinctions between the two. For simple visual perceptual decisions, performance improvement through perceptual learning is usually specific for the stimuli similar to those used in training, and do not fully generalize to other stimuli when the tasks are difficult (Ahissar and Hochstein, 1997; Green and Bavelier, 2003). Practice on more complex tasks, however, may improve performance in other tasks (Green and Bavelier, 2003). Unlike the SAT, the perceptual learning process can be automatic, without conscious insights of the task. For example, motion discrimination improves as participants were exposed to subliminal motion stimuli when performing an motion-irrelevant task (Watanabe et al., 2001). The specificity, generalizability, and implicit nature of perceptual learning indicate changes in early sensory processing as well as top–down influences during the learning process (Gilbert et al., 2001; Furmanski et al., 2004; Yang and Maunsell, 2004; Fahle, 2005; Bao et al., 2010; Zhang and Kourtzi, 2010; Zhang et al., 2010).

The cognitive processes underpin SAT and perceptual learning have previously been investigated by using the drift-diffusion model (DDM) (Stone, 1960; Ratcliff, 1978). The DDM belongs to a large family of decision-making models, namely sequential sampling models (Wald, 1947; Lehmann, 1959; Stone, 1960; Link, 1975; Link and Heath, 1975; Townsend and Ashby, 1983; Luce, 1986; Ratcliff and Smith, 2004; Smith and Ratcliff, 2004; Bogacz et al., 2006). These models assume that information supporting decisions is represented by a stream of noisy observations over time, and conceptualize decision-making as an information accumulation process: momentary evidence is accumulated over time, which reduce the noise in the evidence and hereby facilitate more accurate decisions. The sequential sampling models have been proven successful in providing a close fit to response accuracy and response time (RT) distributions (e.g., Ratcliff and Rouder, 1998), and are consistent with the identification of putative neural accumulators in the cortex from neurophysiological (Kim and Shadlen, 1999; Shadlen and Newsome, 2001; Roitman and Shadlen, 2002; Schall, 2002; Mazurek et al., 2003; Huk and Shadlen, 2005; Hanks et al., 2006; Gold and Shadlen, 2007) and neuroimage studies (Ploran et al., 2007; Heekeren et al., 2008; Ho et al., 2009; Kayser et al., 2010; Zhang et al., 2012).

The DDM is one of the most prominent sequential sampling models for two-choice decisions. It has been applied to a number of perceptual and cognitive tasks, including memory retrieval (Ratcliff, 1978), lexical decisions (Ratcliff et al., 2004; Wagenmakers et al., 2008), visual discrimination (Ratcliff, 2002; Palmer et al., 2005), and categorization (Nosofsky and Palmeri, 1997). The model implies a single accumulator integrating the sample evidence according to a stochastic diffusion process, until the accumulated evidence reaches one of the two decision boundaries, corresponding to the two choice alternatives. As such the model decomposes behavioral data into four parameters mapped on to latent psychological processes (**Figure 1**): boundary separation a for response caution, drift rate v for speed of accumulation, starting point z for *a priori* response bias, and non-decision time T_{er} for stimulus encoding and response execution latencies (Ratcliff and McKoon, 2008; Wagenmakers, 2009). Trial-to-trial variability in model parameters can be included to improve the

FIGURE 1 | Examples of trajectories of the drift-diffusion model. Two decision boundaries (0 and *a*) represent the "leftward" and "rightward" decisions in the motion discrimination task. The drift rate *v* represents mean sensory evidence per unit of time. The magnitude of *v* is determined by the quality of the evidence. A positive *v* (as shown in the figure) indicates that the upper boundary is the correct choice. The diffusion process starts at a starting point between the two boundaries (denoted as a proportion of *a* by *z*) until the accumulated evidence reaches one of two boundaries. If the correct boundary is hit (blue sample paths), the model makes a correct decision. Because of noise, the model may sometime hit the incorrect boundary (red sample path). The predicted response time (RT) is the sum of the duration of the diffusion process and the non-decision time T_{er}.

model fits to experimental data (Laming, 1968; Ratcliff, 1978; Ratcliff et al., 1999; Ratcliff and Tuerlinckx, 2002).

Behavioral changes in SAT and perceptual learning can be explained by different parameter changes in the DDM. The SAT can be simply quantified by the separation of the two decision boundaries. When response speed is emphasized, the distance between decision boundaries is decreased. This reduces the amount of accumulated evidence prior to a decision (i.e., faster RT) and increase the change of hitting the wrong decision boundary (i.e., lower accuracy). When accuracy is emphasized, the distance between decision boundaries is increased and the model predicts slower RT and higher accuracy, because more evidence need to be accumulated prior to a decision. It has indeed been shown that emphasizing decision speed or accuracy leads to changes in the boundary separation (Ratcliff and Rouder, 2000). A few recent studies have also applied the DDM to perceptual learning and identified two separate learning mechanisms (Dutilh et al., 2011, 2009; Petrov et al., 2011). First, training and practice are associated with an increase in the drift rate, leading to higher accuracy and faster RT (Dutilh et al., 2009; Wagenmakers, 2009). The drift rate change is consistent with most learning theories that the quality of sensory processing improves during training (Ahissar and Hochstein, 2004). Second, perceptual learning has been shown to decrease the non-decision time, which may be due to an increase in familiarity with the stimuli and task after training (Dutilh et al., 2011, 2009; Petrov et al., 2011).

However, two important issues remain unsolved. First, although previous research proposed that emphasizing speed or accuracy influence only the boundary separation (Ratcliff and Rouder, 1998; Wagenmakers et al., 2008), recent studies showed

that speed-accuracy instructions affect two other model parameters: drift rate (Vandekerckhove et al., 2011; Rae et al., in press) and non-decision time (Osman et al., 2000; Rinkenauer et al., 2004; Voss et al., 2004; Mulder et al., 2010, 2013). Therefore, it is necessary to examine whether other model parameters are indeed affected by speed emphasis or accuracy emphasis instructions.

Second, previous studies of the SAT and perceptual learning have been largely independent, partly because of the different time scale on which the two effects operate. However, since speed-accuracy instructions and learning can affect the same decision parameters, it is necessary to study these two different task conditions in a single experiment. Here we test the intriguing hypothesis that the SAT be efficiently manipulated over the course of learning a new task. One might establish a stable trade-off between speed and accuracy throughout learning, according to the task instructions. Alternatively, the effects of speed-accuracy instructions in a new task may be different from that in the same task after substantial practice.

The current study examined changes in decision performance and underlying cognitive mechanisms when SAT was manipulated throughout the course of learning. During multiple training sessions, participants learned to perform a coherent motion discrimination task under speed or accuracy emphasis (**Figure 2A**). Speed-accuracy instructions efficiently modulated participants' behavior between short blocks of trials across all sessions and training gradually improves performance specific to the trained directions. By fitting the DDM using Bayesian parameter estimation approach, we quantified the influence of speed-accuracy instructions and learning on the model parameters. Emphasizing decision accuracy rather than speed was related to increased boundary separation, drift rate and non-decision time at the beginning of training. In contrast, the emphasis on accuracy was only related to increased boundary separation after training. Furthermore, faster and more accurate decisions after learning are mainly due to a decrease in boundary separation and an increase in drift rate. Our results demonstrate that decision-making processes are differentially influenced by speed-accuracy instructions and training at different time scales and different stages of learning.

MATERIALS AND METHODS

PARTICIPANTS

Six adults (four females) between the age of 21–35 years (mean age, 25.50 years) participated in the experiment. All participants were right handed with normal hearing and normal or corrected-to-normal vision, and none reported a history of significant neurological or psychiatric illness. None had previous experience with the task. All participants signed a written informed consent before starting the experiment. The study was approved by the Cambridge Psychology Research Ethics Committee.

APPARATUS

The experiment was conducted in a darkened testing room. Each participant's head rested in a chinrest to stabilize the head position and control viewing distance. A computer (Dell Optiplex 745) controlled stimulus delivery and recorded behavioral responses. Visual stimulus was presented on a 21-inch CRT
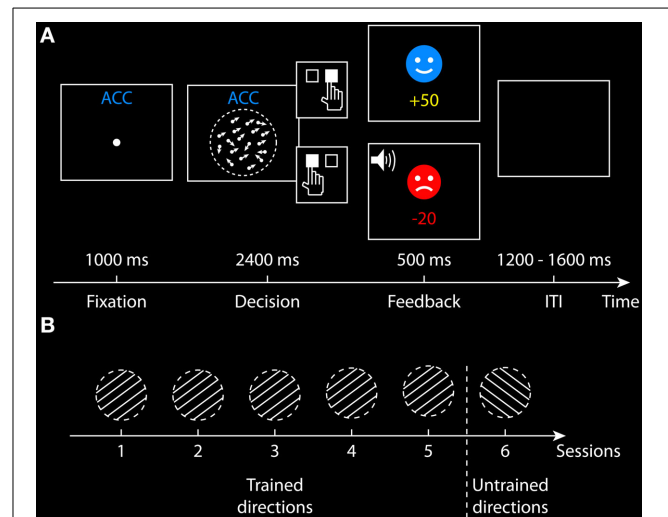
FIGURE 2 | Behavioral paradigm. (A) Structure of a single trial in the accuracy condition. A fixation point was presented for 1000 ms. The random dot kinematogram was then presented for a maximum of 2400 ms, during which participants made a binary decision on whether the coherent motion direction is leftward or rightward by pressing one of the two response buttons. For a correct response, a smiley face was presented for 500 ms and 50 points was credited. For an incorrect response, a sad face was presented and 20 points was lost, together with an auditory feedback. The payoff in the speed condition was slightly different (see section Task and Procedurefor more details). The intertial interval (ITI) was randomized between 1200 and 1600 ms. **(B)** Training procedure across six sessions. In the first five sessions, half of the participants trained at two directions (30 and 210°), and the other half trained at two different directions (150 and 330°). In the sixth session, all participants performed the task at two new directions that were not presented in their first five sessions (i.e., untrained directions).

monitor (Dell P1130) with a resolution of 1024 by 768 pixels and a refresh rate of 85 Hz, located 47.50 cm in front of the participants. Participants' responses were collected from a two-button response box. The experiment was written in Matlab 7.8 (The MathWorks, Natick, USA) and used the Psychophysics Toolbox 3 extensions (Brainard, 1997).

STIMULI

The stimuli were random-dot kinematograms displayed within a central invisible circular aperture (12° diameter) on a black background (100% contrast). Dot density was 16.53 dots per deg^2 per s and the minimum distance between any two dots in each frame was 0.48°. Each dot was white and subtended a visual angle of 0.12° at the screen center. The motion stimulus was formed by interleaving three uncorrelated sequences of dot positions at a rate of 85 frames/s, which was similar to those described elsewhere (Britten et al., 1993; Shadlen and Newsome, 2001; Roitman and Shadlen, 2002; Pilly and Seitz, 2009). To introduce coherent motion information, in each frame a fixed proportion (10.71%) of the dots was replotted at an appropriate spatial displacement in the direction of motion (10°/s velocity), relative to their positions three frames earlier, and the rest of the dots were replotted at random locations within the aperture. For example, three uncorrelated sets of dots were plotted in the first three frames.

A proportion of dots (i.e., the signal dots) in frame 1 moved in frame 4 with spatial displacements, and then a proportion of dots in frame 2 moved in frame 5, and so on. Signal dots that moved outside the aperture were wrapped around from the opposite direction of motion to conserve dot density and avoid attention cues along edges. The coherent dot motion in each trial was in one of four non-cardinal directions (30, 150, 210, and 330°).

TASK AND PROCEDURE

All participants completed six behavioral sessions conducted on different days. Participants performed a two-alternative forced-choice task in all sessions, deciding whether the coherent motion direction of the random-dot stimulus is leftward (toward 150 or 210°) or rightward (toward 30 or 330°) (**Figure 2A**). Participants responded by pressing the left button (for leftward decisions) or the right button (for rightward decisions) on the response box with their right index and middle fingers. In the first five sessions, the random-dot stimulus was always presented at two possible directions along a line (e.g., 30 and 210°), which referred to as the trained directions. In the sixth session, the stimulus was only presented at the other two new directions (e.g., 150 and 330°), which referred to as the untrained directions. One-half of the participants were trained at 30 and 210° directions and the other half of the participants were trained at the 150 and 330° directions in their first five sessions (**Figure 2B**).

Each experiment session comprised 672 trials, which were divided into 12 blocks of 56 trials. Each block had 50% leftwards motion trials and 50% rightwards motion trials at a randomized order. Participants took short breaks between blocks. The speed-accuracy manipulation was introduced at the block level: each session comprised of 6 accuracy blocks and 6 speed blocks. The first block of each session was always an accuracy block, and the order of the accuracy/speed instructions in the rest of the blocks were randomized across sessions and participants. At the beginning of an accuracy block, the text instruction "*Be accurate this time*" was presented on the screen in blue (RGB = 5, 137, 255), indicating that the participants should respond as accurate as possible. At the beginning of a speed block, the text instruction "*Be fast this time*" was presented in red (RGB = 255, 2, 2), indicating that the participants should respond as fast as possible. To ensure participants could easily identify the task instructions during the experiment, a text cue was presented at the top center of the screen throughout each block: "ACC" in blue (RGB = 5, 137, 255) for accuracy blocks, and "SPD" in red (RGB = 255, 2, 2) for speed blocks. Before the first and the 29th trials of each block, four parallel gray lines (RGB = 100, 100, 100, 0.05° thick, 4° apart) were presented within the circular aperture for 2000 ms, indicating the two possible motion directions in the current block (30 and 210°, or 150 and 330°). Before the first session, each participant was familiarized with the task during a short practice run comprising 16 trials for the accuracy condition and 16 trials for the speed condition, during which the proportion of coherently moving dots was set at a high level of 80%.

Each trial began with the presentation of a fixation point (0.12° diameter) at the center of the screen, which was illuminated for 1000 ms, followed by the random-dot stimulus onset. The stimulus was presented for a maximum of 2400 ms, during which the participants were instructed to perform the motion discrimination task under accuracy or speed emphasis. The random-dot stimulus disappeared as soon as a response was made, or the maximum duration was reached. The RT on each trial was measured from the stimulus onset until the participant made a response. Feedback was given 100 ms after the stimulus offset, followed by an intertrial interval randomized between 1200 and 1600 ms (**Figure 2A**).

To help the participants engage in the task and effectively adjust their decision processes to the speed-accuracy instructions, three types of feedback were given in the forms of texts, auditory beeps (tone with frequency of 600 Hz and duration of 0.15 s), and bonus points (see Petrov et al., 2011; Mulder et al., 2013 for similar multi-session designs using bonus points). If the participant failed to respond within 2200 ms or responded within 100 ms, a red warning message "Too slow!" or "Too fast!" was presented for a prolonged period (1500 ms) together with a beep, and the participant lost 50 points. In the accuracy condition, if the participant made a correct response, a smiley face was presented for 500 ms and 50 bonus points were credited. For an incorrect response, a sad face was presented for 500 ms and a beep where given, and the participant lost 20 points. In the speed condition, when the participants failed to respond within a time limit, a red text "Too slow!" and a beep was given and the participant lost 20 points. No further feedback about the accuracy of the participants' responses was given (i.e., they would also lose 20 points for a correct but overtime response). For each session and each participant, the time limit for the speed condition was defined as the 40% quantile of the RT distribution from the participant's first accuracy block in that specific session (see Mulder et al., 2013 for another way of defining participant-specific time limit). If participant's response was within the time limit, the same type of feedback was given for correct and incorrect responses as in the accuracy condition, but the participant would only lose 10 points for an incorrect response (i.e., fewer penalties for errors when instructing speeded responses). Participants started with zero bonus point at the beginning of each session and the cumulative bonus points were displayed at the bottom of the screen throughout the session.

DATA PROCESSING AND ANALYSIS

To eliminate fast guesses, trials with RT faster than 100 ms were removed from further analysis. Trials without a valid response within 2200 ms after the random-dot stimulus onset were also removed. The discarded trials only accounted for 0.3% of all trials. Decision accuracies (proportion of correct responses) and mean RTs from each session were entered into two separate repeated-measures ANOVAs for group analyses, with task conditions (accuracy and speed instructions) and sessions as factors.

Randomization tests were used to examine the statistical significance at the single-subject level (Edgington, 1995; Coolican, 2009). For example, to test whether a single participant had different RT between the speed and accuracy conditions, we first estimated the mean RT separately from each block in each session of the participant, resulting in RT samples from 36 speed blocks and 36 accuracy blocks. The observed RT difference between the two task conditions was quantified by the sample *t*-value

(mean difference between the data from the speed emphasis and accuracy emphasis conditions divided by the standard error of the difference). If the null hypothesis is true, there is no difference between task conditions, and the samples are exchangeable between conditions. We therefore generated a null distribution of the test statistic from 100,000 permutations, with the condition label randomly shuffled in each permutation. The permutation p-value was then calculated as the proportion of the randomized samples with the test statistic exceeded the observed test statistic. The same randomization procedure was applied to test the learning effects between sessions (**Table 1**).

HIERARCHICAL DRIFT-DIFFUSION MODEL

A full version of the DDM was fitted to each participant's accuracy and RT distribution. The model consists of seven parameters (Ratcliff and McKoon, 2008; Wagenmakers, 2009). (1) Boundary separation a ($a > 0$). (2) Mean drift rate v. (3) Mean response bias z as a proportion of boundary separation ($0 < z < 1$), which gives the starting point of the diffusion process relative to the two boundaries ($z * a$). Thus, values of $z > 0.5$ indicate an *a priori* bias toward the upper boundary (right button press) and values of $z < 0.5$ indicate a bias toward the lower boundary (left button press). (4) Mean non-decision time T_{er}. (5) Normally distributed trial-by-trial variability in drift rate s_v. (6) Uniformly distributed trial-by-trial variability in response bias s_z. (7) Uniformly distributed trial-by-trial variability in non-decision time s_t. The model predicts a binary choice as whether the upper or the lower boundary is reached, and predicts the observed RT as a sum of the decision time (i.e., the latency for the accumulator reaching one of the boundaries) and the non-decision time.

We used the hierarchical drift-diffusion model toolbox to fit the data (Wiecki et al., 2013). The hierarchical extension of the DDM assumes that the model parameters for individual participants are random samples drawn from group-level distributions, and uses Bayesian statistical methods to simultaneously estimate all parameters at both the group level and the individual-subject level (Vandekerckhove et al., 2011). The Bayesian approach for parameter estimation has two advantages. First, the Bayesian approach is more robust in recovering model parameters when less data is available (Matzke et al., 2013; Wiecki et al., 2013). Second, Bayesian estimation generates joint posterior distributions of all model parameters, given the observed experimental data. The posterior parameter distribution provides not only a point estimate, but also uncertainty of the estimate, and can be straightforwardly applied for Bayesian inference (Gelman et al., 2004). For example, let $P_{\text{Post}|\text{Data}}(a_{\text{accuracy}})$ and $P_{\text{Post}|\text{Data}}(a_{\text{speed}})$ be the marginal posteriors for the boundary separation from the accuracy and speed conditions. To test whether the boundary separation in the accuracy condition is larger than that in the speed condition, we can directly calculate the probability that the difference between the two parameters is larger than zero $P_{\text{Post}|\text{Data}}(a_{\text{accuracy}} - a_{\text{speed}} > 0)$ from the posterior distributions, and a high probability indicates strong evidence in favor of the testing hypothesis.

Performance differences between speed-accuracy conditions and between sessions suggest changes in one or more model parameters across task conditions and sessions. We therefore examined seven variants of the DDM with different parameter constrains between the two task conditions. The seven models differed on whether the boundary separation a, the drift rate v, the non-decision time T_{er}, or a combination of the three parameters varied between the accuracy and speed conditions (**Figure 4**). In all the models, the four key parameters (a, v, T_{er}, and z) were allowed to vary between sessions and were estimated at both individual-subject level and group level. The trial-by-trial variability parameters (s_v, s_t, and s_z) were shared between sessions and were estimated only at the group level, because it has been shown that the DDM with variability parameters fixed across multiple sessions provided a better explanation of the data (Liu and Watanabe, 2012). Similar to previous studies, the response bias parameter was set to vary between sessions but was invariant between task conditions (Mulder et al., 2013).

For each model, we generated 15,000 samples from the joint posterior distribution of all model parameters by using Markov chain Monte Carlo methods (Gamerman and Lopes, 2006) and discarded the first 5000 samples as burn-in (see Wiecki et al., 2013 for a more detailed description of the procedure). The convergence of the Markov chains were assessed using Geweke statistic (Gelman and Rubin, 1992). Parameter estimates in all models were converged after 15,000 samples.

RESULTS

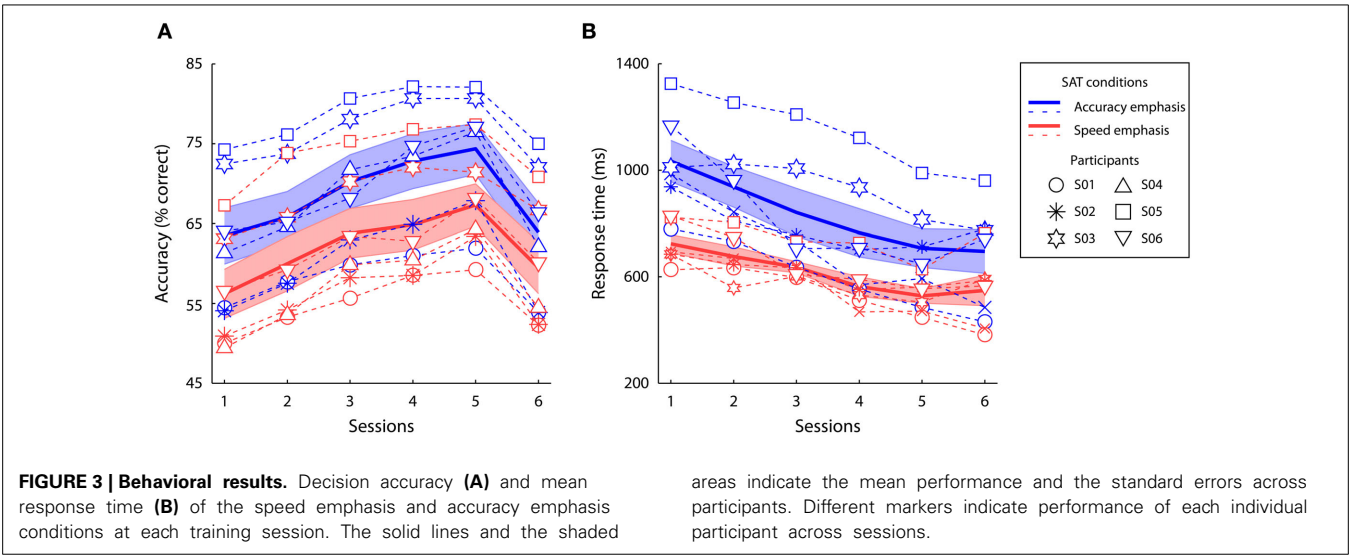### SPEED-ACCURACY TRADEOFF AND LEARNING EFFECTS ON BEHAVIORAL PERFORMANCE

Participants' performance under the accuracy and speed conditions was quantified by the mean decision accuracy and mean RT in each session (**Figure 3**). A two-way repeated-measures ANOVA showed a significant main effect of speed-accuracy instructions [accuracy: $F_{(1, 5)} = 27.57$, $p < 0.01$, partial $\eta^2 = 0.85$; RT: $F_{(1, 5)} = 17.56$, $p < 0.01$, partial $\eta^2 = 0.78$], a significant main effect of session [accuracy: $F_{(5, 25)} = 67.48$, $p < 0.0001$, partial $\eta^2 = 0.93$; RT: $F_{(5, 25)} = 22.82$, $p < 0.0001$, partial $\eta^2 = 0.82$], and a significant interaction between speed-accuracy manipulation and session [accuracy: $F_{(5, 25)} = 4.78$, $p < 0.01$, partial $\eta^2 = 0.49$; RT: $F_{(5, 5)} = 5.08$, $p < 0.01$, partial $\eta^2 = 0.50$]. In each session, the participants had higher accuracy ($p < 0.05$ in all sessions, Wilcoxon signed ranks test) and faster RT ($p < 0.05$ in all sessions, Wilcoxon Signed Ranks Test) in the accuracy condition than in the speed condition. Therefore, throughout training, the participants could effectively trade speed for accuracy (and *vice versa*) as instructed.

During the first five training sessions, behavior performance at the trained directions gradually improved, as shown by a significant linear increase of accuracy [$F_{(1, 5)} = 102.07$, $p < 0.0001$, partial $\eta^2 = 0.95$] and a linear decrease of RT [$F_{(1, 5)} = 53.37$, $p < 0.001$, partial $\eta^2 = 0.91$] over training. To examine whether the behavioral improvement at the trained directions can be generalized to another direction, we compared participants' performance between the 5th session (i.e., the last session at the trained directions) and the 6th session (i.e., untrained directions after training). The learning effect on decision accuracy was specific to individual participants' trained directions, as the accuracy was significantly lower at the untrained directions than the trained directions [$F_{(1, 5)} = 73.56$, $p < 0.0001$, partial $\eta^2 = $

Table 1 | Results of single-subject randomization tests.

Participant	SAT effects				Learning effects				Learning generalization			
	Accuracy (ACC > SPD)		RT (ACC > SPD)		Accuracy (session 1–5)		RT (session 1–5)		Accuracy (session 5–6)		RT (session 5–6)	
	t	p	t	p	t	p	t	p	t	p	t	p
S01	1.92	0.0600	2.05	0.0398	−2.82	0.0090	3.84	<0.0001	2.23	0.0348	3.14	0.0057
S02	1.85	0.0695	3.46	0.0007	−3.89	0.0012	4.57	<0.0001	3.90	0.0008	2.83	0.0102
S03	3.38	0.0015	9.73	<0.0001	−1.74	0.0963	2.45	0.0237	1.28	0.2121	−0.37	0.7150
S04	4.35	<0.0001	5.06	<0.0001	−3.44	0.0028	2.43	0.0194	2.31	0.0336	−0.53	0.6114
S05	2.40	0.0200	10.89	<0.0001	2.40	0.0258	2.73	0.0132	1.70	0.1093	−0.69	0.5077
S06	3.99	0.0002	4.37	<0.0001	−3.52	0.0018	6.50	<0.0001	2.78	0.0110	−1.20	0.2508

The SAT effects compared the behavioral performance between accuracy and speed conditions across all sessions. The learning effects compared the performance between session 1 and 5. The learning generalization effects compared the accuracy and RT between session 5 and 6 (i.e., performance at the untrained directions). Differences between conditions were quantified by sample t-values. Each p-value was obtained from 100,000 permutations of data samples (see section Data Processing and Analysis for details).

FIGURE 3 | Behavioral results. Decision accuracy **(A)** and mean response time **(B)** of the speed emphasis and accuracy emphasis conditions at each training session. The solid lines and the shaded areas indicate the mean performance and the standard errors across participants. Different markers indicate performance of each individual participant across sessions.
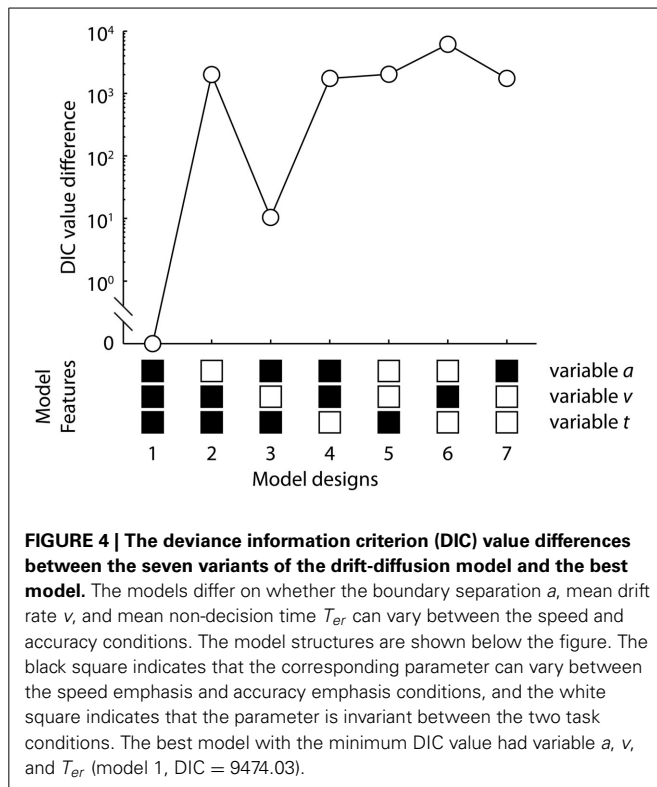
0.94]. Further, the learning effect on decision speed generalized across directions, as the RT at the untrained directions did not significantly differ to the trained directions after training [$F_{(1, 5)} = 0.03$, $p = 0.87$, partial $\eta^2 = 0.01$], but much faster than that in the first session [$F_{(1, 5)} = 35.94$, $p < 0.01$, partial $\eta^2 = 0.88$].

These results indicate strong group effects of speed-accuracy instructions and learning in perceptual decisions. Since the experiment collected substantial amount of data from individual participants, it is effective to further examine whether each individual's performance is consistent with the group effects above (Coolican, 2009; Barnett et al., 2012). We therefore conducted single-subject randomization tests (Bulté and Onghena, 2008, see section Data Processing and Analysis for details), estimating the main effects of task instructions across all sessions, the effect of learning, and generalization between trained and untrained directions for each participant (**Table 1**). Four participants had significantly higher decision accuracy and slower RT across sessions when instructed to trade speed for accuracy, with a trend

effect in the accuracy in two participants (S01 and S02 in **Table 1**). After training, significant improvements in both accuracy and RT were observed in five out of six participants, except one participant (S03) who had faster RT but no significant accuracy change after training. Four participants had significantly lower accuracies at the untrained directions than the trained directions after training. These analyses suggested that the single-subject data are largely consistent with the group inferences.

HIERARCHICAL DRIFT-DIFFUSION MODEL FOR SPEED-ACCURACY TRADEOFF AND LEARNING

To examine which model parameters account for the effects of speed-accuracy instructions during learning, we considered seven variants of the hierarchical DDM, varying systematically in constraints on whether three model parameters (a, v, and T_{er}) were invariant or varied across the task conditions. We used a Bayesian parameter estimation procedure to draw samples from the joint posterior distributions of all the parameters in the hierarchical

FIGURE 4 | The deviance information criterion (DIC) value differences between the seven variants of the drift-diffusion model and the best model. The models differ on whether the boundary separation a, mean drift rate v, and mean non-decision time T_{er} can vary between the speed and accuracy conditions. The model structures are shown below the figure. The black square indicates that the corresponding parameter can vary between the speed emphasis and accuracy emphasis conditions, and the white square indicates that the parameter is invariant between the two task conditions. The best model with the minimum DIC value had variable a, v, and T_{er} (model 1, DIC = 9474.03).

DDM (Vandekerckhove et al., 2011; Wiecki et al., 2013). The posterior samples represents parameter estimates and their uncertainties after having observed the data (i.e., response and RT distributions) (Gelman et al., 2004). Model fits were assessed by comparing each model's deviance information criterion (DIC) value (Spiegelhalter et al., 2002), which has a degree of penalty for additional free model parameters.

The best model (the one with the lowest DIC value) to describe the data across task conditions, sessions and participants allows the boundary separation a, mean drift rate v, and mean non-decision time T_{er} all to vary between speed and accuracy conditions (model 1 in **Figure 4**). The second best model had varied a and T_{er} but invariant v between SAT conditions, which had a DIC value 10.37 larger than the best model (model 3 in **Figure 4**). The model with only varied v but invariant a and T_{er} (model 6 in **Figure 4**) provided the worst fit among the seven models. Thus, changes in the mean drift rate are less likely to significantly account for the observed speed-accuracy effects. In later analysis, we focused on the best model with the minimum DIC value[1].

To evaluate the overall model fit, we generated posterior model predictions of the best model by simulate the same amount of predicted data as observed in the experiment using posterior

---

[1]Conventionally, a DIC difference of more than 10 indicates that the evidence supporting the best model is substantial (Burnham and Anderson, 2002). Because the difference of DIC values between the best and the second best model is close to this criterion, we repeated the same analysis on parameters estimates as in section hierarchical drift-diffusion model analyses for the second best model. The parameter changes between task conditions and sessions remain significant in the second best model.

estimates of the model parameters. There was very good agreement between the observed data and the model predictions across conditions and sessions (**Figure 5**).

HIERARCHICAL DRIFT-DIFFUSION MODEL ANALYSES

The hierarchical DDM incorporates parameters estimates (a, v, T_{er}, and z) at the individual-subject level and population estimates of these parameters at the group level (Wiecki et al., 2013). We used two complementary approaches to determine the effects of speed-accuracy instructions and learning on the model parameters. First, for each parameter at the individual-subject level, the mean of its posterior distribution was used as a point estimate for group analysis. Second, for each group-level parameter, the mean and the standard deviation of its posterior distribution were used to quantify group-level measures and estimation uncertainties (**Figure 6**). We also used the group-level posteriors to compare two parameters in Bayesian methodology (Lindley, 1965; Berger and Bayarri, 2004; Kruschke, 2010, see section Data Processing and Analysis for details). For simplicity, below we used p to refer to classical frequentist p-value from ANOVA, and $P_{P|D}$ to refer to the proportion of the posteriors supporting the testing hypothesis at the group level.
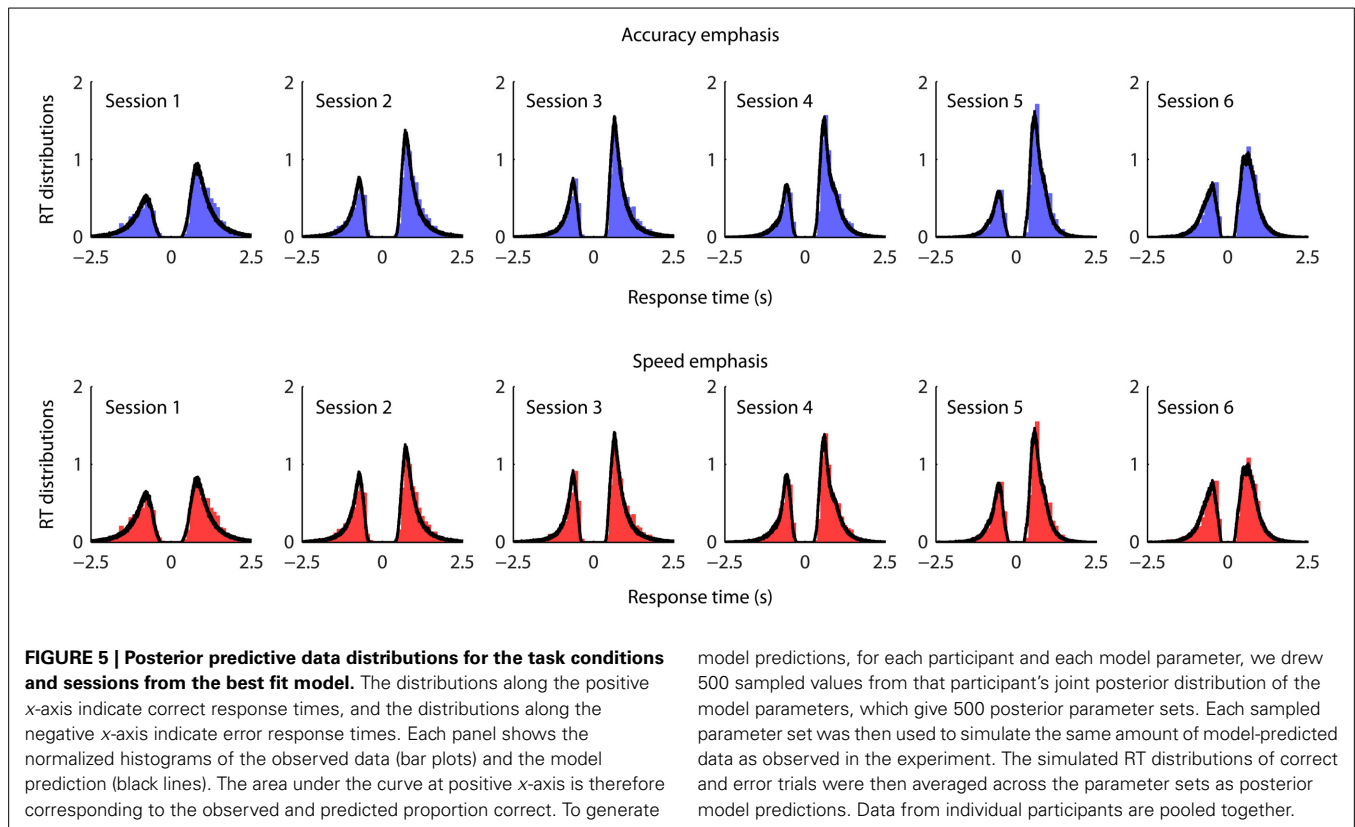
Boundary separation

Figure 6A showed the posterior mean and standard deviation of the boundary separation for each task condition and session. The boundary separation was significantly larger in the accuracy conditions than in the speed conditions [$F_{(1, 5)} = 16.21$, $p < 0.01$, partial $\eta^2 = 0.76$, $P_{P|D} = 0.95$]. *Post-hoc* tests showed significant differences between SAT conditions in all sessions ($p < 0.05$, Wilcoxon signed ranks test, $P_{P|D} > 0.93$). The interaction between the SAT condition and session is not significant [$F_{(5, 25)} = 0.34$, $p = 0.89$ partial $\eta^2 = 0.06$], suggesting similar extent of the speed-accuracy effect on boundary separation across sessions.

There is a main effect of session [$F_{(5, 25)} = 7.83$, $p < 0.001$, partial $\eta^2 = 0.61$]. Learning at the trained directions gradually decreases boundary separation, as supported by a linear effect in the first five sessions [$F_{(1, 5)} = 15.17$, $p < 0.05$, partial $\eta^2 = 0.75$]. Boundary separation at untrained directions after learning (session 6) is lower than that at the first session [$F_{(1, 5)} = 9.41$, $p < 0.05$, partial $\eta^2 = 0.65$, $P_{P|D} = 0.98$], but similar to the parameter at the trained directions after learning (session 5) [$F_{(1, 5)} = 1.68$, $p = 0.25$, partial $\eta^2 = 0.25$, $P_{P|D} = 0.37$]. Therefore, the learning effect on boundary separation generalized between trained and untrained directions.

Drift rate

The mean drift rate (**Figure 6B**) did not significantly differ between SAT conditions across all sessions [$F_{(1, 5)} = 2.93$, $p = 0.15$, partial $\eta^2 = 0.37$, $P_{P|D} = 0.76$], consistent with our model comparison result that the mean drift rate is not the main factor in explaining the effects of speed-accuracy instructions. Interestingly, there was a marginal interaction effect between task conditions and sessions before and after training (sessions 1 and 5) [$F_{(5, 25)} = 6.14$, $p = 0.06$, partial $\eta^2 = 0.55$], which is mainly driven by the higher mean drift rate in the accuracy condition

FIGURE 5 | Posterior predictive data distributions for the task conditions and sessions from the best fit model. The distributions along the positive *x*-axis indicate correct response times, and the distributions along the negative *x*-axis indicate error response times. Each panel shows the normalized histograms of the observed data (bar plots) and the model prediction (black lines). The area under the curve at positive *x*-axis is therefore corresponding to the observed and predicted proportion correct. To generate model predictions, for each participant and each model parameter, we drew 500 sampled values from that participant's joint posterior distribution of the model parameters, which give 500 posterior parameter sets. Each sampled parameter set was then used to simulate the same amount of model-predicted data as observed in the experiment. The simulated RT distributions of correct and error trials were then averaged across the parameter sets as posterior model predictions. Data from individual participants are pooled together.

than the speed condition in the first session ($p < 0.05$, Wilcoxon signed ranks test, $P_{P|D} = 0.86$).

The main effect of session on the mean drift rate was significant [$F_{(5, 25)} = 118.50$, $p < 0.00001$, partial $\eta^2 = 0.96$], with a linear increase in the first five sessions at the trained directions [$F_{(1, 5)} = 350.98$, $p < 0.00001$, partial η^2]. The drift rate at the untrained directions was lower than that at the trained directions after learning [$F_{(1, 5)} = 217.53$, $p < 0.00001$, partial $\eta^2 = 0.98$, $P_{P|D} \approx 1$], consistent with the observed data that improvements in accuracy did not transfer to the untrained directions after learning.

Non-decision time

The non-decision time (**Figure 6C**) was larger in the accurate condition than in the speed condition [$F_{(1, 5)} = 8.21$, $p < 0.05$, partial $\eta^2 = 0.62$, $P_{P|D} = 0.89$]. Pairwise comparison within each session indicates that the effects of speed-accuracy instructions were significant in the first three sessions ($p < 0.05$, Wilcoxon signed ranks test, $P_{P|D} > 0.91$) but not in the last three sessions ($p > 0.08$, Wilcoxon signed ranks test, $P_{P|D} < 0.80$). No significant effect of session was observed [$F_{(5, 25)} = 1.57$, $p = 0.21$, partial $\eta^2 = 0.24$], but there is an interaction between task conditions and sessions before and after training [$F_{(1, 5)} = 6.83$, $p < 0.05$, partial $\eta^2 = 0.58$]. These results suggest that the speed-accuracy instructions affect the non-decision time at a larger extent at the beginning of training.

Response bias

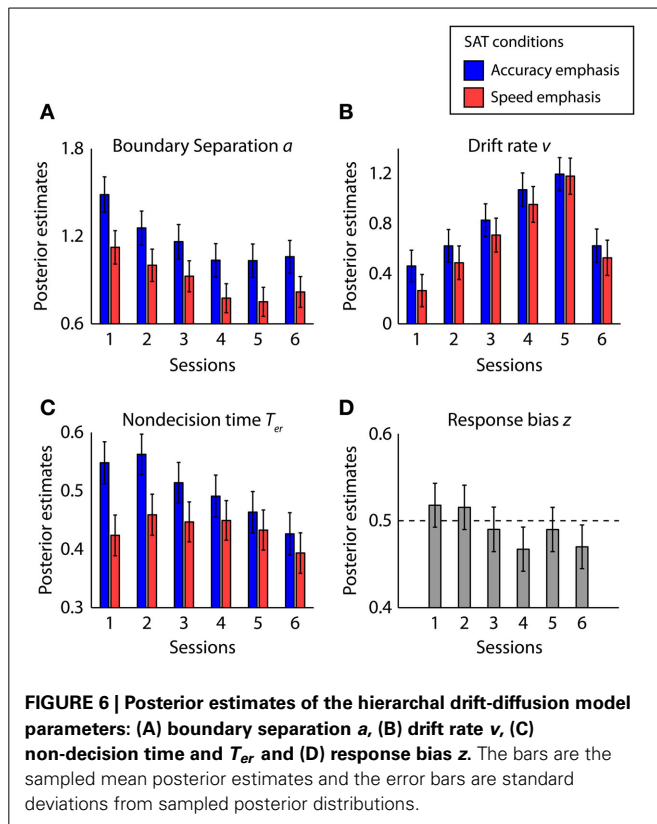The posterior estimates of the response bias were close to 0.5 in all sessions (**Figure 6D**) and a repeated-measures ANOVA showed no effect of sessions [$F_{(5, 25)} = 0.78$, $p = 0.58$, partial $\eta^2 = 0.13$]. Therefore, there was no significant bias toward any of the two responses or change of biases across sessions.

DISCUSSION

This study examined how the two widely observed phenomenon, SAT and perceptual learning, differentially shape decision-making processes over different timescales and stages of learning. Speed emphasis or accuracy emphasis, in a coherent motion discrimination task, rapidly modulated participant's behavior between short blocks of trials (fast and error-prone or slow and accurate). This tradeoff between speed and accuracy was consistent throughout training and generalized between trained and untrained directions. The model analysis suggested that accuracy emphasis, compared with speed emphasis, not only increases the total amount of evidence required to render a decision (i.e., boundary separation), but also increases the quality of the evidence being accumulated (i.e., drift rate) and the latencies on stimulus encoding and motor preparation (i.e., non-decision time). Importantly, the effect of speed-accuracy instructions on boundary separation was significant across multiple sessions, but the effect on drift rate and non-decision time was significant only at the beginning of training.

One common assumption often made is that speed-accuracy instruction influences only the boundary separation. This selective influence assumption was largely accommodated by the ability of the constrained DDM with only varied boundaries to adequately fit behavioral data under SAT manipulations (Ratcliff and Rouder, 1998; Wagenmakers et al., 2008). However, such an

FIGURE 6 | Posterior estimates of the hierarchal drift-diffusion model parameters: (A) boundary separation *a*, (B) drift rate *v*, (C) non-decision time and T_{er} and (D) response bias *z*. The bars are the sampled mean posterior estimates and the error bars are standard deviations from sampled posterior distributions.

approach cannot rule out possible influence of speed-accuracy instructions on other model parameters. Recent studies have considered more flexible models and identified the speed-accuracy effects on drift rate and non-decision time. By reanalyzing the data from Ratcliff and Rouder (1998), Vandekerckhove et al. (2011) suggested that the SAT is better described by changes in both drift rate and boundary separation than changes in boundary alone, with larger drift rate and boundary separation under accuracy emphasis. Similarly, Rae et al. (in press) reported that a constrained model with invariant drift rate between speed emphasis and accuracy emphasis conditions would underpredict the observed decision accuracy difference between the SAT conditions, which we also noticed from simulations of the inferior model (Model 3 in **Figure 4**). Rae et al. (in press) also reported larger drift rate change between speed-accuracy instructions in more difficult tasks than easier tasks. Interestingly, this is consistent with our result of significant drift rate change only in the first session, because the same task is relatively difficult for participants at the beginning of their training. Furthermore, studies using the DDM with variable non-decision time between different speed-accuracy conditions suggested decreased non-decision time when response speed is emphasized (Voss et al., 2004; Mulder et al., 2010, 2013). Therefore, emphasizing speed or accuracy affects multiple processes, not only the total amount of evidence needed for making a decision.

We found different effects of speed-accuracy instructions on the model parameters over the course of learning. For a difficult and unfamiliar task, emphasizing accuracy resulted in increased

boundary separation, drift rate, and non-decision time. Once the participants learned the task after substantial training, the effect of speed-accuracy instructions was evident only on boundary separation. These findings confirmed a substantial role of boundary separation in response to speed-accuracy instructions (Ratcliff and Rouder, 1998; Wagenmakers et al., 2008; Starns and Ratcliff, 2014) throughout learning and generalized between trained and untrained stimulus features. The influence of speed-accuracy instructions on the other two DDM parameters is not intuitive, because unlike boundary separation, changing drift rate or non-decision time itself cannot describe an inverse relationship between decision error and RT as observed in SAT: increasing drift rate results in lower decision errors but shorter RT, and increasing non-decision time results in longer RT but no change in accuracy (Ratcliff and McKoon, 2008).

Nevertheless, several possible hypotheses may explain why learning influences the drift rate and non-decision time in response to speed-accuracy instructions. First, Rae et al. (in press) proposed that the quality of information extracted from the environment improves over the course of a single decision, and the rates of the changes are identical in both speed and accuracy emphasis conditions. Since the RT is smaller when response speed is emphasized, the drift rate estimated from the speed condition is largely based on the quality of information extracted early after stimulus onset, which would be systematically lower than the information quality later in a trial (i.e., as in the accuracy condition). Second, drift rate has been linked to the allocation of attention on the task (Schmiedek et al., 2007). It is possible that speed-accuracy instructions have impacts on the balance of attentional resources allocated between the decision process and other cognitive processes. For example, speed emphasis may facilitate the monitoring of elapsed time within a trial, which limits the attentional resources for extracting information for decision-making. Third, Rinkenauer et al. (2004) examined the SAT effects on lateralized readiness potentials (Leuthold et al., 1996; Eimer, 1998; Masaki et al., 2004) and observed decreased intervals between response-locked lateralized readiness potential onset and motor responses under speed emphasis (see Osman et al., 2000 for similar results). Since lateralized readiness potential intervals refer to the duration of motor processes after a decision being made, the findings from the electrophysiological data posit a role of speed-accuracy instructions on both decision and post-decision processes. This further supports our findings of decreased non-decision time under speed emphasis, because response execution is often considered an important component described by non-decision time in the DDM (Ratcliff and McKoon, 2008). However, it is not immediately clear why the SAT effects on drift rate and non-decision making are more evident at the beginning of training. An active account is that participants change their decision strategy after they become proficient about the procedure and the task (e.g., Adini et al., 2004). In other words, participants may learn to integrate information across larger periods of the stimulus presentation, decreasing the time spent on processes outside of decision-making and hence improving performance. Or, in a more passive account, because the task becomes much easier after training, there is only a limited capacity to improve on the accuracy and RT, which in turn

limits the influence of speed-accuracy instructions on the model parameters other than boundary separation. Future investigations on how learning underpins the SAT at various task difficulty levels are necessary.

Our results demonstrated distinct perceptual learning mechanisms with different properties. As expected, training with feedback led to gradual improvements in decision accuracy and speed. The learning effect on accuracy was specific to the trained directions (Liu and Weinshall, 2000), but the improvement on RT partially generalized to untrained directions after training. Unlike most previous perceptual learning studies, which have focused only on decision accuracy but ignored decision speed (e.g., Fahle and Poggio, 2002; Dosher and Lu, 2007), we used the DDM to provide a mechanistic interpretation of both accuracy and speed improvements during learning (see Dutilh et al., 2009, 2011; Petrov et al., 2011; Liu and Watanabe, 2012 for similar approaches). Drift rate increased over training and the increase was specific to the trained directions, compatible with the theory that sensory processing is enhanced after learning (Karni and Sagi, 1991; Gilbert et al., 2001). This is also consistent with neurophysiological evidence that improved behavioral performance over training is accompanied by changes in sensory-driven responses of neurons in areas associated with perceptual decisions (Law and Gold, 2008). Boundary separation decreased over training and did not significantly differ between trained and untrained directions after training. Therefore, after substantial training of two motion directions, less accumulated evidence is required to discriminate coherent motion between two novel directions, even though the quality of extracted information from novel stimulus (e.g., drift rate for untrained directions) is lower. These findings further confirmed previous studies showing the learning effect on drift rate and boundary separation (Petrov et al., 2011; Liu and Watanabe, 2012).

The current study highlighted the benefits of using Bayesian methods to implement the DDM with the recently proposed hierarchical extension (Vandekerckhove et al., 2011; Wiecki et al., 2013). The hierarchical DDM is powerful in recovering model parameters with limited observed data (e.g., Jahfari et al., 2013). This feature is particularly important for the current study, because data from different training sessions need to be considered separately. One major advantage of using Bayesian methods for parameter estimation is the practicality of the obtained posterior parameter distributions. As we demonstrated in the current study, the posterior distributions can either be used to provide point estimates for classical frequentist inference, or can be directly used for Bayesian inference at both individual and group levels.

Two issues require further consideration. First, the drift-diffusion model is only an exemplar model of a large family of sequential sampling models (Ratcliff and Smith, 2004; Smith and Ratcliff, 2004; Bogacz et al., 2006; Zhang, 2012), and there are also simplified accumulator models omitting the noise in momentary evidence (Brown and Heathcote, 2005, 2008). These models mainly differ in how evidence supporting different alternatives is accumulated over time. It is of theoretical interest to explore whether our findings depend on the specific structure of the models we used. For example, one recent study showed similar influence of speed-accuracy instructions on model parameters in the DDM and in an accumulator model (Rae et al., in press). Second, we used a combination of bonuses and warning messages to help participants engage in the task, which is similar to early studies using a payoff matrix with criterion time (Fitts, 1966; Pachella and Pew, 1968) This design has been proven to be efficient in modulating behavior (Dutilh et al., 2009; Petrov et al., 2011). However, it is possible that participants would adopt a different decision strategy if the feedback or payoff is changed (e.g., the ratio of correct and error bonuses, see Simen et al., 2006, 2009; Bogacz et al., 2010; Balci et al., 2011).

In summary, we showed that the influence of speed-accuracy instructions cannot be attributed to a single change in decision boundary, but also relates to changes in other parameters that are relevant to the decision-making process and depends on the stage of learning. Future research on this topic should therefore take into account the complexity of individual's response to speed-accuracy instructions.

REFERENCES

Adini, Y., Wilkonsky, A., Haspel, R., Tsodyks, M., and Sagi, D. (2004). Perceptual learning in contrast discrimination: the effect of contrast uncertainty. *J. Vis.* 4, 993–1005. doi: 10.1167/4.12.2

Ahissar, M., and Hochstein, S. (1997). Task difficulty and the specificity of perceptual learning. *Nature* 387, 401–406. doi: 10.1038/387401a0

Ahissar, M., and Hochstein, S. (2004). The reverse hierarchy theory of visual perceptual learning. *Trends Cogn. Sci.* 8, 457–464. doi: 10.1016/j.tics.2004.08.011

Balci, F., Simen, P., Niyogi, R., Saxe, A., Hughes, J. A., Holmes, P., et al. (2011). Acquisition of decision making criteria: reward rate ultimately beats accuracy. *Atten. Percept. Psychophys.* 73, 640–657. doi: 10.3758/s13414-010-0049-7

Bao, M., Yang, L., Rios, C., He, B., and Engel, S. A. (2010). Perceptual learning increases the strength of the earliest signals in visual cortex. *J. Neurosci.* 30, 15080–15084. doi: 10.1523/JNEUROSCI.5703-09.2010

Barnett, S. D., Heinemann, A. W., Libin, A., Houts, A. C., Gassaway, J., Sen-Gupta, S., et al. (2012). Small N designs for rehabilitation research. *J. Rehabil. Res. Dev.* 49, 175–186. doi: 10.1682/JRRD.2010.12.0242

Beersma, B., Hollenbeck, J. R., Humphrey, S. E., Moon, H., Conlon, D. E., and Ilgen, D. R. (2003). Cooperation, competition, and team performance: toward a contingency approach. *Acad. Manag. J.* 46, 572–590. doi: 10.2307/30040650

Berger, J. O., and Bayarri, M. J. (2004). The Interplay of Bayesian and frequentist analysis. *Stat. Sci.* 19, 58–80. doi: 10.1214/088342304000000116

Blumen, H. M., Gazes, Y., Habeck, C., Kumar, A., Steffener, J., Rakitin, B. C., et al. (2011). Neural networks associated with the speed-accuracy tradeoff: evidence from the response signal method. *Behav. Brain Res.* 224, 397–402. doi: 10.1016/j.bbr.2011.06.004

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., and Cohen, J. D. (2006). The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol. Rev.* 113, 700–765. doi: 10.1037/0033-295X.113.4.700

Bogacz, R., Hu, P. T., Holmes, P. J., and Cohen, J. D. (2010). Do humans produce the speed-accuracy trade-off that maximizes reward rate? *Q. J. Exp. Psychol.* 63, 863–891. doi: 10.1080/17470210903091643

Brainard, D. H. (1997). The psychophysics toolbox. *Spat. Vis.* 10, 433–436. doi: 10.1163/156856897X00357

Britten, K. H., Shadlen, M. N., Newsome, W. T., and Movshon, J. A. (1993). Responses of neurons in macaque MT to stochastic motion signals. *Vis. Neurosci.* 10, 1157–1169. doi: 10.1017/S0952523800010269

Brown, S., and Heathcote, A. (2005). A ballistic model of choice response time. *Psychol. Rev.* 112, 117–128. doi: 10.1037/0033-295X.112.1.117

Brown, S. D., and Heathcote, A. (2008). The simplest complete model of choice response time: linear ballistic accumulation. *Cogn. Psychol.* 57, 153–178. doi: 10.1016/j.cogpsych.2007.12.002

Bulté, I., and Onghena, P. (2008). An R package for single-case randomization tests. *Behav. Res. Methods* 40, 467–478. doi: 10.3758/BRM.40.2.467

Burnham, K., and Anderson, D. (2002). *Model Selection and Multimodel Inference: a Practical Information-Theoretic Approach, 2nd Edn.* New York, NY: Springer-Verlag.

Chittka, L., Dyer, A. G., Bock, F., and Dornhaus, A. (2003). Psychophysics: bees trade off foraging speed for accuracy. *Nature* 424:388. doi: 10.1038/424388a

Chittka, L., Skorupski, P., and Raine, N. E. (2009). Speed-accuracy trade-offs in animal decision making. *Trends Ecol. Evol.* 24, 400–407. doi: 10.1016/j.tree.2009.02.010

Coolican, H. (2009). *Research Methods and Statistics in Psychology, 5th Edn.* London: Hodder Education.

Dosher, B. A., and Lu, Z.-L. (2007). The functional form of performance improvements in perceptual learning: learning rates and transfer. *Psychol. Sci.* 18, 531–539. doi: 10.1111/j.1467-9280.2007.01934.x

Dutilh, G., Krypotos, A.-M., and Wagenmakers, E.-J. (2011). Task-related versus stimulus-specific practice. *Exp. Psychol.* 58, 434–442. doi: 10.1027/1618-3169/a000111

Dutilh, G., Vandekerckhove, J., Tuerlinckx, F., and Wagenmakers, E.-J. (2009). A diffusion model decomposition of the practice effect. *Psychon. Bull. Rev.* 16, 1026–1036. doi: 10.3758/16.6.1026

Edgington, E. (1995). *Randomization Tests, 4th Edn.* New York, NY: Marcel Dekker.

Eimer, M. (1998). The lateralized readiness potential as an on-line measure of central response activation processes. *Behav. Res. Methods Instrum. Comput.* 30, 146–156. doi: 10.3758/BF03209424

Fahle, M. (2005). Perceptual learning: specificity versus generalization. *Curr. Opin. Neurobiol.* 15, 154–160. doi: 10.1016/j.conb.2005.03.010

Fahle, M., and Poggio, T. (2002). *Perceptual Learning.* Cambridge: MIT Press.

Fitts, P. M. (1966). Cognitive aspects of information processing: III. Set for speed versus accuracy. *J. Exp. Psychol.* 71, 849–857. doi: 10.1037/h0023232

Forstmann, B. U., Dutilh, G., Brown, S., Neumann, J., von Cramon, D. Y., Ridderinkhof, K. R., et al. (2008). Striatum and pre-SMA facilitate decision-making under time pressure. *Proc. Natl. Acad. Sci. U.S.A.* 105, 17538–17542. doi: 10.1073/pnas.0805903105

Franks, N. R., Dornhaus, A., Fitzsimmons, J. P., and Stevens, M. (2003). Speed versus accuracy in collective decision making. *Philos. Trans. R. Soc. B Biol. Sci.* 270, 2457–2463. doi: 10.1098/rspb.2003.2527

Furmanski, C. S., Schluppeck, D., and Engel, S. A. (2004). Learning strengthens the response of primary visual cortex to simple patterns. *Curr. Biol.* 14, 573–578. doi: 10.1016/j.cub.2004.03.032

Gamerman, D., and Lopes, H. F. (2006). *Markov Chain Monte Carlo: Stochastic Simulation for Bayesian Inference. 2nd Edn.* London: Taylor and Francis.

Gelman, A., Carlin, J., Stern, H., and Rubin, D. (2004). *Bayesian Data Analysis, 2nd Edn.* Boca Raton, FL: Chapman & Hall/CRC Press.

Gelman, A., and Rubin, D. B. (1992). Inference from iterative simulation using multiple sequences. *Stat. Sci.* 7, 457–472. doi: 10.1214/ss/1177011136

Gilbert, C. D., Sigman, M., and Crist, R. E. (2001). The neural basis of perceptual learning. *Neuron* 31, 681–697. doi: 10.1016/S0896-6273(01)00424-X

Gold, J. I., and Shadlen, M. N. (2007). The neural basis of decision making. *Annu. Rev. Neurosci.* 30, 535–574. doi: 10.1146/annurev.neuro.29.051605.113038

Green, C. S., and Bavelier, D. (2003). Action video game modifies visual selective attention. *Nature* 423, 534–537. doi: 10.1038/nature01647

Hanks, T. D., Ditterich, J., and Shadlen, M. N. (2006). Microstimulation of macaque area LIP affects decision-making in a motion discrimination task. *Nat. Neurosci.* 9, 682–689. doi: 10.1038/nn1683

Heathcote, A., Brown, S., and Mewhort, D. J. K. (2000). The power law repealed: the case for an exponential law of practice. *Psychon. Bull. Rev.* 7, 185–207. doi: 10.3758/BF03212979

Heekeren, H. R., Marrett, S., and Ungerleider, L. G. (2008). The neural systems that mediate human perceptual decision making. *Nat. Rev. Neurosci.* 9, 467–479. doi: 10.1038/nrn2374

Heitz, R. P., and Schall, J. D. (2012). Neural mechanisms of speed-accuracy tradeoff. *Neuron* 76, 616–628. doi: 10.1016/j.neuron.2012.08.030

Ho, T. C., Brown, S., and Serences, J. T. (2009). Domain general mechanisms of perceptual decision making in human cortex. *J. Neurosci.* 29, 8675–8687. doi: 10.1523/JNEUROSCI.5984-08.2009

Huk, A. C., and Shadlen, M. N. (2005). Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. *J. Neurosci.* 25, 10420–10436. doi: 10.1523/JNEUROSCI.4684-04.2005

Ivanoff, J., Branning, P., and Marois, R. (2008). fMRI evidence for a dual process account of the speed-accuracy tradeoff in decision-making. *PLoS ONE* 3:e2635. doi: 10.1371/journal.pone.0002635

Jahfari, S., Ridderinkhof, K. R., and Scholte, H. S. (2013). Spatial frequency information modulates response inhibition and decision-making processes. *PLoS ONE* 8:e76467. doi: 10.1371/journal.pone.0076467

Karni, A., and Sagi, D. (1991). Where practice makes perfect in texture discrimination: evidence for primary visual cortex plasticity. *Proc. Natl. Acad. Sci. U.S.A.* 88, 4966–4970. doi: 10.1073/pnas.88.11.4966

Kayser, A. S., Erickson, D. T., Buchsbaum, B. R., and D'Esposito, M. (2010). Neural representations of relevant and irrelevant features in perceptual decision making. *J. Neurosci.* 30, 15778–15789. doi: 10.1523/JNEUROSCI.3163-10.2010

Kim, J. N., and Shadlen, M. N. (1999). Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nat. Neurosci.* 2, 176–185. doi: 10.1038/5739

Kruschke, J. (2010). *Doing Bayesian Data Analysis: A Tutorial Introduction with R and BUGS.* San Diego, CA: Academic Press/Elsevier.

Laming, D. R. J. (1968). *Information Theory of Choice-Reaction Times.* Oxford, UK: Academic Press.

Law, C.-T., and Gold, J. I. (2008). Neural correlates of perceptual learning in a sensory-motor, but not a sensory, cortical area. *Nat. Neurosci.* 11, 505–513. doi: 10.1038/nn2070

Lehmann, E. (1959). *Testing Statistical Hypotheses.* New York, NY: Wiley.

Leuthold, H., Sommer, W., and Ulrich, R. (1996). Partial advance information and response preparation: inferences from the lateralized readiness potential. *J. Exp. Psychol. Gen.* 125, 307–323. doi: 10.1037/0096-3445.125.3.307

Lindley, D. V. (1965). *Introduction to Probability and Statistics from a Bayesian Viewpoint.* Cambridge: Cambridge University Press. doi: 10.1017/CBO9780511662973

Li, W., Piëch, V., and Gilbert, C. D. (2004). Perceptual learning and top-down influences in primary visual cortex. *Nat. Neurosci.* 7, 651–657. doi: 10.1038/nn1255

Link, S. W. (1975). The relative judgment theory of two choice response time. *J. Math. Psychol.* 12, 114–135. doi: 10.1016/0022-2496(75)90053-X

Link, S. W., and Heath, R. A. (1975). A sequential theory of psychological discrimination. *Psychometrika* 40, 77–105. doi: 10.1007/BF02291481

Liu, C. C., and Watanabe, T. (2012). Accounting for speed-accuracy tradeoff in perceptual learning. *Vision Res.* 61, 107–114. doi: 10.1016/j.visres.2011.09.007

Liu, Z., and Weinshall, D. (2000). Mechanisms of generalization in perceptual learning. *Vision Res.* 40, 97–109. doi: 10.1016/S0042-6989(99)00140-6

Logan, G. D. (1992). Shapes of reaction-time distributions and shapes of learning curves: a test of the instance theory of automaticity. *J. Exp. Psychol. Learn. Mem. Cogn.* 18, 883–914. doi: 10.1037/0278-7393.18.5.883

Luce, R. D. (1986). *Response Times: Their Role in Inferring Elementary Mental Organization.* New York, NY: Oxford University Press.

Marshall, J. A. R., Dornhaus, A., Franks, N. R., and Kovacs, T. (2006). Noise, cost and speed-accuracy trade-offs: decision-making in a decentralized system. *J. R. Soc. Interface* 3, 243–254. doi: 10.1098/rsif.2005.0075

Masaki, H., Wild-Wall, N., Sangals, J., and Sommer, W. (2004). The functional locus of the lateralized readiness potential. *Psychophysiology* 41, 220–230. doi: 10.1111/j.1469-8986.2004.00150.x

Matzke, D., Dolan, C. V., Logan, G. D., Brown, S. D., and Wagenmakers, E.-J. (2013). Bayesian parametric estimation of stop-signal reaction time distributions. *J. Exp. Psychol. Gen.* 142, 1047–1073. doi: 10.1037/a0030543

Mazurek, M. E., Roitman, J. D., Ditterich, J., and Shadlen, M. N. (2003). A role for neural integrators in perceptual decision making. *Cereb. Cortex* 13, 1257–1269. doi: 10.1093/cercor/bhg097

Mulder, M. J., Bos, D., Weusten, J. M. H., van Belle, J., van Dijk, S. C., Simen, P., et al. (2010). Basic impairments in regulating the speed-accuracy tradeoff predict symptoms of attention-deficit/hyperactivity disorder. *Biol. Psychiatry* 68, 1114–1119. doi: 10.1016/j.biopsych.2010.07.031

Mulder, M. J., Keuken, M. C., van Maanen, L., Boekel, W., Forstmann, B. U., and Wagenmakers, E.-J. (2013). The speed and accuracy of perceptual decisions in a random-tone pitch task. *Atten. Percept. Psychophys.* 75, 1048–1058. doi: 10.3758/s13414-013-0447-8

Nosofsky, R. M., and Palmeri, T. J. (1997). An exemplar-based random walk model of speeded classification. *Psychol. Rev.* 104, 266–300. doi: 10.1037/0033-295X.104.2.266

Osman, A., Lou, L., Muller-Gethmann, H., Rinkenauer, G., Mattes, S., and Ulrich, R. (2000). Mechanisms of speed–accuracy tradeoff: evidence from covert motor processes. *Biol. Psychol.* 51, 173–199. doi: 10.1016/S0301-0511(99)00045-9

Pachella, R. G., and Pew, R. W. (1968). Speed-accuracy tradeoff in reaction time: effect of discrete criterion times. *J. Exp. Psychol.* 76, 19–24. doi: 10.1037/h0021275

Palmer, J., Huk, A. C., and Shadlen, M. N. (2005). The effect of stimulus strength on the speed and accuracy of a perceptual decision. *J. Vis.* 5, 376–404. doi: 10.1167/5.5.1

Petrov, A. A., Van Horn, N. M., and Ratcliff, R. (2011). Dissociable perceptual-learning mechanisms revealed by diffusion-model analysis. *Psychon. Bull. Rev.* 18, 490–497. doi: 10.3758/s13423-011-0079-8

Pilly, P. K., and Seitz, A. R. (2009). What a difference a parameter makes: a psychophysical comparison of random dot motion algorithms. *Vision Res.* 49, 1599–1612. doi: 10.1016/j.visres.2009.03.019

Ploran, E. J., Nelson, S. M., Velanova, K., Donaldson, D. I., Petersen, S. E., and Wheeler, M. E. (2007). Evidence accumulation and the moment of recognition: dissociating perceptual recognition processes using fMRI. *J. Neurosci.* 27, 11912–11924. doi: 10.1523/JNEUROSCI.3522-07.2007

Rae, B., Heathcote, A., Donkin, C., Averell, L., and Brown, S. (in press). The hare and the tortoise: emphasizing speed can change the evidence used to make decisions. *J. Exp. Psychol. Learn. Mem. Cogn.*

Ratcliff, R. (1978). A theory of memory retrieval. *Psychol. Rev.* 85, 59–108. doi: 10.1037/0033-295X.85.2.59

Ratcliff, R. (2002). A diffusion model account of response time and accuracy in a brightness discrimination task: fitting real data and failing to fit fake but plausible data. *Psychon. Bull. Rev.* 9, 278–291. doi: 10.3758/BF03196283

Ratcliff, R., Gomez, P., and McKoon, G. (2004). A diffusion model account of the lexical decision task. *Psychol. Rev.* 111, 159–182. doi: 10.1037/0033-295X.111.1.159

Ratcliff, R., and McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. *Neural Comput.* 20, 873–922. doi: 10.1162/neco.2008.12-06-420

Ratcliff, R., and Rouder, J. N. (1998). Modeling response times for two-choice decisions. *Psychol. Sci.* 9, 347–356. doi: 10.1111/1467-9280.00067

Ratcliff, R., and Rouder, J. N. (2000). A diffusion model account of masking in two-choice letter identification. *J. Exp. Psychol. Hum. Percept. Perform.* 26, 127–140. doi: 10.1037/0096-1523.26.1.127

Ratcliff, R., and Smith, P. L. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychol. Rev.* 111, 333–367. doi: 10.1037/0033-295X.111.2.333

Ratcliff, R., and Tuerlinckx, F. (2002). Estimating parameters of the diffusion model: approaches to dealing with contaminant reaction times and parameter variability. *Psychon. Bull. Rev.* 9, 438–481. doi: 10.3758/BF03196302

Ratcliff, R., Van Zandt, T., and McKoon, G. (1999). Connectionist and diffusion models of reaction time. *Psychol. Rev.* 106, 261–300. doi: 10.1037/0033-295X.106.2.261

Rinkenauer, G., Osman, A., Ulrich, R., Muller-Gethmann, H., and Mattes, S. (2004). On the locus of speed-accuracy trade-off in reaction time: inferences from the lateralized readiness potential. *J. Exp. Psychol. Gen.* 133, 261–282. doi: 10.1037/0096-3445.133.2.261

Roitman, J. D., and Shadlen, M. N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J. Neurosci.* 22, 9475–9489.

Schall, J. D. (2002). The neural selection and control of saccades by the frontal eye field. *Philos. Trans. R. Soc. B Biol. Sci.* 357, 1073–1082. doi: 10.1098/rstb.2002.1098

Schmiedek, F., Oberauer, K., Wilhelm, O., Süss, H.-M., and Wittmann, W. W. (2007). Individual differences in components of reaction time distributions and their relations to working memory and intelligence. *J. Exp. Psychol. Gen.* 136, 414–429. doi: 10.1037/0096-3445.136.3.414

Schouten, J. F., and Bekker, J. A. M. (1967). Reaction time and accuracy. *Acta Psychol. (Amst.)* 27, 143–153. doi: 10.1016/0001-6918(67)90054-6

Shadlen, M. N., and Newsome, W. T. (2001). Neural basis of a perceptual decision in the parietal cortex (Area LIP) of the rhesus monkey. *J. Neurophysiol.* 86, 1916–1936.

Simen, P., Cohen, J. D., and Holmes, P. (2006). Rapid decision threshold modulation by reward rate in a neural network. *Neural Netw.* 19, 1013–1026. doi: 10.1016/j.neunet.2006.05.038

Simen, P., Contreras, D., Buck, C., Hu, P., Holmes, P., and Cohen, J. D. (2009). Reward rate optimization in two-alternative decision making: empirical tests of theoretical predictions. *J. Exp. Psychol. Hum. Percept. Perform.* 35, 1865–1897. doi: 10.1037/a0016926

Smith, P. L., and Ratcliff, R. (2004). Psychology and neurobiology of simple decisions. *Trends Neurosci.* 27, 161–168. doi: 10.1016/j.tins.2004.01.006

Spiegelhalter, D. J., Best, N. G., Carlin, B. P., and van der Linde, A. (2002). Bayesian measures of model complexity and fit. *J. R. Stat. Soc. B Stat. Methodol.* 64, 583–639. doi: 10.1111/1467-9868.00353

Starns, J. J., and Ratcliff, R. (2014). Validating the unequal-variance assumption in recognition memory using response time distributions instead of ROC functions: a diffusion model analysis. *J. Mem. Lang.* 70, 36–52. doi: 10.1016/j.jml.2013.09.005

Stone, M. (1960). Models for choice-reaction time. *Psychometrika* 25, 251–260. doi: 10.1007/BF02289729

Townsend, J. T., and Ashby, F. (1983). *The Stochastic Modeling of Elementary Psychological Processes*. Cambridge, MA: Cambridge University Press.

Trobalon, J. B., Chamizo, V. D., and Mackintosh, N. J. (1992). Role of context in perceptual learning in maze discriminations. *Q. J. Exp. Psychol. B.* 44, 57–73.

Uchida, N., and Mainen, Z. F. (2003). Speed and accuracy of olfactory discrimination in the rat. *Nat. Neurosci.* 6, 1224–1229. doi: 10.1038/nn1142

Vandekerckhove, J., Tuerlinckx, F., and Lee, M. D. (2011). Hierarchical diffusion models for two-choice response times. *Psychol. Methods* 16, 44–62. doi: 10.1037/a0021765

Van Veen, V., Krug, M. K., and Carter, C. S. (2008). The neural and computational basis of controlled speed-accuracy tradeoff during task performance. *J. Cogn. Neurosci.* 20, 1952–1965. doi: 10.1162/jocn.2008.20146

Voss, A., Rothermund, K., and Voss, J. (2004). Interpreting the parameters of the diffusion model: an empirical validation. *Mem. Cognit.* 32, 1206–1220. doi: 10.3758/BF03196893

Wagenmakers, E.-J. (2009). Methodological and empirical developments for the Ratcliff diffusion model of response times and accuracy. *Eur. J. Cogn. Psychol.* 21, 641–671. doi: 10.1080/09541440802205067

Wagenmakers, E.-J., Ratcliff, R., Gomez, P., and McKoon, G. (2008). A diffusion model account of criterion shifts in the lexical decision task. *J. Mem. Lang.* 58, 140–159. doi: 10.1016/j.jml.2007.04.006

Wald, A. (1947). *Sequential Analysis*. New York, NY: Wiley.

Watanabe, T., Náñez, J. E., and Sasaki, Y. (2001). Perceptual learning without perception. *Nature* 413, 844–848. doi: 10.1038/35101601

Wickelgren, W. A. (1977). Speed-accuracy tradeoff and information processing dynamics. *Acta Psychol. (Amst.)* 41, 67–85. doi: 10.1016/0001-6918(77)90012-9

Wiecki, T. V, Sofer, I., and Frank, M. J. (2013). HDDM: hierarchical bayesian estimation of the drift-diffusion model in python. *Front. Neuroinform.* 7:14. doi: 10.3389/fninf.2013.00014

Wylie, S. A., van den Wildenberg, W. P. M., Ridderinkhof, K. R., Bashore, T. R., Powell, V. D., Manning, C. A., et al. (2009). The effect of speed-accuracy strategy on response interference control in Parkinson's disease. *Neuropsychologia* 47, 1844–1853. doi: 10.1016/j.neuropsychologia.2009.02.025

Yang, T., and Maunsell, J. H. R. (2004). The effect of perceptual learning on neuronal responses in monkey visual area V4. *J. Neurosci.* 24, 1617–1626. doi: 10.1523/JNEUROSCI.4442-03.2004

Zhang, J. (2012). The effects of evidence bounds on decision-making: theoretical and empirical developments. *Front. Psychol.* 3:263. doi: 10.3389/fpsyg.2012.00263

Zhang, J., Hughes, L. E., and Rowe, J. B. (2012). Selection and inhibition mechanisms for human voluntary action decisions. *Neuroimage* 63, 392–402. doi: 10.1016/j.neuroimage.2012.06.058

Zhang, J., and Kourtzi, Z. (2010). Learning-dependent plasticity with and without training in the human brain. *Proc. Natl. Acad. Sci. U.S.A.* 107, 13503–13508. doi: 10.1073/pnas.1002506107

Zhang, J., Meeson, A., Welchman, A. E., and Kourtzi, Z. (2010). Learning alters the tuning of functional magnetic resonance imaging patterns for visual forms. *J. Neurosci.* 30, 14127–14133. doi: 10.1523/JNEUROSCI.2204-10.2010

frontiers in
NEUROSCIENCE

Learning to maximize reward rate: a model based on semi-Markov decision processes

Arash Khodadadi*, Pegah Fakhari and Jerome R. Busemeyer

Department of Psychological and Brain Sciences, Indiana University, Bloomington, IN, USA

When animals have to make a number of decisions during a limited time interval, they face a fundamental problem: how much time they should spend on each decision in order to achieve the maximum possible total outcome. Deliberating more on one decision usually leads to more outcome but less time will remain for other decisions. In the framework of sequential sampling models, the question is how animals learn to set their decision threshold such that the total expected outcome achieved during a limited time is maximized. The aim of this paper is to provide a theoretical framework for answering this question. To this end, we consider an experimental design in which each trial can come from one of the several possible "conditions." A condition specifies the difficulty of the trial, the reward, the penalty and so on. We show that to maximize the expected reward during a limited time, the subject should set a separate value of decision threshold for each condition. We propose a model of learning the optimal value of decision thresholds based on the theory of semi-Markov decision processes (SMDP). In our model, the experimental environment is modeled as an SMDP with each "condition" being a "state" and the value of decision thresholds being the "actions" taken in those states. The problem of finding the optimal decision thresholds then is cast as the stochastic optimal control problem of taking actions in each state in the corresponding SMDP such that the average reward rate is maximized. Our model utilizes a biologically plausible learning algorithm to solve this problem. The simulation results show that at the beginning of learning the model choses high values of decision threshold which lead to sub-optimal performance. With experience, however, the model learns to lower the value of decision thresholds till finally it finds the optimal values.

Keywords: semi-Markov decision process, average reward rate maximization, speed-accuracy trade-off, reinforcement learning, sequential sampling models, diffusion process, decision threshold

1. INTRODUCTION

In many problems that animals and humans encounter, the quality of a desired outcome that they can achieve depends on the amount of a resource they spent. For example, one can pay more money (resource) to buy a more stylish (higher quality) coat (desired outcome). If the resource is limited (which is almost always the case), the animal or human should decide how much of the resource she is willing to spend on obtaining one outcome. By spending more of the resource on an outcome the quality increases but less would be left for other outcomes. A rational animal or human, then, should decide how to allocate the resource for obtaining each outcome to maximize the total amount of obtained outcome. That is, she should find out what resource allocation maximizes *outcome per unit of the resource*.

One interesting example of a situation in which the subject should trade a resource with the quality of the outcome is perceptual decision making in which the subject should detect a noisy stimulus and choose a proper response based on it. Because of the noise in the stimulus, to make more accurate responses the subject should spend more time to detect the stimulus. Since faster responses are less accurate, the subject should trade between the amount of time (resource) and the accuracy (which

determines the quality of the outcome). This leads to the so-called speed-accuracy tradeoff (SAT).

In the past few decades, computational modeling has been a popular method for investigating the mechanisms underlying perceptual decision making. A large class of models of perceptual decision making, called sequential sampling models, assume that the subject sequentially samples from the stimulus (Link and Heath, 1975; Townsend and Ashby, 1983; Luce, 1986; Smith and Vickers, 1988; Busemeyer and Townsend, 1993; Smith, 2000; Usher and McClelland, 2001; Ratcliff and Smith, 2004). These samples are noisy and so the decision cannot be made based on a single sample. These models propose that the subject responds whenever the accumulated evidence favoring one of the responses exceeds a specific value called the *decision threshold*. This way, these models separate the perceptual process from the decisional process. The evidence accumulation models the perceptual process and is assumed to be affected by the physical stimulus. The decisional process is modeled by the decision threshold and is assumed to be controlled by the subject. Higher values of the decision threshold mean that more information is needed for making a decision and so the decisions will be more accurate. However, accumulating more information takes more time and

so decisions will be slower. Thus, the SAT is explained in sequential sampling models by changes in the decision threshold. This feature of sequential sampling models has motivated a large body of research on the SAT phenomena. A standard experimental method of investigating this phenomena is to vary the emphasis on speed or accuracy in the task instructions. Sequential sampling models predict that the subjects will choose lower decision threshold in the speed condition in comparison to the accuracy conditions. This prediction has been confirmed in many studies (Ratcliff, 1978; Luce, 1986; Ratcliff, 2002; Palmer et al., 2005; Ratcliff and McKoon, 2007; Ivanoff et al., 2008; Wagenmakers et al., 2008; Bogacz et al., 2010; Forstmann et al., 2010).

Although these results show that subjects choose different values of decision threshold in response to varying the task's instructions, they do not specify what value of the decision threshold should be chosen in each condition. In other words, the results of theses studies do not provide a normative account of the SAT phenomena. The rationality notion explained above, however, suggests a possible way to provide such an explanation: if the total time of the task is fixed, a rational subject should balance between her speed and accuracy such that the total outcome obtained during the whole task is maximized. Spending more time on one trial results in less remaining time for the other trials, meaning the subject experiences fewer trials in the task. However, by spending more time on one trial the subject can increase the chance of responding correctly.

This experimental design was first suggested by Gold and Shadlen (2002). They considered a perceptual decision making task in which the total time of the task is fixed and so the total number of trials that the subject can experience depends on the average time she spends on each trial. Also, the subject receives a reward after each correct response and a penalty after each incorrect response. They proposed that a rational subject sets her decision threshold such that the expected total outcome (sum of rewards and penalties) would be maximized. Because the total time of the task is limited and fixed, this is equivalent to maximizing the expected outcome per unit time, or the *average reward rate*.

Bogacz et al. (2006) further investigated the properties of the average reward rate as a function of the task parameters (e.g., reward, penalty, stimulus salience and so on) and the parameters of a class of sequential sampling models. Specifically, they derived the relationship between the task parameters and the optimal value of the decision threshold in the experimental design of Gold and Shadlen. More recently, Simen et al. (2009) and Balci et al. (2011) conducted a series of experiments to see if human subjects can achieve the optimal performance in this experimental design. The results of these studies showed that after extensive training in tasks similar to what was proposed by Gold and Shadlen (2002), human subjects could learn to set the decision threshold at values close to optimal.

Knowing that subjects can learn to behave optimally, the next question would be how the brain learns the optimal threshold. The aim of this paper is to propose a computational framework to answer this question. To this end, we consider a more general experimental design than the design of Gold and Shadlen. In this design, instead of having one condition, trials in a block

can come from one of several possible conditions and so the subject should set different decision thresholds for different conditions to achieve the maximum average reward rate (section 2). We then show that this experiment can be modeled as a stochastic process, specifically a *semi-Markov decision process* (section 4). Learning the optimal decision threshold will be framed as an optimal control problem in this stochastic environment. We then propose a biologically plausible model that can solve this problem (section 5). In the final section of this paper, we test the performance of our model in learning the optimal value of the decision threshold in different experiments (section 6).

2. COMPUTATIONAL METHODS

Our model is developed to account for a more general experimental design than what was used in previous research on optimal SAT. To the best of our knowledge, Simen et al. (2009) conducted the first experimental study to investigate if human subjects can learn the optimal value of the decision threshold. To contrast their experimental design with the one that is considered in this paper, here we briefly explain experiment 1 of Simen et al. (2009).

The stimulus in each trial of this experiment was the well-known random-dot kinematogram. This stimulus consists of a number of dots, some of them moving coherently toward the left or toward the right, while other dots move randomly. The subjects' task is to decide in each trial if the net direction of motion is toward the left or right. The salience of the stimulus is determined by the percentage of dots that are moving coherently. Other task parameters were the reward that the subject receives after each correct response and the response-stimulus interval (RSI), the time between subject's response and the presentation of the next stimulus. Each session of the experiment consisted of 12 blocks (the number of blocks was more than 12, but here we just consider those that are relevant to our explanation). The blocks' duration was fixed (4 min) and so the number of trials in each block depended on how much time the subject spent on each trial.

Based on Gold and Shadlen's hypothesis (Gold and Shadlen, 2002), because the blocks' duration is fixed, a rational subject will try to balance her speed and accuracy such that the average reward rate is maximized. Since the total reward is the sum of the reward for each block, maximizing the total average reward rate is equivalent to maximizing it in each individual block. In experiment 1 of Simen et al. (2009), the stimulus salience and reward were held constant. RSI was held constant within each block, but manipulated across blocks. Clearly, the average reward rate is a function of RSI, since the longer the delay between the trials, the fewer trials can be experienced within a block. In addition, (Bogacz et al., 2006) showed that if the subjects' performance in this experiment is modeled in the sequential sampling framework, the optimal value of the decision threshold is a function of RSI. Therefore, to achieve the optimal performance in each block (and so maximize the total average reward rate) subjects have to set different decision thresholds for different blocks, dependent on the RSI .

Although the optimal value of the decision threshold in a block depends on the RSI in that block, it does not depend on the RSI in other blocks. In other words, to maximize the average reward rate in one block, the subject does not need to know what are the values of RSI in other blocks. Therefore, the subject can set the

value of the decision threshold in each block with a specific value of RSI, independent of other blocks with different RSIs. This is the main difference between this design and the design we consider in the current paper.

Here, we consider a more general design in which to achieve the optimal performance the subject should consider all conditions together and the optimal decision threshold for one condition depends on all other conditions in the task. As an example, consider two conditions: RSI= 500 ms and RSI= 1000 ms. In the previous design, there would be two types of blocks: in one type the RSI of all trials is 500 ms while the RSI of the trials in the other type of blocks is 1000 ms. In our design, however, trials with RSI= 500 ms and RSI= 1000 ms are all intermixed. In other words, there is no manipulation across blocks. Crucially, a cue associated with each RSI value is presented at the beginning of each trial. For example, in the task set-up shown in **Figure 1**, in trials with RSI= 500 ms a red cross-hair is presented as the cue while in trials with RSI= 1000 the cue is a blue cross-hair. As seen in this figure, the cue is followed by the random dots stimulus. The blocks' duration is fixed and so a rational subject should maximize the average reward rate.

Since a cue associated with the RSI of the trial is presented before the presentation of the stimulus, we assume that subjects can set different values of decision threshold for each value of RSI. In other words, the subject can associate different value of decision threshold to each cue. Thus, like the design in Simen et al. (2009), the average reward rate will be a function of the two decision thresholds. The crucial difference is that in their design the two decision thresholds are independent of each other (since RSI is only manipulated across blocks), whereas in our design the optimal threshold for one value of RSI depends on the value of the other RSI.

The reason for this dependency in our design can be conceived intuitively by noting that because the blocks' duration is fixed,

every second that the subject spends on one trial, she is actually losing the opportunity to spend that time on other trials. If the other trials on average lead to higher reward, it is better to spend less time on the current trial. By being faster in one type of trials, the accuracy decreases and so the subject will lose more rewards in those trials. However, if other types of trials are "rewarding enough" it may be worth it to be fast and inaccurate in those trials which lead to less reward. This means that to set the decision threshold in each condition, the subject should consider all other conditions in the task. In the next section, we derive a formal expression of the average reward rate in our design and investigate its properties in more detail.

3. AVERAGE REWARD RATE

In this section, we investigate the properties of the average reward rate as a function of the task parameters. We first state the formula for the average reward rate in the experimental design explained above. To see how this function is related to the decision threshold in different conditions, we then explain a variant of sequential sampling models called *independent race model* and show how the decision making process is modeled in this framework. Finally, we see some examples of the average reward rate for different task parameters.

3.1. AVERAGE REWARD RATE AS A FUNCTION OF TASK PARAMETERS

In section 2, we explained the experimental design with an example in which only one task parameter (RSI) was manipulated. Before deriving the formula for the average reward rate, we should explain the experimental design in more detail.

In the experimental design, there are several blocks with a fixed duration. Each trial in each block comes from one of the N_c possible "conditions" with each condition specifying the task parameters in that trial. Each trial is drawn from a given condition C_i with probability P_i. As explained before, a cue presented at
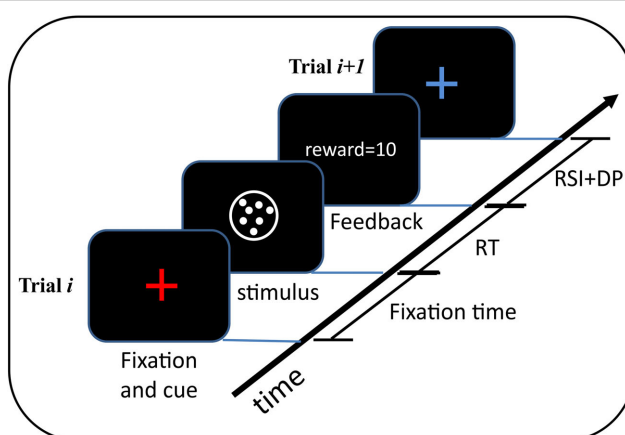
	Condition 1	Condition 2
Associated color	red	blue
Salience	10%	10%
Reward	1	1
Penalty	0	0
DP	0 sec	0 sec
RSI	500 ms	1 sec

FIGURE 1 | An example of the experimental design. In this example, each trial can come from one of two conditions with equal probability. A colored cross-hair presented at the beginning of each trial indicates the condition of the upcoming trial. After the presentation of the cue, the stimulus appears and remains on the screen till the subject responds. After responding, the subject receives feedback. The time between subject's response and the beginning of the next trial is determined by the delay penalty (DP) and response-stimulus interval (RSI). The table on the right of the figure shows the cue-condition association along with the task parameters in each condition. As seen, the two conditions differ only in the value of the RSI. See **Table 1** for a description of the task parameters. In this table, a separate parameter for the fixation time is not considered and instead it is considered as a part of the RSI.

the beginning of each trial indicates which condition this trial is coming from. For example, in **Figure 1** there are two conditions and each trial can come from one of them with equal probability. In this figure, all task parameters are the same in these two conditions except the RSI.

The subject receives a reward after each correct response and a penalty after incorrect responses. Also, there is a delay penalty after each incorrect response. This is the time that the subject should wait in addition to the RSI when the response is incorrect. The task parameters and their notations are specified in **Table 1**.

The average reward rate is defined as the average reward divided by the average time that it takes to obtain the reward. In our experimental design , since the subject can choose different decision thresholds for different conditions, the average time and average reward will be different in different conditions . The average reward in the task, then, is the weighted sum of the average reward in each condition with the weights being the probability of each condition presented in the task. The average time is computed in the same way. The average reward rate, then, can be expressed as follows:

$$\bar{R} = \frac{\sum_{i=1}^{N_c} P_i \cdot \left[r_i^C \cdot P_i^C + r_i^I \cdot (1 - P_i^C) \right]}{\sum_{i=1}^{N_c} P_i \cdot \left[\bar{T}_i^C \cdot P_i^C + (\bar{T}_i^I + T_i^{DP}) \cdot (1 - P_i^C) + T_i^{RSI} + T^{ND} \right]} \tag{1}$$

Among all these parameters, the subject can only control the probability of correct P_c^i, mean correct reaction time $\bar{T}_i^C$ and mean incorrect reaction time $\bar{T}_i^I$ in each condition i, by adjusting her decision threshold in each condition. All other parameters are controlled by the experimenter. The sequential sampling models specify the relationship between the decision threshold and mean reaction time and probability correct. In the next section we explain this relationship.

3.2. DIFFUSION PROCESS MODEL OF PERCEPTUAL DECISION MAKING

As explained before, the sequential sampling models assume that the subject accumulates noisy information favoring each response and she will respond as soon as the evidence favoring one of the responses reaches a decision threshold. Several models have been proposed based on different assumptions about the accumulation process and the decision process (see Ratcliff and Smith, 2004 for

Table 1 | Task parameters in the experiment.

Parameter	Description
N_c	Number of conditions
P_i	Probability of condition i happening
P_i^C	Probability of being correct in condition i
r_i^C	Reward in condition i
r_i^I	Penalty in condition i
$\bar{T}_i^C$	Mean correct reaction time in condition i
$\bar{T}_i^I$	Mean incorrect reaction time in condition i
T_i^{DP}	Delay penalty after incorrect responses in condition i
T_i^{RSI}	Response-stimulus interval in condition i
T^{ND}	Non-decision time which is assumed to be independent of the condition

a comprehensive review of different sequential sampling models). Although different models make different predictions about a subject's performance, most of these models can fulfil the purpose of this paper. In this paper, we consider an independent race model in which the information favoring each response is accumulated in a separate accumulator. This model assumes that the subject responds as soon as the accumulated information in one of the accumulators reaches its decision threshold. The accumulated information in each accumulator is modeled as a diffusion process. In this model, the information is sampled and accumulated in continuous time. A diffusion process X is specified by the stochastic differential equation:

$$dX = \mu \cdot dt + \sigma \cdot dB \tag{2}$$

The parameter μ is called the drift coefficient and determines the mean of the process X (It can be shown that $E[X(t)] = \mu \cdot t$ (see for example Smith, 2000)). This parameter is assumed to be proportional to the stimulus salience. The parameter σ is the diffusion coefficient and specifies the amount of noise in the samples. The process dB specifies the increments of a zero-mean Gaussian process.
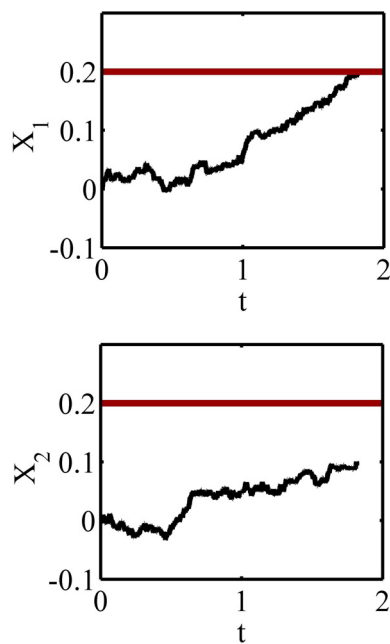
Consider the random-dot kinematogram task explained in section 2. The corresponding independent race model of this task consists of two accumulators that one of them accumulates information favoring the "right" response while the other accumulates information favoring the "left" response. Each of these accumulators is a diffusion process with one decision threshold (see **Figure 2**). Thus, the parameters of the model are the drift coefficients μ_i, the diffusion coefficients σ_i and the decision thresholds a_i, where the subscript $i = 1, 2$ denotes the ith accumulator. For sake of simplicity, we assume that $\sigma_1 = \sigma_2 = \sigma$ and $a_1 = a_2 = a$. In this paper, we do not distinguish between the right and left responses, and instead assume that accumulator 1 corresponds to the correct response and accumulator 2 corresponds to the incorrect response. The probability of giving a correct response, as well as the probability density functions for the correct and incorrect reaction times, are expressed in Supplementary Material. These functions specify the relationship between the average reward rate function in Equation1 and the decision thresholds in the different conditions, and so all parameters being fixed, one can plot $\bar{R}$ as a function of the decision thresholds. Several examples are investigated in the next section.

3.3. SOME EXAMPLES

In this section, we investigate the properties of the average reward rate function in Equation1 with three examples. In the first example, we consider an experiment similar to the experimental design used previously (Simen et al., 2009; Balci et al., 2011). As explained above, in this case the subject has to set only one decision threshold for each block. The average reward rate $\bar{R}$ as a function of the decision threshold for the task parameters given in **Table 2** and different values of T^{RSI} is shown in **Figure 3**. As can be seen in this figure, for all values of T^{RSI} there is one value of the decision threshold a that maximizes the average reward rate. The properties of the average reward rate function for another sequential sampling model, called the drift diffusion model, have
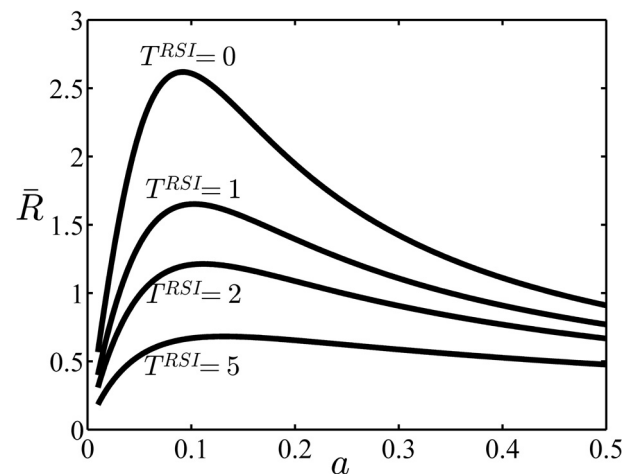
FIGURE 2 | A sample path of the accumulated information in the two accumulators of the race model. The black paths are the accumulated information and the thick red lines are the decision thresholds. The accumulated information in the accumulator 1 reaches the decision threshold at about 1.8 s and before the accumulator 2 and so the response 1 will be selected.

FIGURE 3 | The average reward rate in example 1. In this example, no cue is presented at the beginning of trials and so it is assumed that the subject sets one decision threshold for all trials. In this figure, the average reward rate is plotted as a function of this decision threshold, (denoted as a in the figure), and for different values of the parameter T^{RSI}. Other parameters are given in **Table 2**.

Table 2 | Task parameters used in the first example.

Parameter	μ_1	μ_2	σ	T^{ND}	r^C	r^I	T^{DP}
Value	0.1	0	0.05	0.5	5	−5	5

Table 3 | Task parameters used in the second example.

Parameter	μ_1^1	μ_1^2	μ_2^1	μ_2^2	σ	T^{ND}	r_1^C	r_1^I	r_2^C	r_2^I	T^{DP}	T^{RSI}
Value	0.1	0	0.15	0	0.07	1	2	−1	2	−5	2	1

been investigated thoroughly before (Bogacz et al., 2006; Simen et al., 2006, 2009; Balci et al., 2011). Specifically, it has been shown that this function is uni-modal in the whole parameter space. Our simulations, not reported here, showed that this is also the case for the independent race model used here.
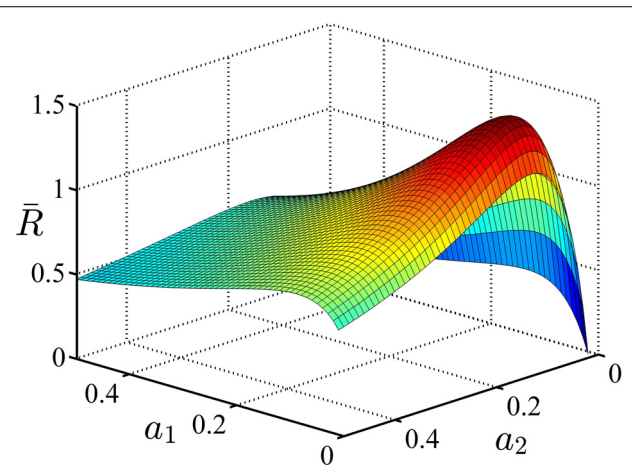
In the second example, we consider an experiment similar to what was shown in **Figure 1**. In this experiment, each trial could come from one of the two conditions with equal probability. The task parameters are given in **Table 3**. In this table, μ_i^j is the drift coefficient of accumulator j in condition i. As explained before, in this experiment the subject can set separate decision thresholds for different conditions. The average reward rate as a function of the decision threshold in condition 1, a_1, and condition 2, a_2, is shown in **Figure 4**. Although we do not prove it here, our simulations suggest that this function is also uni-modal over the whole parameter space and so there is one pair of the decision thresholds that maximize it. In **Figure 4**, the average reward rate is maximized when $a_1 = 0.06$ and $a_2 = 0.11$. As can be seen in **Table 3**, in both conditions, the reward for the correct response is $r_j^C = 2$ but the punishments are different. In condition 2, the punishment is greater and because of that, the subject might ponder more in

FIGURE 4 | The average reward rate in example 2. In this example, there are two conditions and each trials starts with presentation of a cue associated with these conditions. It is assumed that the subject sets different decision thresholds for each condition. In this figure, the average reward rate is plotted as a function of these decision thresholds. a_1 denotes the decision threshold in condition 1 and a_2 denotes the decision threshold in condition 2. The task parameters are given in **Table 3**.

those trials and so the optimal decision threshold for condition 2 is greater than condition 1.

In the last example of this section, we examine how the optimal decision threshold in each condition varies when the difficulty

of one of the conditions changes. Again, we consider an experiment with two conditions. The task parameters are given in **Table 4**. As can be seen in this table, all parameters of the two conditions are the same except the reward after correct responses which is 1 in condition 1 and 5 in condition 2. Here, we want to see how the optimal values of the decision thresholds change when the salience level of condition 1, μ_1^1, changes. The optimal value of the decision thresholds for several values of μ_1^1 in the interval $[0.05, 0.2]$ is plotted in the top panel of **Figure 5**. When the salience levels in the two conditions are equal, that is $\mu_1^1 = \mu_2^1 = 0.05$ and $\mu_1^2 = \mu_2^2 = 0$, the optimal value of the decision threshold in condition 2 is larger than condition 1 ($a_1^{opt} = 0.033$ and $a_2^{opt} = 0.083$). This is because each correct response in condition 2 leads to higher value of reward and so it is worth it to set a higher decision threshold for this condition and so on average spent more time on this condition than condition 1 and make more correct responses. However, as the salience level of condition 1 increases and this condition becomes easier than condition 2, the optimal decision threshold in condition 1 increases while it decreases for condition 2. To investigate this situation more, the probability of giving a correct response and mean time spent in each condition when the optimal decision thresholds are recruited are shown in the left and right panels at the bottom of **Figure 5**, respectively. As seen, by increasing μ_1^1, the optimal decision thresholds change in a way that the probability of correct response increases for condition 1 and decreases for condition 2. The mean time spent in each condition shows a more complicated pattern. In conclusion, even when the task parameters of only one condition change, the subject should adjust her speed and accuracy in all conditions to maximize the global average reward rate.

4. A STOCHASTIC PROCESS MODEL OF THE EXPERIMENT

The main aim of this paper is to propose a computational model of how subjects learn the optimal decision thresholds in the experimental design explained before. The proposed model is based on well-known reinforcement learning algorithms previously used to model optimal action selection and decision making in animals and humans (Barto, 1995; Montague et al., 1996; Schultz et al., 1997; Sutton and Barto, 1998). In these algorithms, the learning problem is formulated as an optimal control problem in a stochastic environment. One critical step in modeling in this framework is to specify the environment corresponding to the problem in hand. In this section, we show how our experimental design can be modeled as a stochastic process called a semi-Markov decision process. In what follows, we first explain Markov decision processes and then discuss how semi-Markov decision processes generalize them to continuous time problems. Finally, we show how our problem can be cast as a semi-Markov decision process.
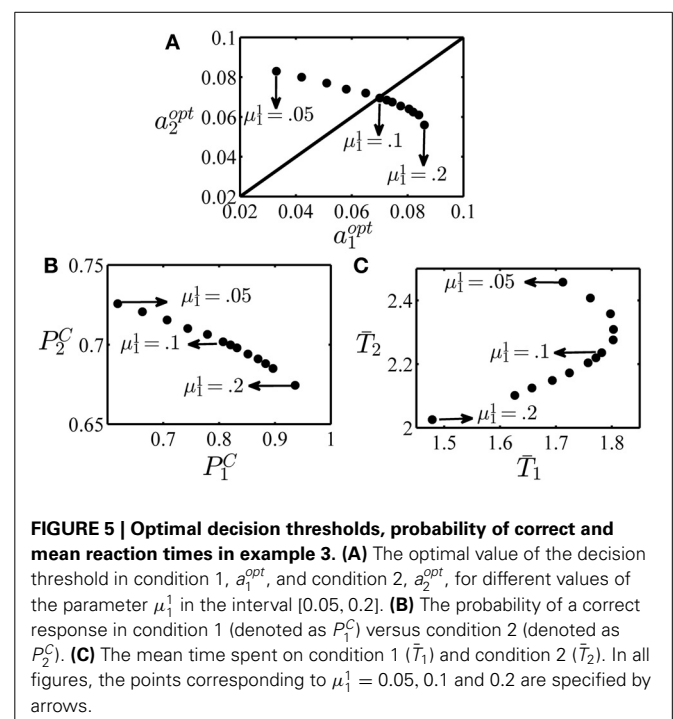
4.1. MARKOV DECISION PROCESS

In a Markov decision process an agent (e.g. an animal or a robot) is interacting with a stochastic environment. The environment consists of N states $S = \{s^1, \cdots, s^N\}$ and at each time step k it is in one of these states, say $s_k = s^i$. At each time step, the agent can choose an action from the set of M possible actions $A = \{a^1, \cdots, a^M\}$. After taking action a_k the environment transfers to a new state s_{k+1} with probability $\mathbf{T}_{ij}^u(k) = \Pr(s_{k+1} = s^j | s_k = s^i, a_k = a^u)$ and the agent receives a probabilistic reward $r_k = r$ with probability $\mathbf{R}_{ij}^u(r, k) = \Pr(r_k = r | s_k = s^i, s_{k+1} = s^j, a_k = a^u)$. The important aspect of these functions is that they possess the Markov property. That is, the transition probability $\mathbf{T}_{ij}^u$ and reward probability $\mathbf{R}_{ij}^u$ only depend on the state at time k and the action a_k and not the whole history of states and actions $\{s^1, a^1, \cdots, s^k, a^k\}$. More formally:

$$\Pr\left(s_{k+1} = s^j | s_k, a_k, \cdots, s_1, a_1\right) = \Pr\left(s_{k+1} = s^j | s_k, a_k\right) (3)$$

$$\Pr\left(r_k = r | s_{k+1}, s_k, a_k, \cdots, s_1, a_1\right) = \Pr\left(r_k = r | s_{k+1}, s_k, a_k\right) (4)$$

This structure is called a Markov decision process (MDP). In short, an MDP consists of a 4-tuple $\langle S, A, \mathbf{T}, \mathbf{R} \rangle$ such that $\mathbf{T}$ and $\mathbf{R}$ possess the Markov property. The state and action spaces in an MDP can be continuous.

The agent's goal in an MDP is to find the optimal policy. A policy $\pi : S \times A \rightarrow [0, 1]$ is a function that maps a state-action pair (s, a) to the probability of selecting action a in state s. To define the optimal policy we need a notion of optimality. This notion can be formalized based on the agent's desire to maximize a function of received rewards called *return*. One popular form of the return function used in many applications of reinforcement

FIGURE 5 | Optimal decision thresholds, probability of correct and mean reaction times in example 3. (A) The optimal value of the decision threshold in condition 1, a_1^{opt}, and condition 2, a_2^{opt}, for different values of the parameter μ_1^1 in the interval $[0.05, 0.2]$. **(B)** The probability of a correct response in condition 1 (denoted as P_1^C) versus condition 2 (denoted as P_2^C). **(C)** The mean time spent on condition 1 ($\bar{T}_1$) and condition 2 ($\bar{T}_2$). In all figures, the points corresponding to $\mu_1^1 = 0.05, 0.1$ and 0.2 are specified by arrows.

Table 4 | Task parameters used in the third example.

Parameter	μ_1^2	μ_2^1	μ_2^2	σ	T^{ND}	r_1^C	r_1^I	r_2^C	r_2^I	T^{DP}	T^{RSI}
Value	0	0.05	0	0.07	0.5	1	-2	5	-2	2	0.5

learning is *expected discounted sum of future reward*:

$$\mathrm{E}\left[\sum_{k=0}^{\infty} \gamma^k r_k\right] \tag{5}$$

where the operator E denotes expectation over all trials. The parameter γ is called the discounting factor and determines the relative weighting of immediate versus later rewards. The optimal policy will maximize this return. One reason for popularity of this return in the literature of reinforcement learning is that, as we will see in section 5.1, it will lead to a set of recursive equations for finding the optimal policy. Some of the psychologically more plausible returns (e.g., hyperbolic discounting) do not possess this property (for a fuller discussion see Daw, 2003, section 2.1.4).

Based on this notion of return, the value of state s^i at time step k under the policy π is defined as:

$$V_\pi\left(s^i, k\right) = \mathrm{E}\left[\sum_{j=k}^{\infty} \gamma^{j-k} r_j | s_k = s^i, \pi\right] \tag{6}$$

This function is called the *state value function* and is the expected discounted sum of rewards that the agent expect to receive given that the state at time step k is $s_k = s^i$ and the agent will choose actions based on policy π afterwards. It is easy to show that an optimal policy that maximizes the return 6 will also maximize the state value functions for all time steps and states. Thus a policy π^* is optimal if:

$$V_{\pi^*}(s, k) \geq V_\pi(s, k) \quad \text{for all } s \in S \text{ and } k \text{ and all policies } \pi \tag{7}$$

that is, if the state value functions under that policy are greater than those under any other policy.

4.2. SEMI-MARKOV DECISION PROCESS

In an MDP the state transitions occur at discrete time steps. Semi-Markov decision processes (SMDPs), generalize MDPs by allowing the state transitions to occur in continuous irregular times. In this framework, after the agent takes action a in state s, the environment will remain in state s for time d and then transits to the next state and the agent receives the reward r. The dwell time d is a random variable with probability density $D(d; s, a)$. An SMDP is specified by the 5-tuple $\langle S, A, \mathbf{T}, \mathbf{R}, D \rangle$ where $\mathbf{T}$, $\mathbf{R}$ and D possess the Markov property. This process is called semi-Markov because the transition from one state to another not only depends on the current state and action but also on the time elapsed since the action has been taken.

Since D is a function of action a, the dwell time in each state depends on the agent's policy. This means that in an SMDP, in addition to the total reward, the total time to achieve that reward depends on the policy. Thus, it is reasonable to define the optimal policy based on a return that takes both reward and dwell time into account. This makes the average reward rate an appealing choice for the return in an SMDP. Assuming that the rewards are delivered only after each transition (and not during the dwell

time) the average reward rate of an SMDP starting at state s^i under the policy π is defined as follows (Das et al., 1999):

$$\bar{R}^\pi\left(s^i\right) = \lim_{N \to \infty} \frac{\mathrm{E}\left[\sum_{k=0}^{N} r_k | s_0 = s^i, \pi\right]}{\mathrm{E}\left[\sum_{k=0}^{N} d_k | s_0 = s^i, \pi\right]} \tag{8}$$

The state value functions are defined accordingly. The optimal policy of an SMDP, then, maximizes the average reward rate.

4.3. AN SMDP MODEL OF THE EXPERIMENT

In this section, we show how the experimental design explained in section 2 can be modeled as an SMDP. Each SMDP is specified by the five-tuple $\langle S, A, \mathbf{T}, \mathbf{R}, D \rangle$ and so to explain our model we should show how these functions correspond to different components of the experiment and mechanisms of subjects' decision making. Before formally defining each of the components of the model, we explain them using the examples in section 3.3. The SMDP corresponding to the second example in section 3.3 is shown in **Figure 6**. There are two conditions in this example. In our model, we assume that each condition corresponds to one state of the environment and so for this example the corresponding SMDP has two states. After presentation of each cue at the beginning of each trial, the environment transits to one of these states. The dwell time in an SMDP is defined as the time between the transition from one state to another. Since state transitions in the model occur at the beginning of each trial (by the presentation of the cue), the dwell time is the time between the presentation of a cue in one trial and the time of the presentation of the cue in the next trial. The critical aspect of the model is the way that we define actions in the corresponding SMDP. In our experimental design, the main concern is the relationship between the decision threshold in each condition and the average reward rate. On the other hand, in an SMDP in each state the agent tries to take actions that maximize the average reward rate. This suggests a plausible choice for actions in the corresponding SMDP: the action in each state is
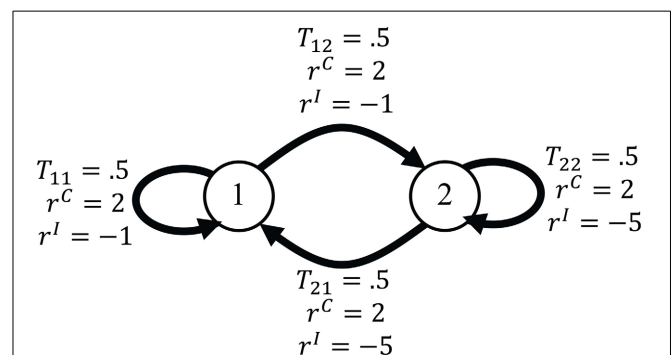
FIGURE 6 | The SMDP corresponding to the second example in section 3.3. In the example each trial could come from one of two conditions and so the corresponding SMDP has two states. The states are shown by circles. The arrows show transitions between states. The probability of each transition and the reward that the agent receives after each transition are written on the arrows.

the decision threshold of the information accumulators. In this way, by learning the optimal policy in the SMDP, the subject is actually learning the optimal value of the decision threshold for each condition. The decision threshold affects both the reaction time and the accuracy and so in the corresponding SMDP, actions affect both the reward probabilities and the dwell times. Since the decision thresholds can be any positive value, the action space in each state is the continuous space of all positive numbers.

The transition from one state to another is determined by the probability that a trial comes from a specific condition. In the example we are considering here, this probability is 0.5 for each condition. The reward that subject receives after each correct and incorrect response depends on the condition. In **Figure 6**, these quantities are shown on the arrows that indicate transitions between states.

Based on this description, the functions $\langle S, A, \mathbf{T}, \mathbf{R}, D \rangle$ can be specified as follows:

The state space S: the state space is the discrete set of all possible conditions in the experiment, that is $S = \{C_1, \cdots, C_{N_c}\}$.

The transition probability function $\mathbf{T}$: For sake of simplicity, in this paper we only consider experiments in which the probability of each trial coming from a specific condition does not depend on either subject's response or the condition presented in the previous trial. As we explained, this probability is instead determined by the experimenter. Thus, the transition probability function is defined as follows:

$$\mathbf{T}_{ij}^{u}(k) = \Pr\left(s_{k+1} = C_j | s_k = C_i, a_k = a^u\right) = \Pr\left(s_{k+1} = C_j\right) = P_j \tag{9}$$

where C_j denotes the j^{th} condition which corresponds to the j^{th} state of the environment.

The reward probability function $\mathbf{R}$: As it was explained before, the reward that subject receives after responding in each trial, depends on the condition and the subject's response. Specifically, in condition i the subject receives reward r_i^C for each correct response and r_i^I for each incorrect response. Therefore, the probability of receiving a reward r in each condition depends on subject's accuracy and so her decision threshold in that condition. Formally, the reward probability function is defined as follows:

$$\mathbf{R}_{ij}^{u}(r, k) = \Pr\left(r_k = r | s_k = C_i, s_{k+1} = C_j, a_k = a^u\right)$$

$$= \begin{cases} P_i^C & \text{if } r = r_i^C, \\ 1 - P_i^C & \text{if } r = r_i^I, \\ 0 & \text{otherwise} \end{cases} \tag{10}$$

The dwell time probability density function D: The dwell time in an SMDP is the time that it takes between transition from one state to another. In our model, this time is the sum of four parts: non-decision time, response time, delay penalty and RSI. The response time is the time between the presentation of the stimulus and the time that the first accumulator hits its decision threshold. The delay penalty depends on the subject's response and the trial condition. The probability density of the dwell time for each condition is a function of the subject's decision threshold

and the task parameters in that condition. The mean dwell time in condition i is:

$$\mathrm{E}[d] = \bar{T}_i^C \cdot P_i^C + \left(\bar{T}_i^I + T_i^{DP}\right) \cdot \left(1 - P_i^C\right) + T_i^{RSI} + T^{ND} \tag{11}$$

The action space A: As it was explained above, the action space in each state is the space of all positive real numbers. The policy $\pi(s_k = C_i, a_k = a)$ is the probability density function that specifies the likelihood of setting the value a as the decision threshold when the cue associated with condition C_i is presented.

5. MODEL

So far, we have shown how our experimental design can be modeled as an SMDP. Following Gold and Shadlen (2002) we speculated that a rational subject learns to balance her speed and accuracy in each condition such that the average reward rate is maximized. The question, then, is how the subject learns this optimal behavior. In this section, we propose a normative model of learning the optimal SAT in our experimental design. In the SMDP framework proposed above, the problem of optimal SAT is equivalent to the problem of learning the optimal policy that maximizes the average reward rate. Fortunately, the problem of learning the optimal policy in an MDP and SMDP has been investigated thoroughly in the machine learning and computer science literature. Specifically, the reinforcement learning (RL) algorithms provide a mechanism for learning the optimal policy without any knowledge about the dynamic of the environment and only by experiencing it (Bertsekas and Tsitsiklis, 1996; Sutton and Barto, 1998). In an SMDP, the dynamic of the environment is determined by the functions $\mathbf{R}$, $\mathbf{T}$ and D. The RL algorithms assume that the agent does not know these functions and can only observe noisy samples from them. This property makes these algorithms appropriate for our problem: here, the subject can only observe the reward, reaction time and condition in each trial and she should learn the optimal value of decision threshold based on these observations. Another appealing feature of these algorithms is that they provide a biologically plausible account of learning. It has been shown that the pattern of fluctuations in the firing of dopaminergic neurons in ventral tegmental area and surrounding neurons in tasks that involve prediction of reward, resembles a signal called the *temporal difference error*, which plays a central role in the RL algorithms (Montague et al., 1996; Schultz et al., 1997).

In this section, we first explain the temporal difference learning method and then propose a model that uses this method to solve the optimal SAT problem.

5.1. TEMPORAL DIFFERENCE LEARNING

As we explained before, the optimal policy in MDPs and SMDPs is defined as the policy that maximizes the state values for all states (see inequality 7). Therefore, to find the optimal policy a learning algorithm should be able to compute the state values for a given policy. For return function 5, the state values are defined in Equation 6. This equation can be written in a recursive form:

$$V_\pi(s_k) = \mathrm{E}\left[r_k + \gamma V_\pi\left(s_{k+1}\right)\right] \tag{12}$$

This equation is known as the Bellman equation. The expectation on the right hand side of this equation is taken with respect to all possible actions and states and so depends on the functions **T** and **R**. Notice that the Bellman equation provides one equation for each state in the state space and so can be considered as a system of equations. If the functions **T** and **R** are known, dynamic programming methods can be used to solve this system of equations efficiently (Bertsekas and Tsitsiklis, 1996). However, in many situations (including our problem) the agent does not know these functions. The temporal difference (TD) leaning method provides a simple and efficient solution to this problem. In its simplest form, the TD learning method uses an estimate of the difference between the two sides of Bellman equation 12 to learn the state value functions. This estimate is called the temporal difference error and is defined as follows:

$$\delta_k = r_k + \gamma \hat{V}_\pi^k(s^j) - \hat{V}_\pi^k(s^i) \tag{13}$$

where $\hat{V}_\pi^k(s)$ is the agent's estimate of the value of state s at time step k, s^i and s^j are the state of the environment at steps k and $k+1$ and r_k is the one step reward that the agent earned by going from s^i to s^j. The agent then updates its estimate of the value of state s^i using this error signal:

$$\hat{V}_\pi^{k+1}(s^i) = \hat{V}_\pi^k(s^i) + \alpha_c \cdot \delta_k \tag{14}$$

where α_c is the learning rate.

Das et al. (1999) showed that for an SMDP with the average reward rate return defined in Equation 8, the TD error signal should be computed as follows:

$$\delta_k = r_k - \hat{\rho}_k \cdot d_k + \hat{V}_\pi^k(s^j) - \hat{V}_\pi^k(s^i) \tag{15}$$
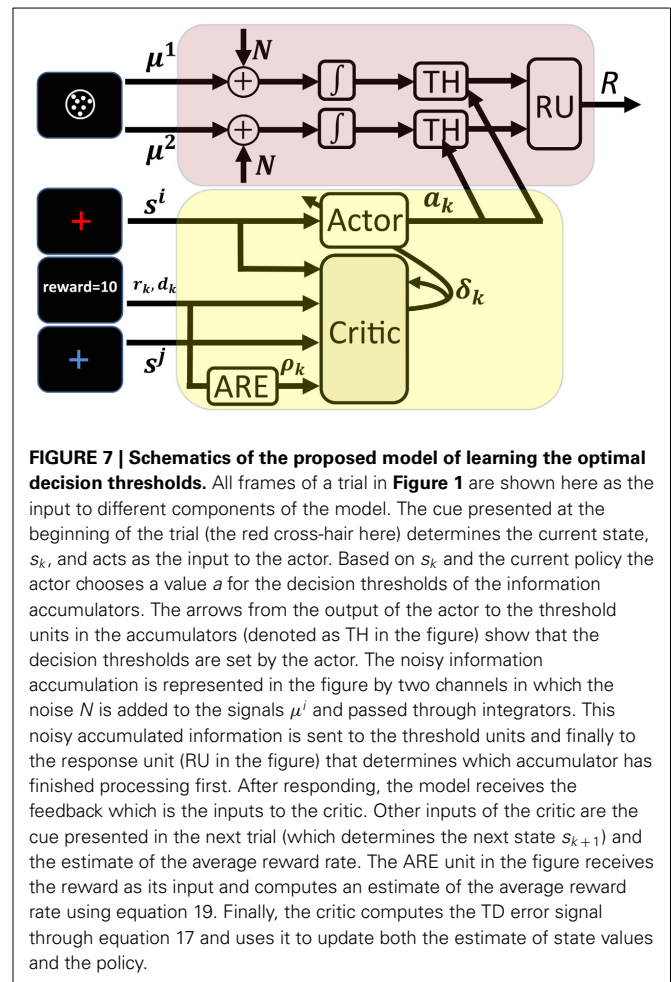
where $\hat{\rho}_k$ is an estimate of the average reward rate defined in Equation 8 at time step k (in the next section, we explain how this estimate can be computed).

Equation 15 together with the update Equation 14 provide an algorithm for learning the state values for a given policy π in an SMDP. However, they do not specify how the optimal policy can be learned. In the next section, we explain a method for learning the optimal policy based on the TD learning algorithm.

5.2. A MODEL OF LEARNING OPTIMAL DECISION THRESHOLDS

In this section, we propose a model for learning the optimal decision thresholds. This model is based on the TD learning algorithm. The schematics of the model is shown in **Figure 7**. The frames shown on the left of this figure as the inputs to different parts of the model, are exactly those that were shown in **Figure 1** (these are the frames shown in one trial of the task).

The model consists of two units: an *information accumulation unit* and an *actor-critic unit*. The information accumulation unit is responsible for processing the stimulus and selecting the appropriate response in each trial. The stimulus presented in each trial is considered as the input to this unit and the selected response (denoted as R in the figure) is the output of the unit. This unit is an independent race model. As we discussed in section 3.2, in the independent race model, the correct and incorrect responses

FIGURE 7 | Schematics of the proposed model of learning the optimal decision thresholds. All frames of a trial in **Figure 1** are shown here as the input to different components of the model. The cue presented at the beginning of the trial (the red cross-hair here) determines the current state, s_k, and acts as the input to the actor. Based on s_k and the current policy the actor chooses a value a for the decision thresholds of the information accumulators. The arrows from the output of the actor to the threshold units (denoted as TH in the figure) show that the decision thresholds are set by the actor. The noisy information accumulation is represented in the figure by two channels in which the noise N is added to the signals μ^i and passed through integrators. This noisy accumulated information is sent to the threshold units and finally to the response unit (RU in the figure) that determines which accumulator has finished processing first. After responding, the model receives the feedback which is the inputs to the critic. Other inputs of the critic are the cue presented in the next trial (which determines the next state s_{k+1}) and the estimate of the average reward rate. The ARE unit in the figure receives the reward as its input and computes an estimate of the average reward rate using equation 19. Finally, the critic computes the TD error signal through equation 17 and uses it to update both the estimate of state values and the policy.

each have a separate accumulator (modeled as a diffusion process) and assumes that whenever the accumulated information reaches one of the thresholds (TH in the figure), the subject will respond. The unit named RU in the figure, simply determines which of the accumulators has won the race and so determines the response.

The speed and accuracy are controlled by the value of the decision threshold in the information accumulation unit. The value of this parameter is set by the other unit of the model, the actor-critic unit (in the figure, the arrows from the output of the actor-critic unit to the threshold units are intended to show this) and so this unit is responsible for learning to solve the optimal SAT problem. Actor-critic architecture is one the most popular TD learning algorithms. Specifically, among several TD based algorithms proposed in the RL literature, the actor-critic algorithm has received significant attention in the computational neuroscience literature. This is because different components of this model mapped nicely to the anatomy of basal ganglia, a brain circuit know to be involved in many motor and cognitive functions (Barto, 1995; Doya, 2000; Frank, 2005, 2006; Bogacz and Larsen, 2011).

The actor-critic architecture, as its name suggests, consists of two units: an actor unit and a critic unit. The actor unit has a representation of the current policy and in each state selects an action based on this policy. In the SMDP model of our

experimental design, the state is determined by the cue presented at the beginning of a trial and the action is the value of the decision threshold for that trial. Thus, in our model the cue presented at the beginning of a trial is the input to the actor unit while its output is the decision threshold for that trial (see **Figure 7**). In each trial, after the presentation of the cue the actor sets the decision thresholds of the information accumulation units. After the presentation of the stimulus, the information accumulation unit selects a response based on which of its accumulators has reached its threshold sooner. Based on the selected response and the trial condition, a reward is presented to the subject and the next trial starts after a while.

At the moment that the cue of the next trial is presented, the critic unit plays its role. The role of this unit, as its name implies, is to criticize the action taken by the actor in a trial. In our model, the critic evaluates if the chosen decision threshold in a trial leads to better or worse than expected performance in that trial. The critic does this by computing the TD error in that trial. To see how the TD error can be employed to evaluate the actions taken by the actor, let us consider an experimental design with only one condition (like the first example in section 3.3). In this situation, $s^j = s^i = C_1$ and so the terms $\hat{V}^k_\pi(s^j)$ and $\hat{V}^k_\pi(s^i)$ in Equation 15 cancel out each other and the TD error is reduced to $\delta_k = r_k - \hat{\rho}_k \cdot d_k$. If the TD error δ_k is positive for a trial, it means that the amount of reward received in that trial, r_k, exceeded the cost spent on that trial which is $\hat{\rho}_k \cdot d_k$. The term $\hat{\rho}_k \cdot d_k$ is considered as the cost spent on the trial because when the subject is spending time d_k to receive the reward r_k, she is actually losing the opportunity of spending this time on other trials that on average lead to $\hat{\rho}_k$ of reward per unit of time. Thus, if the actor has chosen a value for the decision threshold for a trial and the TD error for that trial is positive, it means that decision threshold has led to better than expected performance. Similarly, negative TD error means worse than expected performance. In actor-critic algorithm, this feature of the TD error is used to improve the policy: if in a trial the actor takes an action that leads to positive TD error, the probability of taking that action next time increases. Similarly, the probability of taking actions that lead to negative TD error decreases. This way the policy is improved (at least probabilistically) after each trial till it finally converges to the optimal policy.

In the general experimental design in which there could be more than one condition, to compute the TD error, the critic needs to have an estimate of the state values. Thus, in addition to improving the policy, the TD error computed by the critic is used to estimate these values using Equation 14 (in **Figure 7** the arrow that goes from the output of the critic back to it is intended to show this).

In sum, in our model, after the presentation of the cue in the next trial [which is equivalent to transition to the new state in the corresponding SMDP (see section 4.3)], the critic computes the TD error which is used both to improve the policy and estimate the state values.

To calculate the TD error using Equation 15, the critic needs to know the current state s^i, the next state s^j, the reward r_k, the dwell time d_k and the estimate of the average reward rate $\hat{\rho}_k$. Thus, in **Figure 7** the frames corresponding to the current trial (the red cross-hair), the reward, the next trial (the blue cross-hair) and

also the output of the average reward rate estimator unit (ARE in the figure) are shown as the inputs to the critic unit.

So far, we have explained the role of different units of the model in processing the stimulus and selecting response, setting the decision thresholds and improving the policy. To complete the model, three issues should be addressed and the rest of this section is devoted to them: first, how the policy is represented in the actor, second, how the policy can be improved using the TD error and third, how the average reward rate can be estimated.

The first two problems are tightly related and so we explained them first. When the actor-critic algorithm is used in discrete action space problems, the policy can be represented as a probability mass function with one probability value for each action in the action space. When an action is taken by the actor, the probability of taking it in the next trial will be increased or decreased proportional to the TD error, and the probability of taking other actions will be normalized accordingly. However, when the action space is continuous (as is the case in our model), the policy takes the form of a probability density function, and so updating it using the method for discrete action spaces is not feasible anymore. Another problem associated with continuous action spaces is action selection: even if we are able to fully specify the policy $\pi(\cdot)$, then how should the actions be selected? Several methods have been proposed to address these problems. The comparison between these methods is outside the scope of this paper. Here, we use a slight modification of a simple algorithm proposed by Gullapalli (1990). In this algorithm, the policy is represented by a Gaussian distribution. For a 1-dimensional action space the policy takes the following form:

$$\pi\left(s_k = s^i, a_k = a\right) = \frac{1}{\sqrt{2\pi} \cdot v_i} \cdot \exp\left(-\frac{(a - m_i)^2}{2(v_i)^2}\right) \quad (16)$$

where $\pi(s_k = s^i, a_k = a)$ specifies the likelihood of choosing the value a as the action in state s^i, and m_i and v_i are the mean and standard deviation of the Gaussian distribution representing the policy in state s^i. An advantage of representing the policy as a parametric distribution is that during learning, we just need to update the parameters. In other words, the problem of updating the policy reduces to the problem of updating its parameters. The Gaussian distribution in Equation 16 has only two parameters: m and v. Gullapalli (1990) suggested the following updating rule for the parameter m:

$$m_i(k + 1) \leftarrow m_i(k) + \alpha_m \cdot \triangle m_i(k)$$

$$\triangle m_i(k) = \delta_k \cdot \left(\frac{a_k - m_i(k)}{(v_i(k))^2}\right) \quad (17)$$

In Gullapalli's algorithm v is also updated but for sake of simplicity we do not consider updating it here (v remains constant during learning).

In the actor-critic implementation of this algorithm, in each state the actor draws a sample from the distribution Equation 16 and takes it as the action. After receiving reward the critic computes the TD error signal. This signal is then used to update the policy parameter using Equation 17. For appropriate choice

of learning rates, this algorithm will eventually converge to the optimal policy.

The parameter v can be considered as the exploration-exploitation parameter. For small values of this parameter, the Gaussian distribution is highly concentrated around its mean, m, and so most of the actions (which are random samples from this distribution) will be close to the mean. In this case, the algorithm cannot explore the action space enough. On the other hand, for large values of v, many of the actions will be exploratory. In this case, even if the algorithm finds the optimal value of m, many of the selected actions will still be suboptimal. One way to balance between exploration and exploitation is to start the algorithm with large values of v and decrease its value gradually during learning.

Now, we turn to the third problem mentioned above. In Equation 15, $\hat{\rho}_k$ is the estimate of the global average reward rate. This signal is estimated by a linear filter named ARE in **Figure 7**. Before explaining how this unit works, we should clarify a point. Both the actor and the critic units work at discrete time steps. Specifically, although we can implement them as continuous time systems, they only do their computations at either the beginning or at the end of a trial. Therefore, all signals computed in these two units are indexed by the discrete index k. However, to estimate the average reward rate, the ARE unit needs to work in continuous time and so its input and output are functions of time t. Specifically, the input to this unit is the signal $U(t) = \sum_k r_k \cdot \underline{\delta}(t - t_k)$ where $\underline{\delta}(\cdot)$ is the Dirac delta function, r_k is the reward received in the k^{th} trial and t_k is the time at which this reward was received. It is assumed that the rewards are delivered at the time of state transitions and so $t_k = \sum_{j=1}^{k} d_j$ with d_j being the dwell time in the j^{th} trial. The signal $U(t)$ is the train of impulses created by delivery of rewards. The ARE unit acts as a linear filter on its input and so its output is computed as follows:

$$\frac{d\hat{\rho}(t)}{dt} = -\alpha_{\hat{\rho}} \cdot \hat{\rho}(t) + U(t) \qquad (18)$$

The signal $\hat{\rho}(t)$ is the estimate of the average reward rate at time t. To compute the TD error, the critic uses the value of this signal at the end of each trial, that is at times t_k. It is easy to show that:

$$\hat{\rho}_k = \hat{\rho}(t_k) = \left(\hat{\rho}_{k-1} + \alpha_{\hat{\rho}} \cdot r_{k-1}\right) \cdot e^{-\alpha_{\hat{\rho}} \cdot d_k} \qquad (19)$$

In sum, in one trial of the experiment different components of the model work in the following way: after the presentation of the cue, the actor draws a random sample from the Gaussian distribution in Eqaution 16 and sets the thresholds of the accumulators equal to this value. After the presentation of the stimulus, the accumulators race till one of them reaches its threshold and selects a response. Based on the response and the trial condition a reward is delivered and the next trial starts. After the presentation of the cue at the beginning of the next trial, the critic computes the TD error using Equation 15. This error signal is then used to update both the estimate of state values (using Equation 14) and the policy (using Equation 17). The actor then selects a new threshold for the new trial based on the presented cue.

6. SIMULATION RESULTS

In this section, we analyze the performance of the proposed model in two simulations. The first simulation corresponds to the first example given in section 3.3 in which no cue is presented at the beginning of trials and so there is only one state in the corresponding SMDP. The second simulation corresponds to the second example in that section in which there are two trial conditions and each of them is associated with a specific cue presented at the beginning of each trial and so the corresponding SMDP has two states.

6.1. SIMULATION 1: ONE CONDITION WITH NO CUE

In the experimental design considered in this simulation, no cue is presented at the beginning of trials. There could be one or more than one conditions in the task but all conditions are intermixed and the subject does not know the number of conditions in the task. As explained before, because there is no cue, the subject will treat all the conditions the same and so even if there are more than one condition she will set one decision threshold for all trials. For this simulation, we use the parameters values given in **Table 2** with $T^{RSI} = 0$.

We use the Gaussian policy specified in Equation 16 with $v^2 = 2.25 \times 10^{-4}$. All other parameters being fixed, the average reward rate will be only a function of the mean of the Gaussian, m. The average reward rate as a function of m is shown in **Figure 8**. Supplementary Material explains how this function is computed.

The maximum reward rate is equal to $RR_{max} = 2.5812$ which is obtained when $m = 0.095$. Here, we investigate the performance of the model in learning this optimal value of m.

The results of 20 simulations of the model are captured in **Figure 9**. The learning rates for these simulations are: $\alpha_m = 0.0002$, $\alpha_{\hat{\rho}} = 0.001$. In the top panels of this figure, thin gray lines correspond to the performance in different simulations and the dark thick line is the average among all simulations. The thick red lines show the optimal values. The left top panel shows the
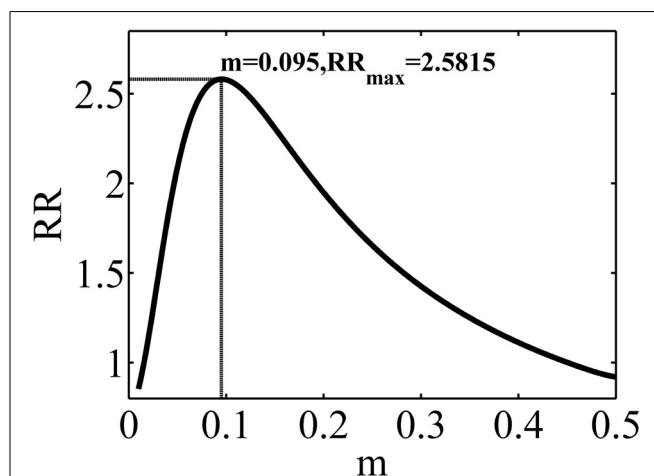
FIGURE 8 | The average reward rate in the experiment considered in simulation 1, as a function of *m*, the mean of the Gaussian policy. This function is computed using the method described in Supplementary Material.
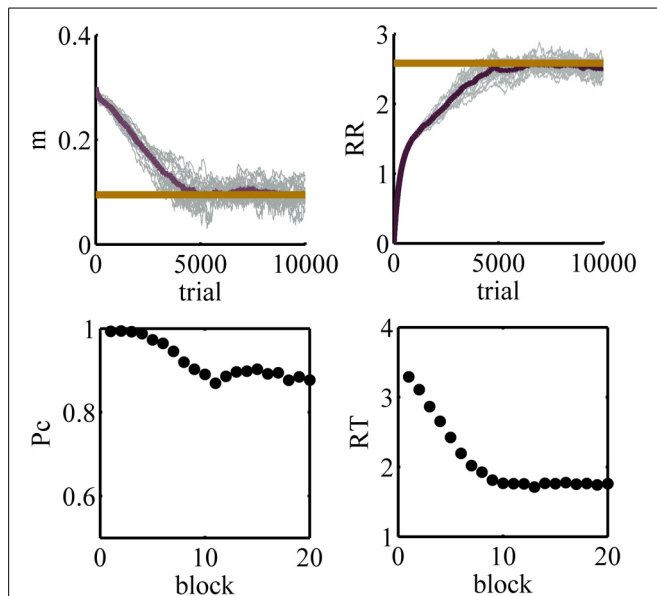
FIGURE 9 | Results of simulation 1. Top left: the value of the parameter *m* as a function of trial number during learning. **Top right:** the estimated average reward rate, $\hat{\rho}_k$, as a function of trial. In the top panels, the tick red line shows the optimal value. **Bottom left:** the accuracy as a function of block. **Bottom right:** the mean reaction time as a function of block. In the bottom panels, a block is defined as 500 trials. As learning progresses, the mean of the policy decreases and so the subject chooses lower decision thresholds more often. This leads to lower accuracy but faster responses.

FIGURE 10 | The average reward rate as a function of the mean of the Gaussian policy in state 1,m_1, and in state 2, m_2. This function is computed using the method described in Supplementary Materials.

value of the parameter *m* as a function of trial and the right top panel shows the estimated average reward rate ($\hat{\rho}_k$ in Equation 19) as a function of trial. The left and right bottom panels show the accuracy and the mean reaction time averaged over all 20 simulations, respectively, as functions of block number where a block is defined arbitrarily as 500 trials.

At the beginning of learning, the value of the parameter *m* is high and so high values of the decision threshold are chosen more often. Thus, at this stage of learning the accuracy is high and the mean reaction times are also high. In other words, the model is too much conservative and so it cannot achieve the maximum average reward rate. Throughout learning, the model gradually learns to lower the value of the parameter *m*. Finally, at about trial 5800 or so the model finds the optimal value of this parameter. As can be seen in the right top panel, the initial estimation of average reward rate is zero and during learning it approaches the optimal value and finally asymptotes at the optimal value.

The learning of the model may seem slow. This raises the question of whether human subjects are also so slow or if they can find the optimal threshold faster. In a recent study, Balci et al. (2011) investigated human subjects' performance in an experimental design similar to what was used in simulation 1. Their results show that, on average, subjects achieve the optimal performance after about 10 sessions of training (Figures 3a and 8 of Balci et al., 2011). Thus, the learning speed of our model is close to human subjects. Then, the next question is why both the algorithm and the subjects learn slowly. The main reason for this slow learning is a high amount of noise in the function that the agent
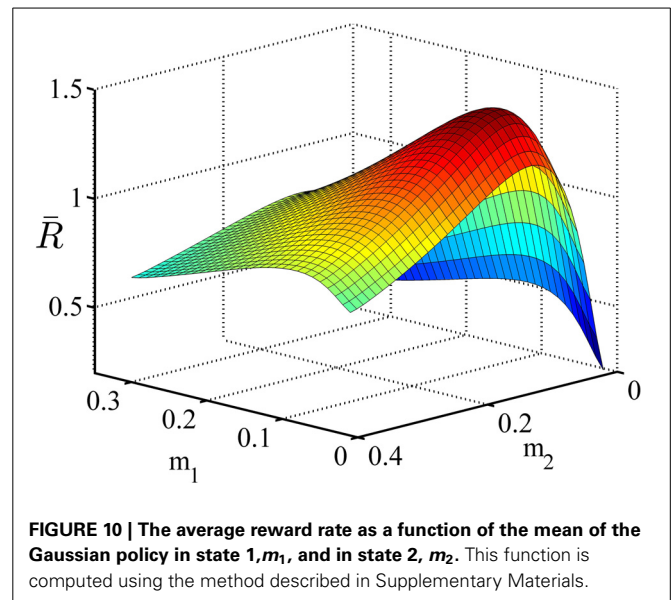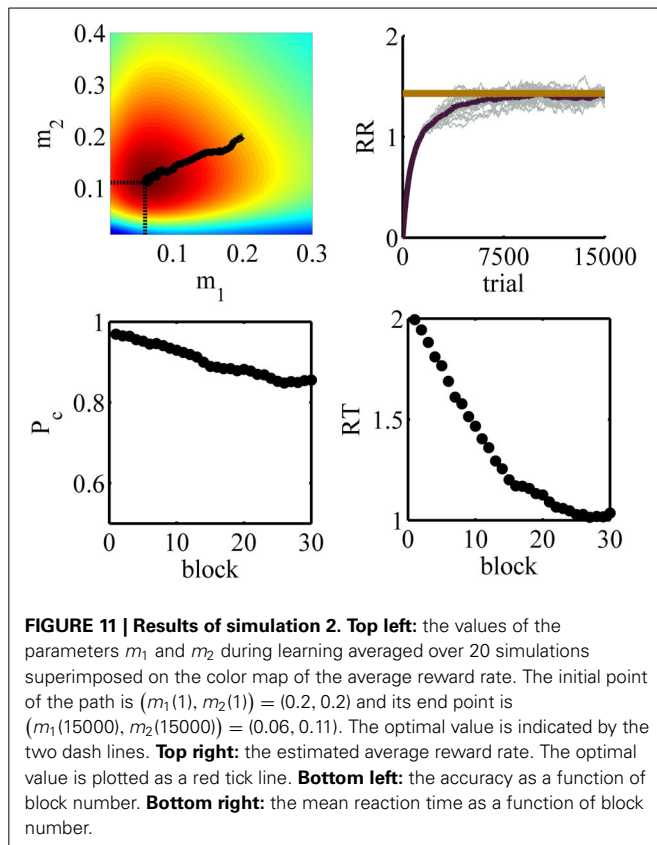
is trying to maximize. For a single value of the decision threshold, the variance of the reaction time could be high. Also, the accuracy could be a value significantly less than one. Therefore, even if the subject keeps her threshold at a fixed value for several trials, the samples of the average reward rate obtained in each trial would be very noisy and it takes a long time before the subject can achieve a reliable estimate of it.

6.2. SIMULATION 2: TWO CONDITIONS WITH CUE

In this simulation, we analyze the performance of the model in the experimental design explained in example 2 of section 3.3. We used the parameters given in **Table 3**. The policy in each state is represented by a Gaussian distribution with $v_1^2 = v_2^2 = 0.015$. The average reward rate as a function of the means of these distributions is shown in **Figure 10**. We used the method explained in Supplementary Material to plot this function. As we see in this figure, the maximum reward rate equals $RR_{max} = 1.43$ which is obtained when $m_1 = 0.07$ and $m_2 = 0.12$.

The model was simulated in this task for 20 times. The learning rate parameters in the model were $\alpha_m = 0.0005$, $\alpha_c = 0.075$ and $\alpha_{\hat{\rho}} = 0.001$. We have chosen a larger learning rate for the critic than the actor to make sure that the critic learns faster. This is because the critic provides the TD error signal necessary for updating both the critic and the actor and so the policy should not be updated a lot before the critic learns the state values.

The simulation results are depicted in **Figure 11**. The panels in this figure correspond to those in **Figure 9**. The top right panel shows the estimated average reward as a function of trial number. As seen, this function asymptotes in an optimal value after about 8000 trials or so. The top right panel shows the upper view of the average reward rate function shown in **Figure 10**. The black tick path superimposed on this figure shows the points $(m_1(k), m_2(k))$ averaged over 20 simulations, with $k = 1, \cdots , 15000$ being the trial number. This curve shows that at the beginning of learning the agent sets the mean of its policy in condition one and two at $m_1 = 0.2$ and $m_2 = 0.2$, respectively $((m_1(1), m_2(1)) =$

FIGURE 11 | Results of simulation 2. Top left: the values of the parameters m_1 and m_2 during learning averaged over 20 simulations superimposed on the color map of the average reward rate. The initial point of the path is $(m_1(1), m_2(1)) = (0.2, 0.2)$ and its end point is $(m_1(15000), m_2(15000)) = (0.06, 0.11)$. The optimal value is indicated by the two dash lines. **Top right:** the estimated average reward rate. The optimal value is plotted as a red tick line. **Bottom left:** the accuracy as a function of block number. **Bottom right:** the mean reaction time as a function of block number.

$(0.2, 0.2)$ is the starting point of the black path). With extensive learning, the agent learns the optimal value of the thresholds $((m_1(15000), m_2(15000)) = (0.06, 0.11)$ is the end point of the black path). An interesting point is that this curve shows that on average the learning algorithm takes the shortest path from the starting values to the optimal values of the parameters. Finally, similar to simulation 1, the bottom panels show that the algorithm learns to choose less conservative values of the decision threshold which leads to less accurate but faster responses.

In the previous two simulations, the initial values of the thresholds were higher than the optimal value. We also performed another simulation with the same parameters as simulation 1 but with a lower than optimal initial value of the threshold. Due to space limitation, we do not present the full details of this simulation. It suffices to mention here that the model could learn the optimal value of the threshold in this situation and its performance was at the same level of simulation 1.

7. DISCUSSION

In this paper, we suggested a theoretical framework to answer the question of how animals learn to set the decision threshold to maximize the average reward rate. We considered an experimental design in which trials from different conditions are intermixed. A cue associated with each condition is presented at the beginning of each trial and indicates which condition this trial comes from. We derived the expression for the average reward rate in this experiment and investigated the properties of this function and showed that to achieve the optimal average reward rate, the

subject has to set different decision thresholds for different conditions. We, then, proposed an SMDP model of the experiment in which each condition is modeled as a state of the environment, decision thresholds are actions and the time spent on each trial is the dwell time in each state. In this way, the problem of learning the optimal decision thresholds becomes the problem of learning the optimal action in each state of the corresponding SMDP. Finally, we proposed a model to solve this problem. In the proposed model, an independent race architecture is responsible for processing the stimulus and selecting responses while an actor-critic architecture learns the optimal value of the decision thresholds.

In the first set of simulation, we considered an experiment in which there is no cue at the beginning of each trial and so there is only one state in the corresponding SMDP. Simen et al. (2006) have proposed a model for learning the optimal decision threshold in this situation. In their model, the decision threshold at time t is $a(t) = \max (0, a_{max} - w \cdot r(t))$ where $r(t)$ is the current estimation of the reward rate. To assure that the algorithm converges to the optimal value of $a(t)$, the two constants a_{max} and w should be chosen such that the line $a_{max} - w \cdot a$ passes through the maximum of the function $R(a)$, the reward rate as a function of threshold a (see **Figure 3** for an example of this function). Notice that for each set of task parameters the function $R(a)$ would be different and so different values of a_{max} and w will assure the convergence of the algorithm. In the simulations, the authors assumed that the subjects have learned the optimal values of these parameters through practice under different trial conditions. The model then predicts fast adaptation of the decision threshold for a well-trained subject. The model proposed in this paper explains slow learning of the optimal decision threshold for an untrained subject. Further research is necessary to see how the two approaches can be combined to develop a model of both slow learning of untrained subjects and fast threshold adaptation of well-trained subjects.

In our simulations, we assumed that the drift coefficients remain constant during learning. As a result, when the threshold decreases the accuracy also decreases [see bottom left panels of **Figures 9**, **11**]. However, this may not always be the case. For example, in Balci et al. (2011) the estimated drift coefficient increased with practice. The effect of this increase in drift coefficient and the decrease in the threshold was such that the subjects' accuracy remained constant while the reaction time decreased with practice. It would be interesting to investigate the behavior of the proposed model in this situation and more generally when the task parameters change during learning. It can be imagined, though, that as long as the task parameters do not change very quickly, the model will still be able to learn the optimal thresholds.

The SMDP framework has been utilized previously in the animal learning literature to model the rate of responding in free-operant tasks (Niv, 2007; Niv et al., 2007). A rat placed in an operant chamber can choose to perform one of the several possible actions (nose poking, lever pressing, etc.). In addition, the rat may choose to perform different actions at different rates. Faster responding has the possible benefit of obtaining more reward but it is also associated with higher costs (e.g., energy cost, cognitive load and so on). Niv (2007); Niv et al. (2007) proposed a

normative account of how fast each action should be taken to achieve an optimal balance between the benefits of behaving fast and its costs. In this model, taking each action incurs a rate-dependent cost and it is assumed that the rat is trying to maximize the average reward rate. Like Niv's model, the model we proposed in this paper learns to optimally balance between the benefits and costs of acting fast. The benefit of acting fast in both models is to be able to experience more trials. The cost in our framework, however, is having less accuracy. This cost is due to the constraints that the sequential sampling model imposes to the relationship between the speed and accuracy which are in turn imposed by the noise in the stimulus.

One feature of our model is that it suggests a way to integrate the theories of perceptual decision making and reinforcement learning. Traditionally, these theories have been developed separately (see Bogacz and Gurney, 2007 and Bogacz and Larsen, 2011 for a discussion of this matter). Theories of perceptual decision making deal with situations in which the subject should process a noisy stimulus and select the appropriate response based on a known stimulus-response mapping. The reinforcement learning theories, on the other hand, deal with situations in which the stimulus is easily detectable but the subject should learn to take optimal actions in response to each stimulus. In our model, the cue presented in each trial is the easily detectable stimulus for the reinforcement learning (the actor-critic) unit. The role of this unit is to learn the optimal mapping between the cues and the decision thresholds which form the action space in the corresponding SMDP. The noisy stimulus in each trial, on the other hand, is processed by the independent race unit. By this division of labor, the model benefits from the strengths of both sets of theories.

In this respect, our model is in line with the recent effort in integrating these two sets of theories (Bogacz and Gurney, 2007; Dayan and Daw, 2008; Law and Gold, 2009; Rao, 2010; Bogacz and Larsen, 2011; Shenoy and Yu, 2011; Ratcliff and Frank, 2012). Bogacz and Larsen (2011) proposed a computational model of basal ganglia that is capable of learning the optimal stimulus-response mapping when the stimulus is noisy. This model is basically an actor-critic architecture in which the actor is a variant of the sequential sampling models. The critic, crudely speaking, provides the error signal necessary for learning the weights between the sensory units and the information accumulators in the actor. By learning these weights, the model learns the correct stimulus-response mapping. However, because the model is developed in the Markov decision process framework, it cannot solve the problem of optimal balance between speed and accuracy.

Rao (2010) proposed a model in which the perceptual decision making problem was cast as action selection in a partially observable Markov decision process (POMDP). In this model, each stimulus is considered as a state of the environment. The subject, however, does not know the state and instead can only make noisy observations of it at discrete time steps. The subject starts with a prior belief about the state and after each observation she updates her belief using the Bayes rule. At each time step, based on her current belief about the state, the subject can either choose one of the responses or make another observation. Using temporal difference learning, the model learns the optimal mapping between the current belief and these actions. Since the model was not developed to solve the optimal speed-accuracy problem, the cost of time is considered as an arbitrary constant in the model (for example -1 for each time-step that the response is not selected). In contrast, in our model the cost of time is proportional to the average reward rate which in turn depends on the task parameters in all conditions of the task. Further research is needed to compare the two models.

One question that remains open is how the brain performs average reward reinforcement learning. The essential part of this algorithm is the computation of the temporal difference error and this in turn needs an estimate of the average reward rate. Thus, the important question is how the brain estimates the average reward rate. Niv et al. (2007) suggested that the average reward rate is coded as tonic dopamine in the brain. This suggestion is based on the observation that higher levels of tonic dopamine is associated with higher response rate and vice versa (see e.g., Salamone and Correa, 2002). One interesting future line of research, then, would be to use model based fMRI technique to investigate the relationship between the average reward rate signal computed in our model and tonic activity (the activity before the delivery of reward) of brain areas previously associated with reward prediction (e.g., striatum).

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: http://www.frontiersin.org/journal/10.3389/fnins.2014.00101/abstract

REFERENCES

Balci, F., Simen, P., Niyogi, R., Saxe, A., Hughes, J. A., Holmes, P., et al. (2011). Acquisition of decision making criteria: reward rate ultimately beats accuracy. *Attent. Percept. Psychophys.* 73, 640–657. doi: 10.3758/s13414-010-0049-7

Barto, A. G. (1995). "Adaptive critics and the basal ganglia," in *Models of Information Processing in the Basal Ganglia*, eds J. Houk, J. L. Davis, and D. G. Beiser, (Cambridge, MA: MIT Press), 215–232.

Bertsekas, D. P., and Tsitsiklis, J. N. (1996). *Neuro-Dynamic Programming*. Belmont, MA: Athena Scientific.

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., and Cohen, J. D. (2006). The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol. Rev.* 113, 700–765. doi: 10.1037/0033-295X.113.4.700

Bogacz, R., and Gurney, K. (2007). The basal ganglia and cortex implement optimal decision making between alternative actions. *Neural Comput.* 19, 442–477. doi: 10.1162/neco.2007.19.2.442

Bogacz, R., and Larsen, T. (2011). Integration of reinforcement learning and optimal decision-making theories of the basal ganglia. *Neural Comput.* 23, 817–851. doi: 10.1162/NECO_a_00103

Bogacz, R., Wagenmakers, E.-J., Forstmann, B. U., and Nieuwenhuis, S. (2010). The neural basis of the speed-accuracy tradeoff. *Trends Neurosci.* 33, 10–16. doi: 10.1016/j.tins.2009.09.002

Busemeyer, J. R., and Townsend, J. T. (1993). Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. *Psychol. Rev.* 100, 432–459. doi: 10.1037/0033-295X.100.3.432

Cox, D. R., and Miller, H. D. (1965). *The Theory of Stochastic Processes*. London: Methuen.

Das, T. K., Gosavi, A., Mahadevan, S., and Marchalleck, N. (1999). Solving semi-markov decision problems using average reward reinforcement learning. *Manag. Sci.* 45, 560–574. doi: 10.1287/mnsc.45.4.560

Daw, N. D. (2003). *Reinforcement Learning Models of the Dopamine System and Their Behavioral Implications*. Unpublished Doctoral dissertation, Carnegie Mellon University.

Dayan, P., and Daw, N. D. (2008). Decision theory, reinforcement learning, and the brain. *Cogn. Affect. Behav. Neurosci.* 8, 429–453. doi: 10.3758/CABN.8.4.429

Doya, K. (2000). Complementary roles of basal ganglia and cerebellum in learning and motor control. *Curr. Opin. Neurobiol.* 10, 732–739. doi: 10.1016/S0959-4388(00)00153-7

Forstmann, B. U., Anwander, A., Schfer, A., Neumann, J., Brown, S., Wagenmakers, E.-J., et al. (2010). Cortico-striatal connections predict control over speed and accuracy in perceptual decision making. *Proc. Natl. Acad. Sci. U.S.A.* 107, 15916–15920. doi: 10.1073/pnas.1004932107

Frank, M. J. (2005). Dynamic dopamine modulation in the basal ganglia: a neurocomputational account of cognitive deficits in medicated and nonmedicated parkinsonism. *J. Cogn. Neurosci.* 17, 51–72. doi: 10.1162/0898929052880093

Frank, M. J. (2006). Hold your horses: a dynamic computational role for the subthalamic nucleus in decision making. *Neural Netw.* 19, 1120–1136. doi: 10.1016/j.neunet.2006.03.006

Gold, J. I., and Shadlen, M. N. (2002). Banburismus and the brain: decoding the relationship between sensory stimuli, decisions, and reward. *Neuron* 36, 299–308. doi: 10.1016/S0896-6273(02)00971-6

Gullapalli, V. (1990). A stochastic reinforcement learning algorithm for learning real-valued functions. *Neural Netw.* 3, 671–692. doi: 10.1016/0893-6080(90)90056-Q

Ivanoff, J., Branning, P., and Marois, R. (2008). fMRI evidence for a dual process account of the speed-accuracy tradeoff in decision-making. *PLoS ONE* 3(7):e2635. doi: 10.1371/journal.pone.0002635

Law, C.-T., and Gold, J. I. (2009). Reinforcement learning can account for associative and perceptual learning on a visual-decision task. *Nat. Neurosci.* 12, 655–663. doi: 10.1038/nn.2304

Link, S. W., and Heath, R. A. (1975). A sequential theory of psychological discrimination. *Psychometrika* 40, 77–105. doi: 10.1007/BF02291481

Luce, R. D. (1986). *Response Times: Their Role in Inferring Elementary Mental Organization.* New York, NY: Oxford University Press.

Montague, P. R., Dayan, P., and Sejnowski, T. J. (1996). A framework for mesencephalic dopamine systems based on predictive hebbian learning. *J. Neurosci.* 16, 1936–1947.

Niv, Y. (2007). *The Effects of Motivation on Habitual Instrumental Behavior.* Unpublished Doctoral dissertation, The Hebrew University of Jerusalem.

Niv, Y., Daw, N. D., Joel, D., and Dayan, P. (2007). Tonic dopamine: opportunity costs and the control of response vigor. *Psychopharmacology* 191, 507–520. doi: 10.1007/s00213-006-0502-4

Palmer, J., Huk, A. C., and Shadlen, M. N. (2005). The effect of stimulus strength on the speed and accuracy of a perceptual decision. *J. Vis.* 5:1. doi: 10.1167/5.5.1

Rao, R. P. N. (2010). Decision making under uncertainty: a neural model based on partially observable markov decision processes. *Front. Comput. Neurosci.* 4:146. doi: 10.3389/fncom.2010.00146

Ratcliff, R. (1978). A theory of memory retrieval. *Psychol. Review* 85, 59–108. doi: 10.1037/0033-295X.85.2.59

Ratcliff, R. (2002). A diffusion model account of response time and accuracy in a brightness discrimination task: fitting real data and failing to fit fake but plausible data. *Psychon. Bullet. Review* 9, 278–291. doi: 10.3758/BF03196283

Ratcliff, R., and Frank, M. J. (2012). Reinforcement-based decision making in corticostriatal circuits: mutual constraints by neurocomputational and diffusion models. *Neural Comput.* 24, 1186–1229. doi: 10.1162/NECO_a_00270

Ratcliff, R., and McKoon, G. (2007). The diffusion decision model: theory and data for two-choice decision tasks. *Neural Comput.* 20, 873–922. doi: 10.1162/neco.2008.12-06-420

Ratcliff, R., and Smith, P. L. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychol. Rev.* 111, 333–367. doi: 10.1037/0033-295X.111.2.333

Salamone, J. D., and Correa, M. (2002). Motivational views of reinforcement: implications for understanding the behavioral functions of nucleus accumbens dopamine. *Behav. Brain Res.* 137, 3–25. doi: 10.1016/S0166-4328(02)00282-6

Schultz, W., Dayan, P., and Montague, P. R. (1997). A neural substrate of prediction and reward. *Science* 275, 1593–1599. doi: 10.1126/science.275.5306.1593

Shenoy, P., and Yu, A. J. (2011). Rational decision-making in inhibitory control. *Front. Hum. Neurosci.* 5:48. doi: 10.3389/fnhum.2011.00048

Simen, P., Cohen, J. D., and Holmes, P. (2006). Rapid decision threshold modulation by reward rate in a neural network. *Neural Netw.* 19, 1013–1026. doi: 10.1016/j.neunet.2006.05.038

Simen, P., Contreras, D., Buck, C., Hu, P., Holmes, P., and Cohen, J. D. (2009). Reward rate optimization in two-alternative decision making: empirical tests of theoretical predictions. *J. Exp. Psychol. Hum. Percept. Perform.* 35, 1865–1897. doi: 10.1037/a0016926

Smith, P. L. (2000). Stochastic dynamic models of response time and accuracy: a foundational primer. *J. Math. Psychol.* 44, 408–463. doi: 10.1006/jmps.1999.1260

Smith, P. L., and Vickers, D. (1988). The accumulator model of two-choice discrimination. *J. Math. Psychol.* 32, 135–168. doi: 10.1016/0022-2496(88)90043-0

Sutton, R., and Barto, A. G. (1998). *Reinforcement Learning: An Introduction (Adaptive Computation and Machine Learning).* Cambridge, MA:MIT Press.

Townsend, J. T., and Ashby, F. G. (1983). *The Stochastic Modeling of Elementary Psychological Processes.* Cambridge: Cambridge University Press.

Usher, M., and McClelland, J. L. (2001). The time course of perceptual choice: the leaky, competing accumulator model. *Psychol. Rev.* 108, 550–592. doi: 10.1037/0033-295X.108.3.550

Wagenmakers, E.-J., Ratcliff, R., Gomez, P., and McKoon, G. (2008). A diffusion model account of criterion shifts in the lexical decision task. *J. Mem. Lang.* 58, 140–159. doi: 10.1016/j.jml.2007.04.006

Speed accuracy trade-off under response deadlines

Hakan Karşılar[1], Patrick Simen[2], Samantha Papadakis[2] and Fuat Balcı[1]*

[1] *Department of Psychology, Koç University, Istanbul, Turkey*
[2] *Department of Neuroscience, Oberlin College, Oberlin, OH, USA*

Perceptual decision making has been successfully modeled as a process of evidence accumulation up to a threshold. In order to maximize the rewards earned for correct responses in tasks with response deadlines, participants should collapse decision thresholds dynamically during each trial so that a decision is reached before the deadline. This strategy ensures on-time responding, though at the cost of reduced accuracy, since slower decisions are based on lower thresholds and less net evidence later in a trial (compared to a constant threshold). Frazier and Yu (2008) showed that the normative rate of threshold reduction depends on deadline delays and on participants' uncertainty about these delays. Participants should start collapsing decision thresholds earlier when making decisions under shorter deadlines (for a given level of timing uncertainty) or when timing uncertainty is higher (for a given deadline). We tested these predictions using human participants in a random dot motion discrimination task. Each participant was tested in free-response, short deadline (800 ms), and long deadline conditions (1000 ms). Contrary to optimal-performance predictions, the resulting empirical function relating accuracy to response time (RT) in deadline conditions did not decline to chance level near the deadline; nor did the slight decline we typically observed relate to measures of endogenous timing uncertainty. Further, although this function did decline slightly with increasing RT, the decline was explainable by the best-fitting parameterization of Ratcliff's diffusion model (Ratcliff, 1978), whose parameters are constant within trials. Our findings suggest that at the very least, typical decision durations are too short for participants to adapt decision parameters within trials.

Keywords: response deadlines, optimality, speed-accuracy, timing uncertainty, decision making
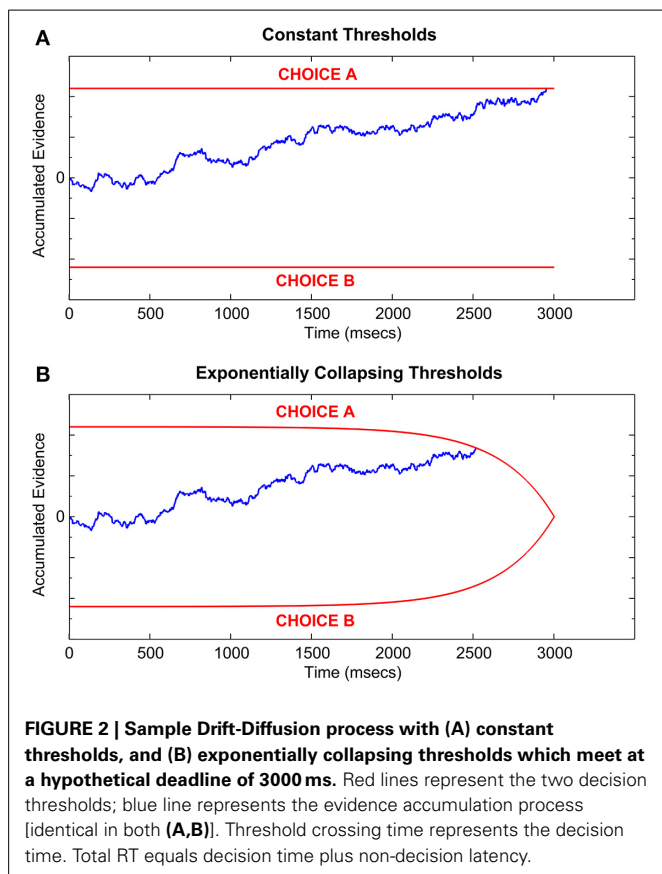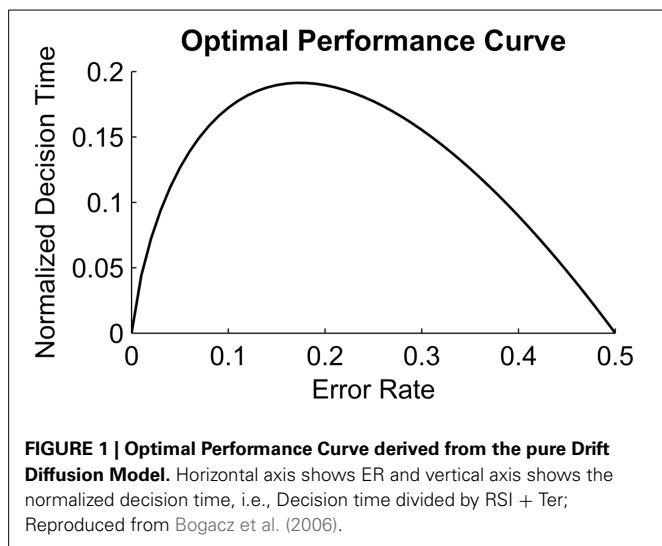
INTRODUCTION

Noisy evidence accumulation models such as the drift-diffusion model (DDM, Ratcliff, 1978, 1981, 1985, 1988, 2002) have successfully explained accuracy and RT patterns in two-alternative forced choice (2AFC) perceptual decision tasks. The DDM has also been useful in defining an optimality-based benchmark for decision making. For instance, Bogacz et al. (2006) formulated a parameter-free optimal performance curve (OPC; **Figure 1**) relating the DDM's decision speed to its accuracy in a class of 2AFC tasks. Specifically, on tasks in which the signal-to-noise ratio (SNR) stays constant within a test block and within trials, the two stimulus types are equally likely and participants are free to wait as long as they wish prior to responding. The OPC prescribes an optimal normalized decision time (DT) for a given level of accuracy in order to maximize the expected reward rate (RR) in such free-response paradigms. If the signal quality is very high, then little evidence needs to be accumulated to achieve high accuracy; conversely if there is no signal in the environment (necessarily yielding an error rate around 0.5), the decision maker should accumulate little or no evidence before making a choice. In this way, the participant can maximize the number of decisions made (trials generated) in a fixed amount of test duration. However, when the SNR is at an intermediate level, the optimal decision strategy requires accumulating more evidence

(and thus generating fewer trials) for maximizing the RR; the maximum decision time is associated with accuracy levels of roughly 0.8. Note that the OPC for 2AFC tasks was defined based on the assumptions of the reduced DDM analyzed by Bogacz et al. (2006), which lacks the between-trial variability of the core parameters found in Ratcliff's DDM.

Inherent in the formulation of the OPC is a trade-off between speed and accuracy of decisions (SAT; Wickelgren, 1977), which posits that fast responses suffer from less evidence accumulation and are thus less accurate, whereas slower responses benefit from more evidence accumulation resulting in higher accuracy at the cost of time. In formal decision making models such as the DDM, SAT is represented by a threshold parameter that determines how much evidence is accumulated in favor of each hypothesis in a 2AFC task (**Figure 2**). A higher threshold requires more evidence accumulation and thus underlies a slower response, on average, whereas a lower threshold leads to a faster response at the expense of an increased chance of errors due to noisy evidence accumulation (e.g., Ratcliff and McKoon, 2008). Research shows that, with extensive training, participants can maximize their RR by setting the optimal threshold, which defines the optimal trade-off between the speed and accuracy of their decisions (e.g., Simen et al., 2009; Balci et al., 2011b). However, behavioral studies testing for optimality in 2AFC paradigms typically do not enforce

FIGURE 1 | Optimal Performance Curve derived from the pure Drift
Diffusion Model. Horizontal axis shows ER and vertical axis shows the
normalized decision time, i.e., Decision time divided by RSI + Ter;
Reproduced from Bogacz et al. (2006).

FIGURE 2 | Sample Drift-Diffusion process with (A) constant
thresholds, and (B) exponentially collapsing thresholds which meet at
a hypothetical deadline of 3000 ms. Red lines represent the two decision
thresholds; blue line represents the evidence accumulation process
[identical in both (A,B)]. Threshold crossing time represents the decision
time. Total RT equals decision time plus non-decision latency.

hard time constraints on the decision making process (e.g., Feng
et al., 2009; Simen et al., 2009; Bogacz et al., 2010; Starns and
Ratcliff, 2010, 2012; Balci et al., 2011b), which provides a the-
oretically infinite (in reality limited by the test block duration)
amount of time to the participant before a decision must be
made.

Decisions in real life scenarios rarely enjoy such temporal
luxury for gathering evidence, but instead often need to be

terminated before a pre-specified deadline, after which no reward
can be earned (e.g., in class exams). Optimal behavior in such
settings requires the decision maker to collapse decision thresh-
olds as the deadline approaches, such that they meet when the
deadline is reached, in order to secure at least a 50% chance of
earning a reward, as opposed to a 0% chance if responding late.
In this regard, see Frazier and Yu (2008), who analyzed optimal
threshold collapse for a loss function that linearly combines an
indicator of on-time, accurate responding, the RT itself, and a
penalty for late responding. This loss function is closely related,
but not identical, to an objective function equaling the RR. As
such, the notion of time-dependent collapsing thresholds (or sim-
ilarly, time-dependent inflation of evidence accumulation rates)
has received a great deal of attention in the decision making lit-
erature (Luce, 1986; Rao, 2010; Drugowitsch et al., 2012; Thura
et al., 2012).

Two interesting hypotheses emerge from this formulation.
First, a higher level of endogenous timing uncertainty (for a
fixed deadline) requires an earlier threshold collapse, along with
a lower rate of decline (see Frazier and Yu, 2008; Figures 2A,B).
Within this formulation, endogenous timing uncertainty refers
to the trial-to-trial variability in a participant's estimates of time
intervals (Buhusi and Meck, 2005). Second, for a given level of
timing uncertainty, threshold collapsing should begin earlier for
a shorter deadline. Balci et al. (2011a) tested these previously
untested predictions in a pilot study but found little evidence
of collapsing thresholds; however, their design might not have
been optimized to investigate these predictions that might have
obscured signs of threshold collapse (e.g., not terminating the
RDM stimulus at the deadline). This study tests these predic-
tions more rigorously, and thereby elucidates the extent to which
optimal behavior in 2AFC is achievable when reward maxi-
mization entails within-trial modulation of decision thresholds.
Additionally, we aim to investigate the extent to which, if at
all, participants are successful in factoring their level of timing
uncertainties into their threshold modulation.

In order to formally define the optimal 2AFC behavior,
whether under response deadlines or not, we need mathemat-
ical models which can accurately describe accuracy along with
RT in 2AFC tasks by relying on various psychomechanistic com-
ponents underlying a complete decision making process. One
such model is the above-mentioned DDM, which conceptualizes
decision making as a bounded, noisy, evidence accumulation pro-
cess (Figure 2) in the form of an ongoing computation of the
current log-likelihood ratio of the two hypotheses under con-
sideration (Stone, 1960). At its core, the DDM is a continuous
version of the Sequential Probability Ratio Test (SPRT), which
is a statistical procedure for minimizing the number of sam-
ples necessary to decide between two hypotheses with a given
mean accuracy, as well as maximizing the likelihood of arriv-
ing at the correct hypothesis for any given number of samples
(Wald and Wolfowitz, 1948). In the formulation of the DDM,
the step time between the samples accumulated in an SPRT
becomes infinitesimal, resulting in a continuous random walk,
where the duration from the start of the evidence accumula-
tion until a threshold crossing represents the decision time (see
Stone, 1960).

The drift-diffusion process is defined by the stochastic differential equation:

$$dx = Adt + cdW \qquad (1)$$

Here, as in Bogacz et al. (2006), *x* denotes the difference between the evidence supporting two different alternatives at time *t*, *Adt* represents the average increase in *x* during the interval *dt*, and *cdW* is Gaussian distributed white noise with mean 0 and variance $c^2 dt$ (Ratcliff and Smith, 2004). When *x* crosses one of the two decision thresholds (one above the starting point, and one below it) a decision is made. This threshold crossing time represents the decision time. Within this formulation the drift rate *A* represents the average rate of the evidence accumulation, and is the slope of this random walk process. On the other hand, the noise component explains the random fluctuations in the same process and accounts for the fact that a given SNR can lead to correct decisions in some trials and errors in some others. This model is now referred to as the pure DDM (**Figure 2**; see Bogacz et al., 2006). It uses RT and accuracy data in order to describe decision performance by quantifying drift rate (*v*; rate of evidence accumulation), boundary separation (*a*; decision threshold), non-decision related latency (T_{er}), and starting point (*z*) parameters. In a more generalized version, three parameters of the DDM (*v*, *z*, and T_{er}) were made variable on a trial-by-trial basis, mainly to allow for fitting data with unequal average RT for correct and incorrect responses (Ratcliff and Rouder, 1998) and is appropriately named the extended DDM (see Bogacz et al., 2006).

The DDM has been successful in explaining RT and accuracy data in various psychophysical studies (see Voss et al., 2013 for a review) including recognition memory (Ratcliff, 1978; McKoon and Ratcliff, 2012), brightness discrimination (Ratcliff, 2002), color discrimination (Spaniol et al., 2011), and even the classification of clinical disorders (Mulder et al., 2010; White et al., 2010). Of greater relevance to this study, however, is the DDM's utilization in prescribing unique threshold parameters for RR-maximizing (i.e., optimal) performance in 2AFC tasks. As mentioned earlier, the theoretical work by Bogacz et al. (2006) has defined a closed-form RR-maximizing function that prescribes a specific average decision time for each error rate (ER), and also defines the OPC. Bogacz et al. (2010) and Simen et al. (2009) have tested the extent to which human participants are optimal in setting RR-maximizing thresholds, and have found that within a single session, thresholds were generally set too high compared to their optimal values. Balci et al. (2011b) have replicated this finding, but have also shown that this accuracy bias diminishes with practice.

Bogacz et al. (2010) and Balci et al. (2011b) argued that suboptimal performance due to favoring accuracy over reward rate (observed in their studies after a limited level of training) might be an adaptive threshold setting bias that takes into account endogenous timing uncertainty. This adaptive bias was attributed to the asymmetry (i.e., lower rate of decline in RR for thresholds higher than the optimal threshold) in the RR curves as a function of decision threshold (Bogacz et al., 2006; Figure 15), which entails that setting the threshold higher than the optimal threshold leads to a higher RR than setting it too low by the same

amount. A more adaptive response under endogenous timing uncertainty therefore entails favoring slower yet more accurate responses (Bogacz et al., 2006; Balci et al., 2011b). Balci et al's (2011a) findings suggest that participants can "monitor" their levels of uncertainty regarding temporal properties of the task, and thereby factor it into the decision process. This proposition is further supported by studies showing that humans and other animals can in fact take normative account of their timing uncertainties at both sub- and supra-second intervals in order to reach optimal performance when they make decisions based on the durations of stimuli/events (e.g., Hudson et al., 2008; Balci et al., 2009; Jazayeri and Shadlen, 2010; Simen et al., 2011; Çavdaroğlu et al., 2014; for a review see Balci et al., 2011a). Overall, these studies suggest that timing uncertainty is instrumental in shaping choice behavior and determining how much reward is earned both in temporal and non-temporal decision-making. The importance of interval timing to perceptual decision making is further emphasized by recent studies proposing possible mechanisms (e.g., gain modulation) by which temporal information processing can modulate speed-accuracy tradeoffs (e.g., Standage et al., 2011, 2013).

Endogenous timing uncertainty becomes even more relevant to optimal choice behavior in 2AFC perceptual decision making when a response deadline is explicitly introduced to the decision process. Such situations are familiar to most organisms in their natural settings, within which contextual temporal properties constantly require arriving at a decision before a stochastic deadline. For instance, correctly identifying when and how long a prey will be available in a hunting ground, as well as which prey to hunt among the alternatives ("Slow but old?," "Young but fast?") are of vital importance for a predator's survival. The optimal predator in its attempt to choose the best option should also require less and less information for arriving at a decision as the time for the prey animals to leave approaches. This strategy ensures that it catches at least one prey, though perhaps not an ideal one, instead of losing all. Moreover, it should engage in this decision process while simultaneously relying on its level of uncertainty regarding how much time it has before a choice must be made. If it is too uncertain about temporal intervals, or the time until the prey animals leave is too short, the predator should start reducing the required level of evidence earlier, and should at worst pick a random prey right before the time to leave, if it still hasn't done so. This hypothetical naturalistic scenario exemplifies the above-mentioned optimal strategy in a situation with a response deadline, which is to collapse the decision threshold such that by the time the deadline is reached, a response of at least 50% accuracy is ensured.
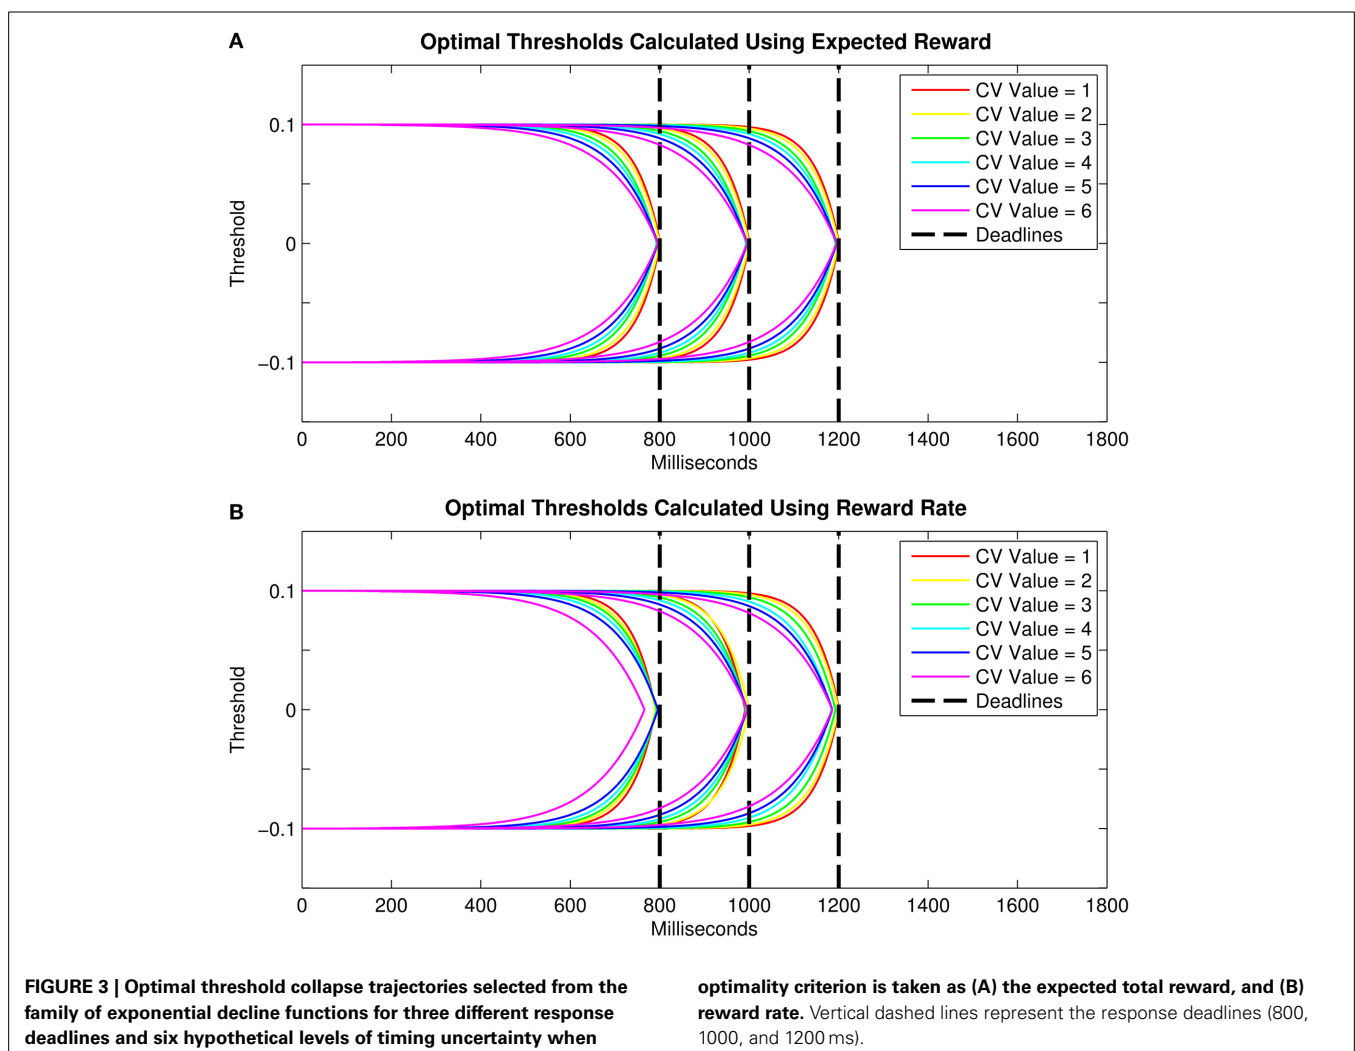
Two main hypotheses emerge under this scenario. First, for a given deadline, higher timing uncertainty makes it necessary to collapse thresholds earlier compared to lower timing uncertainty, so that the deadline is not passed by accident, ultimately resulting in an opportunity cost. Second, for a given timing uncertainty, participants need to start collapsing decision thresholds earlier for shorter deadlines, compared to longer ones. Frazier and Yu (2008) have shown that both predictions should manifest themselves with steady decline in accuracy as time approaches the deadline, which should closely parallel the presumed decline in

decision thresholds. We can quantify this time-dependent decline in thresholds by calculating accuracy levels for RTs bins of a specific size. The resulting curve formed by connecting the accuracy levels in these bins constitutes the conditional accuracy (a.k.a. Micro Speed Accuracy Trade-off) curve (Wickelgren, 1977; Luce, 1986). Since the diffusion process calculates the log-likelihood ratio of the two hypotheses, a particular accuracy level is assured by setting a particular decision threshold. When accuracy data is sorted and binned in this way, this principle should still hold for each individual RT bin. Thus, if the threshold is dynamically set lower in later time bins, then by definition this also prescribes lower accuracy for those bins (Luce, 1986).

Here, we conduct simulations in order to approximate the optimal relationship between threshold collapsing and (1) the deadline duration and (2) the level of endogenous timing uncertainty. For the collapsing thresholds we use two closed-form collapse functions: exponential and linear. **Figure 3** depicts the threshold collapsing functions (assuming exponential collapse functions) that yielded the highest number of rewards for different response deadlines (for a given level of timing uncertainty) and for different levels of endogenous timing uncertainty (for a

given deadline). As predicted by Frazier and Yu (2008), visual inspection of **Figures 3A,B** suggests that reward-maximizing threshold trajectories should nearly meet at the response deadline, and threshold collapsing should start earlier in the trial for shorter deadlines and higher levels of timing uncertainty. Our simulations showed very similar results when RR instead of "reward amount" is taken as an indicator of optimality. These results qualitatively mimicked the analytically derived functions found by Frazier and Yu (2008) for an objective function closely related to RR (see Methods).

To the best of our knowledge, the aforementioned predictions have not been directly tested by employing hard response deadlines (but see Balci et al., 2011a for description of a pilot study). Neither has the relationship of 2AFC behavior under response deadlines been empirically related to the decision maker's level of endogenous timing uncertainty. The present study fills this empirical gap. Finally, we conducted further simulations to determine whether different levels of trial-to-trial variability of the core DDM parameters that might result from the introduction of the response deadlines can explain our data without alluding to dynamic (within-trial) threshold modulation. These simulations

FIGURE 3 | Optimal threshold collapse trajectories selected from the family of exponential decline functions for three different response deadlines and six hypothetical levels of timing uncertainty when optimality criterion is taken as **(A)** the expected total reward, and **(B)** reward rate. Vertical dashed lines represent the response deadlines (800, 1000, and 1200 ms).

were necessary given that it is also possible to observe a reduction in conditional accuracy curves without any corresponding threshold modulation as suggested by Frazier and Yu (2008). Our simulations confirmed this possibility by showing that such declines in accuracy with RT in these conditional accuracy curves can emerge directly from Ratcliff's model without any within-trial collapsing of the threshold, as shown previously (Ratcliff and Rouder, 1998; Ratcliff and McKoon, 2008).

MATERIALS AND METHODS

PARTICIPANTS

Eleven adults (6 males and 5 females), aged between 18 and 24 years ($M = 20$) were recruited through announcements posted online at the daily newsletter of Koç University. One participant (male, aged 24) stopped attending experiments after the first session, and his data were discarded from all analyses. The experiment consisted of eight, daily, one-hour long sessions comprised of two Free Response (FR) sessions, four Deadlined Response (DR) sessions, and two Temporal Reproduction (TR) sessions in that order (see Procedure below). One participant missed a single DR session, and another participant missed the second TR session. The experiment was approved by the Institutional Review Panel for Human Subjects of Koç University and was in accordance with the principles of the Declaration of Helsinki. All participants provided written consent for their participation.

APPARATUS

All stimuli were presented on a 21″ LCD screen on an Apple iMac G4 computer, generated in Matlab using the Psychtoolbox Extension (Brainard, 1997; Pelli, 1997; Kleiner et al., 2007) on the SnowDots framework developed by Joshua Gold at the University of Pennsylvania. Participants sat at a distance of 58–63 cm from the screen, in a dimly lit room and provided their responses using a standard Apple iMac keyboard, and stereo noise-cancelling headphones worn throughout the experiment gave auditory feedback.

STIMULI AND PROCEDURE

Free response dot motion discrimination task

Stimuli were random dot kinematograms (see Gold and Shadlen, 2001; Shadlen and Newsome, 2001). These Random Dot Motion (RDM) stimuli consisted of a circular aperture of randomly moving white dots (3×3 pixels) on a black background, with a diameter of approximately 3 inches, centered on the screen. On each trial, 16% of the dots moved coherently in rightward or leftward direction (0 or 180 degrees respectively). The motion direction was assigned randomly with equal probability. Participants' task was to use the 'Z' or 'M' keys on the keyboard to report the direction of the coherently moving dots. Stimuli stayed on the screen until a response was given, at which moment they were terminated. Trials were separated by a response-to-stimulus (RSI) interval, sampled from a truncated exponential distribution with a mean of 2 s, a lower bound of 1 s, and an upper bound of 5.6 s. Correct responses were followed by an auditory beep indicating positive feedback, whereas no feedback was given for incorrect responses. This method of giving auditory feedback is standard in most 2AFC tasks, and has been shown to

aid acquisition (e.g., Herzog and Fahle, 1997; Seitz et al., 2006), which was also the central purpose of our FR sessions (**Figure 6**). Premature/anticipatory responses (i.e., responses less than 100 ms after the offset of the stimulus) were penalized by a 4 s time-out, following a buzzing sound. Participants earned 2 kurus (approximately 1 cent) per correct response in experimental trials (excluding practice blocks), whereas no punishment was given for incorrect responses. The cumulative number of correct responses was presented on the screen every 10 trials in font size 12 (approximately 0.7 cm height). FR session consisted of a 2-min practice block, followed by eight 5-min test blocks, and a 4-min Signal Detection (SD) block. The data from these SD blocks were not used in this study.

Deadlined dot motion discrimination tasks

The DR sessions consisted of a 2-min practice block with FR trials, followed by one 5-min experimental FR block (same as the one described above), followed by two groups of four DR blocks, each group preceded by a 2-min practice block of the corresponding deadline (see below). Stimulus types and presentation schedules in DR blocks in these DR sessions were identical to those used in FR sessions, except for the assignment of either a short (800 ms) or a long (1000 ms) deadline to every trial in the block. In these DR trials, if the participant failed to respond before the pre-specified fixed deadline, the RDM stimulus disappeared, a buzzing sound was played (indicating a "late response") and no reward was given for that trial. Otherwise, identical to the FR trials, the RDM stimulus disappeared upon a given response and a reward was given for correct responses.

After a 10 s intermission following the above-mentioned single 5-min FR block, and the 2-min practice block of DR trials, four 5-min experimental blocks with the same type of DR trials employing one of the deadlines (i.e., short or long) were presented. These blocks were followed by a 30-s intermission, after which the same order of practice and experimental blocks was presented, this time using the other deadline. Individual blocks were separated by a minimum break of 10-s, after which the participant made a button press to start the following block. The order of deadlines was randomized across the two halves of the eight DR blocks in each session. Identical with the FR sessions, two 2-min SD blocks were presented at the end of each session, and the data from these SD blocks were not used in this study.

These two hard deadlines were chosen based on the data collected from single session pilot testing with only the FR blocks. These data showed that the majority of participants' RTs ranged between 400 and 2500 ms, with a mean of 700 ms. Based on these data we chose two deadlines, the "easy" deadline of 1000 ms (on average 15% of the RTs were longer than 1000 ms; s.e.m. = 4.41) and the more "stringent" deadline of 800 ms (on average 28% of the RTs were longer than 800 ms; s.e.m. = 6.32). This way, we planned to have enough data from trials with RTs near the response deadlines. It can be argued that shorter deadlines might have made the task so difficult as to preclude strategic time-based decision-making. That said, we observed that participants sped up their free response RTs in the deadline blocks (**Figure 7**) and thus the deadline stringency was not as high as we intended during the study design. Nonetheless, the deadlines clearly exerted

an effect on speed and accuracy relative to free responding, as we demonstrate below, and the two deadline durations should have been sufficiently discriminable from each other that a differential effect on behavior could have been expected. The ratio between 800 and 1000 ms constitutes a discriminable difference for humans; given a coefficient of variation (CV; Section Temporal reproduction task – static stimuli) of 0.12, the difference is over two standard deviations for the standard duration of 800 ms (Malapani and Fairhurst, 2002). This CV value is also consistent with earlier data (see Wearden, 2003).

Temporal reproduction task – static stimuli

The TR task consisted of the presentation of a stimulus for a specific duration, after which the participant tried to reproduce the same duration as accurately as possible by holding down the space button. The stimulus used in the first TR task was a 3 × 3 inch green square, placed in the middle of the screen. Each TR trial started with a button press after which the square was presented for a specific duration. The TR session started with a practice block of 9 trials using 3 randomly ordered target durations (i.e., 1.3, 2.3, and 3.3 s) with equal frequency. After the reproduced interval on practice trials, visual feedback was given by placing an approximately 1 cm white vertical line either to the left or right of a red reference line in the middle of the screen, representing the reproduced and given durations, respectively. The offset length of the white line was proportional to the difference between given and reproduced durations, whereas its location (left vs. right) showed under- or over-reproduction, respectively.

Nine 5-min test blocks of three target durations (1, 2.12, and 4.24 s), were presented in pseudo-random order following the practice trials. No feedback was given in test trials. The amount of money earned in each block was a function of the target duration, the average of absolute deviance scores for that block, and a maximum of 2.5 Turkish Liras that could hypothetically be earned with perfect performance (i.e., mean deviance score of 0), calculated using the following formula;

$$Total\ Earnings = Maximum\ Possible\ Amount$$
$$\times \left(1 - Average\ Deviance\ Score/Target\ Time\right)\quad (2)$$

Therefore, a smaller deviance score was required in a block of shorter target durations, compared to a block of longer to be-reproduced durations, in order to earn the same monetary reward.

The total amount earned was shown at the end of each block. Participants' endogenous timing uncertainties were quantified using reproduction data for each duration by dividing the standard deviation of reproduced durations by their mean. This is a statistical procedure for obtaining the CV of a dataset, and is used as an indicator of endogenous timing uncertainty, which is typically constant for different durations within an individual (Gibbon, 1977; Buhusi and Meck, 2005). The CV is an appropriate measure of timing uncertainty since when the CV is known, one can estimate the expected error of the same individual for other intervals (CV times t). Thus, many studies in the interval timing literature use CV as a measure of timing uncertainty (e.g., Gibbon, 1977; Balci et al., 2011a).

Temporal reproduction task–RDM stimuli

These additional TR sessions were identical to the original TR session described above, except for replacing the static green square with a RDM stimulus identical to the one used in FR and DR sessions (i.e., dot motion stimulus with 16% coherence). The purpose in replacing the static stimulus with the RDM stimulus was to replicate as closely as possible the conditions in which the FR and DR sessions took place, since a TR task more similar to these 2AFC tasks could better capture the representation of attentional, as well as temporal, dynamics underlying the decision making process (see Zakay and Block, 1996). This in turn should lead to more accurate estimates of timing performance (i.e., CV) as manifest in the decision task and thus values that are more appropriate for generating threshold collapse predictions in DR sessions. In order to make sure that the motion direction was being attended to, participants were asked to report the direction of motion using the "Z" or "M" keys in 20% of the trials, following the time reproduction. "Total Earnings" (Equation 2) were multiplied by the proportion of accuracy in reporting the direction of motion in each block.

Since the error rate in direction judgments would inevitably decrease the total amount earned in these TR sessions compared to those using the static stimuli, the maximum possible amount that could be earned per block was increased from 2.5 to 3 Turkish Liras. Each TR task (i.e., with static or RDM stimuli) lasted for a single session. The TR testing was shorter than the 2AFC tasks because estimating temporal accuracy and precision does not require as large of a dataset as one needs for the DDM fits and conditional accuracy curves.

DATA ANALYSIS
Quantifying declining accuracy with time

In order to quantify a possible decline in accuracy as time elapsed within trials, accuracy levels were calculated for each 50 ms RT bin, forming the conditional accuracy curves. Bins with less than 4 data points, as well as RTs above 5 s, were excluded from all further analyses. The exclusion criterion for bin size was based on post-hoc analyses of the data, especially for the last two RT bins (i.e., at around the deadline), which generally contained less data points than the ones that corresponded to shorter RTs. Our analysis showed that nine participants had at least 4 data points in the last RT bin in the short deadline condition, whereas this number declined to four participants in the long deadline condition. Since the accuracy at and near the deadline was of high relevance to this study, we set our exclusion criterion to allow for involving these participants' RT data in further analyses. Note that our original choice of the specific response deadlines based on free response RT distributions aimed for more data points to fall in these later bins.

A conditional accuracy curve allows us to determine the RT bin where a decline in accuracy starts, as well as the rate of this decline. In order to define the specific point where the accuracy trend changes, we found the RT bin at which the sum of squared errors of two piece-wise linear fits to data before and after that point (a.k.a. the knot) is minimized. This was achieved by running an algorithm which fits the piece-wise linear functions to data by using each RT bin as a putative knot location where the

first linear function is "latched on" to the second one. Specifically, the algorithm constrains the intercept of the second linear fit to be the last value of the first fit, forming two connected lines. Since the last data point of the first fit affects the fit of the second line by slightly modifying its slope, the algorithm runs in both forward and reverse directions, ensuring that it finds the knot location where the total error of the piece-wise fit is minimal, regardless of which of the two slopes is modified. The purpose of using this algorithm was to quantify the onset (i.e., inflection point), as well as the slope of a possible decline in accuracy with RT separately for two different deadlines. The correlations of these two values (i.e., onset & slope) with timing uncertainty were later calculated (Section Effect of Deadlines on Response Time and Accuracy) in order to test if higher levels of timing uncertainty predicted an earlier onset of decline in accuracy characterized by a lower (as opposed to a steep) negative slope.

Optimal threshold collapse simulations

We conducted simulations in order to approximate the optimal threshold collapsing trajectories for different deadline durations (800, 1000, and 1200 ms) and six linearly increasing levels of endogenous timing uncertainty (i.e., CV), using two different closed-form collapse functions (i.e., exponential and linear). Below we describe the details for the exponential threshold collapse function, but the same procedure applies to the linear collapse function as well. Although our response paradigm employed only two deadline durations (800 and 1000 ms), we have also tested the 1200 ms deadline in these simulations. For the objective function analyzed by Frazier and Yu (2008)–which may approximate but is not identical to RR—analytically optimal collapse functions look much like our exponentials.

In order to find the exponential threshold collapsing trajectory that maximizes the number of rewards for a given deadline and a given timing uncertainty, we first constructed a total of 101 threshold trajectories with 0.01 second increments, separately for each CV value. The following formula was used to construct an exponential curve:

$$a = \left(Asymptote + \left(Starting\ Point\ -\ Asymptote \right) \times e^{(-c*t)} \right) \quad (3)$$

where *Asymptote* was set at 0.1 for the upper threshold, Starting Point was set at 0, c represented the rate of exponential decline (i.e., as a proxy for temporal discriminability), and t is time. The resulting curve was then flipped on its y-axis to construct the upper threshold. This mirror image of the upper threshold was used as the lower threshold (**Figure 3**).

All thresholds collapsed exponentially with time to the starting point of evidence accumulation (**Figures 3A,B**). The upper and lower thresholds with the earliest evaluated collapse onset met well before the shortest deadline (i.e., 800 ms), and the thresholds with the latest evaluated collapse onset met well after the longest deadline (i.e., 1200 ms). The presumed effect of the timing uncertainty was implemented by changing the exponential decay parameter (c; e.g., steeper collapse for higher temporal discriminability due to lower timing uncertainty).

For each response deadline, we defined the optimal threshold trajectory as the one (out of 101 per CV) that yielded the greatest number of rewards out of 10^6 drift-diffusion simulations. In line with our experimental paradigm, in these simulations RTs longer than the deadline duration were not assigned any reward. The drift diffusion processes were simulated based on Equation 1. The drift rate was set to 0.1, the noise coefficient was set to 0.1, the starting point was set to 0 and non-decision time was set to 0. The two decision thresholds were set to -0.1 and 0.1 at trial onset. For simplicity, the core parameters were not allowed to vary between trials. The results of our simulations supported Frazier and Yu's formulation; the optimal thresholds for a given deadline and a given CV were the ones which nearly reached the starting point at the response deadline even with closed-form collapse functions (**Figure 3A**). These simulations also suggested that higher timing uncertainty requires an earlier onset of threshold collapsing, so that the upper and lower decision thresholds are ensured to meet virtually at the deadline.

We have also calculated the optimal threshold collapse trajectories by setting the criterion for optimality as the highest RR instead of the highest amount of expected reward (**Figure 3B**). The RR for each collapse trajectory was calculated by dividing the mean accuracy by the mean RT. In calculating the RR, late responses (i.e., those beyond the deadline) were given a value of 0 for accuracy (i.e., they were counted as error trials). RT was defined as "DT + RSI + T_{er}"for trials with RTs faster than the deadline, and "deadline + RSI" for trials where RTs were slower than the deadline. Using values for the RSI and T_{er} very close to the ones derived from our experimental paradigm, calculated the expected RR for each collapse trajectory and found that, similar to those in **Figure 3A**, optimal thresholds for a given CV were the ones that roughly collapsed to the starting point near the deadline (**Figure 3B**).

Visual inspection of **Figure 4A** shows that the order of the optimal threshold (i.e., the order of a given threshold among the 101 thresholds tested with 0.01 s increments) increases with longer deadlines for a given CV, in addition to decreasing with higher CVs for a given deadline. Additionally, conditional accuracy curves were plotted for the six hypothetical CV levels, separately for the three deadline durations (**Figure 4B**). The level of CV (i.e., the level of endogenous timing uncertainty) was increased or decreased by decreasing or increasing the rate of exponential decline (the c parameter in Equation 3), respectively. Visual inspection of **Figure 4B** suggests that accuracy in our simulations declines with time for all levels of CV. However, contrary to our expectations, accuracy never fully reaches 50% (chance level) in these curves. Both **Figures 4A,B** were constructed based on expected total reward as the optimality criterion.

Finally, **Figure 5A** shows the expected total reward curves for all 101 collapse functions constructed with the lowest and the highest CV levels (out of the six CV levels) for the three deadline durations. Visual inspection of **Figure 5A** suggests that the expected total reward steadily increases with the order of exponentially collapsing thresholds, and sharply declines immediately following the deadline. Additionally, **Figure 5B** shows the mean RTs and expected total rewards
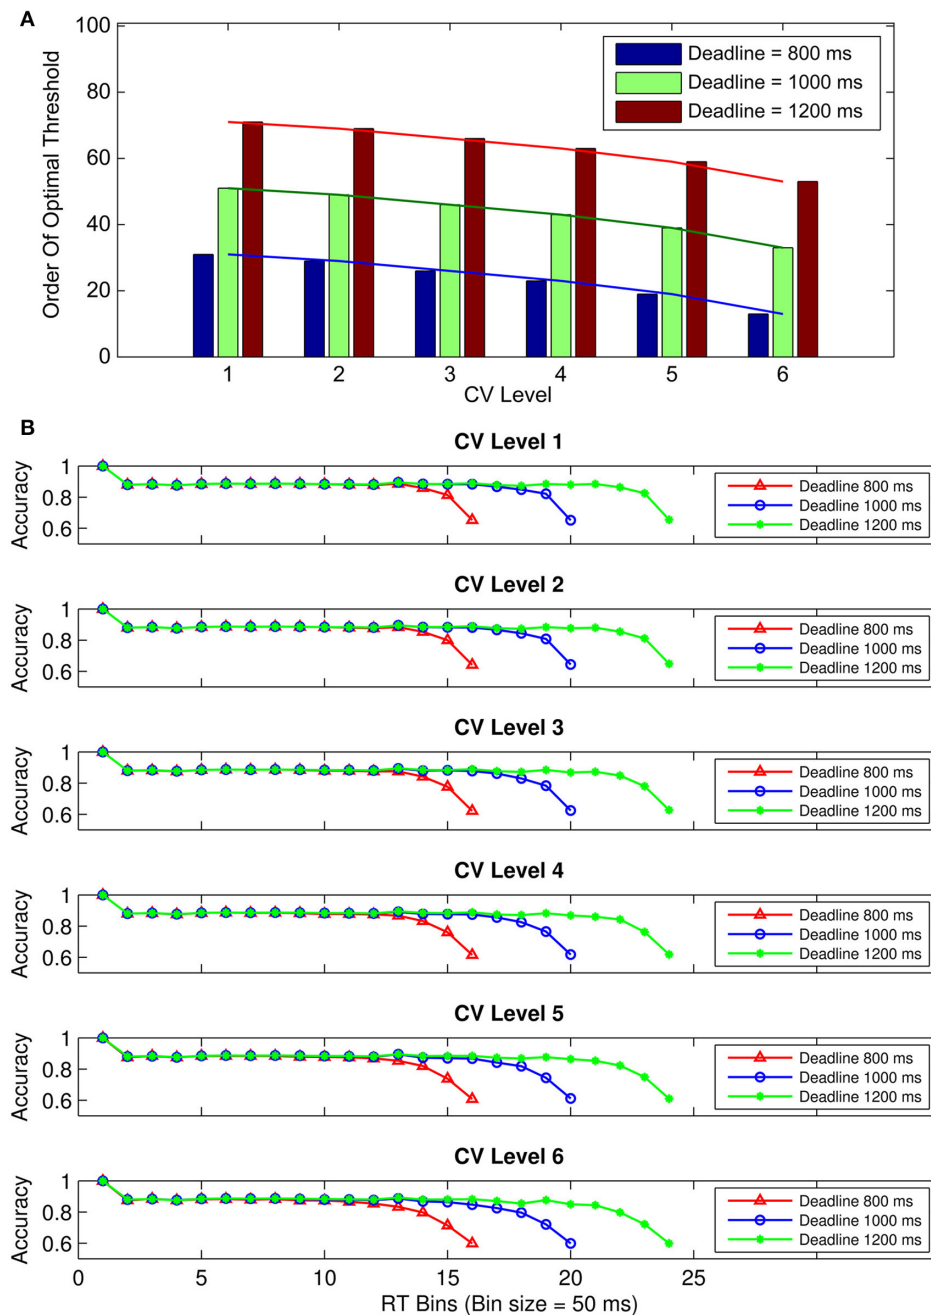
FIGURE 4 | (A) Bar graphs depicting the order of the optimal threshold collapse trajectories (out of 101 thresholds with 0.01 s increments) selected from the family of exponential decline functions for six hypothetical levels of timing uncertainty. Lines connect the bars. **(B)** Conditional accuracy curves for the six CV conditions, shown separately for the three response deadlines. Red lines represent the conditional accuracy curves for the short deadline (800 ms), blue lines for the medium deadline (1000 ms), and green lines for the long deadline (1200 ms). Both **(A,B)** are based on expected total reward as the optimality criterion.

predicted for optimal threshold trajectories as a function of CV, separately for the three deadlines. **Figure 5B** suggests that with increasing timing uncertainty (i.e., CV level), both the mean RT and the expected total reward decline. See Supplementary Material for the linear threshold collapse results.

RESULTS

ACCURACY AND RESPONSE TIME IN THE FREE RESPONSE CONDITIONS

The data from the two FR sessions showed that the participants' error rates declined from a mean of 10% in the first 4 blocks of the first FR session, to a mean of 4.3% in the last 4 blocks of
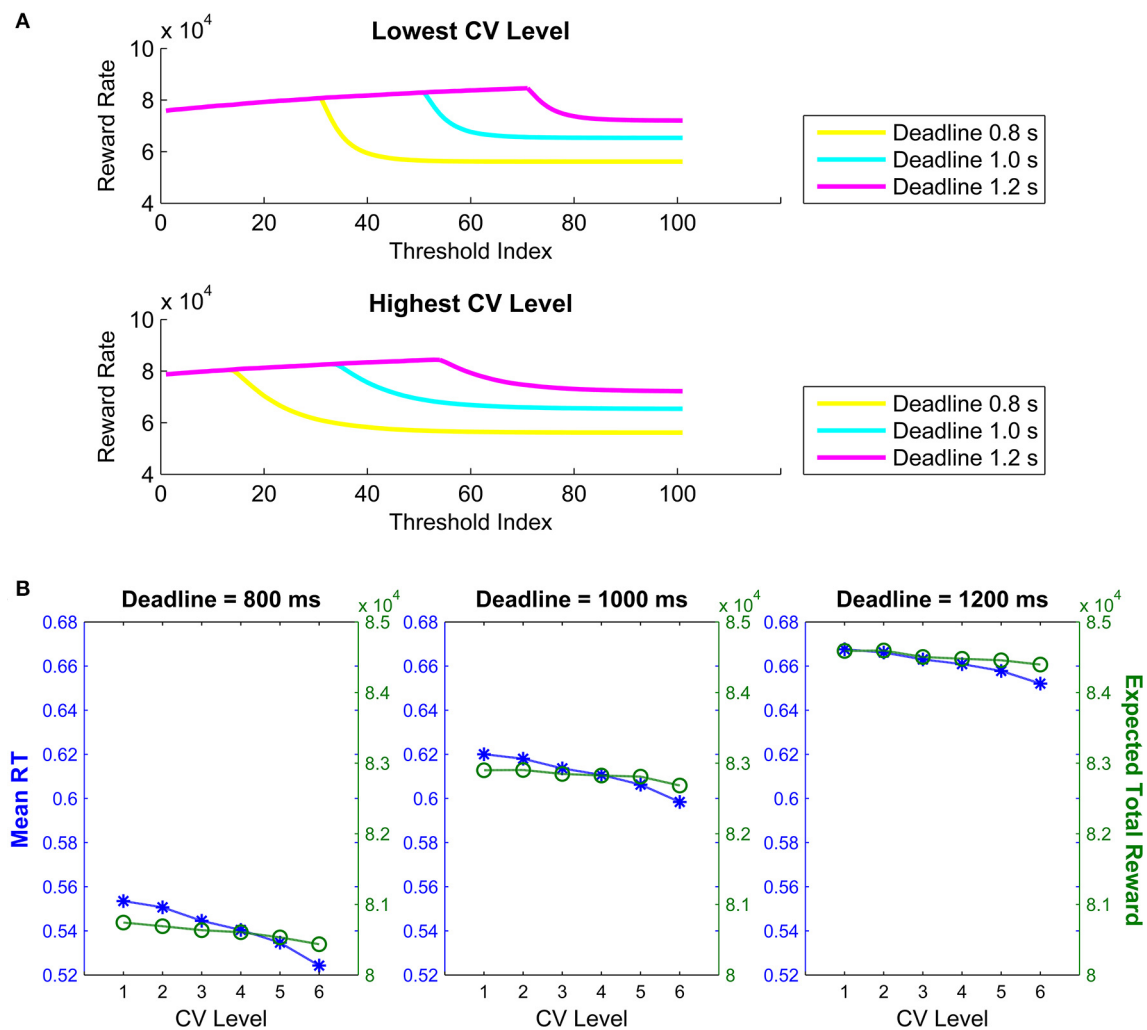
FIGURE 5 | (A) Expected total reward amount for the highest and lowest CV levels as a function of the order of threshold among the 101 thresholds tested (here defined as "Threshold Order"). **(B)** Mean response times and expected total reward amounts as a function of six levels of CV defining six exponential threshold collapse trajectories for the short (800 ms), medium (1000 ms) and long (1200) simulated deadlines. Both **(A,B)** are based on expected total reward as the optimality criterion.

the second FR session [$t_{(9)} = 3.1$, $p < 0.05$; **Figure 6**] suggesting that the FR sessions were successful in training the participants on the RDM discrimination task. Additionally, the RTs showed a similar decline with increasing blocks, with a mean of 0.94 s in the first 4 blocks of the first FR session, to a mean of 0.75 s in the last 4 blocks of the second FR session, however, this difference failed to reach significance ($p > 0.05$). RTs between the first and second halves within the two FR sessions did not differ significantly (both $ps > 0.05$), excluding the potential role of factors such as an increased fatigue or inattention toward the end of a test session.

Figure 7 shows the RT distributions in the FR blocks in FR sessions, FR blocks in DR sessions, and the two deadline blocks in the DR sessions. **Figure 7** shows the plots either of all RTs pooled across participants (**Figure 7A**), or RTs below the short deadline duration (**Figure 7B**). A mean of 844.85 (s.e.m. = 20.1) trials

were completed in FR blocks in FR sessions, whereas this number was 105.2 (s.e.m. = 1.24) in FR blocks in DR sessions, 433.88 (s.e.m. = 2.27) in Short Deadline blocks in DR sessions, and 432.48 (s.e.m. = 2.82) in Long Deadline blocks in DR sessions.

EFFECT OF DEADLINES ON RESPONSE TIME AND ACCURACY

In order to determine whether introducing a deadline for responding was successful in modifying behavior, we first compared the mean RT values obtained by pooling data from both FR sessions, the 4 DR sessions (separately for the short and long deadline conditions), and the single FR blocks presented at the start of each DR session for each participant. A one-way repeated measures ANOVA was conducted to compare the effect of response time limitations on mean RT in four conditions; two free response (i.e., FR blocks in FR sessions and FR blocks in DR sessions) and two deadline (i.e., short & long) conditions.

FIGURE 6 | Mean error rate as a function of FR block. Mean error rates per FR block in the first two sessions (i.e., FR sessions). Blue squares correspond to FR blocks in the first FR session and red circles correspond to FR blocks in the second FR session. Error bars denote the standard error of the mean.
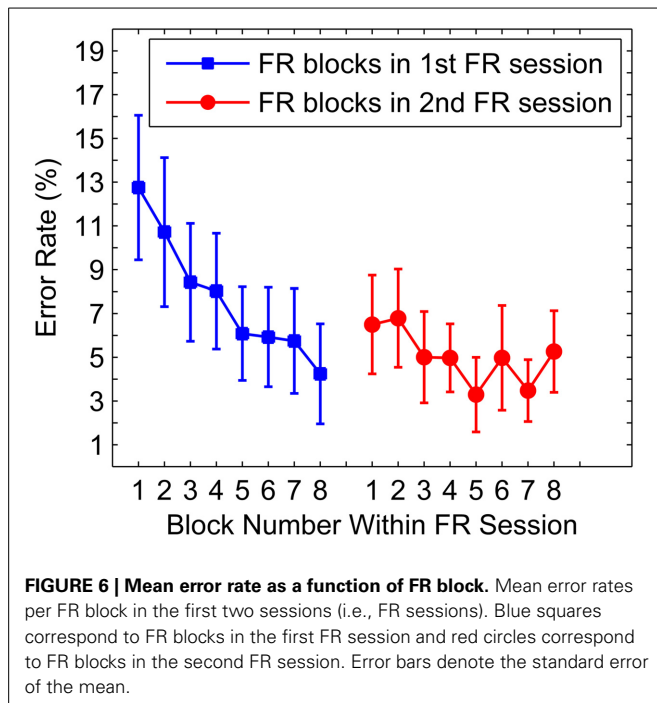
Since response deadlines act as a procedural censoring point for slower RTs, only the RT values up to the short deadline (800 ms) were compared in all conditions. Our analysis indicated a significant effect of different experimental conditions on the RTs, $F_{(3,6)} = 32.78$, $p < 0.001$. Tests of six pair-wise comparisons were conducted using Holm-Bonferroni adjusted alpha levels. These comparisons showed that RTs in FR blocks in FR sessions ($M = 602$ ms) were significantly longer than both the short deadline ($M = 519$ ms, $p < 0.001$) and the long deadline ($M = 525$ ms, $p < 0.001$) conditions, as well as the response times of FR blocks in DR sessions ($M = 548$ ms, $p < 0.001$). The difference between the RTs in the two separate deadline conditions and the FR blocks in DR sessions did not reach significance (both $ps > 0.05$). However, when no correction was applied for multiple comparisons, the mean RT differences between FR blocks in DR sessions and the two separate deadline conditions reached significance (both $ps < 0.05$).

In order to further test if introducing a short vs. long deadline was effective, we compared the number of missed deadlines for each deadline condition. A mean of 1.68% of deadlines were missed in the short deadline condition (s.e.m. = 0.35), whereas this percentage declined to a mean of 0.36% in the long deadline condition (s.e.m. = 0.09). A paired samples t-test revealed that the percentage of missed deadlines was higher for the short deadline condition, compared to long deadline condition $t_{(9)} = 4.5$, $p < 0.001$. In other words, participants as expected were more likely to miss the deadline in the short DL conditions compared to the long DL conditions. The hypothetical percentage of missed deadlines was computed for the RT distributions of the FR blocks in DR sessions by calculating the percentage of the data above the RTs corresponding to the two deadlines separately. A mean of 9.13% of the trials (s.e.m. = 3.02) had RTs above the short

deadline duration (i.e., 800 ms), whereas a mean of 3.26% of the trials (s.e.m. = 1.15) had RTs above the long deadline duration (i.e., 1000 ms). Matched-sample t-tests showed that the percentage of RTs above the short deadline duration in FR blocks in DR sessions was significantly higher compared to the percentage of missed deadlines in the short deadline condition $t_{(9)} = 2.89$, $p < 0.05$. Similarly, the percentage of RTs above the long deadline duration in FR blocks in DR sessions was significantly higher compared to the percentage of missed deadlines in the long deadline condition $t_{(9)} = 2.81$, $p < 0.05$. These results point at the effect of response deadlines on RTs.

An additional One-Way repeated measures ANOVA was conducted to compare the effect of four experimental conditions on overall accuracy, using accuracy data corresponding to RTs below 800 ms (again due to the procedural censoring factor). There was a significant effect of experimental condition on accuracy, $F_{(3,6)} = 22.59$, $p < 0.001$. Tests of six pair-wise comparisons conducted using Holm-Bonferroni adjusted alpha levels revealed that, whereas the accuracy in FR sessions ($M = 0.96$) and FR blocks in DR sessions ($M = 0.94$) did not differ significantly from each other ($p > 0.05$), both accuracy means differed significantly from the short ($M = 0.90$, both $ps < 0.001$) and long deadline ($M = 0.90$, $ps < 0.001$) conditions. Mean accuracy in the two deadline conditions did not differ significantly ($p > 0.05$).

The effect of four experimental conditions on overall accuracy were also compared using all data, without excluding those above 800 ms. There was a significant effect of experimental condition on accuracy, $F_{(3,6)} = 8.07$, $p < 0.001$. Tests of six pair-wise comparisons conducted using Holm-Bonferroni adjusted alpha levels revealed that, whereas the accuracy in FR blocks in FR sessions ($M = 0.94$) and FR blocks in DR sessions ($M = 0.93$) did not differ significantly ($p > 0.05$), mean accuracy in FR blocks in DR sessions differed significantly from both the short ($M = 0.90$) and long deadline ($M = 0.90$) conditions (both $ps < 0.001$). The mean accuracy in the two deadline conditions did not differ significantly either from each other or from the mean accuracy in FR blocks in FR sessions (all $ps > 0.05$). However, when no correction was applied for multiple comparisons, the mean accuracy differences between FR blocks in FR sessions and the two separate deadline conditions reached significance (all $ps < 0.05$).

We analyzed within block RTs in both deadline conditions to verify that inattention/fatigue did not set in toward the end of a 5-min block, possibly resulting in slower RTs toward the end of a block. For this purpose, we first calculated individual participants' mean RTs for each trial order in separate deadlined blocks across all DR sessions, for the two deadline conditions. For instance the mean RT for trial number 14 in the second block of all short deadlined DR sessions was calculated by taking the mean of all RTs corresponding to the 14th trial in the second blocks of the short deadlined DR sessions and so on. For later trials where some blocks did not have RT data due to unequal number of trials per block, mean RT was calculated by using available data only. Given that there were four blocks in each deadline condition per session, this procedure resulted in four sets of mean RTs per participant, which were fit by a linear regression using a least-squares method. It was reasoned that an increase in RTs over the course of a block of trials should manifest itself as a positive slope of a linear fit

FIGURE 7 | Response time distributions for FR blocks in FR sessions, FR blocks in DR sessions, Short DR blocks in DR sessions and Long DL blocks in DR sessions. RT data pooled across participants are plotted either **(A)** without an upper limit, or **(B)** with an upper limit of the short deadline duration (800 ms).
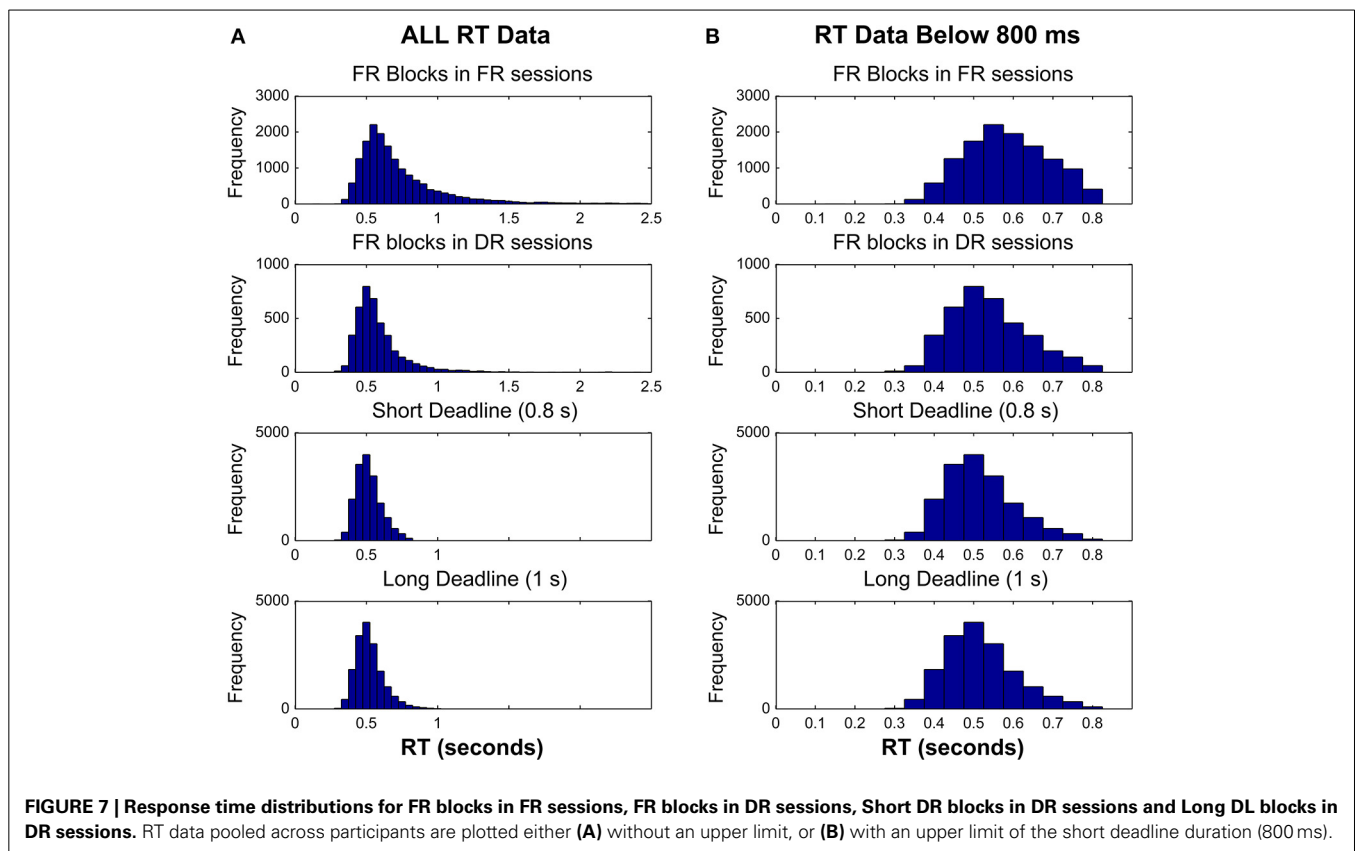
to data. A total of eight one-sample t-tests were conducted (four for each deadline condition) in order to determine whether the slopes of the linear fits were different from 0. None of the slopes were significantly higher or lower compared to the test value of 0 (all $ps > 0.05$), suggesting that RTs did not increase or decrease toward the end of a test block.

Finally, we wanted to see if error trials were more likely to occur in the first half or the second half of a DR block, due to possibly increasing fatigue or inattention. Using the same method described above, we calculated individual participants' mean accuracies in the first and the second halves of each block, separately for the two deadline conditions. Eight paired sample t-tests were conducted to compare accuracy in the two halves of each block in the two deadline conditions (i.e., four t-tests for each condition). None of the differences were significant, suggesting that accuracy did not decline toward the end of a deadlined test block (all $ps > 0.05$).

ACCURACY AT DEADLINE
In order to see if it declined to chance level at the deadline, accuracy in the last 50 ms RT bin was calculated for both deadline conditions. Nine participants had valid data (i.e., more than 4 data points) in this RT bin in the short deadline condition, with a mean accuracy of 78.4% (s.e.m. = 3.6%), whereas 4 participants had data in the last bin in the long deadline condition with a mean of 75.6% (s.e.m. = 5.8%). Of those with valid data in the last bin, no participant's accuracy fell below 63% in the short

deadline condition, whereas the lowest accuracy in the last bin was 60% in the long deadline condition. A Wilcoxon signed ranks test indicated that accuracy in the last RT bin in the short deadline condition ($Mdn = 0.76$) was significantly higher than a hypothetical value of 0.5 ($Z = 45$, $p < 0.05$), whereas this difference did not reach significance for the last RT bin in the long deadline condition ($Mdn = 0.78$, $p > 0.05$).

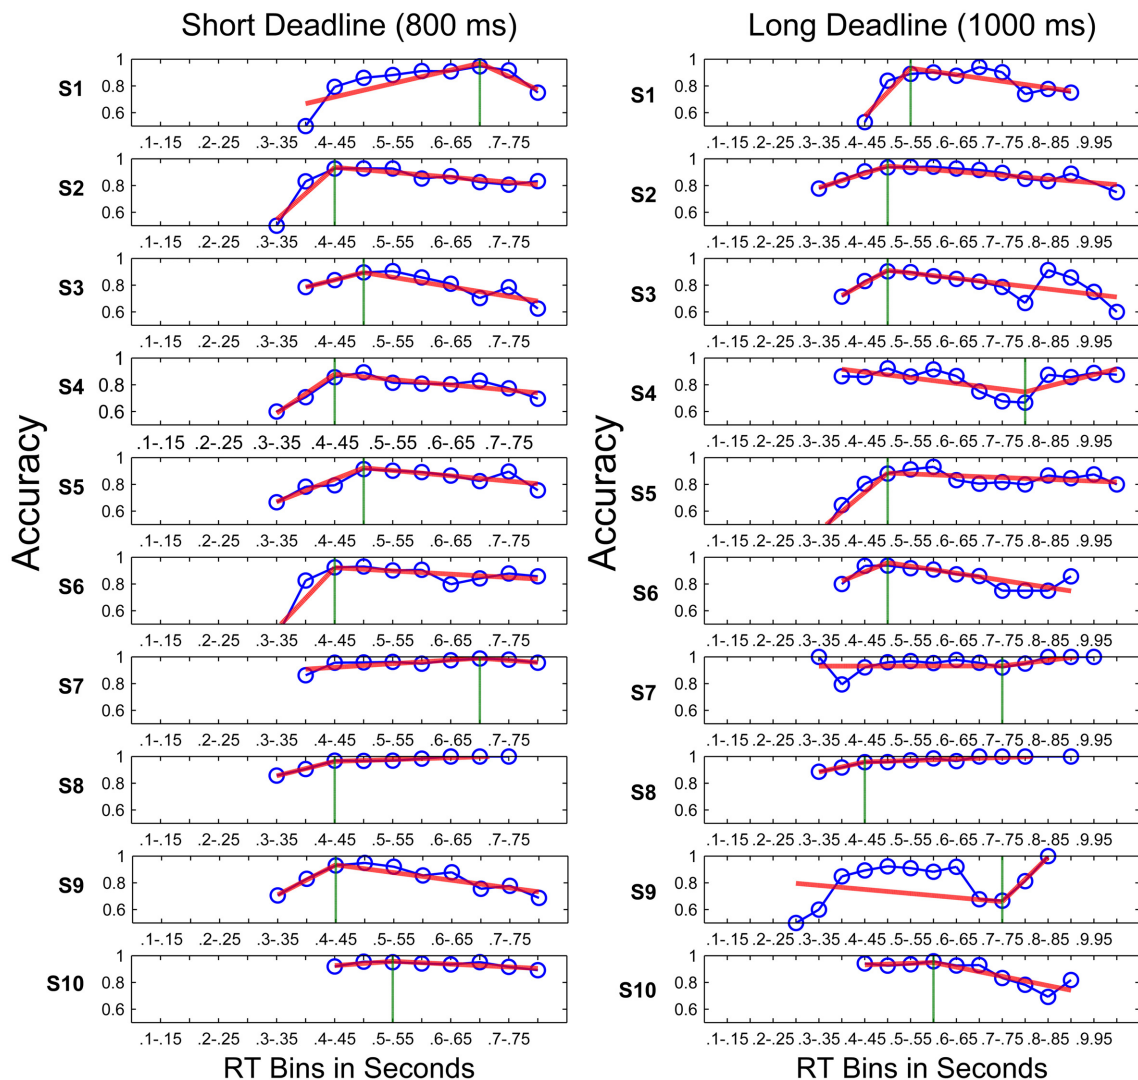### PIECE-WISE LINEAR FITS OF CONDITIONAL ACCURACY CURVES
Figure 8 shows the conditional accuracy curves plotted for each condition by pooling data across participants. The analysis using piece-wise linear fits was also based on each participant's data expressed as conditional accuracy curves (**Figure 9**). The knot locations (defined in terms of RT bins) of the piece-wise linear fits to these data and the slopes of the best fit lines were calculated using the algorithm described in the Methods Section, in order to quantify the onset, as well as the rate of a potential decline in accuracy with time. **Figure 9** shows fits to individual participants' data. A total of 9 out of 10 participants had declining accuracies after the inflection point (i.e., knot location) with time (i.e., negative slope) in the short deadline condition, whereas 6 had declining accuracies after the inflection point in the long deadline condition. Two one sample t-tests were conducted in order to compare the slopes of the second line for the two deadline conditions to the slope of "0" (i.e., no decline in accuracy with time). Although, the slopes in the short deadline condition ($M = -0.3$) differed significantly from 0 [$t_{(9)} = 2.84$,
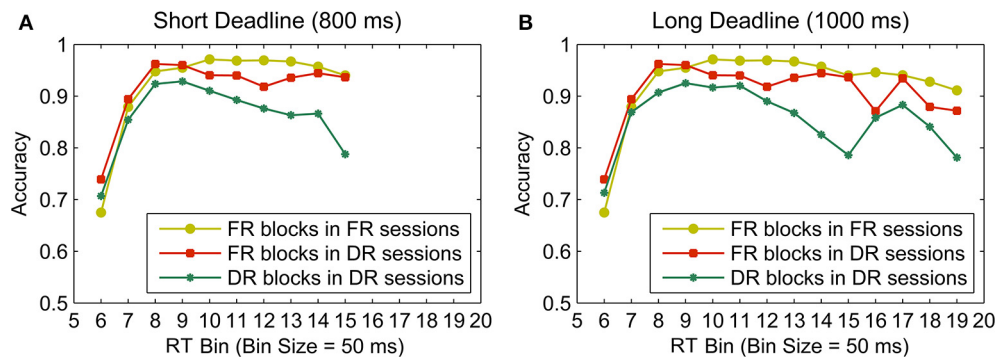
FIGURE 8 | Conditional accuracy curves for the FR blocks in FR sessions, FR blocks in DR sessions, and DR blocks separately for the short (A) and long (B) deadlines (800 and 1000 ms, **respectively).** The conditional accuracy curves for the two FR blocks are identical between two columns up to the 15th RT bin. Data were pooled across participants.

FIGURE 9 | Piece-wise linear fits (red lines) to conditional accuracy for the DR blocks (blue lines with circles) of all participants in the short deadline (left) and long deadline (right) conditions. Vertical green lines indicate inflection points.

$p < 0.05$], this difference failed to reach significance in the long deadline condition ($M = 0.01$, $p > 0.05$). The insignificant difference remained for the long deadline condition when the data from participant 9 with a bad fit were not included in the analysis.

TEMPORAL UNCERTAINTY AND CONDITIONAL ACCURACY CURVES

Coefficient of variation values for each participant were calculated for both TR tasks by taking the average of all CVs for the three target durations (see Methods Section; **Figure 10**). Mean CV values obtained from the first TR task using static stimuli were significantly higher compared to CVs obtained from the second TR task using RDM stimuli [$t_{(9)} = 3.97$, $p < 0.01$], which may reflect a practice effect since the first TR task always used static stimuli or the specific stimulus effect. A potentially significant correlation between RT and CV was examined. Neither of the CV values obtained from the two TR tasks correlated significantly with mean RTs in the FR or DR conditions (all $ps > 0.05$).

While the positive correlation between CVs in the TR task with static stimuli and the knot location of the piece-wise fits to RT data in the short deadline condition reached significance [$r_{(8)} = 0.85$, $p < 0.01$, two-tailed], the same CVs did not correlate with the knot locations in the long deadline condition ($p > 0.05$). Conversely, the CVs obtained in the TR task with dynamic stimuli were positively correlated with the knot location of the piece-wise fits in the long deadline condition [$r_{(7)} = 0.72$, $p < 0.05$, two-tailed], whereas they did not correlate with those knot locations in the short deadline condition ($p > 0.05$). Neither of the CVs correlated with the slopes of the first or second line of the piece-wise linear fits (both $p > 0.05$).

As can be seen in **Figure 9**, participant number 9 had a visibly bad piece-wise linear fit to his/her conditional accuracy curve in the long deadline condition. Therefore, the same correlations were also calculated by excluding this participant's data in the long deadline condition. While the correlation between CVs in the TR task with static stimuli and the knot location of the piece-wise fits in the long deadline condition remained insignificant ($p > 0.05$), the correlation between CVs in the TR task with dynamic stimuli and the knot location in the long deadline condition also failed to reach significance when calculated by excluding this participant's data. Excluding this participant's data also did not result in a significant correlation between CVs and the slopes of the first or second line of the piece-wise linear fits to the long deadline condition (all $ps > 0.05$). None of these results support the optimal performance predictions, since we expected participants with higher CVs to start reducing their accuracy earlier (under the threshold collapsing assumption). If anything we observed the opposite relationship with the CVs in TR task with static stimuli in the short deadline condition, and the CVs in TR task with dynamic stimuli in the long deadline condition. When the data only from the participants with a negative slope in the second line of the piece-wise linear fits were taken into consideration, none of the correlations between either of the CVs and the knot locations, or between the CVs and the slopes of both the first and second line of the piece-wise linear fits reached significance (all $ps > 0.05$).

Even though we had a minimum number of data points per RT bin used in forming the conditional accuracy curves, investigating the declining accuracy using binned RTs may be misleading in the sense that some bin accuracies calculated with fewer yet highly accurate/inaccurate trials may be artificially inflated / deflated. In other words, the binning methodology may fail to accurately represent the dynamics of a declining accuracy with time, since it entails estimating accuracies for a specific time period from the average of sometimes a very limited number of data points. Therefore, we also calculated peak accuracy by taking the cumulative average of accuracy with increasing time (i.e., RT), and correlated the location of these peaks in time with CV values. This was achieved by first sorting RTs for each trial in increasing order and then forming an "accuracy vector" by coding 0 for error trials and 1 for correct trials corresponding to each RT value. Cumulative accuracy was then calculated for each trial by taking the average accuracy of all trials with RTs at and below that trial, which formed a cumulative average accuracy curve. Consistent with the findings reported above, the RTs at which the cumulative average of accuracy peaked did not correlate significantly with the CVs estimated from either TR task (both $ps > 0.05$). These results further supported the above-mentioned results obtained by using the RT binning approach, further suggesting that even if participants collapsed their decision thresholds, they did not take into account their endogenous timing uncertainties.

Finally, in order to see if there was a bias toward over- or underestimating the durations/deadlines additional analyses were conducted. Normalized mean reproduction durations of all participants were first calculated by dividing the mean reproduction duration by the target duration. This was done separately for all three durations (1–2.12–4.24 s) tested in the two TR session types (static or dynamic stimuli). Six one-sample t-tests were conducted using "1" as test value for accurate normalized performance. Only the 1 s test duration in the dynamic stimulus condition ($M = 1.31$, s.e.m. $= 0.0.06$) was systematically

FIGURE 10 | CV values obtained in the temporal reproduction task with two different types of stimuli (red circles for static stimuli & green triangle for dynamic stimuli) for each subject. Subject 4 did not participate in the second TR session with dynamic stimuli.
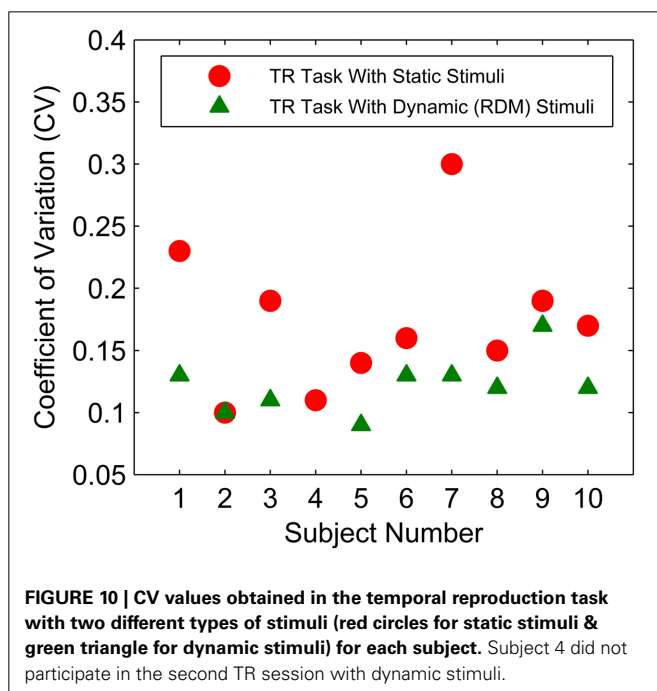
overproduced by the participants [$t_{(8)} = 4.73$, $p < 0.001$], suggesting that subjects tended to underestimate 1 s of dynamic stimulus presentation. This result suggests that if thresholds did in fact collapse with time, this collapse may have started declining later than optimally, since participants were underestimating the deadlines. In order to test this possibility, the correlation between the mean reproduction duration of 1 second (separately in the TR tasks using static & dynamic stimuli), and the knot location, as well as the slope parameter of the conditional accuracy curves was calculated. This procedure was also repeated by excluding the long deadline data of participant ID 9. None of these correlation coefficients reached significance (all ps > 0.05).

DRIFT-DIFFUSION MODEL SIMULATIONS
Since we observed accuracy reduction within trials for some participants in DR sessions, it is important to address whether models with fixed parameters within trials can account for this pattern. Thus, we tested if observed reduction in accuracy as a function of RTs could be due to factors other than collapsing thresholds. For this purpose, individual participants' data between FR blocks in DR sessions were fit by the extended DDM (i.e., allowing for inter-trial variability parameters, *all variability parameters* > 0, and also allowing for starting point bias) using the diffusion model analysis toolbox (DMAT) (Vandekerckhove and Tuerlinckx, 2008). These parameters were then averaged across participants in order to obtain a representative set of parameters that could be used for DDM simulations to follow.

The following mean parameters were obtained; decision boundary (a) = 0.1214, non-decision related delays (T_{er}) = 0.4419, drift rate variability ($Var(v)$) = 0.1922, starting point (z) = 0.0608, starting point variability ($Range(z)$) = 0.0547, non-decision time variability ($Range(T_{er})$) = 0.1668, and drift rate (v) = 0.4447. Data from FR blocks in DR sessions were used instead of FR blocks in FR sessions to estimate the DDM parameters because they represent performance that is closer to steady-state.

Using these DDM parameters, we simulated three sets of 10^6 data points using DMAT's simulation feature, in which either of the threshold (a), drift rate variability ($Var(v)$), or the starting point variability ($Range(z)$) parameters were increased or decreased by 10 and 20% (depending on the condition; see **Figure 11**). Therefore, each set contained five levels of its corresponding parameters. This procedure aimed to investigate if changes unrelated to within-trial threshold collapsing might also lead to decreasing accuracy levels with slower RTs. These specific parameters were chosen for incrementing/decrementing because large/small values of these parameters are known to lead to longer/shorter RTs for incorrect choices (Ratcliff and Rouder, 1998; Ratcliff and McKoon, 2008). Specifically, larger values of threshold and drift rate variability parameters lead to slower error RTs, whereas a larger variability in starting point should present itself as faster RTs for error trials (Ratcliff and Rouder, 1998; Ratcliff and McKoon, 2008). Such response patterns formed by slower responses for error trials compared to correct ones cannot be explained by the pure DDM when it is unbiased toward one threshold over the other (Laming, 1968). Importantly for our purposes, if error trials are slower than correct trials, this pattern

automatically implies a declining conditional accuracy curve. In other words, the decline in accuracy observed in our data may not necessarily be a behavioral manifestation of a collapsing decision threshold (a), but instead may result from changes in the values of the other parameters such as the drift rate variability ($Var(v)$) or an overall reduction in decision threshold (a) that stays constant within a trial. **Figure 11** shows the results of these simulations by plotting accuracies as a function of corresponding RTs (using a bin size of 0.05 s).

Conditional accuracy curves based on simulated data showed a steadily declining accuracy with increasing RT (**Figure 11**). Moreover, although the rate of this decline is higher for a lower threshold parameter, a similarly increasing rate of decline is observed for higher levels of the drift rate variability parameter as well, with no modification of the threshold or any other parameter within a trial. Additionally, increasing or decreasing the starting point variability had no discriminable effect on the rate of decline in accuracy with time. These results suggest that, importantly, decreasing the constant decision threshold (i.e., without the need for within trial modulation) or increasing the variability in drift rate could underlie decreasing accuracy toward a deadline.

DISCUSSION
Many studies using 2AFC tasks have focused on the optimality of decisions in free response paradigms (e.g., Bogacz et al., 2006, 2010; Simen et al., 2009; Starns and Ratcliff, 2010; Balci et al., 2011b). Some of these studies showed that with enough training human participants can optimize the speed-accuracy tradeoff in their decisions by adopting RR-maximizing decision thresholds. When response deadlines are imposed in these tasks, reward maximization instead requires the decision-maker to collapse decision thresholds within a trial such that at the time of deadline, they meet at the starting point of the evidence accumulation process. This is an adaptive process as it secures at least a 50% chance that the reward will be obtained instead of earning nothing if the decision-maker is late. Frazier and Yu (2008) showed the relevance of timing uncertainty to the parameterization of this adaptive within-trial threshold crossing process. Participants with higher timing uncertainty should start collapsing decision thresholds earlier to maximize reward. Thus, reward maximization in these tasks entails factoring timing uncertainty into decisions in a normative fashion.

To this end, previous research has shown that humans and non-human animals are able to take normative account of their endogenous timing uncertainties in both temporal and non-temporal decision making tasks (for review see Balci et al., 2011a). This prediction was tested in the current study by examining conditional accuracy curves and evaluating how their shape depends on deadlines and participants' endogenous timing uncertainty. Although our results showed that accuracy decreased with time toward the deadline for many participants, this rate of decline was much lower than expected from an optimal decision-maker and did not correlate with measured levels of timing uncertainty. In contrast to optimal performance predictions, the timing of the onset of decline in accuracy increased rather than decreased with higher levels of timing uncertainty in the short deadline condition, when this uncertainty was quantified using a static visual
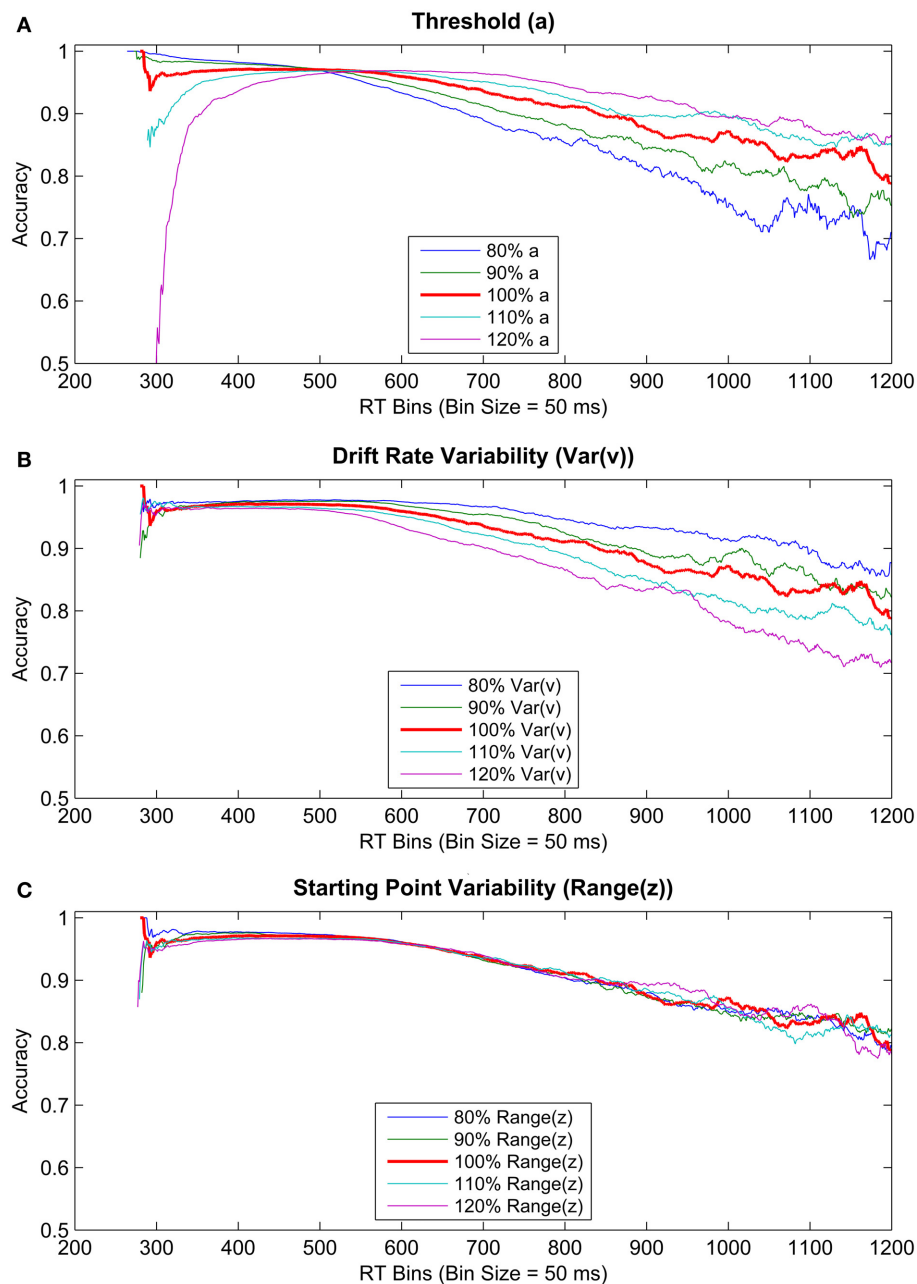
FIGURE 11 | Conditional accuracy curves gathered from simulated data. RT bin size of 50 ms was used for plotting. Three parameters were systematically modified by increasing and decreasing their values by 20 and 10%. The parameters manipulated were **(A)** threshold (top panel), **(B)** drift rate variability (middle panel), and **(C)** starting point variability (bottom panel).

stimulus, and also in the long deadline condition when it was quantified using a dynamic visual stimulus. It is possible that our analytical approach, i.e., using linear fits to accuracy levels of binned RT data, was not sensitive enough to capture such relations and might be vulnerable to artifacts depending on the number of data points included per bin. However, this relation did not hold even when the onset of this decline in accuracy was characterized by the location of peak accuracy levels using a non-binning approach. Overall, these results suggest that there is

no relation between decreasing accuracy and timing uncertainty. Importantly, however, our analyses showed that slopes were less negative in the long deadline condition compared to the short deadline condition, suggesting that interval timing still had an effect on participants' choice behavior.

There are at least three possible explanations for sub-optimal behavior in the deadline blocks. First, participants may have kept favoring accuracy over reward rate throughout the experiment, which has been previously reported (e.g., Maddox and Bohil,

2004; Bogacz et al., 2006, 2010; Balci et al., 2011b). Thus, accuracy bias could have prevented within trial modulation of thresholds to reduce overall error rates. This possibility relies on the implicit assumption that errors are subjectively more costly than missed trials. Second, participants may have started collapsing thresholds later than the optimal case due to underestimation of the deadline. In this case, accuracy would remain above the chance level at the time of response deadline. However, our analyses did not support this possibility. Third, sub-optimal decision making may be caused by mechanistic limitations at the neuronal level which may not allow for within-trial decision threshold modulation, at least for decisions made in less than one second. This is a plausible explanation of our results, given that the cognitive cost (i.e., executive load) of modulating the value of the decision threshold in real-time may outweigh its benefits in terms of increasing the overall reward attained throughout a session. Importantly, participants differed in terms of decreasing and increasing accuracy with time (see **Figure 9**, where some participants' accuracies increased rather than decreased toward the deadline), which could again be explained by individual differences in bias toward accuracy, as opposed to maximizing reward.

Slower RTs on error trials are commonly found in 2AFC research with free responding (Ratcliff and Rouder, 1998; Ratcliff and McKoon, 2008). These patterns can be accounted for by the extended DDM by allowing the drift rate to vary between trials. Drift variability enables the extended DDM to account for slower average error RTs than correct RTs. Inflation of this variability parameter (in addition to decreasing the constant threshold) should therefore produce decreasing accuracy with slower RTs in conditional accuracy curves, even in the absence of collapsing thresholds within a trial. Our simulations confirmed that accuracy can decline steadily with RT without any accompanying threshold collapse. We have shown that, while a concomitantly decreasing threshold parameter yields an additionally higher rate of decline in accuracy, a similar effect is observed by increasing drift rate variability across trials, whereas modifying starting point variability had no such effect. This lack of a visible effect of the starting point parameter on the rate of decline in accuracy with time was expected, given that increasing this parameter results in faster error RTs, which should not necessarily translate into slower error RTs when the same parameter is decreased. Overall, these results suggest that increasing drift rate variability or setting the constant decision threshold to a lower value might be a way to mimic the effect of collapsing thresholds on accuracy without actually collapsing them.

Finally, it is also important to note that a cross-over between faster and slower error responses has been suggested depending on the difficulty of the task (see Luce, 1986). Namely, harder tasks (i.e., higher error rates) have been shown to lead to slower RTs for error trials, whereas participants had faster error RTs in easier tasks (e.g., Ratcliff and Rouder, 1998). It is possible that our task was a relatively easy one, given the low error rates observed (**Figure 6**), the small number of trials in the last RT bin of the conditional accuracy curves (**Figure 9**), and a relatively high estimated drift rate (i.e., 0.4447) (see Section Drift-Diffusion Model Simulations). However, we still observe slower RTs for error trials, as can be seen in **Figure 8**. Therefore, studies using an easier task

still may not observe a more pronounced decline in accuracy with time, but this remains an open question.

Future studies should increase the cost of missing a deadline by explicitly adding a penalty. Under such payoff structures, one might be more likely to observe threshold collapsing. However, note that in these cases the optimal threshold collapse trajectories will also change (possibly meeting prior to the response deadline) due to the explicit penalty for late responses. Additionally, speed-accuracy tradeoff functions in tasks that use response signal methodology do not exhibit reduction in accuracy with increasing lags (e.g., Wickelgren, 1977). On the other hand, in our free response paradigm, such decline in accuracy was apparent in conditional accuracy curves. Response signal paradigms typically employ a single signal (or a series of equally distributed signals) after which the participant is instructed to respond as soon as possible, ensuring that there are no fast guesses, in addition to making within trial strategic manipulation of decision making parameters harder (Heitz, 2014). This difficulty is due to the fact that, by the time the response signal is given, subjects need to make a choice using the already accumulated (and potentially partial) evidence. This approach contrasts with the one we have used in a number of ways. First, subjects do not necessarily need to keep track of the time to respond in response signal tasks, whereas in our experimental design, participants needed to constantly rely on endogenous markers of the passage of time in order to maximize reward, which is likely more taxing in terms of information processing throughout the decision process. In turn, the relatively higher amount of cognitive resources available to the decision maker in the response signal paradigm might present itself as lower variability in drift rate, which as we showed can underlie declining accuracy with time. Secondly, the response signal paradigm allows post-signal accumulation of evidence to a certain extent, whereas our methodology does not permit it at all. As a result, one might expect that, even if participants were able to modulate thresholds within a trial (which we show here to not be the case), giving the chance to accumulate more evidence after a response signal might obscure a decline in accuracy with slower RTs. Further empirical work is needed to elucidate the possible sources of these differences between the two experimental paradigms, although the similarity of the implementation of SAT by decision makers has been questioned due to fundamental differences in the two approaches (see Heitz, 2014).

Overall, our empirical results do not support the optimal performance predictions regarding within-trial collapsing of thresholds under response deadlines. A slight decline in accuracy was observed for decisions made near the response deadlines; however, this decline never reached chance level, which is predicted by optimal threshold collapse. Moreover, the observed decline in accuracy was not related to the level of endogenous timing uncertainty in the expected direction, and it could be accounted for by DDM parameters that are constant within trials.

AUTHOR NOTE

ACKNOWLEDGMENTS

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: http://journal.frontiersin.org/Journal/10.3389/fnins.2014.00248/abstract

REFERENCES

Balci, F., Freestone, D., Simen, P., deSouza, L., Cohen, J. D., and Holmes, P. (2011a). Optimal temporal risk assessment. *Front. Integr. Neurosci.* 5:56. doi: 10.3389/fnint.2011.00056

Balci, F., Gallistel, C. R., Allen, B. D., Frank, K., Gibson, J., and Brunner, D. (2009). Acquisition of peak responding: what is learned? *Behav. Processes* 80, 67–75. doi: 10.1016/j.beproc.2008.09.010

Balci, F., Simen, P., Niyogi, R., Saxe, A., Hughes, J., Holmes, P., et al. (2011b). Acquisition of decision-making criteria: reward rate ultimately beats accuracy. *Atten. Percept. Psychophys.* 73, 640–657. doi: 10.3758/s13414-010-0049-7

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., and Cohen, J. D. (2006). The physics of optimal decision making: a formal analysis of models of performance in two alternative forced choice tasks. *Psychol. Rev.* 113, 700–765. doi: 10.1037/0033-295X.113.4.700

Bogacz, R., Hu, P. T., Holmes, P., and Cohen, J. D. (2010). Do humans produce the speed-accuracy tradeoff that maximizes reward rate?. *Q. J. Exp. Psychol.* 63, 863–891. doi: 10.1080/17470210903091643

Brainard, D. H. (1997). The psychophysics toolbox. *Spat. Vis.* 10, 433–436. doi: 10.1163/156856897X00357

Buhusi, C. V., and Meck, W. H. (2005). What makes us tick? Functional and neural mechanisms of interval timing. *Nat. Rev. Neurosci.* 6, 755–765. doi: 10.1038/nrn1764

Çavdaroğlu, B., Zeki, M., and Balci, F. (2014). Time-based reward maximization. *Philos. Trans. R. Soc. B Biol. Sci.* 369:20120461. doi: 10.1098/rstb.2012.0461

Drugowitsch, J., Moreno-Bote, R., Churchland, A. K., Shadlen, M. N., and Pouget, A. (2012). The cost of accumulating evidence in perceptual decision making. *J. Neurosci.* 32, 3612–3628. doi: 10.1523/JNEUROSCI.4010-11.2012

Feng, S., Holmes, P., Rorie, A., and Newsome, W. T. (2009). Can monkeys choose optimally when faced with noisy stimuli and unequal rewards? *PLoS Comput. Biol.* 5:e1000284. doi: 10.1371/journal.pcbi.1000284

Gibbon, J. (1977). Scalar expectancy theory and Weber's law in animal timing. *Psychol. Rev.* 84, 279–325. doi: 10.1037/0033-295X.84.3.279

Gold, J. I., and Shadlen, M. N. (2001). Neural computations that underlie decisions about sensory stimuli. *Trends Cogn. Sci.* 5, 10–16. doi: 10.1016/S1364-6613(00)01567-9

Heitz, R. P. (2014). The speed-accuracy tradeoff: history, physiology, methodology, and behavior. *Front. Neurosci.* 8:150. doi: 10.3389/fnins.2014.00150

Herzog, M. H., and Fahle, M. (1997). The role of feedback in learning a vernier discrimination task. *Vis. Res.* 37, 2133–2141. doi: 10.1016/S0042-6989(97)00043-6

Hudson, T. E., Maloney, L. T., and Landy, M. S. (2008). Optimal compensation for temporal uncertainty in movement planning. *PLoS Comput. Biol.* 4:e1000130. doi: 10.1371/journal.pcbi.1000130

Jazayeri, M., and Shadlen, M. N. (2010). Temporal context calibrates interval timing. *Nat. Neurosci.* 13, 1020–1026. doi: 10.1038/nn.2590

Kleiner, M., Brainard, D., and Pelli, D. (2007). What's new in Psychtoolbox-3? *Perception* 36, ECVP, Abstract Supplement. Available online at: http://www.perceptionweb.com/ecvp07/ecvp07.pdf

Laming, D. R. J. (1968). *Information Theory of Choice Reaction Time*. New York, NY: Wiley.

Luce, R. D. (1986). *Response Times*. New York, NY: Oxford University Press.

Maddox, W. T., and Bohil, C. J. (2004). Probability matching, accuracy maximization, and a test of the optimal classifier's independence assumption in perceptual categorization. *Percept. Psychophys.* 66, 104–118. doi: 10.3758/BF03194865

Malapani, C., and Fairhurst, S. (2002). Scalar timing in animals and humans. *Learn. Motiv.* 33, 156–176. doi: 10.1006/lmot.2001.1105

McKoon, G., and Ratcliff, R. (2012). Aging and IQ effects on associative recognition and priming in item recognition. *J. Mem. Lang.* 66, 416–437. doi: 10.1016/j.jml.2011.12.001

Mulder, M. J., Bos, D., Weusten, J. M. H., Van Belle, J., Van Dijk, S. C., Simen, P., et al. (2010). Basic impairments in regulating the speed-accuracy tradeoff predict symptoms of attention-deficit/hyperactivity disorder. *Biol. Psychiatry* 68, 1114–1119. doi: 10.1016/j.biopsych.2010.07.031

Pelli, D. G. (1997). The videotoolbox software for visual psychophysics: transforming numbers into movies. *Spat. Vis.* 10, 437–442. doi: 10.1163/156856897X00366

Frazier, P., and Yu, A. J. (2008).Sequential hypothesis testing under stochastic deadlines. *Adv. Neural Inf. Process. Syst.* 20, 465–472.

Rao, R. P. N. (2010). Decision making under uncertainty: a neural model based on partially observable Markov decision processes. *Front. Comput. Neurosci.* 4:146. doi: 10.3389/fncom.2010.00146

Ratcliff, R. (1978). A theory of memory retrieval. *Psychol. Rev.* 85, 59–108. doi: 10.1037/0033-295X.85.2.59

Ratcliff, R. (1981). A theory of order relations in perceptual matching. *Psychol. Rev.* 88, 552–572. doi: 10.1037/0033-295X.88.6.552

Ratcliff, R. (1985). Theoretical interpretations of speed and accuracy of positive and negative responses. *Psychol. Rev.* 92, 212–225. doi: 10.1037/0033-295X.92.2.212

Ratcliff, R. (1988). Continuous versus discrete information processing: modeling the accumulation of partial information. *Psychol. Rev.* 95, 238–255. doi: 10.1037/0033-295X.95.2.238

Ratcliff, R. (2002). A diffusion model account of reaction time and accuracy in a brightness discrimination task: fitting real data and failing to fit fake but plausible data. *Psychon. Bull. Rev.* 9, 278–291. doi: 10.3758/BF03196283

Ratcliff, R., and McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. *Neural Comput.* 20, 873–922. doi: 10.1162/neco.2008.12-06-420

Ratcliff, R., and Rouder, J. N. (1998). Modeling response times for two-choice decisions. *Psychol. Sci.* 9, 347–356. doi: 10.1111/1467-9280.00067

Ratcliff, R., and Smith, P. L. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychol. Rev.* 111, 333–367. doi: 10.1037/0033-295X.111.2.333

Seitz, A. R., Nanez, J. E. Sr., Holloway, S., Tsushima, Y., and Watanabe, T. (2006). Two cases requiring external reinforcement in perceptual learning. *J. Vis.* 6, 966–973. doi: 10.1167/6.9.9

Shadlen, M. N., and Newsome, W. T. (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J. Neurophysiol.* 86, 1916–1935.

Simen, P., Balci, F., deSouza, L., Holmes, P., and Cohen, J. D. (2011). A model of interval timing by neural integration. *J. Neurosci.* 31, 9238–9253. doi: 10.1523/JNEUROSCI.3121-10.2011

Simen, P., Contreras, D., Buck, C., Hu, P., Holmes, P., and Cohen, J. D. (2009). Reward-rate optimization in two-alternative decision making: empirical tests of theoretical predictions. *J. Exp. Psychol. Hum. Percept. Perform.* 35, 1865–1897. doi: 10.1037/a0016926

Spaniol, J., Voss, A., Bowen, H. J., and Grady, C. L. (2011). Motivational incentives modulate age differences in visual perception. *Psychol. Aging* 26, 932–939. doi: 10.1037/a0023297

Standage, D., You, H., Wang, D.-H., and 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. *Front. Comput. Neurosci.* 5:7. doi: 10.3389/fncom.2011.00007

Standage, D., You, H., Wang, D.-H., and Dorris, M. C. (2013). Trading speed and accuracy by coding time: a coupled-circuit cortical model. *PLoS Comput. Biol.* 9:e1003021. doi: 10.1371/journal.pcbi.1003021

Starns, J. J., and Ratcliff, R. (2010). The effects of aging on the speed-accuracy compromise: boundary optimality in the diffusion model. *Psychol. Aging* 25, 377–390. doi: 10.1037/a0018022

Starns, J. J., and Ratcliff, R. (2012). Age-related differences in diffusion model boundary optimality with both trial-limited and time-limited tasks. *Psychon. Bull. Rev.* 19, 139–145. doi: 10.3758/s13423-011-0189-3

Stone, M. (1960). Models for choice-reaction time. *Psychometrika* 25, 251–260. doi: 10.1007/BF02289729

Thura, D., Beauregard-Racine, J., Fradet, C.-W., and Cisek, P. (2012). Decision making by urgency gating: theory and experimental support. *J. Neurophysiol.* 108, 2912–2930. doi: 10.1152/jn.01071.2011

Vandekerckhove, J., and Tuerlinckx, F. (2008). Diffusion model analysis with MATLAB: a DMAT primer. *Behav. Res. Methods* 40, 61–72. doi: 10.3758/BRM.40.1.61

Voss, A., Nagler, M., and Lerche, V. (2013). Diffusion models in experimental psychology: a practical introduction. *Exp. Psychol.* 60, 385–402. doi: 10.1027/1618-3169/a000218

Wald, A., and Wolfowitz, J. (1948). Optimum character of the sequential probability ratio test. *Ann. Math. Stat.* 19, 326–339. doi: 10.1214/aoms/1177730197

Wearden, J. H. (2003). "Applying the scalar timing model to human time psychology: progress and challenges" in *Time and Mind II,* ed H. Helfrich (Göttingen: Hogrefe & Huber), 21-39.

White, C. N., Ratcliff, R., Vasey, M. W., and McKoon, G. (2010). Using diffusion models to understand clinical disorders. *J. Math. Psychol.* 54, 39–52. doi: 10.1016/j.jmp.2010.01.004

Wickelgren, W. A. (1977). Speed-accuracy tradeoff and information processing dynamics. *Acta Psychologica* 41, 67–85. doi: 10.1016/0001-6918(77)90012-9

Zakay, D., and Block, R. A. (1996). "The role of attention in time estimation processes," in *Time, Internal Clocks and Movement,* eds M. A. Pastor and J. Artieda (Amsterdam: Elsevier Science), 143–164.

A comparative study of drift diffusion and linear ballistic accumulator models in a reward maximization perceptual choice task

Stephanie Goldfarb[1]*, Naomi E. Leonard[2,3,4], Patrick Simen[5], Carlos H. Caicedo-Núñez[2] and Philip Holmes[2,3,4]

[1] HRL Laboratories, Center for Neural and Emergent Systems, Malibu, CA, USA
[2] Mechanical and Aerospace Engineering, Princeton University, Princeton, NJ, USA
[3] Program in Applied and Computational Mathematics, Princeton University, Princeton, NJ, USA
[4] Princeton Neuroscience Institute, Princeton University, Princeton, NJ, USA
[5] Neuroscience Department, Oberlin College, Oberlin, OH, USA

We present new findings that distinguish drift diffusion models (DDMs) from the linear ballistic accumulator (LBA) model as descriptions of human behavior in a two-alternative forced-choice reward maximization (Rmax) task. Previous comparisons have not considered Rmax tasks, and differences identified between the models' predictions have centered on practice effects. Unlike the parameter-free optimal performance curves of the pure DDM, the extended DDM and LBA predict families of curves depending on their additional parameters, and those of the LBA show significant differences from the DDMs, especially for poorly discriminable stimuli that incur high error rates. Moreover, fits to behavior reveal that the LBA and DDM provide different interpretations of behavior as stimulus discriminability increases. Trends for threshold setting (caution) in the DDMs are consistent between fits, while in the corresponding LBA fits, thresholds interact with distributions of starting points in a complex manner that depends upon parameter constraints. Our results suggest that reinterpretation of LBA parameters may be necessary in modeling the Rmax paradigm.

Keywords: drift diffusion model, linear ballistic accumulator model, reward maximization, optimal performance theory

1. INTRODUCTION

Among the many models proposed to describe decision tasks, leaky competing accumulators (LCAs) (Usher and McClelland, 2001) and drift diffusion models (DDMs) e.g., Ratcliff and Rouder (1998) have been especially prominent. More recently the linear ballistic accumulator (LBA) (Brown and Heathcote, 2008) was introduced as a conceptually simpler alternative to DDMs. All these models employ drift terms that describe mean rates of evidence accumulation, thresholds that signal decision times when crossed, and sources of variability, either within or across trials. All have been validated against particular behavioral data, but since they differ in structure, number of parameters, and the manner in which variability enters, they may suggest different processing mechanisms [although the DDM can be derived from the LCA under certain conditions (Bogacz et al., 2006)]. It is therefore of interest to compare their accounts of given data sets.

The comparative study of Donkin et al. (2011) revealed few differences between the abilities of the LBA and DDM to fit and predict behavioral data. However, an earlier comparison of DDM fits to simulated data from LBA, DDM, and LCA found that DDM and LCA parameters correlated in a one-to-one manner, but those of LBA and DDM did not (van Ravenzwaaij and Oberauer, 2009). Subsequently, differences in drift rates, non-decision times and

caution parameters were found in many-parameter fits of practice effects (Heathcote and Hayes, 2012), but these differences were not connected to optimal theories of performance in perceptual choice tasks. LBA fits were not included in a substantial recent paper (Teodorescu and Usher, 2013) that compared several race and LCA models. Nor have the LBA and DDMs been compared for reward maximization (Rmax) tasks in which participants have learned strategies and apply task-based knowledge to optimize performance.

In Rmax tasks participants are instructed to adopt a strategy that yields maximum rewards, and are given a fixed time interval to complete each block of trials, during which they may attempt the task as many times as they wish, as detailed in section 2.3. Task difficulty is held constant within a block but varied between blocks. In the two-alternative forced-choice (2AFC) task from which data is analyzed here, visual moving dots stimuli were used and task difficulty was adjusted via motion coherence (Balci et al., 2011). Depending on difficulty, a participant may attempt the task few times, slowly and cautiously, or she may work faster but more carelessly. The 2AFC Rmax task performance and model fits have been tested against the DDM, and DDM fits have been shown to describe a speed-accuracy tradeoff quite close to that of high performing participants (Bogacz et al., 2006; Simen et al.,

2009; Bogacz et al., 2010; Balci et al., 2011). However, Rmax task performance and fits have not previously been compared across models.

Here we compare the LBA, which represents evidence in favor of two or more options, with the pure and extended (Ratcliff, e.g., Ratcliff and Rouder, 1998) DDMs, which assess differences of evidence between options. In the LBA two drift rates, believed to be correlated with neural activity (e.g., Gold and Shadlen, 2000, 2001; Gold et al., 2008), represent preferences for each of the two options; in the DDMs, a single drift rate represents the difference between these preferences.

For both the DDMs and the LBA models, thresholds, also called caution parameters (Donkin et al., 2011), set a level of accumulated activity at which a decision is made. Caution is key in setting the speed-accuracy tradeoff: high caution implies low speed and high accuracy and low caution implies high speed and low accuracy (Bogacz et al., 2006; Brown and Heathcote, 2008; Balci et al., 2011). For example, caution can explain the relatively slow response times of elderly individuals (Ratcliff et al., 2004). Caution can be experimentally manipulated by adjusting task difficulty from block to block, and optimal values of caution can be determined analytically for the pure DDM and numerically for the extended DDM and the LBA, as shown below in section 2.2.

An important difference between the DDMs and the LBA models is the treatment of variability. In the DDMs variability enters as additive Gaussian noise in the evidence accumulation dynamics during each individual trial. In the extended DDM, there are additional trial-to-trial variabilities in the starting point, in the drift rate of evidence accumulation, and in the non-decision time. In contrast, there is no additive noise in the LBA during individual trials, as implied by the adjectives "linear" and "ballistic." Instead there is only trial-to-trial variability in the starting points and in the drift rates. Nonetheless, the LBA models can capture much of the same behavior as the extended DDM, and they do so with fewer parameters (Brown and Heathcote, 2008; Donkin et al., 2011).

Direct numerical comparisons of the role of caution in the models are straightforward. Parameters can be fit to participant behavior at each difficulty level for each model. For a given parameter set, changes in speed and accuracy of responses as caution is varied can be computed, and thus the value of the caution

parameter that yields the highest reward rate can be determined for each difficulty level. An optimal speed-accuracy tradeoff for each model and participant can then be derived, assuming that caution is varied from one difficulty condition to the next.

The models can also be evaluated and compared by examining their predictions of optimal performance. For the pure DDM, a unique parameter-free Optimal Performance Curve (OPC) describes the relationship between error rate (ER) and a normalized decision time (DT), independent of model parameters (Bogacz et al., 2006). Parameterized families of OPCs may also be determined for the extended DDM, and as the values of the additional parameters (variances in drift rate, in starting point, and in non-decision time) become small, these curves approach that of the pure DDM (Bogacz et al., 2006). Like the extended DDM, the LBA does not predict unique OPCs, and our analysis of the LBA reveals a critical interaction between thresholds and variance in starting points.
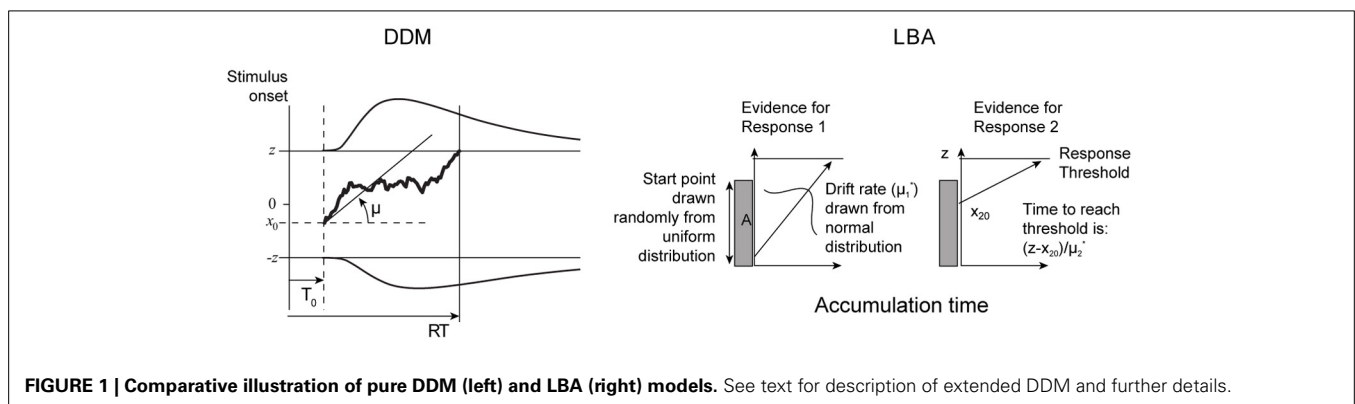
While the DDMs and the LBA can reproduce key aspects of behaviors, the DDM fits suggest that participants are on average *least* cautious on the most difficult tasks, in which the optimal strategy is random guessing. In contrast, an LBA fit indicates that participants are on average *more* cautious on the most difficult tasks, and that they reduce variance in starting points as difficulty decreases. These differences between the DDM and the LBA predictions highlight the role of diffusive noise within trials in the DDMs, which is sacrificed for simplicity in the LBA model.

The paper is structured as follows. In section 2 we discuss our methods, describing the LBA model, the pure and extended DDMs, and parameter-fitting procedures. We review optimal performance theory for the pure and extended DDMs, develop analogous results for the LBA, propose an adapted LBA that decouples starting point from thresholds and compare model performances in the limit of large noise. Section 3 describes our results primarily in terms of parameter fits across subjects, and section 4 contains a discussion and directions for future work. Details of fits to individual participants are provided in the Supplementary Material.

2. MATERIALS AND METHODS
2.1. DRIFT DIFFUSION AND LINEAR BALLISTIC ACCUMULATION
In this section we describe the pure and extended (Ratcliff) DDMs and the LBA models. The two models are illustrated in **Figure 1**.

FIGURE 1 | Comparative illustration of pure DDM (left) and LBA (right) models. See text for description of extended DDM and further details.

2.1.1. Pure drift diffusion model

In the pure DDM the difference in evidence for the two choices evolves according to the following scalar equation:

$$dx = \mu\,dt + \sigma\,dW, \; x(0) = x_0, \tag{1}$$

where $x(t)$ is the aggregate evidence at time t, μ is the drift rate, σ is the diffusion rate, and dt and dW represent time and Wiener noise increments, respectively. Evidence accumulates noisily from the starting point $x(0) = x_0$ at time $t = 0$ to the first time $t = T$ at which $x(T) = +z$ or $-z$. Without loss of generality, we assume that $\mu \geq 0$ (Bogacz et al., 2006). Thus, the two thresholds, $+z$ and $-z$, respectively, correspond to selecting the correct and incorrect choices. We will refer to z interchangeably as the *threshold* or the *caution* (parameter); z can take any non-negative value.

Only four parameters are required to predict DT for Equation (1): μ, x_0, σ, and z, and there are closed form analytical expressions for mean DT and ER (Bogacz et al., 2006):

$$\mathrm{DT} = \frac{z}{\mu}\tanh\left(\frac{z\mu}{\sigma^2}\right) + \frac{2z}{\mu}\cdot\left(\frac{1 - \exp\left(\frac{2x_0\mu}{\sigma^2}\right)}{\exp\left(\frac{2z\mu}{\sigma^2}\right) - \exp\left(-\frac{2z\mu}{\sigma^2}\right)}\right) - \frac{x_0}{\mu}, \tag{2}$$

$$\mathrm{ER} = \frac{1}{1 + \exp\left(\frac{2z\mu}{\sigma^2}\right)} - \left(\frac{1 - \exp\left(\frac{2x_0\mu}{\sigma^2}\right)}{\exp\left(\frac{2z\mu}{\sigma^2}\right) - \exp\left(-\frac{2z\mu}{\sigma^2}\right)}\right). \tag{3}$$

In addition, the pure DDM is augmented by a non-decision time parameter, T_0, corresponding to non-decision processes. The estimated reaction time (RT) is the sum of the decision and non-decision times: RT = DT + T_0.

Although our data were derived from unbiased stimuli, we allow non-zero starting points in order to make direct comparisons among all the models, since extended DDMs and LBAs use ranges of starting points.

2.1.2. Extended drift diffusion model

In the extended (Ratcliff) DDM, evidence accumulation in each trial is governed by the same process as in the pure DDM, but with added variability in starting points, drift rate, and non-decision time:

$$dx = \mu^*dt + \sigma\,dW, \; x(0) = x_0^*, \tag{4}$$

where μ^*, σ, z, x_0^*, and T_0^*, respectively represent the drift rate, diffusion rate, threshold, starting point, and non-decision time for a given trial. Evidence accumulation proceeds from the starting point $x(0) = x_0$ at time $t = 0$ to the first time $t = T$ at which $x(T) = +z$ or $-z$. For each trial, μ^* is selected from $\mathcal{N}\left(\mu, s_\mu^2\right)$, x_0^* is selected from $\mathcal{U}\left(x_0 - \frac{s_{x_0}}{2}, x_0 + \frac{s_{x_0}}{2}\right)$, and T_0^* is selected from $\mathcal{N}\left(T_0, s_{T_0}^2\right)$, where $\mathcal{N}$ and $\mathcal{U}$ respectively denote Gaussian (normal) and uniform distributions, and μ, s_μ, s_{x_0}, T_0, s_{T_0} are all non-negative constant scalars.

The additional variability in the model parameters from trial to trial augments the model's descriptive power. In particular, the extended DDM, unlike the pure DDM, can predict different

RT distributions for correct and error trials, even with unbiased mean starting points. Prior work has suggested that the parameters new to the extended DDM sufficiently extend the descriptive capabilities of the DDM to merit the additional modeling cost (Ratcliff and Rouder, 1998; Ratcliff and Smith, 2004; Balci et al., 2011). However, analytical expressions for DT and ER analogous to Equations (2, 3) do not exist for the extended DDM. The extended DDM is frequently called the Ratcliff DDM due to a large body of work by Ratcliff to characterize it. Here the threshold z can assume any non-negative value outside the range of starting points (Tuerlinckx, 2004).

2.1.3. Linear ballistic accumulator model

The LBA model is conceptually simple and yet can provide rich descriptions of behavior, rivaling those of the extended DDM (Brown and Heathcote, 2008; Donkin et al., 2011). Evidence $x_i(t)$ for each of two or more choices accumulates *linearly* and *ballistically* in time t from $x_i(0) = x_{i0}*$ toward a threshold z at a drift rate μ_i^*:

$$x_i(t) = x_{i0}^* + \mu_i^* t, \; i = 1, 2, \ldots, N. \tag{5}$$

As in the extended DDM, parameters may vary from trial to trial. Here μ_i^* is selected from $\mathcal{N}\left[\mu_i, s^2\right]$ and x_{i0}^* from $\mathcal{U}[0, A]$ on each trial. The parameter $A > 0$ defines the maximum value that any starting point x_{i0}^* may assume. The accumulator $x_i(t)$ that is first to reach the threshold z is selected. In prior work, A has been restricted to lie below z, i.e., $A \leq z$ (Brown and Heathcote, 2008; Donkin et al., 2011). A non-decision time T_0 is also included. While drift rates generally differ for each accumulator ($\mu_i \neq \mu_j$), the remaining parameters A, z, s, T_0 are common to all accumulators.

Closed form expressions describing the LBA model's behavior were derived in Brown and Heathcote (2008). The cumulative distribution function (CDF), $F_i(t)$, and the probability density function (PDF), $f_i(t)$, can be written in terms of the LBA parameters for individual accumulators:

$$F_i(t) = 1 + \frac{z - A - t\mu_i}{A}\Phi\left(\frac{z - A - t\mu_i}{ts}\right)$$
$$- \frac{z - t\mu_i}{A}\Phi\left(\frac{z - t\mu_i}{ts}\right) + \frac{ts}{A}\phi\left(\frac{z - A - t\mu_i}{ts}\right)$$
$$- \frac{ts}{A}\phi\left(\frac{z - t\mu_i}{ts}\right), \tag{6}$$

$$f_i(t) = \frac{1}{A}\left[-\mu_i\Phi\left(\frac{z - A - t\mu_i}{ts}\right) + s\phi\left(\frac{z - A - t\mu_i}{ts}\right)\right.$$
$$\left. + \mu_i\Phi\left(\frac{z - t\mu_i}{ts}\right) - s\phi\left(\frac{z - t\mu_i}{ts}\right)\right], \tag{7}$$

where $\Phi(\,\cdot\,)$ is the CDF and $\phi(\,\cdot\,)$ the PDF for the normal distribution with zero mean and unit variance. See Brown and Heathcote (2008, Supplementary Material) for the derivations of Equations (6, 7).

To determine mean first passage times for competing accumulations, we use the *defective PDF*, denoted $\mathrm{PDF}_i(t)$; unlike the

standard PDF, the defective PDF generally integrates to a value between 0 and 1. $\text{PDF}_i(t)$ describes the likelihood that accumulator $x_i(t)$ reaches the threshold *provided that no other accumulator has already done so*:

$$\text{PDF}_i(t) = f_i(t) \prod_{j \neq i} \left(1 - F_j(t)\right). \quad (8)$$

Because drift rates μ_i^* are selected from a normal distribution, in some cases all μ_i^*'s are negative. When this happens, the model produces an infinite decision time, and no response is given. Thus to compare LBA responses to those predicted by the two DDMs, which are finite on every trial, we consider only simulated LBA trials that yield a finite response time, i.e., that have at least one accumulator with a positive drift rate on that trial. To do so we modify the expressions above by a normalization factor $\alpha(\mu_1, \ldots, \mu_N, s) = 1 - \prod_{i=1}^N \Phi\left(-\frac{\mu_i}{s}\right)$, which is the probability that no accumulator reaches threshold. This follows since $\Phi\left(-\frac{\mu_i}{s}\right)$ is the probability that the ith accumulator has $\mu_i^* < 0$. The normalized defective probability density functions $p_i(t)$ are given in Brown and Heathcote (2008) as

$$p_i(t) = \frac{\text{PDF}_i(t)}{\alpha(\mu_1, \ldots, \mu_N, s)}. \quad (9)$$

For a two choice task, we therefore have

$$p_1(t) = \frac{\text{PDF}_1(t)}{1 - \Phi\left(-\frac{\mu_1}{s}\right)\Phi\left(-\frac{\mu_2}{s}\right)}, \quad (10)$$

$$p_2(t) = \frac{\text{PDF}_2(t)}{1 - \Phi\left(-\frac{\mu_1}{s}\right)\Phi\left(-\frac{\mu_2}{s}\right)}, \quad (11)$$

with $\int_0^\infty (p_1(t) + p_2(t))dt = 1$. The expressions for DT and ER are

$$\text{DT} = \int_0^\infty t(p_1(t) + p_2(t))dt, \quad (12)$$

$$\text{ER} = \int_0^\infty p_2(t)dt. \quad (13)$$

Following the convention adopted for the DDM, we shall assume that $\mu_1 \geq \mu_2$, so that $p_1(t)$ and $p_2(t)$ represent correct and incorrect responses, respectively, and the corresponding DTs may be written as

$$\text{DT}_{\text{correct}} = \frac{1}{1 - \text{ER}} \int_0^\infty t p_1(t)dt, \quad (14)$$

$$\text{DT}_{\text{error}} = \frac{1}{\text{ER}} \int_0^\infty t p_2(t)dt. \quad (15)$$

We also normalize the sum of the mean drift rates: $\mu_1 + \mu_2 = 1$. For the LBA described in the literature (Brown and Heathcote, 2008; Donkin et al., 2011), the threshold must not fall within the range of starting points, i.e., we must have $z \geq A$. The LBA, unlike the two DDMs, therefore almost always predicts non-zero DTs. Implications of this in determining an optimal speed-accuracy tradeoff are discussed in the next section.

2.2. OPTIMAL PERFORMANCES

As in Bogacz et al. (2006), we define optimal performance as a strategy that maximizes the Reward Rate (RR):

$$\text{RR} = \frac{1 - \text{ER}}{\text{DT} + T_0 + \text{RSI}}, \quad (16)$$
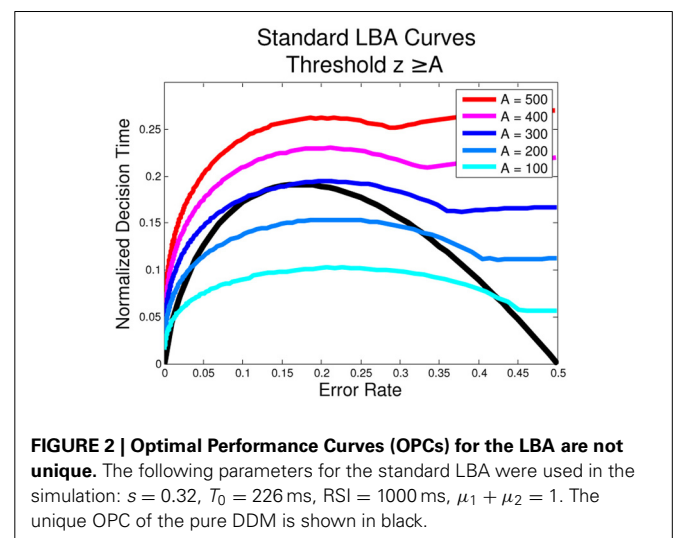
where RSI denotes the response to stimulus interval (see section 2.3 below). To assess performance, we seek a relationship between the behavioral measures ER and DT that yields the maximum RR for a given decision making model. This Optimal Performance Curve (OPC) (Bogacz et al., 2006) relates normalized DT to ER, where the former is defined as $\frac{\text{DT}}{\text{D}_{\text{tot}}}$ with $\text{D}_{\text{tot}} = T_0 + \text{RSI}$. We now describe OPCs for the DDM, extended DDM, and LBA.

2.2.1. Optimal performance under the pure DDM is described by a unique curve

The pure DDM has a unique, parameter-free OPC, defined by

$$\frac{\text{DT}}{\text{D}_{\text{tot}}} = \left(\frac{1}{\text{ER} \log \frac{1 - \text{ER}}{\text{ER}}} + \frac{1}{1 - 2\text{ER}}\right)^{-1}, \quad (17)$$

which is derived by finding the threshold that maximizes RR for a given task difficulty (Bogacz et al., 2006). This function is shown in solid black in **Figure 2** below. Its general shape can be intuitively explained by noting that for very noisy stimuli, prolonged evidence accumulation cannot improve much over random choices, so at the righthand end optimal thresholds approach zero, $\text{DT} \to 0$ and $\text{ER} \to 0.5$. Alternatively, very easy tasks require little accumulation to achieve high accuracy, so DTs are also small at the left, but $\text{ER} \to 0$. For each intermediate difficulty level and a given $\text{D}_{\text{tot}} = T_0 + \text{RSI}$, there is a unique optimal threshold with associated DT and ER between 0 and 0.5 that maximizes RR, thus defining the curve. All other thresholds, associated with faster or slower RTs, yield smaller net rewards at that task condition. Going from right to left, RRs rise as difficulty

FIGURE 2 | Optimal Performance Curves (OPCs) for the LBA are not unique. The following parameters for the standard LBA were used in the simulation: $s = 0.32$, $T_0 = 226$ ms, RSI = 1000 ms, $\mu_1 + \mu_2 = 1$. The unique OPC of the pure DDM is shown in black.

decreases, but it is important to recognize that any point on the OPC corresponds to maximum RR for a specific task condition. See Bogacz et al. (2006) and Zacksenhouse et al. (2010) for further discussion and illustrations of the OPC.

2.2.2. Optimal performance under the extended DDM is not uniquely defined

The extended DDM has families of OPCs rather than a unique OPC, as in the pure DDM. In the extended DDM, variability in starting points precludes the possibility of trials with a DT = 0. However, for low values of variance parameters in the extended DDM, the OPCs for the extended DDM approach the OPC for the pure DDM. For a sample OPC for the extended DDM see (Bogacz et al., 2006, Figure 14). To compute such curves, all parameters except drift rate and threshold are fixed. Then the threshold which optimizes RR is computed for each drift rate and used to determine ER and normalized DT. Further details can be found in Bogacz et al. (2006).

2.2.3. Optimal performance under the LBA is not uniquely defined

The LBA expressions of Equations (6–11) are complicated, and simple analytical expressions of their OPC families do not seem possible. Instead we approximate them numerically. To do this, we fix T_0, RSI, s, set $\mu_1 + \mu_2 = 1$ and choose A. We then calculate ER and DT for each $\mu_1 \in [0.5, 1]$. From these we estimate the optimal z and find the corresponding ER and DT, producing a point on the OPC for the selected A value. We find that a different OPC is generated for each value of A, as shown in **Figure 2**, i.e., there is no unique OPC for the LBA.

This is consistent with the observation that, for $\mu_1 = \mu_2 = 0.5$ (equal evidence for both options), different choices of A will affect the DT but not the ER. This is because the expected accuracy will be exactly 0.5 and no greater accuracy may be realized or information accumulated over time. It follows that the optimal solution is the lowest possible threshold and therefore the shortest possible DT.

For the pure DDM and the extended DDM with zero starting point variance ($s_x = 0$), the threshold parameter, z, can go to 0, and likewise the DT. However, in the LBA, the threshold must lie at or above the top of the range of starting points, i.e., $z \geq A$. Since this is a key source of variability, in general $A > 0$. Moreover, the lowest (and optimal) threshold for $\mu_1 = \mu_2 = \mu$ is therefore $z = A$, which gives DT > 0. The OPC curve plotting $\frac{DT}{DT_{tot}}$ varies with the value of A as shown in **Figure 2**; the smooth portions of each curve correspond to $z > A$ (on the left) and $z = A$ (on the right). That is, when the task is difficult and the ER is close to 0.5, the threshold, z, is as small as it can be. Moreover, unlike the OPC for the DDM, the OPCs for the LBA terminate at ER = 0.5 with finite normalized DTs. As $A \to 0$, the normalized DTs approach 0 for all ERs.

2.2.4. An adapted LBA allows fast responses at high error rates

The above analysis prompts us to define an *adapted version of the LBA*, in which thresholds can take values in the range of starting points, i.e., $z \geq A$ is relaxed to $z \geq 0$. If the starting point in one or both accumulators is greater than z, then DT = 0 and the accumulator with the higher starting point is selected for that trial.

The mean error rate and decision time, ER_a and DT_a, for the adapted DT with $z < A$ are defined accordingly:

$$ER_a(A, z, \mu_1, \mu_2, s) = \frac{1}{2} \cdot \frac{A - z}{A} + \frac{z}{A} \cdot ER(z, z, \mu_1, \mu_2, s), \quad (18)$$

$$DT_a(A, z, \mu_1, \mu_2, s) = \frac{z}{A} \cdot DT(z, z, \mu_1, \mu_2, s), \quad (19)$$

where ER and DT are calculated as in Equations (12–15) for the standard LBA.

Instantaneous decisions occurring for starting points above z can be seen as representing a prior resolution to respond as fast as possible, as is optimal for entirely noisy (zero-coherence) stimuli, see the OPC of **Figure 2**. In this case RT = T_0. Hence the adapted LBA can produce bimodal RT distributions, with a delta function at T_0 for trials with starting points above z and a second peak at longer RTs due to those starting below z.

Numerically-derived OPCs for the adapted LBA are also non-unique. As stimuli become less discriminable and ERs approach 0.5, the best values of z are those that minimize DT. Hence $z_{opt} \to 0$ as $\mu_1 \to \mu_2$, leading to many rapid responses. For $\mu_1 >> \mu_2$, we also find $z_{opt} \to 0$ and DT $\to 0$, but with ER $\to 0$ due to fast drift toward correct choices. As A varies this produces a family of OPCs with portions on the left similar to those of **Figure 2**, but approaching DT = 0 as ER $\to 0.5$. Also, as $A \to 0$, $z_{opt} \to 0$ for $\mu_1 \neq \mu_2$, as for the standard LBA.

2.2.5. Noise scales differently in the pure DDM and standard LBA

Poorly discriminable stimuli correspond to low signal-to-noise ratios μ/σ in the pure DDM, and may also correspond to variability in drift rates and starting points in the extended DDM. The LBA has trial-to-trial variability ("noise") in drift rates and starting points but lacks additive noise in individual trials. We now compare noise scaling in the two models in the case of equal mean evidence for both alternatives, represented as $\mu_1 = \mu_2$ in the LBA model and as $\mu = 0$ in the DDM. We show that the DDM and LBA behave differently as noise increases.

We first consider approximations of DT and ER about the point $\frac{1}{\sigma^2} = 0$ (analogous to $\mu = 0$) for the DDM. Taking μ and z fixed and strictly positive, and expanding the expressions (2, 3) with respect to the small parameter $\frac{1}{\sigma^2}$ in Taylor series, we obtain

$$DT\left(\frac{1}{\sigma^2}\right) = z^2\left(\frac{1}{\sigma^2}\right) - \frac{\mu^2 z^4}{3}\left(\frac{1}{\sigma^6}\right) + O\left(\frac{\mu^4 z^6}{\sigma^{10}}\right) \text{ and} \quad (20)$$

$$ER\left(\frac{1}{\sigma^2}\right) = \frac{1}{2} - \frac{z\mu}{2}\left(\frac{1}{\sigma^2}\right) + \frac{z^3 \mu^3}{6}\left(\frac{1}{\sigma^6}\right) + O\left(\frac{z^5 \mu^5}{\sigma^{10}}\right). \quad (21)$$

Here ER is $O(1)$, and DT is $O\left(\frac{1}{\sigma^2}\right)$. Thus, ER scales differently with high noise σ as compared to DT. In particular, as $\sigma \to \infty$, ER$\to 0.5$ and DT$\to 0$. Intuitively, large noise pushes the accumulation process rapidly toward one of the boundaries.

We now consider scaling of ER and DT with noise in the standard LBA, and show that large noise generally leads to non-zero DTs and always leads to non-zero mean DTs, given non-zero thresholds. For non-discriminable stimuli $\mu_1 = \mu_2 = \frac{1}{2}$ the LBA has two sources of noise or variability: in drift rates, s, and in

starting points, A. In all cases, since $\mu_1 = \mu_2$ and drift is the only source of bias, $ER = 0.5$. We note that the mean DT can be 0 if and only if $A = 0$, due to the constraint that $z \geq A$, so that $A \neq 0$ implies a non-zero distance to accumulate to threshold. To see this, first suppose that $s = 0$ and then allow s to increase, producing a distribution of drift rates centered around $\mu_i = \frac{1}{2}$. In fact for $s = 0$ the RT distribution is

$$f(t) = \begin{cases} 0, & t < \frac{z-A}{\mu} \\ \frac{2\mu Z}{A^2}\left(1 - \frac{\mu t}{z}\right), & t \in \left[\frac{z-A}{\mu}, \frac{z}{\mu}\right], \\ 0, & t > \frac{z}{\mu} \end{cases} \quad (22)$$

with mean $DT = (3z - 2A)/3\mu$, and this value will change continuously with s. It follows that for $\mu_1 = \mu_2$ and the minimum $z^* = A$, DT increases with A. This behavior for $\mu_1 = \mu_2$ and high A is quite different from that of the optimized pure DDM, in which large noise implies that mean DT $\to 0$.

2.3. REWARD MAXIMIZATION EXPERIMENT

In order to compare model fits, we reanalyze human behavioral data from a free response motion discrimination task previously presented in Balci et al. (2011). Participants ($n = 17$, 6 male) were asked to discriminate the direction of displays of moving dots on a computer screen and instructed to maximize their rewards. Task difficulty was adjusted via motion coherence determined by the fraction of dots moving leftward or rightward while the rest moved randomly.

Stimuli were viewed at ≈ 60 cm from the CRT monitor. The participant indicated motion direction by pressing a key on a standard keyboard: leftward with the "Z" key and rightward with the "M" key. Leftward and rightward stimuli were presented with equal probabilities. Premature responses, either anticipatory or with RTs of less than 100 ms, were penalized with a buzzing sound and a 4 s timeout period. When participants did not respond prematurely, RSIs were selected from a truncated exponential distribution with mean of 1 s. Experiments were conducted at a Macintosh computer, using the Psychophysics Toolbox (Brainard, 1997).

Each participant completed at least 13 daily sessions of 60 min total duration. The first four of these sessions involved training and practice without monetary reward. In each of the remaining sessions, participants completed five blocks with motion stimuli presented at a different coherence in each block (0, 4, 8, 16, and 32%, randomized across participants); participants earned $0.02 for each correct response. Performance improved markedly over the first 5 sessions and for certain participants continued to improve until session 9. Here we only use data from sessions 10–13.

After completing the motion discrimination task, participants performed an interval timing task and a signal detection task. The signal detection task was the same as above, except that participants were instructed to respond merely to motion onset, regardless of direction. In one block they were instructed to press the "M" key, and in the other, the "Z" key, again receiving $0.02 for each correct (non-anticipatory) response. The signal detection data was used to compute non-decision times as described in the

following section. Interval timing data is not used here, so we do not describe that task. For more details, see Balci et al. (2011).

2.4. DATA FITTING PROCEDURES

Fits were performed to participant data for the two DDMs and the LBA using published toolboxes for the models in Matlab (Tuerlinckx, 2004) and R (Donkin et al., 2009), respectively. Fits to the LBA model required some modifications to the standard LBA code, as outlined in Donkin et al. (2009).

Data were separated by difficulty and fits were computed for individual participants over all five difficulty conditions. Multiple fits were performed for each condition and participant, first varying only drift rate and caution (threshold) with difficulty level for the pure DDM, and then also varying the range of starting points, while the remaining parameters were held constant. The DDMs and LBA were fit separately to RT distributions for correct and error trials in each condition using five quantiles (10%, 30%, 50%, 70%, 90%). The same data and partitioning were used for both model fit toolboxes. The toolboxes fit non-decision times T_0, so that mean RTs and DTs can be computed for all models.

Empirical non-decision times were also estimated for each participant from their mean RTs for the fastest 25% of signal detection trials, as in Balci (Personal Communication, 2011) and Balci et al. (2011). These non-decision times were only used in computing normalized mean DTs for the human data shown below in **Figures 3**, **4** and Supplementary Figure S4; normalized mean DTs for the models were derived from the model fit toolboxes.
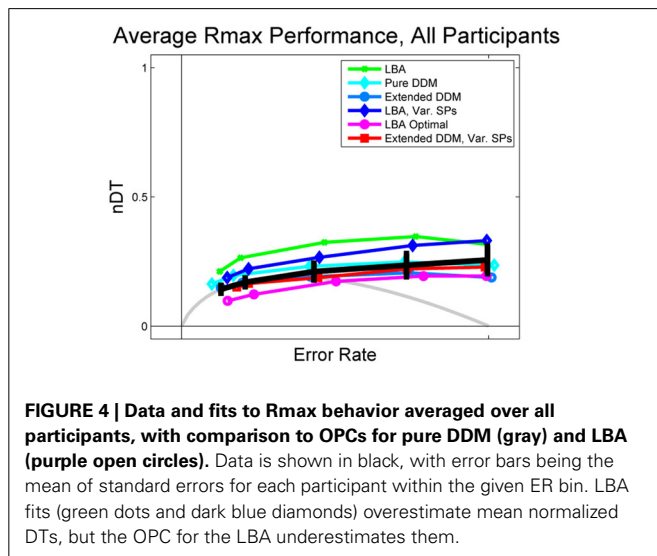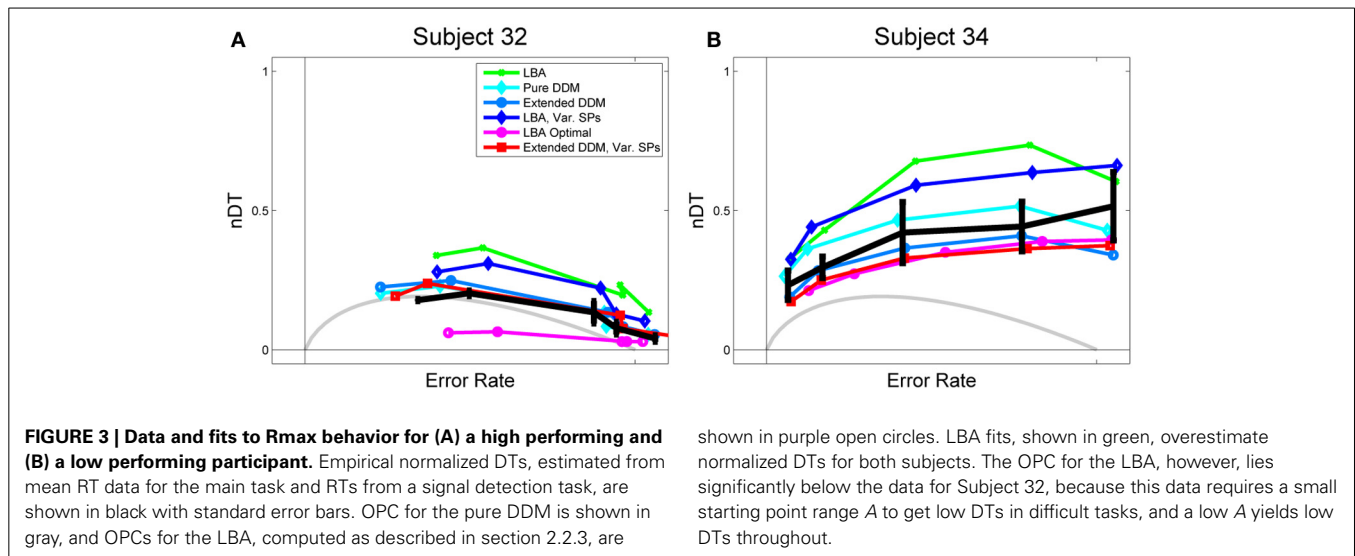
Both the DMAT (DDM) and LBA toolboxes allow the user to constrain some parameters to constant values while others are allowed to vary across conditions. The DMAT toolbox does this by using a combination of a system of matrix equations similar to those in general linear models, coupled with post-processing to remove outliers (Vandekerckhove and Tuerlinckx, 2007, 2008). The LBA toolbox uses the quantile maximum probability estimator method (Heathcote et al., 2004) to estimate PDFs for correct and error trials, which are then used to select model parameters.

The qualities of the resulting model fits were then assessed by comparing their predicted mean RTs and ERs with data for each condition and participant, using the Akaike, Corrected Akaike, and Bayesian Information Criteria (AIC, AIC_c, and BIC) (Akaike, 1974, 1980, 1981). Finally, in addition to each participant's actual performance, a theoretically optimal performance for each difficulty condition was estimated by allowing the caution parameter to vary freely while holding all other parameters fixed at that participant's fitted values. The optimal value of caution is defined as that yielding the highest possible reward rate, given constant (fitted) values for the remaining parameters.

3. RESULTS

To compare the properties of the DDM and LBA in fitting Rmax task data, we fitted the following models to data for each participant:

- A pure DDM, in which thresholds z and drift rates μ vary among coherence conditions (13 parameters: x_0, σ^2, T_0, 5 z's and 5 μ's).

FIGURE 3 | Data and fits to Rmax behavior for (A) a high performing and (B) a low performing participant. Empirical normalized DTs, estimated from mean RT data for the main task and RTs from a signal detection task, are shown in black with standard error bars. OPC for the pure DDM is shown in gray, and OPCs for the LBA, computed as described in section 2.2.3, are shown in purple open circles. LBA fits, shown in green, overestimate normalized DTs for both subjects. The OPC for the LBA, however, lies significantly below the data for Subject 32, because this data requires a small starting point range A to get low DTs in difficult tasks, and a low A yields low DTs throughout.

FIGURE 4 | Data and fits to Rmax behavior averaged over all participants, with comparison to OPCs for pure DDM (gray) and LBA (purple open circles). Data is shown in black, with error bars being the mean of standard errors for each participant within the given ER bin. LBA fits (green dots and dark blue diamonds) overestimate mean normalized DTs, but the OPC for the LBA underestimates them.

Table 1 | DDM and LBA model fit comparisons, with fit quality defined by the match to mean ER and RT data.

Model	Total parameters	AIC	AIC$_c$	BIC	R^2
Pure DDM	13	63.57	−1.43	93.50	0.98
Extended DDM	16	66.37	20.66	103.21	0.98
Extended DDM with var. SPs	20	66.76	30.40	112.81	0.99
LBA	13	80.44	15.44	110.37	0.95
LBA with var. SPs	17	66.18	23.68	105.32	0.99

Lower AIC, AIC$_c$, and BIC scores indicate better fits. See text for discussion.

- An extended DDM, in which thresholds z and mean drift rates μ^* vary among coherence conditions (16 parameters: x_0, σ^2, T_0, 5 z's, 5 μ's and s_x, s_{T_0}, s_μ).
- A second extended DDM, in which thresholds z, mean drift rates μ^* and ranges s_x of starting points (SPs) vary among coherence conditions (20 parameters: x_0, σ^2, T_0, 5 z's, 5 μ's, 5 s_x's, s_{T_0} and s_μ).
- A LBA model, in which thresholds z and mean drift rates μ_i vary among coherence conditions (13 parameters, with $\mu_1 + \mu_2 = 1$: A, s, T_0, 5 z's and 5 μ_1's).
- A second LBA model, in which thresholds z, mean drift rates μ_i and ranges A of starting points (SPs) vary among coherence conditions (17 parameters, with $\mu_1 + \mu_2 = 1$: 5 A's, s, T_0, 5 z's and 5 μ_1's).

Note that, for increased flexibility, σ is not set to 1.

AIC, AIC$_c$, and BIC scores for each participant and model were computed based on mean RT and ER data and model predictions. These were then averaged over all participants and conditions to determine mean scores for each model. All three metrics reward goodness of fit while penalizing extra parameters; lower scores are desirable and negative values are possible (Akaike, 1974, 1980, 1981). The scores, along with mean values of the correlation coefficient R^2, are summarized in **Table 1**. Figures S1 and S2 in the Supplementary Material show individual fits to RT and ER data. According to AIC, AIC$_c$, and BIC, the pure DDM provides the best overall fit to mean RT and ER data, but fits for each participant and condition are quite good for all five models.

Lower AIC, AIC$_c$, and BIC scores for the pure DDM are due to the fact that it has fewer parameters than all other models used here except the standard LBA, and the pure DDM predicts mean ER and RT data well. The extended DDM and LBA with variable starting points achieve slightly higher mean R^2 values than the pure DDM, indicating better fits to mean behavior when additional parameters are included. Comparing the pure DDM fit to the extended DDMs over RT *distributions* using the DDM toolbox, the extended DDM ($\chi^2 = 168.35$, $p < 0.05$) and extended DDM with variable starting points ($\chi^2 = 408.94$, $p < 0.001$)

yield superior deviance scores (Chernoff and Lehmann, 1954). However, fitting distributions is problematic because individual participant data sets separated by coherence condition are relatively small, and assessment of Rmax performance requires mean RTs and ERs.

Table 1 also shows that allowing the range of starting points to vary in the LBA yields better fits according to AIC and BIC but not AIC_c, and that such variability in the extended DDM does not improve fits according to any of these metrics.

Figure 3 shows normalized DTs for representative high and low performing participants (Subjects 32 and 34) and Figure 4 shows the DTs averaged over all participants. Individual fits appear in Figure S4 of the Supplementary Material. The data is shown in solid black with error bars, and model fits are superimposed in curves of various colors with different markers: the LBA in green dots, the pure DDM in light blue with diamonds, the extended DDM in dark blue with circles, the LBA with varied starting point ranges in dark blue with diamonds, and the extended DDM with varied starting point ranges in red with squares. The OPC for the LBA is indicated in purple with circles, and the OPC for the DDM in light gray.

Figure 3 illustrates a key difference between high and low performing participants. For the former (Figure 3A), the fits all trend downward and mean normalized DTs decrease as ERs increase, as they do for the OPC for the DDM. For the latter (Figure 3B), this trend is reversed and all fits diverge from the OPC for the DDM. The pure and extended DDM fits lie close to the data in both cases. The OPC for the LBA generally predicts the smallest normalized DTs, but it lies far below the data at intermediate ERs for the high performing participant.

As with individual subjects who do not perform at the highest level (Figure 3B), the average behavior shown in Figure 4 diverges from the OPC for the pure DDM as ERs increase. The difference between the empirical normalized mean DTs of Figure 4 and the OPC for the pure DDM is a good predictor of overall RR for each of the coherence conditions ($R^2 = 0.53$, $p < 0.001$).

The LBA models and the pure DDM overestimate normalized DTs and the extended DDMs slightly underestimate them. This is due in part to some subjectivity in the estimation of non-decision time T_0. For example, the LBA tends to fit smaller T_0 values than do the DDMs, and thus the LBA yields larger DTs (Donkin et al., 2011). However, while the LBA fits lie above the data curve, the OPC for the LBA lies below it, especially at intermediate ERs. This holds for many high performing participants (see Figure 3A and Supplementary Figure S4). For such individuals the range of starting points is small (Supplementary Figure S6), so that thresholds can be small without a major sacrifice in accuracy. We investigate how starting point ranges depend on coherence below (Figure 6A).

To better understand differences among the five models, we next consider mean parameter fits for the caution parameter, i.e., threshold, z. For each participant, coherence condition, and DDM fit, we calculated a threshold, and then averaged over individual threshold values for a given model and coherence. Thresholds for individual participants and coherence levels appear in Figure S5 of the Supplementary Material.

Figure 5A shows that mean threshold values increase with coherence for all three DDMs. Allowing starting point ranges to vary with coherence in the extended DDM increased thresholds at the two lowest coherence levels, and this variation with coherence over all thresholds is significant in the extended DDM [$F_{(4, 64)} = 73.72$, $p < 0.01$, $\eta^2 = 0.82$]. Extended DDM fits to the same data set in Balci et al. (2011) suggested approximately equal evidence for both variable and constant thresholds, because of differences in outlier treatment and in fitting algorithm options.

Figure 5B illustrates a similar analysis for the two LBA model fits. The fit to the LBA model with starting point range fixed over coherences also indicates that caution increases with coherence, but allowing starting point ranges to vary with coherence reverses this trend. A within-groups ANOVA on parameter values for the LBA with and without variability in starting point ranges shows that the main effect of LBA model type [$F_{(1, 16)} = 5.62$, $p < 0.05$,
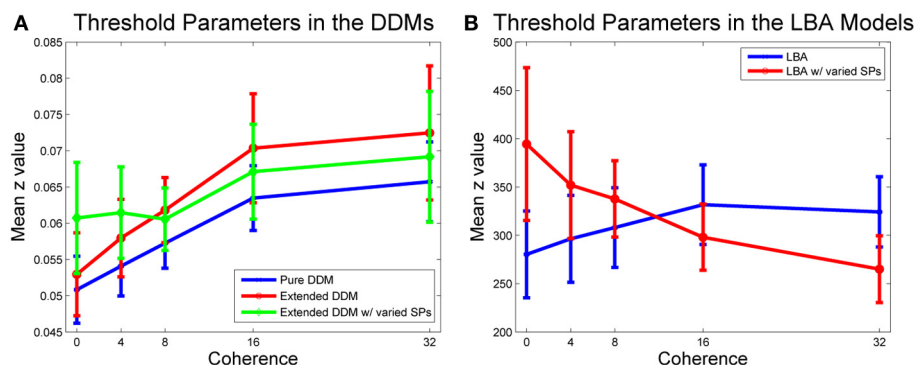
FIGURE 5 | Mean threshold values versus coherence. (A) Caution increases with coherence in all DDM fits. Pure (blue) and first extended DDM (red) fits do not differ significantly, but difficulty condition is significant [$F_{(4, 64)} = 3.91$, $p < 0.01$, $\eta^2 = 0.16$]. Allowing the starting point range to vary in the second extended DDM (green) does not improve fits by AIC/BIC, but compared with the other two DDM fits, main effects of model [$F_{(2, 32)} = 3.99$, $p < 0.05$, $\eta^2 = 0.02$] and interaction of model type and

difficulty condition [$F_{(8, 128)} = 2.06$, $p < 0.01$, $\eta^2 = 0.01$] are significant. (B) Caution increases with coherence in the first LBA fit, and decreases with coherence in the second LBA fit, with variable starting point range. A within-groups ANOVA shows that main effect of model type [$F_{(1, 16)} = 5.62$, $p < 0.05$, $\eta^2 = 0.01$] and interaction of model type and difficulty condition [$F_{(4, 64)} = 4.29$, $p < 0.01$, $\eta^2 = 0.08$] are both significant. Bars indicate standard errors for $n = 17$ subjects.

$\eta^2 = 0.01$] and the interaction of LBA model type and difficulty condition [$F_{(4, 64)} = 4.29$, $p < 0.01$, $\eta^2 = 0.08$] are both significant.

We next consider the role of starting point variability. Mean values of starting point ranges, averaged across all participants, are shown in **Figure 6A** for all models. The two models in which variability is allowed exhibit similar monotonically decreasing starting point ranges as coherence increases. Analogous data for individual participants appears in Figure S6 of the Supplementary Material, illustrating substantial variability among participants.

Figure 6B compares the DDM and LBA estimates of mean drift rates. Here LBA drift values are reduced by subtracting $\frac{1}{2}$ from μ_1, so that $\mu = 0$ corresponds to zero coherence in both models, allowing direct comparisons. All models predict a monotonically increasing mean drift rate with increasing coherence. The effect of coherence on drift parameters is significant for all models with a large effect size [$F_{(4, 64)} = 118.80$, $p < 0.001$, $\eta^2 = 0.76$]. Interaction of model and condition type is also significant, but effect size is modest [$F_{(16, 256)} = 5.48$, $p < 0.001$, $\eta^2 = 0.11$]. The effect of model type on drift is also significant [$F_{(4, 64)} = 7.79$, $p < 0.001$, $\eta^2 = 0.10$]. Drift rate estimates for individual participants appear in Figure S7 of the Supplementary Material, showing more uniformity than the starting point ranges of Supplementary Figure S6. Thus, estimates of task difficulty are in general agreement across all models.

4. DISCUSSION

In this paper we compare accounts of behavior provided by fitting DDM and LBA models to behavioral data from a 2AFC Rmax task. Adjustments of thresholds, equivalent to caution, are known to be central to DDM descriptions of Rmax behavior. For example, participants may adjust their thresholds to best suit each difficulty condition (Bogacz et al., 2006, 2010; Balci et al., 2011), or pick a single threshold that works well, albeit suboptimally, across multiple difficulty levels (Balci et al., 2011).
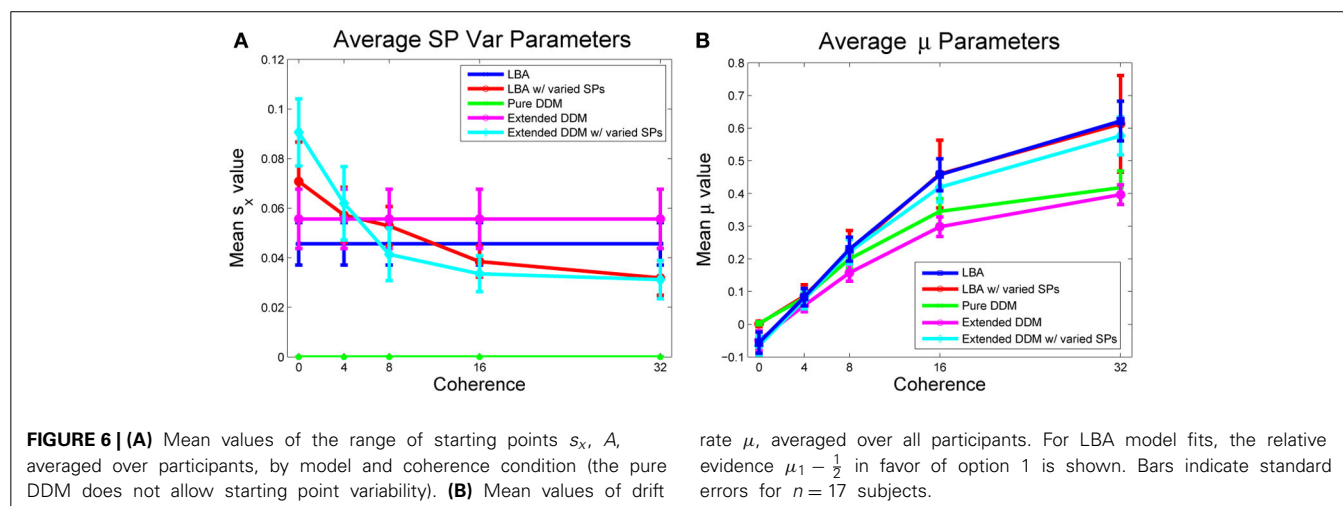
We first showed that, while the optimal performance curve (OPC) for the pure DDM is unique and parameter free, OPCs for the LBA are non-unique (**Figure 2**), like those for the extended DDM. Moreover, for a given parameter set, optimal behavior in

the LBA is at least partially determined by the range of starting points, A. If $A > 0$, fast responses at near signal detection speed are impossible because of thresholds $z \geq A$, and if $A \approx 0$, the quality of fits is limited. Lacking diffusive noise during trials, the LBA requires variable starting points as well as variable drift rates to produce a range of DTs and corresponding estimated RT. Allowing A to vary with task difficulty yields significantly better fits, and this parameter variability is consequently critical to the success of the LBA. We also proposed an adapted LBA, in which thresholds can lie below A, which we intend to analyze and fit to data in the future.

We then fitted five models to an Rmax data set: a pure DDM, an extended DDM, an extended DDM allowing starting point ranges that vary with task difficulty, a standard LBA, and an LBA that allows starting point ranges to vary with task difficulty. For consistency, we employed the LBA parameterization of Brown and Heathcote (2008) and Donkin et al. (2011) to parallel that of the extended DDM. We found that DDMs yielded somewhat better fits with a single starting point range. The AIC and BIC criteria indicated improved fits for the LBA with varying starting point ranges, although AIC$_c$ did not (**Table 1**).

In all three DDMs and the LBA with a common starting point range, participants modestly increased caution with coherence. In contrast, allowing starting point ranges to vary in the LBA predicted that participants *decreased* caution with coherence (compare **Figures 5A,B**). However, starting point ranges decreased with coherence in both models that allowed variability (**Figure 6A**). The LBA models and the extended DDMs require thresholds to lie at or above starting point ranges, but the additional source of within-trial randomness in the DDM may allow smaller starting point ranges, and hence yield better fits for difficult stimuli. Thresholds can be arbitrarily low for the pure DDM, since it has a single starting point. All models agree that drift rates increase with coherence (**Figure 6B**).

Critically, then, the Rmax task reveals that the DDMs and the LBA with varying starting point ranges provide fundamentally different accounts of behavior. In the DDMs, increased participant caution accounts for changes in behavior as stimulus coherence increases from 0% to 32%. According to the LBA model

FIGURE 6 | (A) Mean values of the range of starting points s_x, A, averaged over participants, by model and coherence condition (the pure DDM does not allow starting point variability). **(B)** Mean values of drift

rate μ, averaged over all participants. For LBA model fits, the relative evidence $\mu_1 - \frac{1}{2}$ in favor of option 1 is shown. Bars indicate standard errors for $n = 17$ subjects.

with variable starting point range, participants instead decrease caution and simultaneously reduce their range of starting points as coherence increases. Consequently, mean accumulation distances can still decrease with coherence in the LBA, in spite of smaller starting points. The corresponding RT and ER data therefore remain comparable between LBA and DDM fits. However, interpretations of the role of caution in these fits are inconsistent, adding to earlier findings that LBA parameters may not correlate straightforwardly with those of the DDM (van Ravenzwaaij and Oberauer, 2009; Heathcote and Hayes, 2012).

Our results raise the broad question of model design and selection, and suggest several directions for future work. While good overall, neither the LBA nor DDM accounts of behavior are perfect. The pure DDM is analytically tractable and predicts a unique, parameter-free OPC against which Rmax task performances can be assessed (Bogacz et al., 2006), but it can fail to fit RT distributions, especially when correct and error trials are separated (Ratcliff and Smith, 2004). Additional freedom in extended DDMs with variable drift rates and starting points across trials allows good fits, but defies analytical description and produces multiparameter families of OPCs. The LBA, incorporating trial-to-trial variability but omitting diffusive noise within trials, is almost as simple and tractable as the pure DDM, but it yields families of OPCs in which the allowed range of starting points plays a central and apparently complex role. In contrast, shifting the unique starting point of the pure DDM makes clear predictions regarding biased stimuli or incentivized rewards (Feng et al., 2009; Simen et al., 2009; Rorie et al., 2010; Gao et al., 2011).

Future work might re-adjust our interpretation of LBA parameters. For example, one might assume that the starting point range is controlled in tandem with thresholds. Indeed, a recent study suggests that this range narrows with practice (Heathcote and Hayes, 2012). The adapted LBA introduced in section 2.2.4 may provide accounts of average behavior that are more consistent with those provided by DDM fits, and, as noted there, it can produce bimodal RT distributions such as those sometimes observed for low stimulus discriminability, e.g., Simen et al. (2009) and Balci et al. (2011). Additional numerical and theoretical analyses are also of interest. For example, following Teodorescu and Usher (2013) one could compare models such as the Leaky Competing Accumulator as well as optimal Bayesian accounts of behavior with LBAs and DDMs.

ACKNOWLEDGMENTS

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: http://www.frontiersin.org/journal/10.3389/fnins.2014.00148/abstract

REFERENCES

Akaike, H. (1974). A new look at the statistical model identification. *IEEE Trans. Autom. Control* 19, 716–723. doi: 10.1109/TAC.1974.1100705

Akaike, H. (1980). Likelihood and the Bayes procedure. *Trabajos de Estadística y de Investigación Operativa* 31, 143–166. doi: 10.1007/BF02888350

Akaike, H. (1981). Likelihood of a model and information criteria. *J. Econometrics* 16, 3–14. doi: 10.1016/0304-4076(81)90071-3

Balci, F., Simen, P., Niyogi, R., Saxe, A., Hughes, J., Holmes, P., et al. (2011). Acquisition of decision making criteria: reward rate ultimately beats accuracy. *Atten. Percept. Psychophys.* 73, 640–657. doi: 10.3758/s13414-010-0049-7

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., and Cohen, J. (2006). The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol. Rev.* 113, 700–765. doi: 10.1037/0033-295X.113.4.700

Bogacz, R., Hu, P., Holmes, P., and Cohen, J. (2010). Do humans produce the speed-accuracy trade-off that maximizes reward rate? *Q. J. Exp. Psychol.* 63, 863–891. doi: 10.1080/17470210903091643

Brainard, D. (1997). The psychophysics toolbox. *Spat. Vis.* 10, 433–436. doi: 10.1163/156856897X00357

Brown, S., and Heathcote, A. (2008). The simplest complete model of choice response time: linear ballistic accumulation. *Cogn. Psychol.* 57, 153–178. doi: 10.1016/j.cogpsych.2007.12.002

Chernoff, H., and Lehmann, E. (1954). The use of maximum likelihood estimates in chi-square tests for goodness of fit. *Annal. Math. Stat.* 25, 579–586. doi: 10.1214/aoms/1177728726

Donkin, C., Averell, L., Brown, S., and Heathcote, A. (2009). Getting more from accuracy and response time data: methods for fitting the linear ballistic accumulator. *Behav. Res. Methods* 41, 1095–1110. doi: 10.3758/BRM.41.4.1095

Donkin, C., Brown, S., Heathcote, A., and Wagenmakers, E. (2011). Diffusion versus linear ballistic accumulation: different models but the same conclusions about psychological processes? *Psychon. Bull. Rev.* 18, 1–9. doi: 10.3758/s13423-010-0022-4

Feng, S., Holmes, P., Rorie, A., and Newsome, W. (2009). Can monkeys choose optimally when faced with noisy stimuli and unequal rewards? *PLoS Comput. Biol.* 5:e1000284. doi: 10.1371/journal.pcbi.1000284

Gao, J., Tortell, R., and McClelland, J. (2011). Dynamic integration of reward and stimulus information in perceptual decision-making. *PLoS ONE* 6:e16749. doi: 10.1371/journal.pone.0016749

Gold, J., Law, C., Connolly, P., and Bennur, S. (2008). The relative influences of priors and sensory evidence on an oculomotor decision variable during perceptual learning. *J. Neurophysiol.* 100, 2653–2668. doi: 10.1152/jn.90629.2008

Gold, J., and Shadlen, M. (2000). Representation of a perceptual decision in developing oculomotor commands. *Nature* 404, 390–394. doi: 10.1038/35006062

Gold, J., and Shadlen, M. (2001). Neural computations that underlie decisions about sensory stimuli. *Trends Cogn. Sci.* 5, 10–16. doi: 10.1016/S1364-6613(00)01567-9

Heathcote, A., Brown, S., and Cousineau, D. (2004). QMPE: estimating lognormal, Wald, and Weibull RT distributions with a parameter-dependent lower bound. *Behav. Res. Methods* 36, 277–290. doi: 10.3758/BF03195574

Heathcote, A., and Hayes, B. (2012). Diffusion versus linear ballistic accumulation: different models for response time with different conclusions about psychological mechanisms? *Can. J. Exp. Psychol.* 66, 125–136. doi: 10.1037/a0028189

Ratcliff, R., and Rouder, J. (1998). Modeling response times for two-choice decisions. *Psychol. Sci.* 9, 347–356. doi: 10.1111/1467-9280.00067

Ratcliff, R., and Smith, P. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychol. Rev.* 111:333. doi: 10.1037/0033-295X.111.2.333

Ratcliff, R., Thapar, A., Gomez, P., and McKoon, G. (2004). A diffusion model analysis of the effects of aging in the lexical-decision task. *Psychol. Aging* 19, 278–289. doi: 10.1037/0882-7974.19.2.278

Rorie, A., Gao, J., McClelland, J., and Newsome, W. (2010). Integration of sensory and reward information during perceptual decision-making in lateral intraparietal cortex (LIP) of the macaque monkey. *PLoS ONE* 5:e9308. doi: 10.1371/journal.pone.0009308

Simen, P., Contreras, D., Buck, C., Hu, P., Holmes, P., and Cohen, J. (2009). Reward rate optimization in two-alternative decision making: Empirical tests of theoretical predictions. *J. Exp. Psychol. Hum. Percept. Perform.* 35, 1865–1897. doi: 10.1037/a0030776

Teodorescu, A., and Usher, M. (2013). Disentangling decision models: from independence to competition. *Psychol. Rev.* 120, 1–38. doi: 10.1037/a0030776

Tuerlinckx, F. (2004). The efficient computation of the cumulative distribution and probability density function in the diffusion model. *Behav. Res. Methods* 36, 702–716. doi: 10.3758/BF03206552

Usher, M., and McClelland, J. (2001). The time course of perceptual choice: the leaky, competing accumulator model. *Psychol. Rev.* 108, 550–592. doi: 10.1037/0033-295X.108.3.550

van Ravenzwaaij, D., and Oberauer, K. (2009). How to use the diffusion model: parameter recovery of three methods: EZ, fast-dm, and DMAT. *J. Math. Psychol.* 53, 463–473. doi: 10.1016/j.jmp.2009.09.004

Vandekerckhove, J., and Tuerlinckx, F. (2007). Fitting the Ratcliff diffusion model to experimental data. *Psychon. Bull. Rev.* 14, 1011–1026. doi: 10.3758/BF03193087

Vandekerckhove, J., and Tuerlinckx, F. (2008). Diffusion model analysis with MATLAB: a DMAT primer. *Behav. Res. Methods* 40, 61–72. doi: 10.3758/BRM.40.1.61

Zacksenhouse, M., Bogacz, R., and Holmes, P. (2010). Robust versus optimal strategies for two-alternative forced choice tasks. *J. Math. Psychol.* 54, 230–246. doi: 10.1016/j.jmp.2009.12.004

frontiers in NEUROSCIENCE

When natural selection should optimize speed-accuracy trade-offs

Angelo Pirrone[1,2], Tom Stafford[1] and James A. R. Marshall[2,3]*

[1] Department of Psychology, University of Sheffield, Sheffield, UK
[2] Kroto Research Institute, University of Sheffield, Sheffield, UK
[3] Department of Computer Science, University of Sheffield, Sheffield, UK
*Correspondence: james.marshall@sheffield.ac.uk

1. INTRODUCTION

In psychology and neuroscience, and in other disciplines studying decision-making mechanisms, it is often assumed that optimal decision-making means statistical optimality. This is attractive because statistically optimal decision procedures are known, can be simply implemented in biologically-plausible models, and because such models have been shown to give good fits to behavioural as well as neural data. Here we question when statistical optimality is the kind of optimality we should expect natural selection to aim towards, by considering what kinds of loss function should be optimised under different behavioural scenarios. In laboratory settings subjects are often rewarded only on making a correct choice, so optimisation of a zero-one loss function is appropriate, and this is achieved by implementing a statistically-optimal decision procedure that gives the best compromise between speed and accuracy of decision-making. Many naturalistic decisions may also be described by such a loss function; however others, such as selecting food items of potentially different value, appear to be different since the animal is rewarded by the value of the item it chooses regardless of whether it was the best available. We argue that most naturalistic decisions are value-based. Mechanisms that optimise speed-accuracy trade-offs need to be parameterised, using information about the decision problem, in order to deal with value-based decision-making. Mechanisms for value-sensitive decision-making have been described,

however, which adaptively change between decision-making strategies without the need for continual re-parameterisation.

2. SPEED-ACCURACY TRADE-OFFS

It is usually assumed that decision-makers have to decide to be either fast or accurate. When speed is important mistakes are more frequent, while when accuracy is needed decisions are slower. This obvious problem is defined as the speed-accuracy trade-off and is a distinctive feature of many types of decision making (Wickelgren, 1977).

The speed-accuracy trade-off can be explained within the theoretical framework of sequential sampling models of decision making that have been shown to fit behavioral and neural data from human and animal choice tasks (Ratcliff and Rouder, 2000; Ratcliff et al., 2003, 2004; Ratcliff and Smith, 2004; Busemeyer et al., 2013). In particular, the Drift Diffusion Model (DDM; Ratcliff, 1978) describes choice between two alternatives (see Smith and Ratcliff, 2004; Bogacz et al., 2006; Basten et al., 2010) and recently has been shown also to be quantitatively accurate in describing trinary choices (Krajbich and Rangel, 2011) and value-based choices (Krajbich et al., 2010; Milosavljevic et al., 2010; Krajbich and Rangel, 2011; Krajbich et al., 2012), suggesting that the DDM can be thought of as a unifying computational framework for describing decision making (Basten et al., 2010). Moreover, Bogacz et al. (2006) have demonstrated that several connectionist decision-making models can approximate

the DDM under specific conditions. The DDM is a special case of the statistically-optimal Sequential Probability Ratio Test (SPRT; Wald, 1947; Wald and Wolfowitz, 1948). In the DDM noisy sensory evidence supporting the alternatives is integrated over time until the net evidence in favor of one alternative exceeds a certain positive or negative threshold value, precipitating a decision for the corresponding alternative. These thresholds can be varied to compromise optimally between the average speed and accuracy of decisions.

3. SPEED-VALUE TRADE-OFFS

In situations where decisions are rewarded according to whether they are correct or not, optimizing the speed-accuracy trade-off is sensible. When decisions are rewarded according to the value of the option chosen, however, a different criterion needs to be optimized. This can be illustrated with the simplest case of choosing between two equal value options; here there is no decision accuracy, since choosing either option is "correct." Similarly, there is no difference in average evidence for which of the two options is more valuable, meaning that the SPRT/DDM will only reach a decision by integrating sufficient noise to cross a decision threshold. Thus in this scenario there is no speed-accuracy trade-off to manage; the optimal decision is to choose anything as quickly as possible. The fundamental insight is that for certain decisions, speed-value trade-offs are more appropriate to optimize, rather than speed-accuracy trade-offs.

The SPRT/DDM can be optimized to take account of the value of the alternatives but, as we discuss here, doing so requires knowledge of the decision problem faced. The thresholds for an optimal decision depend on the goals of the decision maker and are task specific. By way of example, one route to accounting for the values associated with different decision outcomes is to minimize an extended version of the Bayes Risk (BR). BR is a linear combination of expected decision delay and expected terminal decision loss, first proposed by Wald and Wolfowitz (1948), and assumes that decision makers seek to minimize a cost function that is the weighted sum of decision times (DTs) and error rate (ERs). This was subsequently extended by Edwards to also account for non-zero rewards for incorrect decisions (Edwards, 1965; Bogacz et al., 2006). Formally Edwards' extension of BR, which implements Wald and Wolfowitz's version as a special case, can be defined as

$$BR_E = c_1 DT + c_2 \begin{pmatrix} ER \\ 1 - ER \end{pmatrix} \quad (1)$$

where c_1 is the cost of observing the stimulus per unit time, while c_2 is a row-vector specifying the payoffs from incorrect and correct choices (Bogacz et al., 2006). If $c_2 = (k\ 0)$, where $k > 0$ is a constant, then Wald and Wolfowitz's original BR is recovered. Several studies demonstrate that, under specific circumstances, subjects choose decision thresholds close to those that minimize BR_E (Busemeyer and Rapoport, 1988; Mozer et al., 2002). Bayes risk is not the only criterion proposed to date that decision-makers might optimize. Bogacz et al. survey alternatives, such as reward-rate, however, these alternatives are all calculated based on decision-accuracy, which requires explicit parameterizations based on the values of correct and incorrect choices (Bogacz et al., 2006). We therefore concentrate our analysis on Bayes risk. Bayes risk can be used to optimize value-sensitive decision-making; for example in a decision between two equal alternatives, each having value v if chosen, we would set the vector $c_2 = (v\ v)$ (e.g., dashed green line in **Figure 1**), thus simplifying Equation (1) above to

$$BR_E = c_1 DT + v. \quad (2)$$
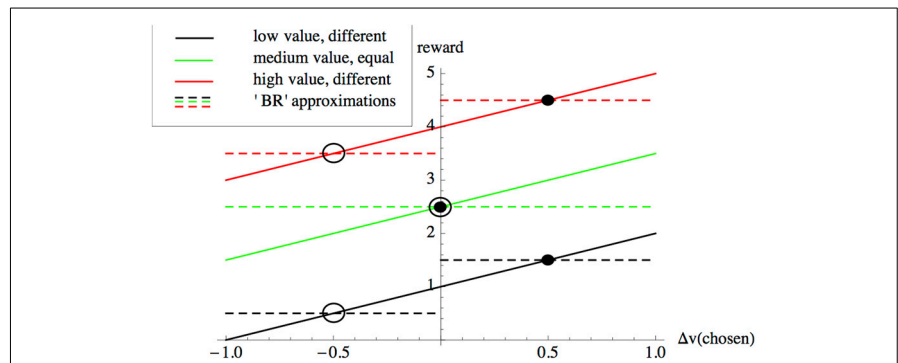
FIGURE 1 | The accuracy-based component of Bayes Risk (BR_E as defined by Equation 1) can be used to approximate a value-based reward scheme. In value-based decisions individuals are rewarded according to the value $|v| + \Delta v$ of the option they choose (solid lines), where $|v|$ is the average value of the alternatives under consideration, and Δv is the deviation from this average of the value of the option chosen by the subject. With knowledge of the values of the alternatives, BR_E can be used to optimize value sensitive decision-making as described in the main text; for example the dashed lines show payoffs used in BR_E for: options having values of 0.5 and 1.5 units (black), options having equal values of 2.5 and 2.5 units (green) and options having values of 3.5 and 4.5 units (red). Intersections between payoffs selected for BR_E (dashed lines) with value-based reward (solid lines of matching colors) correspond to choice scenarios between different-valued options for which BR_E implements reward-by-value of the selected option; these intersections represent choice scenarios involving "poor" (hollow circles) and "good" (filled circles) options having particular values. However, the cost parameters for BR_E need to be recalculated according to the values of the options under consideration; for example, although the difference in the values of the alternatives does not change from the low-value (black) to the high-value (red) scenarios, since their absolute values change the BR_E payoffs need to be recalculated in each case. As described in the text, value-sensitive decision-mechanisms have been described that are able adaptively to deal with a variety of such decision scenarios, without re-parameterizations.

Equation (2) shows us that, intuitively, an optimal decision-maker in our equal-alternatives scenario should minimize decision-time DT, since doing so incurs no penalty as the error rate ER no longer features. However, using Bayes risk in this way requires the values of the alternatives to be known on a case by case basis, as shown in **Figure 1**. Subjects might learn the values of incorrect and correct choices over time, for example when trials are blocked in psychophysical experiments (see Bogacz et al., 2006). However, in the following we argue that in most naturalistic decision scenarios decision-makers will not have this opportunity, and will therefore use other mechanisms that directly optimize speed-value trade-offs, rather than optimizing decisions indirectly via optimization of the speed-accuracy trade-off with an appropriate payoff vector c_2.

4. NATURALISTIC DECISIONS ARE USUALLY VALUE-BASED

We argue that most naturalistic decisions faced by animals, including humans, are value-based, in that the animal is rewarded

according to the value of the option it chooses. Such a view on decision-making is not new to behavioral ecologists, where a long tradition exists of studying behaviors such as mate choice and foraging (Davies et al., 2012) or nest-site selection (Stroeymeyt et al., 2014). Recently many studies have focused on how value and reward are represented and integrated during the decision process (Platt and Glimcher, 1999; Sugrue et al., 2004; Padoa-Schioppa and Assad, 2006; Rangel et al., 2008; Kable and Glimcher, 2009; Krajbich et al., 2010; Philiastides et al., 2010; Hare et al., 2011; Krajbich and Rangel, 2011; Louie and Glimcher, 2012; Tsetsos et al., 2012; Cassey et al., 2013; Towal et al., 2013); however, in psychology and neuroscience, experiments are usually designed such that there is always a correct choice, and only correct choices are rewarded (see Gold and Shadlen, 2003; Bogacz et al., 2006). While studying behavior in psychophysical tasks is beneficial in that it gives a well-controlled decision environment, our point is that only rewarding subjects when they make correct choices may not correspond to the

kind of decisions animals, and their neural circuitry, have typically evolved to deal with. Even in the value-based decision experiments cited above, which are analyzed using the DDM, it is typical to only present subjects with a choice between options known to have *different* values. Moreover, even though some studies have looked at how reward information is integrated (Rorie et al., 2010; Gao et al., 2011), much of this work has not yet focused on the tradeoff between value and speed. While usually in the decision-making literature the optimal behavior is to optimize speed-accuracy trade-offs, and subjects can apparently do this (Busemeyer and Rapoport, 1988; Bogacz et al., 2006), we argue that these scenarios are not representative of many naturalistic settings, and that there is great value in considering how subjects make value-sensitive decisions and how these should be optimized. In the following section we discuss theory that may be useful for this.

At least one important class of naturalistic decisions does require optimization of speed-accuracy trade-offs; these are life-or-death decisions. If we analyze for example the case of an animal attempting to forage while avoiding predators (Trimmer et al., 2008), a slow-but-accurate decision would mean being killed by the predator, a maximal loss. On the other hand if the decision is fast-but-inaccurate the animal would escape even when the stimulus is not a predator, and this would mean losing food. The best strategy for the animal is thus that which optimizes the speed-accuracy trade-off, taking into account the payoffs arising from the different decision outcomes; hence Trimmer et al.'s hypothetical animal is modeled with a single-threshold DDM, with evidence sufficient to cross that single decision threshold leading to the animal taking anti-predator action such as running away.

5. MECHANISMS FOR VALUE-SENSITIVE DECISION-MAKING

Recent modeling work inspired by studying another value-sensitive decision-making system, collective nest-site selection by honeybees (Seeley et al., 2012), has described a very simple mechanism able to adaptively account for the value of different decision outcomes, with minimal parameter tuning (Pais et al.,

2013). This simple model implements a variety of sophisticated decision-making strategies; for example, when equal but low-value alternatives are presented, a decision deadlock is maintained that can be broken should a third, higher-value alternative, be made available. However, if equal-but-high-value alternatives are presented, or sufficient time passes, deadlock is spontaneously and randomly broken (Pais et al., 2013). This is particularly interesting, since the classic DDM is insensitive to the absolute value of the alternatives under consideration, and only integrates the difference in their values. When differences between alternative values are sufficient, the value-sensitive mechanism of Pais et al. becomes closer to a classic DDM, allowing speed-accuracy trade-offs to be managed, although not optimized, through modification of decision thresholds. All of the different behavioral regimes of the model arise without direct parameterizations regarding alternatives' values, simply through the dependence of the model's dynamics on the mean values of inputs to its integrator populations; this allows the model to adaptively respond to different decision scenarios on a trial-by-trial basis, which cannot be achieved in pure DDM models without the decision-maker having access to explicit information on the decision-task at hand. Modifications to DDM-type models have been proposed to deal with trial-by-trial variability such as online estimation of task parameters (Deneve, 2012) or the use of time-dependent change in parameters such as decision-thresholds, urgency signals or asymmetry of inhibition (Ditterich, 2006; Hanks et al., 2011; Drugowitsch et al., 2012; Thura et al., 2012); fundamentally, however, these modifications are still interpreted under the assumption that decision speed vs accuracy is the trade-off to be maximized, unlike the model of Pais et al. (2013) in which the dynamics are naturally interpreted in terms of value vs time trade-offs. Pais et al.'s mechanism also exhibits other characteristics of natural value-discrimination systems, such as Weber's law of just-noticeable difference; interestingly Weber's law arises from the deterministic dynamics of the mechanism rather than from noise processes (Pais et al., 2013) (cf. Deco and Rolls, 2006; Deco et al., 2007). Finally,

it is important to note that the DDM cannot account for the non-linearity that characterizes many decision making dynamics (e.g., food recruitment by social insects; Nicolis and Deneubourg, 1999) while the model of Pais et al. (2013) is non-linear.

6. CONCLUSION

The study of speed-accuracy trade-offs has been tremendously fruitful for psychology, neuroscience and animal behavior, and will doubtless prove fruitful for many years to come. Yet as we have argued here most naturalistic decisions, which animals' brains should have evolved to optimize, are value-based rather than accuracy-based. This leads us to argue that the drift-diffusion model, which optimizes speed-accuracy trade-offs, is not an ideal computational framework to describe value-based decision-making; although it has had some success in describing particular experiments on value-based decision-making, discussed in the section "Speed-Accuracy Trade-Offs," as we have shown here the DDM requires special case-by-case parameterizations to implement true value-based decision-making. We suggest that this limits the generality of the DDM as a unifying framework for all ecologically-relevant decision-making problems. However, recent theory has presented mechanisms that can manage value-sensitive decision problems without the additional informational requirements of the DDM. At the same time, experimental and theoretical psychologists and neuroscientists have started to tackle problems of value-based decision-making. We have presented our arguments for value in terms of animal decision-making, but unicellular organisms and individual cells also make decisions (e.g., Perkins and Swain, 2009; Latty and Beekman, 2011), and value is likely to be similarly important for these. We believe that the evolutionary perspective we have presented here should motivate further research into value-sensitivity and decision-making.

AUTHOR CONTRIBUTIONS

James A. R. Marshall conceived of the paper; James A. R. Marshall, Angelo Pirrone, and Tom Stafford discussed the material; James A. R. Marshall developed

REFERENCES

Basten, U., Biele, G., Heekeren, H. R., and Fiebach, C. J. (2010). How the brain integrates costs and benefits during decision making. *Proc. Natl. Acad. Sci. U.S.A.* 107, 21767–21772. doi: 10.1073/pnas.0908104107

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., and Cohen, J. D. (2006). The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol. Rev.* 113:700. doi: 10.1037/0033-295X.113.4.700

Busemeyer, J., Diederich, A., Heathcote, A., Rangel, A., Rieskamp, J., and Usher, M. (2013). "Sequential sampling models representing a unifying framework of human decision making," in *Symopsium, COGSCI Annual Meeting* (Berlin).

Busemeyer, J. R., and Rapoport, A. (1988). Psychological models of deferred decision making. *J. Math. Psychol.* 32, 91–134. doi: 10.1016/0022-2496(88)90042-9

Cassey, T. C., Evens, D. R., Bogacz, R., Marshall, J. A. R., and Ludwig, C. J. H. (2013). Adaptive sampling of information in perceptual decision-making. *PLoS ONE* 8:e78993. doi: 10.1371/journal.pone.0078993

Davies, N. B., Krebs, J. R., and West, S. A. (2012). *An Introduction to Behavioural Ecology, 4th Edn.* Hoboken, NJ: John Wiley & Sons.

Deco, G., and Rolls, E. T. (2006). Decision-making and weber's law: a neurophysiological model. *Eur. J. Neurosci.* 24, 901–916. doi: 10.1111/j.1460-9568.2006.04940.x

Deco, G., Scarano, L., and Soto-Faraco, S. (2007). Weber's law in decision making: integrating behavioral data in humans with a neurophysiological model. *J. Neurosci.* 27, 11192–11200. doi: 10.1523/JNEUROSCI.1072-07.2007

Deneve, S. (2012). Making decisions with unknown sensory reliability. *Front. Neurosci.* 6:75. doi: 10.3389/fnins.2012.00075

Ditterich, J. (2006). Stochastic models of decisions about motion direction: behavior and physiology. *Neural Netw.* 19, 981–1012. doi: 10.1016/j.neunet.2006.05.042

Drugowitsch, J., Moreno-Bote, R., Churchland, A. K., Shadlen, M. N., and Pouget, A. (2012). The cost of accumulating evidence in perceptual decision making. *J. Neurosci.* 32, 3612–3628. doi: 10.1523/JNEUROSCI.4010-11.2012

Edwards, W. (1965). Optimal strategies for seeking information: models for statistics, choice reaction times, and human information processing. *J. Math. Psychol.* 2, 312–329. doi: 10.1016/0022-2496(65)90007-6

Gao, J., Tortell, R., and McClelland, J. L. (2011). Dynamic integration of reward and stimulus information in perceptual decision-making. *PLoS ONE* 6:e16749. doi: 10.1371/journal.pone.0016749

Gold, J. I., and Shadlen, M. N. (2003). The influence of behavioral context on the representation of a perceptual decision in developing oculomotor commands. *J. Neurosci.* 23, 632–651. doi: 10.1146/annurev.neuro.29.051605.113038

Hanks, T. D., Mazurek, M. E., Kiani, R., Hopp, E., and Shadlen, M. N. (2011). Elapsed decision time affects the weighting of prior probability in a perceptual decision task. *J. Neurosci.* 31, 6339–6352. doi: 10.1523/JNEUROSCI.5613-10.2011

Hare, T. A., Schultz, W., Camerer, C. F., O'Doherty, J. P., and Rangel, A. (2011). Transformation of stimulus value signals into motor commands during simple choice. *Proc. Natl. Acad. Sci. U.S.A.* 108, 18120–18125. doi: 10.1073/pnas.1109322108

Kable, J. W., and Glimcher, P. W. (2009). The neurobiology of decision: consensus and controversy. *Neuron* 63, 733–745. doi: 10.1016/j.neuron.2009.09.003

Krajbich, I., Armel, C., and Rangel, A. (2010). Visual fixations and the computation and comparison of value in simple choice. *Nat. Neurosci.* 13, 1292–1298. doi: 10.1038/nn.2635

Krajbich, I., Lu, D., Camerer, C., and Rangel, A. (2012). The attentional drift-diffusion model extends to simple purchasing decisions. *Front. Psychol.* 3:193. doi: 10.3389/fpsyg.2012.00193

Krajbich, I., and Rangel, A. (2011). Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. *Proc. Natl. Acad. Sci. U.S.A.* 108, 13852–13857. doi: 10.1073/pnas.1101328108

Latty, T., and Beekman, M. (2011). Speed–accuracy trade-offs during foraging decisions in the acellular slime mould physarum polycephalum. *Proc. R. Soc. B Biol. Sci.* 278, 539–545. doi: 10.1098/rspb.2010.1624

Louie, K., and Glimcher, P. W. (2012). Efficient coding and the neural representation of value. *Ann. N.Y. Acad. Sci.* 1251, 13–32. doi: 10.1111/j.1749-6632.2012.06496.x

Milosavljevic, M., Malmaud, J., Huth, A., Koch, C., and Rangel, A. (2010). The drift diffusion model can account for value-based choice response times under high and low time pressure. *Judge. Decis. Making* 5, 437–449. doi: 10.2139/ssrn.1901533

Mozer, M. C., Colagrosso, M. D., and Huber, D. E. (2002). "A rational analysis of cognitive control in a speeded discrimination task," in *Advances in Neural Information Processing Systems*, Vol. 11, eds T. Dietterich, S. Becker, and Z. Ghahramani (Cambridge, MA: MIT Press), 51–57.

Nicolis, S. C., and Deneubourg, J.-L. (1999). Emerging patterns and food recruitment in ants: an analytical study. *J. Theor. Biol.* 198, 575–592. doi: 10.1006/jtbi.1999.0934

Padoa-Schioppa, C., and Assad, J. A. (2006). Neurons in the orbitofrontal cortex encode economic value. *Nature* 441, 223–226. doi: 10.1038/nature04676

Pais, D., Hogan, P. M., Schlegel, T., Franks, N. R., Leonard, N. E., and Marshall, J. A. R. (2013). A mechanism for value-sensitive decision-making. *PLoS ONE* 8:e73216. doi: 10.1371/journal.pone.0073216

Perkins, T. J., and Swain, P. S. (2009). Strategies for cellular decision-making. *Mol. Syst. Biol.* 5, 1–15. doi: 10.1038/msb2009.83

Philiastides, M. G., Biele, G., and Heekeren, H. R. (2010). A mechanistic account of value computation in the human brain. *Proc. Natl. Acad. Sci. U.S.A.* 107, 9430–9435. doi: 10.1073/pnas.1001732107

Platt, M. L., and Glimcher, P. W. (1999). Neural correlates of decision variables in parietal cortex. *Nature* 400, 233–238. doi: 10.1038/22268

Rangel, A., Camerer, C., and Montague, P. R. (2008). A framework for studying the neurobiology of value-based decision making. *Nat. Rev. Neurosci.* 9, 545–556. doi: 10.1038/nrn2357

Ratcliff, R. (1978). A theory of memory retrieval. *Psychol. Rev.* 85:59. doi: 10.1037/0033-295X.85.2.59

Ratcliff, R., Gomez, P., and McKoon, G. (2004). A diffusion model account of the lexical decision task. *Psychol. Rev.* 111:159. doi: 10.1037/0033-295X.111.1.159

Ratcliff, R., and Rouder, J. N. (2000). A diffusion model account of masking in two-choice letter identification. *J. Exp. Psychol.* 26:127. doi: 10.1037/0096-1523.26.1.127

Ratcliff, R., and Smith, P. L. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychol. Rev.* 111:333. doi: 10.1037/0033-295X.111.2.333

Ratcliff, R., Thapar, A., and Mckoon, G. (2003). A diffusion model analysis of the effects of aging on brightness discrimination. *Percept. Psychophys.* 65, 523–535. doi: 10.3758/BF03194580

Rorie, A. E., Gao, J., McClelland, J. L., and Newsome, W. T. (2010). Integration of sensory and reward information during perceptual decision-making in lateral intraparietal cortex (lip) of the macaque monkey. *PLoS ONE* 5:e9308. doi: 10.1371/journal.pone.0009308

Seeley, T. D., Visscher, P. K., Schlegel, T., Hogan, P. M., Franks, N. R., and Marshall, J. A. R. (2012). Stop signals provide cross inhibition in collective decision-making by honeybee swarms. *Science* 335, 108–111. doi: 10.1126/science.1210361

Smith, P. L., and Ratcliff, R. (2004). Psychology and neurobiology of simple decisions. *Trends Neurosci.* 27, 161–168. doi: 10.1016/j.tins.2004.01.006

Stroeymeyt, N., Jordan, C., Mayer, G., Hovsepian, S., Giurfa, M., and Franks, N. (2014). Seasonality in communication and collective decision-making in ants. *Proc. R. Soc. B. Biol. Sci.* 281:20133108. doi: 10.1098/rspb.2013.3108

Sugrue, L. P., Corrado, G. S., and Newsome, W. T. (2004). Matching behavior and the representation of value in the parietal cortex. *Science* 304, 1782–1787. doi: 10.1126/science.1094765

Thura, D., Beauregard-Racine, J., Fradet, C.-W., and Cisek, P. (2012). Decision making by urgency gating: theory and experimental support. *J. Neurophysiol.* 108, 2912–2930. doi: 10.1152/jn.01071.2011

Towal, R. B., Mormann, M., and Koch, C. (2013). Simultaneous modeling of visual saliency and

value computation improves predictions of economic choice. *Proc. Natl. Acad. Sci. U.S.A.* 110, E3858–E3867. doi: 10.1073/pnas.1304429110

Trimmer, P. C., Houston, A. I., Marshall, J. A. R., Bogacz, R., Paul, E. S., Mendl, M. T., et al. (2008). Mammalian choices: combining fast-but-inaccurate and slow-but-accurate decision-making systems. *Proc. R. Soc. B Biol. Sci.* 275, 2353–2361. doi: 10.1098/rspb.2008.0417

Tsetsos, K., Chater, N., and Usher, M. (2012). Salience driven value integration explains decision biases and preference reversal. *Proc. Natl. Acad. Sci. U.S.A.* 109, 9659–9664. doi: 10.1073/pnas.1119 569109

Wald, A. (1947). *Sequential Analysis*. New York, NY: Wiley and Sons.

Wald, A., and Wolfowitz, J. (1948). Optimum character of the sequential probability ratio test. *Ann. Math. Stat.* 19, 326–339. doi: 10.1214/aoms/1177 730197

Wickelgren, W. A. (1977). Speed-accuracy trade-off and information processing dynamics. *Acta Psychol.* 41, 67–85. doi: 10.1016/0001-6918(77) 90012-9

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Neural dynamics implement a flexible decision bound with a fixed firing rate for choice: a model-based hypothesis

Dominic Standage[1]*, Da-Hui Wang[2] and Gunnar Blohm[1]

[1] Department of Biomedical and Molecular Sciences, Queen's University, Kingston, ON, Canada
[2] Department of Systems Science/National Key Laboratory of Cognitive Neuroscience and Learning, Beijing Normal University, Beijing, China

Decisions are faster and less accurate when conditions favor speed, and are slower and more accurate when they favor accuracy. This speed-accuracy trade-off (SAT) can be explained by the principles of bounded integration, where noisy evidence is integrated until it reaches a bound. Higher bounds reduce the impact of noise by increasing integration times, supporting higher accuracy (*vice versa* for speed). These computations are hypothesized to be implemented by feedback inhibition between neural populations selective for the decision alternatives, each of which corresponds to an attractor in the space of network states. Since decision-correlated neural activity typically reaches a fixed rate at the time of commitment to a choice, it has been hypothesized that the neural implementation of the bound is fixed, and that the SAT is supported by a common input to the populations integrating evidence. According to this hypothesis, a stronger common input reduces the difference between a baseline firing rate and a threshold rate for enacting a choice. In simulations of a two-choice decision task, we use a reduced version of a biophysically-based network model (Wong and Wang, 2006) to show that a common input can control the SAT, but that changes to the threshold-baseline difference are epiphenomenal. Rather, the SAT is controlled by changes to network dynamics. A stronger common input decreases the model's effective time constant of integration and changes the shape of the attractor landscape, so the initial state is in a more error-prone position. Thus, a stronger common input reduces decision time and lowers accuracy. The change in dynamics also renders firing rates higher under speed conditions at the time that an ideal observer can make a decision from network activity. The difference between this rate and the baseline rate is actually greater under speed conditions than accuracy conditions, suggesting that the bound is not implemented by firing rates *per se*.

Keywords: speed-accuracy trade-off, neural dynamics, bounded integration, decision threshold, threshold-baseline difference

1. INTRODUCTION

In decision making experiments, subjects make faster, less accurate decisions when conditions favor speed, and make slower, more accurate decisions when conditions favor accuracy (e.g., Bogacz et al., 2010a; Heitz and Schall, 2012). These data describe the speed-accuracy trade-off (SAT) and can be explained by the principles of bounded integration. According to these principles, noisy evidence for the alternatives of a decision is integrated until the running total for one of the alternatives reaches a criterion level. The running total is referred to as a decision variable and the criterion is referred to as the bound. A higher bound allows evidence to be integrated for longer, increasing the percentage of correct decisions. A lower bound has the opposite effect. These abstract models have been invaluable in characterizing the computations underlying decisions and the SAT (see Smith and Ratcliff, 2004; Ratcliff and McKoon, 2008; Bogacz et al., 2010b).

The computations characterized by bounded integration models are hypothesized to be implemented by competitive interactions between neural populations selective for the alternatives of

a decision (Usher and McClelland, 2001; Wang, 2002; Machens et al., 2005; Bogacz et al., 2006; Wong and Wang, 2006; Standage et al., 2011; You and Wang, 2013). According to this widely held hypothesis, temporal integration and competitive interactions are supported by recurrent excitation and feedback inhibition respectively, where each population implements a decision variable and a choice is made when the aggregate firing rate of one of the populations reaches a threshold. This hypothesis is supported by electrophysiological recordings from several cortical areas in non-human primates performing decision tasks, where the spike rates of neurons responsive to the chosen alternative (target-in neurons) increase over several hundreds of milliseconds prior to the animal's choice, and the spike rates of neurons unresponsive to the chosen alternative (target-out neurons) are much lower (e.g., Roitman and Shadlen, 2002; Thomas and Pare, 2007; Bollimunta and Ditterich, 2011; Ding and Gold, 2012).

Under several task paradigms, target-in activity of putative integrator neurons has been shown to reach an approximately fixed rate at the time of commitment to a choice (the choice

threshold), regardless of the speed or accuracy of decisions (Hanes and Schall, 1996; Shadlen and Newsome, 2001; Roitman and Shadlen, 2002; Churchland et al., 2008; Purcell et al., 2010; Ding and Gold, 2012). These data have been interpreted as indicating that the neural implementation of the bound is fixed across conditions emphasizing speed over accuracy or *vice versa* (see Bogacz et al., 2010b). Under the assumption of linear integration, adjusting the starting point of a decision variable is equivalent to adjusting the bound, so it has been hypothesized that subjects trade speed and accuracy by adjusting the "baseline" rate of integrator populations, i.e., the activity on which a decision variable builds (see Bogacz et al., 2010b). According to this hypothesis, the SAT is controlled by a cognitive signal projecting uniformly to all integrator populations, where a stronger (weaker) signal favors speed (accuracy) by decreasing (increasing) the difference between the choice threshold and baseline activity (the threshold-baseline difference). We refer to this possibility as the threshold-baseline hypothesis (a.k.a. the changing-baseline hypothesis, Bogacz et al., 2010b). Several recent neuroimaging (Forstmann et al., 2008; Ivanoff et al., 2008; van Veen et al., 2008; Wenzlaff et al., 2011) and electrophysiological (Heitz and Schall, 2012; Hanks et al., 2014) studies have provided evidence for such a signal, reporting higher baseline (pre-stimulus) activity in decision-correlated cortical areas under speed conditions than accuracy and/or neutral conditions.

Here, we present an alternative hypothesis that does not assume linear integration. As above, we assume that a cognitive signal controls the SAT by projecting uniformly to integrator populations, but the underlying mechanism is grounded in the framework of attractor dynamics (e.g., Machens et al., 2005; Bogacz et al., 2006; Wong and Wang, 2006; Standage et al., 2011; You and Wang, 2013; see Wang, 2008, 2012 for review). According to this framework, integration times are determined by the non-linear dynamics of decision circuitry, where stronger and weaker dynamics furnish shorter and longer integration times respectively (Wong and Wang, 2006; Standage et al., 2011). The SAT can therefore be accomplished by any mechanism that modulates the strength of dynamics within and between neural populations selective for the decision alternatives (see Standage et al., 2014). Spatially non-selective excitation provides just such a mechanism (Salinas and Abbott, 1996), where a stronger (weaker) signal corresponds to speed (accuracy) conditions (Furman and Wang, 2008; Roxin and Ledberg, 2008). Of course, this input also entails higher (lower) baseline activity under speed (accuracy) conditions. In attractor network models, higher (lower) baseline activity will indeed decrease (increase) the threshold-baseline difference, but this decrease (increase) is epiphenomenal. The SAT is supported by the resulting changes to network dynamics.

Below, we use a neurally-derived model (Wong and Wang, 2006) to demonstrate that adjusting the strength of spatially non-selective excitation can control the SAT (Furman and Wang, 2008; Roxin and Ledberg, 2008). We demonstrate that this signal raises (lowers) the baseline activity of integrator populations, consistent with higher (lower) baseline activity under speed (accuracy, neutral) conditions in SAT experiments (Forstmann et al., 2008; Ivanoff et al., 2008; van Veen et al., 2008; Wenzlaff et al., 2011; Heitz and Schall, 2012; Hanks et al., 2014). We use a fixed

choice threshold in the model, so the spatially non-selective signal decreases (increases) the threshold-baseline difference under speed (accuracy) conditions, relative to a neutral condition. We demonstrate that the threshold-baseline difference cannot account for the SAT in the model, since raising (lowering) the threshold to compensate for the higher (lower) baseline activity under the speed (accuracy) condition does not "untrade" speed and accuracy, i.e., reinstating the threshold-baseline difference of the neutral condition does not recover the neutral behavior of the model. Using dynamic systems analysis, we show that a higher (lower) baseline decreases (increases) the effective time constant of integration of the network under speed (accuracy) conditions, accounting for the SAT in a manner consistent with a flexible bound, while also changing the shape of the decision space so as to further decrease (increase) accuracy. Finally, we show that decision-selective firing rates in the model are actually higher (lower) under speed (accuracy) conditions at the time at which an ideal observer can discriminate between the rates of the integrator populations; as is the difference between these rates and the baseline rate (the discrimination-baseline difference). Thus, the discrimination-baseline difference increases under speed conditions and decreases under accuracy conditions, opposite to the principles of the threshold-baseline hypothesis. Our analysis explains these observations.

Our simulations show that under the framework of attractor dynamics, there is no discrepancy between a flexible bound and a fixed choice threshold. The bound—or the difference between the bound and the starting point of a decision variable—is a computational device for controlling the duration of evidence accumulation in abstract models. It can be implemented by the effective time constant of integration of decision circuitry, with corresponding changes to the decision space. This space and its time evolution are emergent properties of network dynamics and are qualitatively different than the synaptic current required to elicit choice behavior.

2. A COMMON INPUT TO INTEGRATORS CONTROLS THE SAT IN AN ATTRACTOR MODEL, BUT NOT BY THE THRESHOLD-BASELINE DIFFERENCE

In their seminal study, Wong and Wang (2006) used analytic methods to reduce a biophysically-based cortical network model (Wang, 2002) to a 2-variable system, tractable for analysis (depicted in **Figure 1A**). They showed that each of the populations selective for the decision alternatives corresponds to a stable state in the space of possible states of network activity, i.e., each population corresponds to an attractor (**Figures 1B,C**). The attractors are separated by an unstable "saddle" steady state with two manifolds: a stable manifold that draws the network toward the saddle point, and an unstable manifold that repels it toward one of the stable attractors (**Figure 1C**). They further calculated the time constants of these two manifolds, showing that the dynamics in the vicinity of the saddle support integration times much longer than the time constants of decay of contributing biophysical processes, such as those of neurons and synapses.

We used Wong and Wang's (2006) model in simulations of a 2-choice random dot motion (RDM) task (Supplementary Material
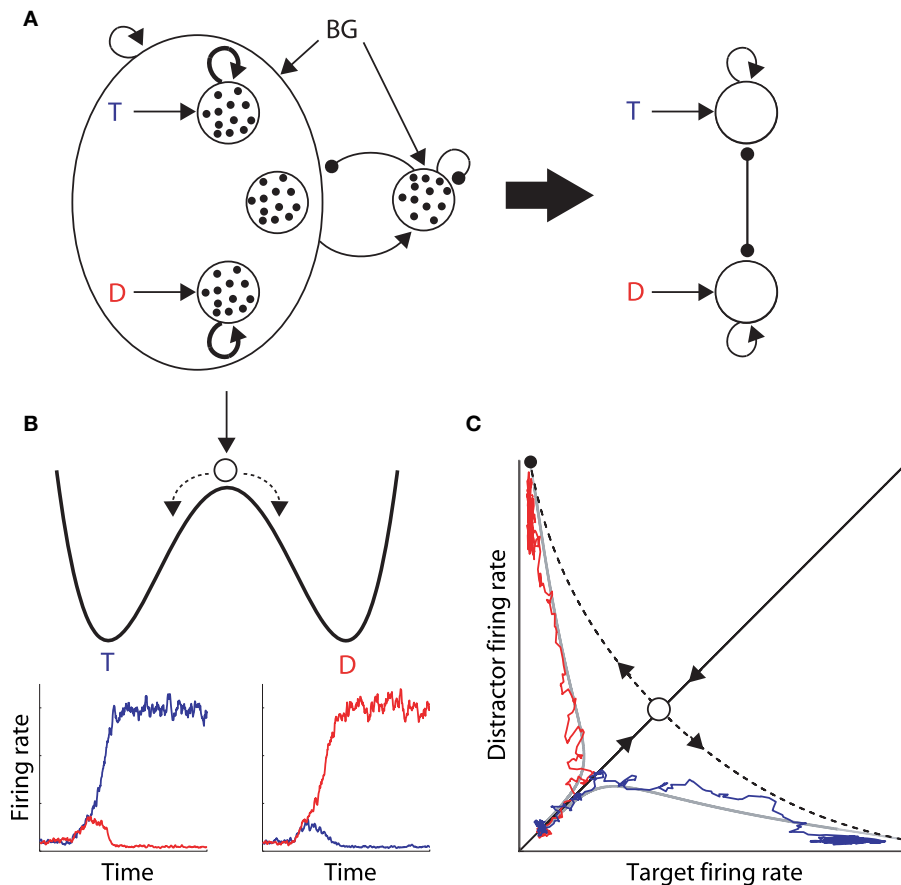
FIGURE 1 | (A) The reduced model by Wong and Wang (2006), approximating a biophysically-based cortical network model (left of the thick arrow) with a 2-variable system (right). The thick arrow depicts the derivation of the latter from the former. The large oval on the left depicts a network of cortical pyramidal neurons. Inside the oval, the three open circles depict the target and distractor populations with selective input T and D respectively, and a population unresponsive to the evidence for either alternative. Looping arcs depict recurrent synapses, which are stronger within each selective population (thicker arcs). All pyramidal neurons excite a common inhibitory pool, which uniformly inhibits all pyramidal neurons. Excitatory and inhibitory synapses are depicted by arrows and closed circles respectively, small black dots depict individual neurons, and BG refers to background input. **(B)** Cartoon depiction of an attractor "energy landscape" for 2-choice decisions, where the energy decreases over time. An unstable steady state (high energy) separates two stable attractors (low energy), corresponding to the target and

distractor stimuli. Conceptually, a ball placed between the two attractors will eventually role one way or the other, depicted by the dashed arrows. The ball enters an attractor basin sooner (later) under speed (accuracy) conditions because the dynamics evolve more quickly (slowly). Below the cartoon, the firing rates of target (blue) and distractor (red) neural populations are plotted over time during two decision trials, corresponding to the ball rolling into the target attractor basin (left) and the distractor attractor basin (right). **(C)** Decision space for two choices. Stable (solid) and unstable (dashed) manifolds of the saddle point (intersection of the manifolds, see text). The system moves toward this state along the stable manifold and is repelled along the unstable manifold. The firing rates of the winning populations in the two decision trials in **(B)** are plotted against each other, superimposed on the decision space, along with two noise-free trajectories (gray) with initial conditions inside each attractor basin. On each trial, the network state moves along the stable manifold before being repelled toward an attractor.

Section 1). We ran 1000 trials for each motion coherence $c \in \{0, 1, 2, 4, 8, 16, 32\}$%, where the motion stimulus was provided for 5s following a 2.5s pre-stimulus interval. We refer to the integrator population receiving the stronger (weaker) stimulus as the target (distractor) population. We modeled speed and accuracy conditions by increasing and decreasing a uniform input to the two populations respectively, relative to a neutral condition. To this end, we adjusted the mean background current I_0, capturing the total input current from upstream neurons other than those encoding motion stimuli. This current therefore subsumes the hypothesized cognitive signal controlling the SAT. Because

the model's parameter values and corresponding dynamics are rigorously described by Wong and Wang (2006), we used the same parameter values here (excepting I_0 and its corresponding standard deviation, see Supplementary Material Section 1).

Unsurprisingly, the spatially non-selective current I_0 produced higher and lower pre-stimulus (baseline) firing rates under speed and accuracy conditions respectively, compared to the neutral condition. Baseline rates can be seen to the left of the vertical line in **Figure 2A** for an example coherence value ($c = 4$%, see Figure caption). The resulting SAT can be seen in **Figures 2B,C**, where the psychometric curve is shifted to the right and left
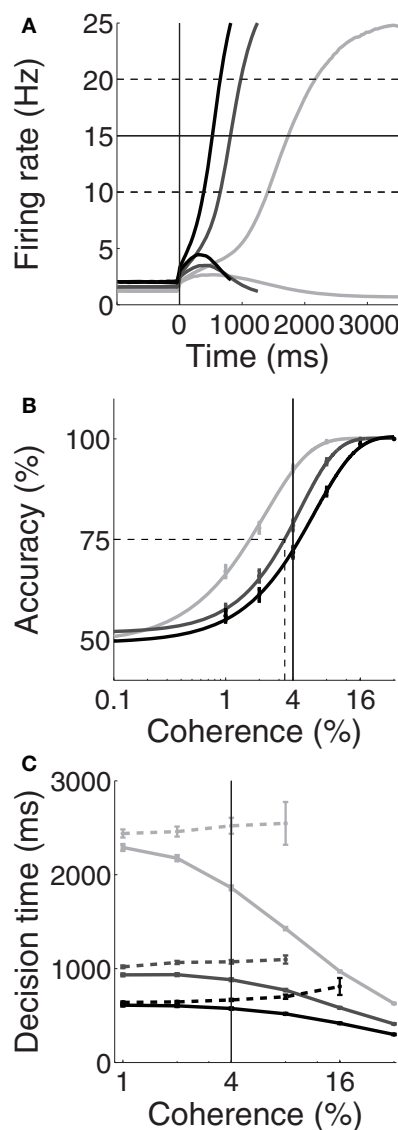
FIGURE 2 | Trading speed and accuracy as a function spatially non-selective input I_0. Simulated neural activity **(A)** and resulting psychometric **(B)** and chronometric **(C)** curves for neutral ($I_0 = 321$pA, medium gray), speed ($I_0 = 325$pA, black) and accuracy ($I_0 = 316$pA, light gray) conditions. **(A)** Trial-averaged firing rates for coherence $c = 4\%$. For each condition, the upper and lower curves show the mean rate over all correct trials for the target and distractor populations respectively. The vertical line at 0ms indicates the time of simulated motion onset. To the left of this line, pre-stimulus/baseline firing rates are higher (lower) under speed (accuracy) conditions compared to the neutral condition. Thus, the threshold-baseline difference is smaller (larger) under speed (accuracy) conditions. The solid horizontal line shows the "default" choice threshold $\theta = 15$Hz used by Wong and Wang (2006). The dashed horizontal lines depict other possible thresholds. **(B)** The percentage of correct trials as a function of coherence. The data are fitted with a Weibull function for each condition. Error bars show standard error. The solid vertical line indicates coherence $c = 4\%$, corresponding to the firing rates in **(A)**. The dotted lines indicate the coherence value at 75% accuracy (see **Figure 3A**). **(C)** Mean decision times over coherence for correct (solid) and error (dashed) trials for each condition. Error bars show standard error. The vertical line indicates coherence $c = 4\%$, corresponding to the firing rates in **(A)**.

under speed and accuracy conditions respectively; and for correct and error trials, mean decisions times are shorter and longer respectively. Thus, **Figure 2** shows that by raising and lowering baseline activity, uniform input to both integrator populations controls the SAT. At first glance, these results appear to support the threshold-baseline hypothesis.

However, the threshold-baseline hypothesis dictates that the speed and accuracy of decisions are determined by the threshold-baseline difference. According to this hypothesis, a fixed threshold-baseline difference will produce uniform decision making performance, regardless of the rate of baseline activity. The threshold-baseline hypothesis therefore requires that any changes to the speed or accuracy of decisions resulting from a change in baseline activity (with a fixed threshold) can be "reversed" by an equal change to the threshold. We therefore increased the threshold under the speed condition by the difference between baseline activity under speed and neutral conditions (Δ_{ns}, the mean difference over the last 1000 ms of the pre-stimulus interval), and we decreased the threshold under the accuracy condition by the difference between baseline activity under neutral and accuracy conditions (Δ_{na}). These adjustments to the threshold did not recover the psychometric and chronometric curves produced under the neutral condition, i.e., the black and light gray curves in **Figures 2B,C** do not overlay the medium gray curves. Denoting the threshold used by Wong and Wang (2006) as θ (vertical line in **Figure 3**), increasing (decreasing) θ by Δ_{ns} (Δ_{na}) under the speed (accuracy) condition has almost no effect on performance. The same is true for any value of the choice threshold above θ. For thresholds below θ, the effect of these adjustments increases with decreasing threshold, but the psychometric (**Figure 3A**) and chronometric (**Figures 3B,C**) curves under speed and accuracy conditions do not come close to overlaying the neutral curves. For the lowest value of the threshold, there is a moderate effect on the psychmetric curves (the difference between the solid and dotted curves for speed and accuracy conditions), but such a low threshold does not allow a firing-rate excursion, so this moderate effect can only be achieved if the model deviates from the neural data on which the threshold-baseline hypothesis is founded, i.e., a fixed rate of target-in activity that is much higher than target-out activity at the time of commitment to a choice (e.g., Shadlen and Newsome, 2001; Roitman and Shadlen, 2002; Thomas and Pare, 2007; Purcell et al., 2010; Bollimunta and Ditterich, 2011; Ding and Gold, 2012). See the Discussion for other issues with such a low threshold. The psychometric and chronometric curves break down for thresholds lower than those in the figure. Note that **Figures 3B,C** show results for coherence values of $c = 1\%$ and $c = 32\%$ respectively. Values in between these extremes yield the same qualitative result. These results demonstrate that the threshold-baseline hypothesis does not account for the SAT under the principles of the attractor framework.

3. THE SAT IS CONTROLLED BY NETWORK DYNAMICS

Returning to **Figure 2A**, the mean firing rates following motion onset (to the right of the vertical line) point to the mechanism by which the spatially non-selective input I_0 controls the SAT in the model. The rate of increase of target activity is higher and
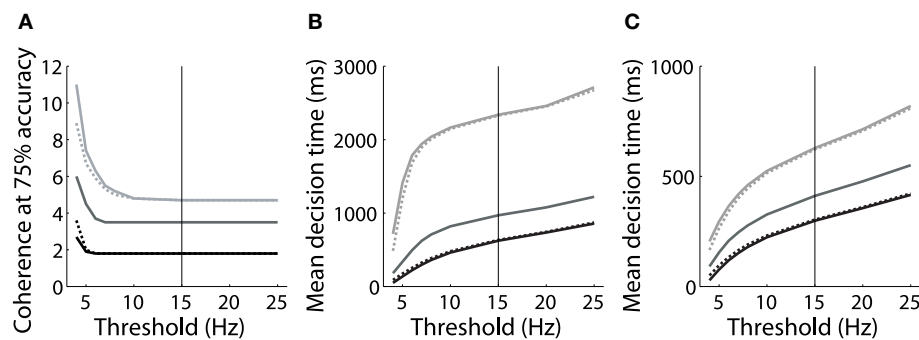
FIGURE 3 | (A) The value of simulated motion coherence at which accuracy equals 75% for a range of choice thresholds under speed (black), neutral (medium gray) and accuracy (light gray) conditions (see dashed lines in **Figure 2B**). Dotted curves show results for simulations in which the threshold was raised (speed condition) and lowered (accuracy condition) by Δ_{ns} and Δ_{na} respectively (see text). Solid curves show results for the corresponding unadjusted threshold. Adjusting the choice threshold makes little difference to accuracy, i.e., the solid and dotted curves are similar for speed and accuracy conditions, and do not overlay the neutral curve. **(B,C)** Mean decision times over all trials for coherence $c = 1\%$ **(B)** and $c = 32\%$ **(C)**. Conventions are the same as in **(A)**. Adjusting the threshold by Δ_{ns} and Δ_{na} makes little difference to decision times, regardless of the threshold chosen. The solid vertical line in each panel indicates the threshold used by Wong and Wang (2006).

lower under speed and accuracy conditions respectively, relative to the neutral condition. The different rates of increase reflect the dynamics furnished by the different values of I_0 under speed, accuracy and neutral conditions. As shown by Wong and Wang (2006), the dynamics in the vicinity of the saddle point determine the length of time the network can integrate evidence, which can be approximated by the time constant of the unstable manifold (the effective time constant of integration, Supplementary Material Section 2). Wong and Wang (2006) calculated this time constant for several values of the strength of recurrent excitation, showing the consequent changes to the speed and accuracy of decisions (see their Figure 11). **Figure 4A** shows these calculations for our changes to I_0. Under speed and accuracy conditions, higher and lower values of I_0 furnish shorter and longer time constants respectively, relative to the neutral condition. Here, it is worth noting that the effective time constant behaves in exactly the same way as the bound of bounded integration models, decreasing (increasing) integration time under speed (accuracy) conditions (**Figure 4A**). Additionally, the shape of the attractor landscape changes with I_0. **Figures 4B–D** show that for a given task difficulty ($c = 4\%$ in the figure), higher values of I_0 push the stable manifold toward the midline at low rates below the saddle point. Since the network approaches the saddle from below (**Figure 1C**) and since errors occur when noise pushes the state of the network over the stable manifold (Wong and Wang, 2006), this re-positioning of the stable manifold further lowers (raises) accuracy under speed (accuracy) conditions. This mechanism is evident in **Figures 4B–D**, in which the solid circle in each panel shows the mean initial state of the network (immediately prior to the onset of evidence). With increasing I_0, the stable manifold moves toward this initial state, which becomes increasingly precarious. Thus, a common input to integrators controls the rate of baseline activity, but the SAT does not result from the consequent changes to the threshold-baseline difference. The SAT results from the changes to network dynamics.

Increasing I_0 not only re-positions the stable manifold, but also re-positions the saddle point, so that both populations fire at higher rates (**Figures 4B–D**). This change in position of the saddle dictates that firing rates will be higher when the network begins its descent into an attractor basin under speed conditions. In other words, firing rates will be higher when decision-selective rates separate from those of the competing population. To confirm this effect, we used signal detection theory to determine when an ideal observer can discriminate target activity from distractor activity in the model under speed, accuracy and neutral conditions (Supplementary Material Section 3). Signal detection theory is commonly used to estimate the time of target selection from neural data (Thompson et al., 1996; Cohen et al., 2009) and assumes that a downstream circuit makes decisions by discriminating the activity of neural populations selective for the alternatives (see Standage and Pare, 2011). Firing rates at the time of discrimination were higher under speed conditions and lower under accuracy conditions (**Figure 5**).

Next, we subtracted the baseline rate under speed, accuracy and neutral conditions from the corresponding rate at discrimination time (the discrimination-baseline difference). The discrimination-baseline difference was larger under speed conditions and smaller under accuracy conditions. Because decisions are over when the firing rates separate, the rate at this time approximates a "decision threshold," as opposed to the choice threshold (see the Discussion). To summarize: the difference between this decision threshold and baseline activity is larger under speed conditions and smaller under accuracy conditions in the model. Thus, stronger (weaker) non-selective input under speed (accuracy) conditions modulates decision-selective firing rates in a manner opposite to the principles of the threshold-baseline hypothesis. We confirmed these findings with an alternative method, in which decision times (and correctness) were determined by the last intersection of target and distractor activity on each trial, i.e., decisions were made when target and distractor activity separated for the final time. The mean rate at the time of separation was higher (lower) under speed (accuracy) conditions, as was the difference between this rate and the baseline rate (not shown). Importantly, our analysis in this section makes
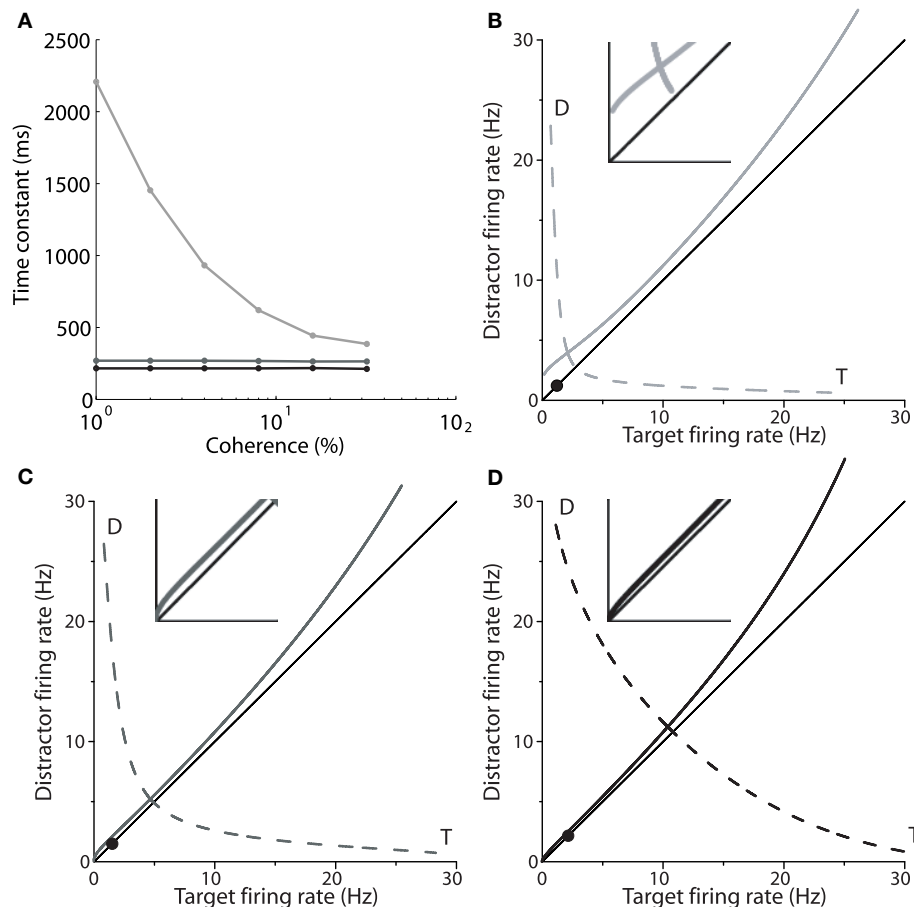
FIGURE 4 | (A) The time constant of the unstable manifold of the saddle point (see **Figure 1**) for speed (black), neutral (medium gray) and accuracy (light gray) conditions, as a function of coherence. The time constant determines the time over which the system is repelled from the saddle toward an attractor corresponding to the target or the distractor (**Figure 1C**; T and D in **B–D**). **(B–D)** Stable (solid) and unstable (dashed) manifolds of the

saddle for the accuracy (**B**, $I_0 = 316$pA), neutral (**C**, $I_0 = 321$pA) and speed (**D**, $I_0 = 325$pA) conditions for coherence $c = 4\%$. At low rates below the saddle, the stable manifold is pushed closer to the midline with increasing I_0, while rates at the saddle increase. Solid circles show the initial state of the network. Insets show close-ups of the stable manifold and the midline at frequencies ≤ 5 Hz.

two predictions for electrophysiological studies of the SAT: (1) target-in and target-out data will separate at higher (lower) rates under speed (accuracy) conditions, and (2) the discrimination-baseline difference will be larger (smaller) under speed (accuracy) conditions.

4. DISCUSSION AND CONCLUSIONS

We have demonstrated that spatially non-selective excitation can control the SAT in an attractor model (**Figures 2B,C**), as shown previously (Furman and Wang, 2008; Roxin and Ledberg, 2008). The non-selective input increases and decreases baseline activity under speed and accuracy conditions respectively (**Figure 2A**), which unavoidably decreases and increases the difference between baseline activity and a fixed choice threshold. The threshold-baseline difference, however, does not control the SAT in the model (**Figure 3**). Rather, an increase (decrease) in non-selective input increases (decreases) the strength of network dynamics, which decreases (increases) the effective time constant of integration (**Figure 4A**) and renders the initial state of the network closer

to (farther from) the stable manifold of the saddle, the crossing of which results in errors (**Figures 4B–D**).

Our findings are consistent with the hypothesis that a cognitive signal controls the SAT by adjusting a uniform input to integrator populations (see Bogacz et al., 2010b; Standage et al., 2014). This hypothesis is supported by neuroimaging (Forstmann et al., 2008; Ivanoff et al., 2008; van Veen et al., 2008; Wenzlaff et al., 2011) and electrophysiological (Heitz and Schall, 2012; Hanks et al., 2014) data from SAT tasks, where pre-stimulus activation has been shown to be higher (lower) under speed (accuracy, neutral) conditions. Like the threshold-baseline hypothesis, our results are consistent with these data. Our results conflict with the threshold-baseline hypothesis because the changes in network dynamics engendered by a uniform input dwarf the corresponding changes to the threshold-baseline difference. A related reason is that the choice threshold is qualitatively different than the bound of bounded integration models. The rate of target-in activity at the time of commitment to a choice has been shown to be considerably higher than the rate at which this activity separates

FIGURE 5 | The mean firing rate of the target population at the time at which an ideal observer can discriminate target activity from distractor activity, calculated across all positive coherence values. The rate is higher (lower) under speed (accuracy) conditions.
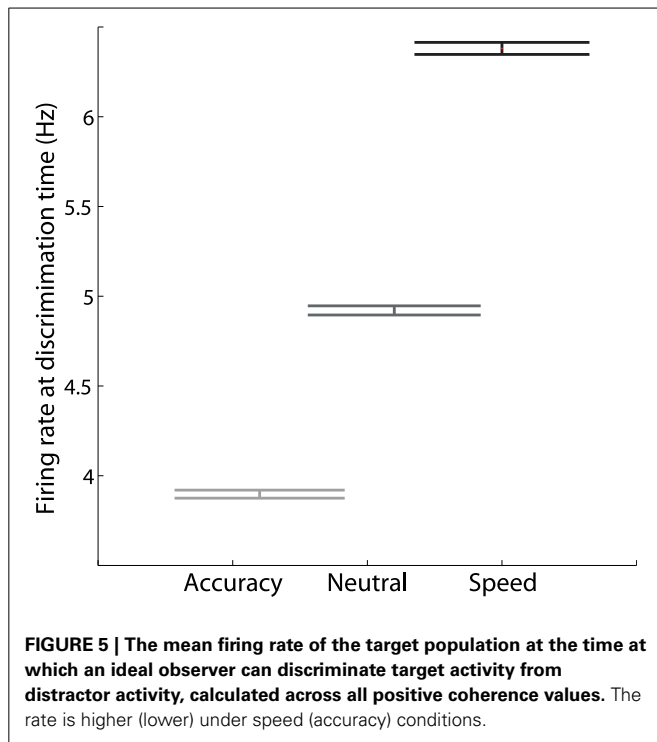
from target-out activity (see e.g., Shadlen and Newsome, 2001; Roitman and Shadlen, 2002; Bollimunta and Ditterich, 2011; Ding and Gold, 2012). Under the framework of attractor dynamics, this excursion of target-in activity corresponds to the repulsion of a decision network from the saddle along its unstable manifold. Thus, these neural data suggest that the choice threshold is much higher than the saddle. As such, changes to the choice threshold will not influence decision accuracy over a broad range, unlike the bound of bounded integration models. This effect is clear in **Figure 2A**. As noted in Section 3, the rate at which target and distractor activity separates can be thought of as a "decision threshold," but our simulations predict that this rate is not fixed across speed and accuracy conditions. Indeed, we predict that it changes (**Figure 5**) in a manner opposite to a flexible bound (e.g., Ratcliff and McKoon, 2008; Bogacz et al., 2010a). Our findings therefore suggest that the bound is not implemented in terms of firing rates *per se*. In this regard, the astute reader may have noticed our use of the term "choice threshold" when referring to decision-selective firing rates at the time of commitment to a choice, as opposed to the more conventional "decision threshold." We believe the latter term is misleading in this context.

There are potential advantages to choice thresholds being higher than decision thresholds. For example, a high choice threshold alleviates the need for fine tuning (Roxin and Ledberg, 2008). Furthermore, the difference between the choice threshold and a decision threshold provides a buffer between decisions and their enactment. This buffer may confer advantages to decision makers. For instance, a high choice threshold gives an upstream decision variable the opportunity to suppress its competitors, that is, the choice is not made until the "winning" integrator

population is firing at a high rate and the losing populations are firing at much lower rates. Thresholds are hypothesized to be implemented by networks with very strong dynamics (Simen, 2012), which are poorly suited to decision making (Standage and Pare, 2011), i.e., they implement an all-or-none response to a critical level of input. If the respective rates of the choice threshold and the decision threshold were similar (a small buffer), then the difference between the decision variables would be smaller when the largest one reaches the choice threshold, increasing the possibility that the thresholding circuit would inadvertently choose the wrong decision variable. Simultaneous electrophysiological recordings from decision circuitry and thresholding circuitry would be informative in this regard. It seems unlikely that target-in activity in one structure would coincide with target-out activity in the other, even infrequently. Another possibility is that thresholding circuitry implements an ideal observer of integrator circuitry, where back-projections from the former to the latter account for the excursion of decision-selective activity prior to choice selection (see Simen, 2012). Under this scenario, bidirectionally-coupled decision circuits would collectively implement both integration and choice, a compelling possibility that warrants further investigation.

Another perspective on the difficulties of equating the difference between the bound and the starting point of a decision variable with the threshold-baseline difference relates to levels of abstraction in models of brain function (Marr, 1982; Trappenberg, 2010). From this perspective, bounded integration models can be considered algorithms that characterize the computations underlying decisions. They have been (and continue to be) invaluable for our understanding of decision processing and the SAT, but it is not necessary to attribute direct biological correlates to each of their parameters. Qualitatively, the effective time constant of integration under speed and accuracy conditions changes in the same manner as the bound (**Figure 4A**) and therefore provides a plausible neural implementation of this abstract term, but the corresponding changes to the attractor basins show that this interpretation may be overly simplistic (**Figures 4B–D**). Note that we do not suggest the twain shall never meet. Far from it, formal equivalence has been shown between different classes of (linear) bounded integration models and the (non-linear) biophysically-based model on which our simulations are based (Bogacz et al., 2006). The constraints under which these models are equivalent define the relationship between decision models at these two levels of abstraction, allowing the systematic consideration of one class in terms of the other. Where earlier work has largely considered the commonalities between classes of model, e.g., the range of parameters under which non-linear, feedback-inhibition models are well-approximated by linear integration models (Usher and McClelland, 2001; Bogacz et al., 2006), we have focused on their differences. In this sense, we have shown what is lost in translation in relation to the SAT, suggesting that caution is warranted when interpreting neural data in terms of models that are purposefully simplified. Note that earlier discussions of the threshold-baseline hypothesis have made it clear that changes to the bound and the starting point of a decision variable are not equivalent in all abstract models (Bogacz et al., 2010b). For more extensive treatment of the constraints of the

threshold-baseline hypothesis in relation to implementation-level models, see Marshall et al. (2012).

It is possible that a different kind of threshold-baseline difference could account for the SAT. If the baseline rate of *thresholding circuitry* were increased (decreased) under speed (accuracy) conditions, then lower rates of integrator activity would be sufficient to elicit choice behavior, i.e., to drive the relevant motor circuitry (see Standage et al., 2014 for review). As such, a cognitive signal controlling the SAT could bypass integrator populations. However, the rates of integrator populations at the time of commitment to a choice would be lower under speed conditions and higher under accuracy conditions, which conflicts with recent electrophysiological recordings from putative integrator neurons showing the opposite profile of activity (Heitz and Schall, 2012). Notably, these data also show higher (lower) baseline rates and a higher (lower) rate of increase under speed (accuracy) conditions, suggesting that speed and accuracy conditions do modulate integrator neurons. These findings are qualitatively reproduced by our simulations (**Figure 2A**).

Finally, we do not suggest that single-circuit attractor models provide a complete picture of decision making. For example, these models produce slower mean decision times on error trials than correct trials because the network state has to cross the unstable manifold (Wong and Wang, 2006; Standage et al., 2011), but error trials are faster than correct trials under some task paradigms (see Smith and Ratcliff, 2004). Such shortcomings point to the need for coupled-circuit models (e.g., Lo and Wang, 2006; Standage et al., 2013). The recent surge in neuroimaging studies of decision making and the SAT represents an important direction in this regard, identifying contributing brain regions and pointing to their respective roles in decision processing (Forstmann et al., 2008; Ivanoff et al., 2008; van Veen et al., 2008; Forstmann et al., 2010; van Maanen et al., 2011; Wenzlaff et al., 2011; Green et al., 2012; Ho et al., 2012). Guided by these data, models of distributed decision circuitry are an exciting direction in decision neuroscience (Frank, 2006; Lo and Wang, 2006; Bogacz and Gurney, 2007). Simulations of the bidirectional coupling between circuits supporting evidence integration and choice may be highly informative about the relationship between decision bounds and choice thresholds.

SUPPLEMENTARY MATERIAL
The Supplementary Material for this article can be found online at: http://www.frontiersin.org/journal/10.3389/fnins.2014.00318/abstract

REFERENCES
Bogacz, R., Brown, E., Moehlis, J., Holmes, P., and Cohen, J. D. (2006). The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol. Rev.* 113, 700–765. doi: 10.1037/0033-295X.113.4.700

Bogacz, R., and Gurney, K. (2007). The basal ganglia and cortex implement optimal decision making between alternative actions. *Neural Comput.* 19, 442–477. doi: 10.1162/neco.2007.19.2.442

Bogacz, R., Hu, P. T., Holmes, P. J., and Cohen, J. D. (2010a). Do humans produce the speed-accuracy tradeoff that maximizes reward rate? *Q. J. Exp. Psychol.* 63, 863–891. doi: 10.1080/17470210903091643

Bogacz, R., Wagenmakers, E.-J., Forstmann, B. U., and Nieuwenhuis, S. (2010b). The neural basis of the speed-accuracy tradeoff. *Trends Neurosci.* 33, 10–16. doi: 10.1016/j.tins.2009.09.002

Bollimunta, A., and Ditterich, J. (2011). Local computation of decision-relevant net sensory evidence in parietal cortex. *Cereb. Cortex* 22, 903–917. doi: 10.1093/cercor/bhr165

Churchland, A. K., Kiani, R., and Shadlen, M. N. (2008). Decision-making with multiple alternatives. *Nat. Neurosci.* 11, 693–702. doi: 10.1038/nn.2123

Cohen, J. Y., Crowder, E. A., Heitz, R. P., Subraveti, C. R., Thompson, K. G., Woodman, G. F., et al. (2009). Cooperation and competition among frontal eye field neurons during visual target selection. *J. Neurosci.* 30, 3227–3238. doi: 10.1523/JNEUROSCI.4600-09.2010

Ding, L., and Gold, J. I. (2012). Neural correlates of perceptual decision making before, during, and after decision commitment in monkey frontal eye field. *Cereb. Cortex* 22, 1052–1067. doi: 10.1093/cercor/bhr178

Forstmann, B. U., Anwander, A., Schafer, A., Neumann, J., Brown, S., Wagenmakers, E.-J., et al. (2010). Cortico-striatal connections predict control over speed and accuracy in perceptual decision making. *Proc. Natl. Acad. Sci. U.S.A.* 107, 15916–15920. doi: 10.1073/pnas.1004932107

Forstmann, B. U., Dutilh, G., Brown, S., Neumann, J., von Cramon, D. Y., Ridderinkhof, K. R., et al. (2008). Striatum and pre-sma facilitate decision-making under time pressure. *Proc. Natl. Acad. Sci. U.S.A.* 105, 17538–17542. doi: 10.1073/pnas.0805903105

Frank, M. J. (2006). Hold your horses: a dynamic computational role for the subthalamic nucleus in decision making. *Neural Netw.* 19, 1120–1136. doi: 10.1016/j.neunet.2006.03.006

Furman, M., and Wang, X.-J. (2008). Similarity effect and optimal control of multiple-choice decision making. *Neuron* 60, 1153–1168. doi: 10.1016/j.neuron.2008.12.003

Green, N., Biele, G. P., and Heekeren, H. R. (2012). Changes in neural connectivity underlie decision threshold modulation for reward maximization. *J. Neurosci.* 32, 14942–14950. doi: 10.1523/JNEUROSCI.0573-12.2012

Hanes, D. P., and Schall, J. D. (1996). Neural control of voluntary movement initiation. *Science* 274, 427–430. doi: 10.1126/science.274.5286.427

Hanks, T., Kiani, R., and Shadlen, M. N. (2014). A neural mechanism of speed-accuracy tradeoff in macaque area lip. *eLife* 3:e02260. doi: 10.7554/eLife.02260

Heitz, R. P., and Schall, J. D. (2012). Neural mechanisms of speed-accuracy tradeoff. *Neuron* 76, 616–628. doi: 10.1016/j.neuron.2012.08.030

Ho, T., Brown, S., van Maanen, L., Forstmann, B. U., Wagenmakers, E.-J., and Serences, J. T. (2012). The optimality of sensory processing during the speed-accuracy tradeoff. *J. Neurosci.* 32, 7992–8003. doi: 10.1523/JNEUROSCI.0340-12.2012

Ivanoff, J., Branning, P., and Marois, R. (2008). fmri evidence for a dual process account of the speed-accuracy tradeoff in decision-making. *Public Library Sci. One* 3, 1–14. doi: 10.1371/journal.pone.0002635

Lo, C.-C., and Wang, X.-J. (2006). Corticobasal ganglia circuit mechanism for a decision threshold in reaction time tasks. *Nat. Neurosci.* 9, 956–963. doi: 10.1038/nn1722

Machens, C. K., Romo, R., and Brody, C. D. (2005). Flexible control of mutual inhibition: a neural model of two-interval discrimination. *Science* 307, 1121–1124. doi: 10.1126/science.1104171

Marr, D. (1982). *Vision.* New York, NY: W. H. Freeman and Company.

Marshall, J. A. R., Bogacz, R., and Gilchrist, I. D. (2012). Consistent implementation of decisions in the brain. *PLoS ONE* 7:e43443. doi: 10.1371/journal.pone.0043443

Purcell, B. A., Heitz, R. P., Cohen, J. Y., Schall, J. D., Logan, G. D., and Palmeri, T. J. (2010). Neurally constrained modeling of perceptual decision making. *Psychol. Rev.* 117, 1113–1143. doi: 10.1037/a0020311

Ratcliff, R., and McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. *Neural Comput.* 20, 873–922. doi: 10.1162/neco.2008.12-06-420

Roitman, J. D., and Shadlen, M. N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J. Neurosci.* 22, 9475–9489.

Roxin, A., and Ledberg, A. (2008). Neurobiological models of two-choice decision making can be reduced to a one-dimensional nonlinear diffusion equation. *Public Library Sci. Comput. Biol.* 4:e1000046. doi: 10.1371/journal.pcbi.1000046

Salinas, E., and Abbott, L. F. (1996). A model of multiplicative neural responses in parietal cortex. *Proc. Natl. Acad. Sci. U.S.A.* 93, 11956–11961. doi: 10.1073/pnas.93.21.11956

Shadlen, M. N., and Newsome, W. T. (2001). Neural basis of a perceptual decision in the parietal cortex (area lip) of the rhesus monkey. *J. Neurophysiol.* 86, 1916–1936.

Simen, P. (2012). Evidence accumulator or decision threshold - which cortical mechanism are we observing? *Front. Psychol.* 3:183. doi: 10.3389/fpsyg.2012.00183

Smith, P. L., and Ratcliff, R. (2004). Psychology and neurobiology of simple decisions. *Trends Neurosci.* 27, 161–168. doi: 10.1016/j.tins.2004.01.006

Standage, D., Blohm, G., and Dorris, M. C. (2014). On the neural implementation of the speed-accuracy trade-off. *Front. Neurosci.* 8:236. doi: 10.3389/fnins.2014.00236

Standage, D., and Pare, M. (2011). Persistent storage capability impairs decision making in a biophysical network model. *Neural Netw.* 24, 1062–1073. doi: 10.1016/j.neunet.2011.05.004

Standage, D., You, H., Wang, D.-H., and 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. *Front. Comput. Neurosci.* 5:7. doi: 10.3389/fncom.2011.00007

Standage, D., You, H., Wang, D.-H., and Dorris, M. C. (2013). Trading speed and accuracy by coding time: a coupled-circuit cortical model. *PLoS Comput. Biol.* 9:e100321. doi: 10.1371/journal.pcbi.1003021

Thomas, N. W. D., and Pare, M. (2007). Temporal processing of saccade targets in parietal cortex area LIP during visual search. *J. Neurophysiol.* 97, 942–947. doi: 10.1152/jn.00413.2006

Thompson, K., Hanes, D., Bichot, N., and Schall, J. (1996). Perceptual and motor processing stages identified in the activity of macaque frontal eye field. *J. Neurophysiol.* 76, 440–455.

Trappenberg, T. P. (2010). *Fundamentals of Computational Neuroscience.* New York, NY: Oxford University Press.

Usher, M., and McClelland, J. L. (2001). On the time course of perceptual choice: The leaky competing accumulator model. *Psychol. Rev.* 108, 550–592. doi: 10.1037/0033-295X.108.3.550

van Maanen, L., Brown, S. D., Eichele, T., Wagenmakers, E.-J., Ho, T., Serences, J., et al. (2011). Neural correlates of trial-to-trial fluctuations in response caution. *J. Neurosci.* 31, 17488–17495. doi: 10.1523/JNEUROSCI.2924-11.2011

van Veen, V., Krug, M. K., and Carter, C. S. (2008). The neural and computational basis of controlled speed-accuracy tradeoff during task performance. *J. Cogn. Neurosci.* 20, 1952–1965. doi: 10.1162/jocn.2008.20146

Wang, X.-J. (2002). Probabilistic decision making by slow reverberation in cortical circuits. *Neuron* 36, 955–968. doi: 10.1016/S0896-6273(02)01092-9

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

Wang, X.-J. (2012). Neural dynamics and circuit mechanisms for decision-making. *Curr. Opin. Neurobiol.* 22, 1039–1046. doi: 10.1016/j.conb.2012.08.006

Wenzlaff, H., Bauer, M., Maess, B., and Heekeren, H. R. (2011). Neural characterization of the speed-accuracy tradeoff in a perceptual decision-making task. *J. Neurosci.* 31, 1256–1266. doi: 10.1523/JNEUROSCI.4000-10.2011

Wong, K.-F., and Wang, X.-J. (2006). A recurrent network mechanism of time integration in perceptual decisions. *J. Neurosci.* 26, 1314–1328. doi: 10.1523/JNEUROSCI.3733-05.2006

You, H., and Wang, D.-H. (2013). Dynamics of multiple-choice decision making. *Neural Comput.* 25, 2108–2145. doi: 10.1162/NECO_a_00473

On the neural implementation of the speed-accuracy trade-off

Dominic Standage[1]*, Gunnar Blohm[1] and Michael C. Dorris[2]

[1] Department of Biomedical and Molecular Sciences, Queen's University, Kingston, ON, Canada
[2] Institute of Neuroscience, Shanghai Institutes for Biological Sciences, Chinese Academy of Sciences, Shanghai, China

Decisions are faster and less accurate when conditions favor speed, and are slower and more accurate when they favor accuracy. This phenomenon is referred to as the speed-accuracy trade-off (SAT). Behavioral studies of the SAT have a long history, and the data from these studies are well characterized within the framework of bounded integration. According to this framework, decision makers accumulate noisy evidence until the running total for one of the alternatives reaches a bound. Lower and higher bounds favor speed and accuracy respectively, each at the expense of the other. Studies addressing the neural implementation of these computations are a recent development in neuroscience. In this review, we describe the experimental and theoretical evidence provided by these studies. We structure the review according to the framework of bounded integration, describing evidence for (1) the modulation of the encoding of evidence under conditions favoring speed or accuracy, (2) the modulation of the integration of encoded evidence, and (3) the modulation of the amount of integrated evidence sufficient to make a choice. We discuss commonalities and differences between the proposed neural mechanisms, some of their assumptions and simplifications, and open questions for future work. We close by offering a unifying hypothesis on the present state of play in this nascent research field.

Keywords: speed-accuracy trade-off, decision making, neural mechanisms of cognition, bounded integration, review

1. INTRODUCTION

The ability to trade-off speed and accuracy against each other is a hallmark of decision making across species and tasks (Chittka et al., 2009; Bogacz et al., 2010a; Heitz and Schall, 2012). For a given task difficulty, decisions are typically faster and less accurate when conditions favor speed, and are slower and more accurate when conditions favor accuracy. Given the near-ubiquity of this behavior in experiments, the speed-accuracy trade-off (SAT) can almost be considered a psychophysical law. It can also be considered a cognitive phenomenon, since it captures a change in strategy toward an ostensibly unchanging task.

The SAT has long been the subject of behavioral experiments (Fitts, 1966; Wickelgren, 1977), but studies addressing its neural basis are a fairly recent development in the field of decision making (Bogacz et al., 2010b). These studies have built on a large body of work on the neural basis of decisions more generally. This work has characterized the computations underlying decisions (see Smith and Ratcliff, 2004; Ratcliff and McKoon, 2008), identified neural correlates of these computations (see Schall, 2001; Gold and Shadlen, 2007; Kable and Glimcher, 2009) and provided mechanistic hypotheses that explain behavioral data in terms of neural data (see Wang, 2008, 2012). This body of work provides a persuasive account of neural decision processing, but does not speak directly to the mechanisms by which decision processing is differentially modulated by conditions favoring speed or accuracy.

In this review, we describe hypotheses on the neural implementation of the SAT. We take a modeling perspective. We classify models according to two general levels of abstraction, sometimes referred to as the algorithmic level and the level of implementation (Marr, 1982). These classes need not be considered discrete, but rather, can be considered as the extreme ends of a continuum. At one end, algorithmic models characterize the computations underlying brain function. At the other end, neural models address the implementation of these computations. In the domain of decision making, analytic studies have shown the assumptions and constraints under which implementation-level models are formally equivalent to algorithmic models, providing a principled foundation for considering the latter in terms of the former (Bogacz et al., 2006). We endeavor to utilize the flexibility and explanatory power of this modeling perspective.

Our review is structured according to the framework of bounded integration. This framework not only provides a set of organizing principles for the review, but provides the background for this collection more generally. Most of the neural and behavioral data we consider were recorded from perceptual decision tasks. We assume that the neural mechanisms underlying perceptual decisions generalize to other kinds of decisions, but the sources of evidence differ according to the decision domain (Gold and Shadlen, 2007). Sensory systems and memory systems provide examples of sources of evidence. We begin by defining the SAT (Section 2). We then describe bounded integration as a

computational framework for characterizing decisions (Section 3), along with a widely held hypothesis on the neural implementation of these computations (Section 3.1). We categorize hypotheses on the SAT according to the major components of the bounded integration framework, describing the evidence for differential modulation of these components under speed and accuracy conditions (Sections 4.1, 4.2, and 4.3). We close with a discussion of the assumptions underlying these hypotheses, the relationship between mechanisms, and some open questions for future research (Section 5).

2. DEFINING THE SPEED-ACCURACY TRADE-OFF

In decision tasks, subjects must determine which decision alternative is favored by the evidence. If the evidence for one alternative is clearly stronger than the evidence for the others, the task is easy. Conversely, if the evidence for each alternative is similar, the task is difficult. Accuracy decreases with task difficulty, while decision times increase, characterizing the common psychometric and chronometric curves respectively (**Figure 1**). Task difficulty therefore imposes a systematic relationship between the speed and accuracy of decisions (see Stone, 2014 in this collection), but these curves do not define the SAT. The SAT refers to changes in the speed and accuracy of decisions *for a given task difficulty*. While many decision tasks manipulate the strength of evidence, this experimental parameter need not vary in SAT experiments.

The SAT captures a control mechanism for decision processing, and can be further distinguished according to the timescale of adjustments to speed and accuracy conditions. Over longer timescales, the SAT may be accomplished by adaptive mechanisms that extract a balance between the speed and accuracy of decisions in order to maximize reward over a block of trials (Gold and Shadlen, 2002; Simen et al., 2006; Furman and Wang, 2008; Standage et al., 2011). This approach has been demonstrated by algorithmic models (Gold and Shadlen, 2002; Bogacz et al., 2010a), biophysically-based neural models (Lo and Wang, 2006; Furman and Wang, 2008) and models in between these levels of abstraction (Simen et al., 2006). In contrast, experimental subjects often learn to respond to speed or accuracy conditions from trial to trial, according to a pre-trial cue (Forstmann et al., 2008; Heitz and Schall, 2012). We point out this difference because we are unaware of any implementation-level models to simulate trial-to-trial switching of response "modes" for speed and accuracy. Since there is an optimal trade-off for *each* condition that depends on its associated reward schedule, it is plausible that long-timescale mechanisms correspond to a learning phase for each response mode; however, it is important to note that switching between speed and accuracy modes necessarily involves additional mechanisms to associate the cues with the appropriate mode, and to switch between modes on cue.

3. THE BOUNDED INTEGRATION FRAMEWORK

Under the bounded integration framework, the evidence for each alternative of a decision is integrated until the running total for one of the alternatives reaches a criterion level. Thus, the bound refers to the criterion and integration refers to the accumulation of evidence. The accumulated evidence for a given alternative is referred to as a decision variable. According to
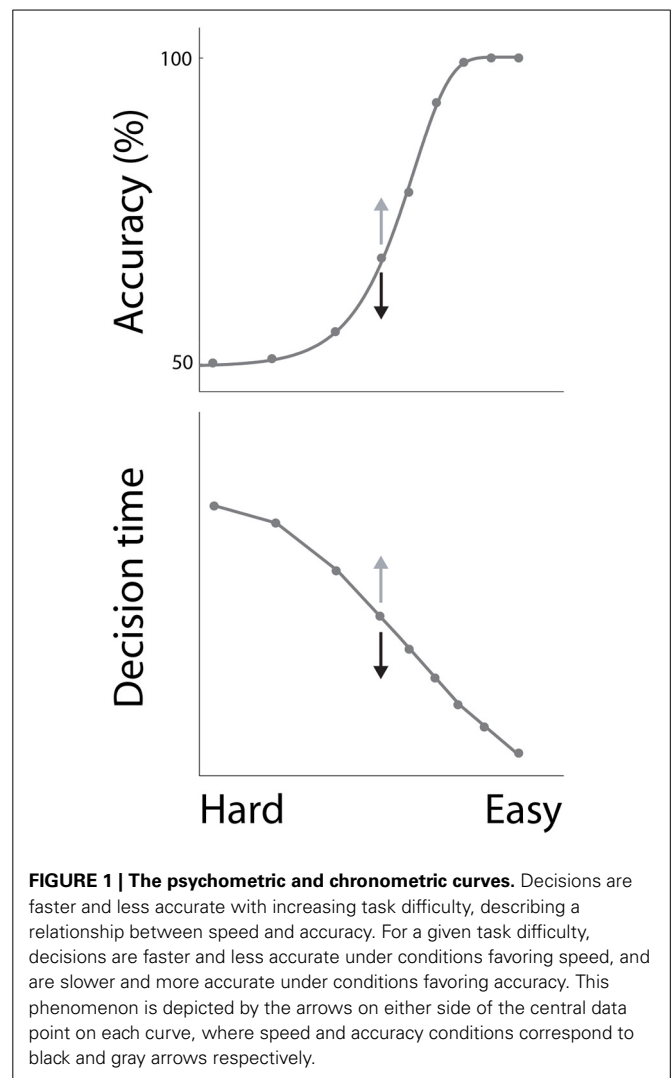
FIGURE 1 | The psychometric and chronometric curves. Decisions are faster and less accurate with increasing task difficulty, describing a relationship between speed and accuracy. For a given task difficulty, decisions are faster and less accurate under conditions favoring speed, and are slower and more accurate under conditions favoring accuracy. This phenomenon is depicted by the arrows on either side of the central data point on each curve, where speed and accuracy conditions correspond to black and gray arrows respectively.

this sequential sampling approach (see Ratcliff and Smith, 2004; Smith and Ratcliff, 2004), integration is necessary because neural processing of the evidence is noisy, as may be the evidence itself. By integrating the evidence over time, an average is computed, so that decisions are not based on moment-to-moment fluctuations in the evidence or its processing. The longer the integration period, the better the average and the higher the probability of identifying the alternative with the most evidence. Clearly, speed and accuracy make conflicting demands under this framework.

Bounded integration subsumes a number of algorithmic models. Most generally, these models can be distinguished according to whether the evidence for each choice is integrated independently from the others, or whether the evidence for each choice serves as evidence against the others. The former are often referred to as race models (**Figure 2A**) and the latter as diffusion models (**Figure 2C**). A flexible approach between these extremes is provided by competing accumulator models (Usher and McClelland, 2001; Bogacz et al., 2007; Purcell et al., 2012), in which decision variables for the respective alternatives are
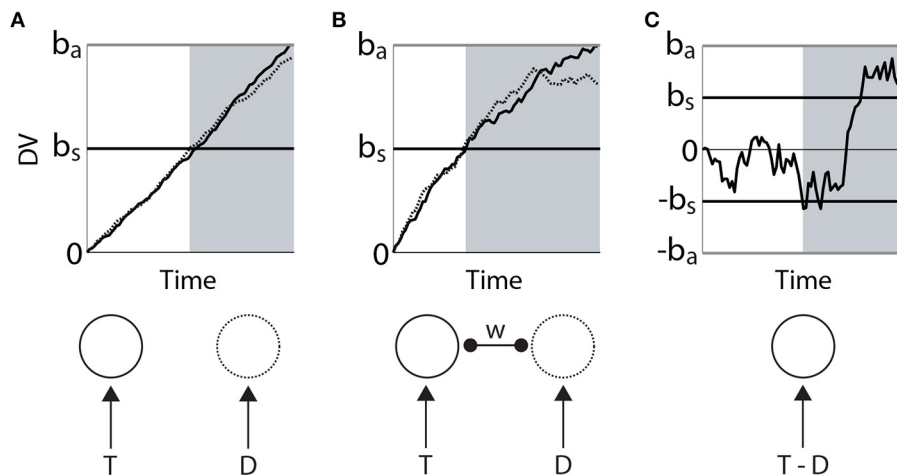
FIGURE 2 | Three classes of bounded integrator model. Each model receives the same two noisy stimuli, one with a higher mean (target T, solid) than the other (distractor D, dotted). Curves in the upper figures correspond to integrators (decision variables DV), depicted in the lower figure. **(A)** Independent race model of a 2-choice decision. The black horizontal line b_s corresponds to a low decision bound, supporting faster decisions that are less likely to identify the target. The gray horizontal line b_a corresponds to a higher bound, favoring the accurate identification of the target at the expense of processing time. The independence of the integrators is depicted in the lower figure. **(B)** Competing accumulator model. The weight (w) of subtraction between the two integrators is depicted in the lower figure. Different values of this weight would yield different curves. **(C)** Drift diffusion model. The decision variable is the integrated difference between the two stimuli. Black (b_s) and gray (b_a) horizontal lines correspond to bounds favoring speed and accuracy respectively. In each panel, the gray shaded region depicts the time of crossing of the lower (speed condition) threshold.

subtracted from one another according to a scaling parameter or *weight* (**Figure 2C**). In 2-choice tasks, the weight of subtraction can effectively (though not always formally) interpolate between the independent race model and the diffusion model, i.e., it controls the strength of competition between accumulators. Moreover, competing accumulators accommodate tasks with any number of choices and they provide an important link between models at the algorithmic level and the implementation level (see the next section). For an intuitive description of the formal relationships between race models, diffusion models and competing accumulators, see Bogacz (2007). For a rigorous mathematical treatment, see Bogacz et al. (2006).

This brief description of bounded integration warrants several technical points. Firstly, integration refers to the accumulation of evidence in continuous time, but for simplicity, we do not distinguish accumulation in discrete time from the continuous-time case. Secondly, the benefits of integration depend on the timescale of noise correlations. Thirdly, we only consider unbiased tasks, in which the bound (or its mean) is the same for each alternative, as is the starting value (or its mean) of each decision variable. Note that "unbiased" does not imply that the mean evidence for each alternative is equal, but rather, the prior probability of each alternative is equal. The framework is readily extended to biased conditions (see Gold and Shadlen, 2001). For a comprehensive description of bounded integration, see Smith and Ratcliff (2004); Bogacz et al. (2006).

3.1. INTERPRETING BOUNDED INTEGRATOR MODELS

As noted in the Introduction, bounded integrator models can be thought of as abstract algorithms that characterize the computations underlying decisions. From this perspective, the terms and parameters of these models are independent of their implementation and do not require explicit neural interpretation. On the other hand, it can be instructive to interpret these parameters in neural terms if they resemble neural activity. As such, the evidence in perceptual decision tasks corresponds to the response by sensory (and sensory-association) neurons to task-relevant stimuli, and decision variables correspond to the activity of downstream neural populations hypothesized to integrate this activity. Accordingly, the starting point of a decision variable is commonly equated with the baseline (pre-trial) level of integrator activity and the bound is commonly equated with the level of this activity at the time of commitment to a choice (see Bogacz et al., 2010b).

There is considerable evidence supporting this general interpretation. For example, in random dot motion (RDM) tasks, subjects are rewarded for identifying the direction of coherent movement of a proportion of randomly moving dots on a computer screen. The coherence of the dots provides the evidence in the task, which can be precisely controlled by the experimenter. Neurons in the medial temporal area (MT) of monkeys are responsive to movement of the dots (Britten et al., 1992, 1993), and in tasks in which monkeys indicate their choices by making an eye-movement to a visual target, neurons in the lateral intraparietal area (LIP) that are responsive to the chosen target (target-in neurons) show buildup activity prior to choice selection (Roitman and Shadlen, 2002; Churchland et al., 2008). Since MT projects to LIP, it is widely believed that neurons in LIP integrate the evidence provided by MT, projecting in turn to the circuitry mediating eye-movements (see Gold and Shadlen, 2007; Shadlen and Kiani, 2013). Note that neural correlates of decision variables in RDM tasks have also been recorded in other cortical areas, e.g., dorsolateral prefrontal cortex (dlPFC)

(Kim and Shadlen, 1999) and the frontal eye fields (FEF) (Ding and Gold, 2012). Similar data have been recorded from these and other brain regions in different task paradigms, described below in relation to SAT experiments. Importantly, electrophysiological recordings from neurons responsive to a visual target that is *not* chosen on a given trial (target-out neurons) typically show a much lower rate of activity than target-in neurons prior to choice selection (e.g., Roitman and Shadlen, 2002; Thomas and Pare, 2007; Bollimunta and Ditterich, 2011; Ding and Gold, 2012). Taken together, increasing activity by target-in neurons and suppressed activity by target-out neurons have been interpreted as revealing competitive interactions between neural decision variables (Usher and McClelland, 2001; Wang, 2002; Albantakis and Deco, 2009; Standage and Pare, 2011). In competing accumulator models, each accumulator can be thought of as a population of neurons responsive to one of the alternatives, where the weight of subtraction corresponds to the strength of inhibition between these populations (**Figure 2B**).

Competing accumulator models can also have parameters governing leakage and recurrent excitation of decision variables, both of which are important for interpreting these models in neural terms. To begin with, neurons leak, e.g., membrane potential and synaptic activation decay. Importantly, the relevant time constants of decay (e.g., the time of decay from maximum to half-maximum) are on the order of tens of milliseconds, whereas perceptual decision times are typically in the range of several hundreds of milliseconds. Thus, the time constants of these currents are not long enough to support temporal integration. Such long integration times are believed to require recurrent excitation (Wang, 2002), provided by synaptic connectivity within a population of excitatory neurons responsive to a given alternative. To provide an idealized example, if the leakage and inhibitory synaptic currents of individual neurons (responding linearly to their inputs) were precisely offset by the strength of recurrent excitation from other neurons in the population, then each neuron would support perfect integration of evidence, limited only by its maximum firing rate. In reality, local-circuit dynamics constrain the length of time each population can support integration, described in the next section.

This neural interpretation of competing accumulator models sets the stage for our consideration of the neural basis of the SAT. In bounded integrator models, we interpret noisy evidence as the response by populations of sensory (and sensory-association) neurons to stimuli in perceptual tasks. We interpret temporal integration as the buildup activity of neural populations receiving projections from sensory neurons. We interpret the starting point of a decision variable as the activity of integrator populations at the time of evidence onset (the baseline rate), and for simplicity, we interpret the bound as the rate of integrator activity at the time of commitment to a choice. We consider another interpretation of the bound in Section 4.2.1.

3.1.1. Attractor dynamics

The time over which competing neural populations can integrate evidence is an emergent property of network dynamics. The relevant dynamics are most easily described for 2-choice decisions, but are applicable to more than two decision alternatives (You and

Wang, 2013). As noted above, when the activity of an integrator population builds up in a 2-choice task, it suppresses the other population by recurrent inhibition. The eventual state of high-rate activity by one population and low-rate activity by the other is an attractor in the space of possible states of the network, and the increase in activity by the "winning" population and the suppression of the losing population (**Figure 3B**) corresponds to a descent into its basin of attraction (**Figure 3C**). The attractors are stable states of the network, that is, the state of the network evolves toward these states for a given set of conditions. Once there, the mean activity of the network is fixed until conditions change, such as the offset of evidence. In the domain of decision making, the "getting there" is the decision process.

The attractors are separated by an *unstable* steady state, toward which the network is drawn with the onset of the evidence, and from which it is repelled toward one of the two attractors (**Figure 3C**). The dynamics in the vicinity of the unstable steady state are slow, supporting temporal integration. The time over which integration is supported is referred to as the *effective time constant* of the network, and corresponds to the rate at which the dynamics evolve near this state. See Wong and Wang (2006) for a thorough description of the dynamics. The crucial point here is that the effective time constant is shorter with stronger recurrent dynamics, limiting the amount of time the network can integrate evidence. Accordingly, moderate dynamics can be considered to support neutral conditions, where stronger and weaker dynamics support speed and accuracy conditions respectively. We refer to local-circuit dynamics with these properties as the "decision regime." We refer to weaker dynamics without these properties as the "leakage regime." In the leakage regime, the effective time constant of the network is similar in principle to the time constant of decay of membrane potential or synaptic activation, though it can be considerably longer. In the decision regime, the effective time constant does not correspond to leakage; rather, it corresponds to an amplification of the decision variable, and is thus qualitatively different than a time constant of decay (see Standage et al., 2011).

4. THREE GENERAL MECHANISTIC HYPOTHESES ON THE SAT

Hypotheses on the neural implementation of the SAT must provide mechanistic explanations for differential decision processing under speed and accuracy conditions. Under the principles of bounded integration, these hypotheses can be grouped into three mutually-compatible classes: modulation of the encoding of evidence, modulation of the integration of encoded evidence, and modulation of the amount of integrated evidence sufficient to make a choice. In principle, each class of hypothesis (and each mechanistic hypothesis in each class) is sufficient to account for the SAT, but we do not favor any one hypothesis over the others. Rather, we believe the SAT is likely to result from the interplay of multiple mechanisms, with different mechanisms (or combinations of mechanisms) playing a greater role in different contexts.

The three general classes of hypothesis provide an intuitive basis for organizing the review, but they also correspond to three successive processing stages of decisions: the encoding of evidence, the integration of encoded evidence, and choosing. Under

FIGURE 3 | (A) A neural implementation of the principles of bounded integration. Neural populations selective for the decision alternatives (T and D) compete via a common inhibitory pool (solid black circle). Arcs with arrows and filled circles depict excitatory and inhibitory synaptic connectivity respectively. Target (T) and distractor (D) stimuli provide stronger and weaker evidence to the integrator populations respectively. **(B)** Competitive interactions between integrator populations lead to an increased spike rate by one population (solid) and a decreased rate by the other (dotted). **(C)** Cartoon depiction of an attractor "energy landscape" supported by the neural model, where the energy decreases over time. An unstable steady state (high energy) separates two attractors (low energy), corresponding to the target and the distractor. The ball depicts the state of the network, which is drawn toward the unstable steady state at stimulus onset (vertical arrow), and from which it is repelled toward one of the "attractor basins" (bent arrows). Descent into the attractor basin corresponds to the firing-rate excursion of the target population in **(B)**, where the vertical line approximates the position of the ball in **(C)**. The evolution of the network state (conceptually, the movement of the ball) is faster (slower) under speed (accuracy) conditions.

FIGURE 4 | Three processing stages for decisions: the encoding of evidence (left), the integration of encoded evidence (middle) and choice selection (right). Evidence-encoding populations **(left)** are responsive to target (T) and distractor (D) stimuli. Weak dynamics prevent integration, depicted by the lack of recurrent connectivity. Evidence-encoding populations project to integrator populations **(middle)**. Feedback connectivity depicts moderately strong dynamics, suitable for temporal integration (corresponding to **Figure 3**). Integrator populations project to thresholding circuitry **(right)**. Thick connectivity depicts very strong dynamics, suitable to an all-or none response to a critical level of input (see Simen, 2012).

the attractor framework, the computational requirements of these stages are supported by weak, moderate and strong local-circuit dynamics respectively. Weak dynamics support the encoding of evidence by "giving way" to their inputs, i.e., the dynamics are dominated by leakage. Moderately strong dynamics furnish a long effective time constant, supporting temporal integration (Section 3.1.1). Strong dynamics furnish a short effective time constant within the decision regime, allowing an all-or-none response to a critical level of input (see Simen, 2012). Thus, the principles of bounded integration are captured by a three-stage neural system, in which evidence-encoding circuitry with weak dynamics projects to integrator circuitry with moderate dynamics, which in turn projects to thresholding circuitry with strong dynamics. This three-stage process is depicted in **Figure 4**.

Finally, it is important to clarify our usage of several terms before proceeding with the review. We define the "correct" alternative as the one for which the evidence has the highest mean, and as suggested in Section 2, we define task difficulty as the difference between the mean of the evidence for the correct alternative and that for the alternative with the next hig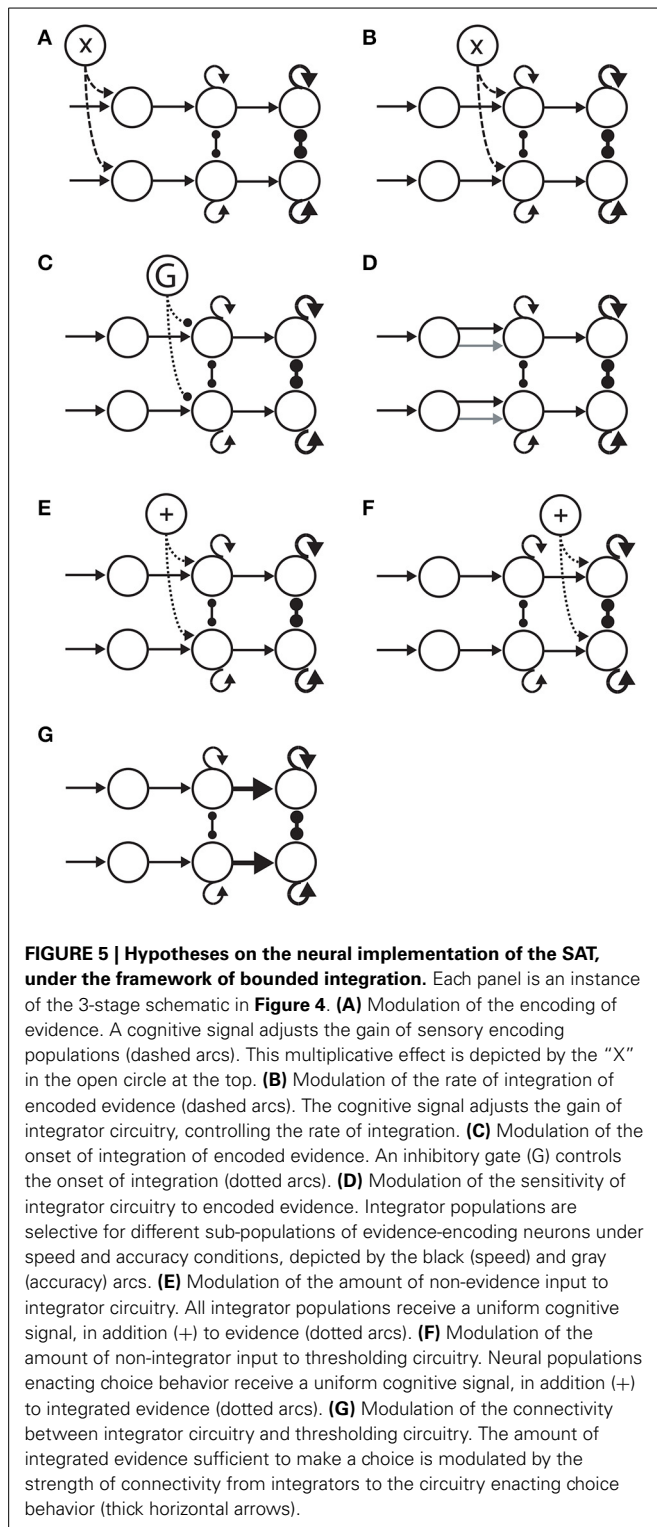hest mean. Task difficulty overlaps with the rate of integration in bounded integrator models, but this overlap depends on model specifics. For example, in race models, increasing the evidence for the correct alternative increases its integration rate (there's more instantaneous input to accumulate) and reduces task difficulty if the evidence for the other alternatives is not increased; however, increasing the evidence for each alternative by the same amount increases the integration rate of each integrator, but does not influence task difficulty. In diffusion models, an increase in the evidence for the correct alternative necessarily decreases task difficulty, unless the
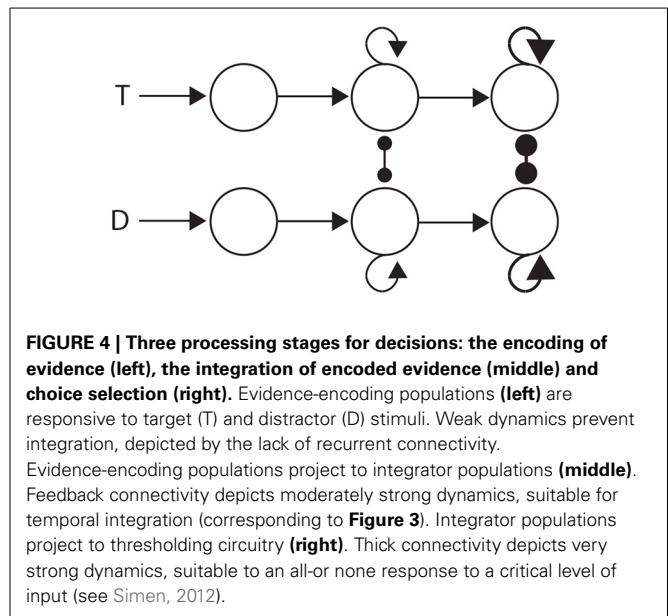
signal-to-noise ratio (SNR) of the evidence is preserved. Here, it is important to remember our definition of the SAT in Section 2: improvements in speed (accuracy) at the expense of accuracy (speed) for a *given* task difficulty. In Section 4.2.1, we describe hypotheses on the neural implementation of the SAT by modulation of the rate of integration. We define the rate of integration as the inverse of the difference between the rate of integrator neurons at the time of commitment to a choice and their baseline rate. These considerations highlight two important points. Firstly, the hypotheses in Section 4.2.1 do not refer to changes in integration rate resulting solely from upstream changes to the encoding of evidence (support for this possibility is described in Section 4.1). Secondly, these hypotheses address the neural mechanisms by which the rate of rise of putative integrator activity is modulated by speed and accuracy conditions, not task difficulty.

4.1. MODULATION OF THE ENCODING OF EVIDENCE

Evidence for the modulation of sensory processing under speed and accuracy conditions (**Figure 5A**) has been shown in a visual search task, in which monkeys were rewarded for making a saccade to a target stimulus, while single-cell activity was recorded from FEF (Heitz and Schall, 2012). A substantial body of electrophysiological data from visual decision tasks indicates that FEF neurons can be classified as visual neurons and movement neurons (Cohen et al., 2010; Purcell et al., 2010). Visual neurons are responsive to task-relevant stimuli, but do not show saccade-related activity, whereas movement neurons show saccade-related activity, but do not respond to stimuli. As such, movement neurons are hypothesized to integrate the evidence encoded by visual neurons (the first and second stages of **Figure 4**), loosely analogous to the hypothesis that LIP neurons integrate the activity of MT neurons in RDM tasks (Section 3.1). In the study by (Heitz and Schall, 2012), the SAT was correlated with multiple adjustments to the activity of both classes of neuron, including the baseline rate of visual neurons (**Figure 6A**), the magnitude

FIGURE 5 | Hypotheses on the neural implementation of the SAT, under the framework of bounded integration. Each panel is an instance of the 3-stage schematic in **Figure 4**. **(A)** Modulation of the encoding of evidence. A cognitive signal adjusts the gain of sensory encoding populations (dashed arcs). This multiplicative effect is depicted by the "X" in the open circle at the top. **(B)** Modulation of the rate of integration of encoded evidence (dashed arcs). The cognitive signal adjusts the gain of integrator circuitry, controlling the rate of integration. **(C)** Modulation of the onset of integration of encoded evidence. An inhibitory gate (G) controls the onset of integration (dotted arcs). **(D)** Modulation of the sensitivity of integrator circuitry to encoded evidence. Integrator populations are selective for different sub-populations of evidence-encoding neurons under speed and accuracy conditions, depicted by the black (speed) and gray (accuracy) arcs. **(E)** Modulation of the amount of non-evidence input to integrator circuitry. All integrator populations receive a uniform cognitive signal, in addition (+) to evidence (dotted arcs). **(F)** Modulation of the amount of non-integrator input to thresholding circuitry. Neural populations enacting choice behavior receive a uniform cognitive signal, in addition (+) to integrated evidence (dotted arcs). **(G)** Modulation of the connectivity between integrator circuitry and thresholding circuitry. The amount of integrated evidence sufficient to make a choice is modulated by the strength of connectivity from integrators to the circuitry enacting choice behavior (thick horizontal arrows).

of their response to stimuli (**Figure 6B**) and the time at which target-in activity can be discriminated from target-out activity (**Figure 6B**). To summarize, the search array was identical across conditions, but the baseline rates and response magnitude of visual neurons were higher, and the time of discrimination

was earlier, under the speed condition, in which the monkeys made faster, less accurate decisions. Conversely, baseline rates and response magnitude were lower, and discrimination was later, under the accuracy condition, in which the monkeys made slower, more accurate decisions.

These data provide strong support for the hypothesis that the modulation of the encoding of evidence contributes to the SAT, but the data alone do not explain the underlying neural mechanism. Gain modulation provides an explanation. The baseline rates of target-in *and* target-out visual neurons were higher (lower) under speed (accuracy) conditions (solid and dashed curves before stimulus onset in **Figure 6B**), suggesting that visual neurons received a common signal, regardless of whether they were encoding evidence for the target or a distractor. Spatially non-selective (global, uniform, diffuse) excitation is an established form of gain modulation in attractor models (Salinas and Abbott, 1996; Furman and Wang, 2008; Standage et al., 2013), so a stronger (weaker) common signal under speed (accuracy) conditions would account for the higher (lower) response magnitude of visual neurons. If the SNR of encoded evidence were unaffected (or lowered) by this signal, then other things being equal, higher-rate activity by visual neurons under the speed condition would be manifest in a higher rate of integration of this activity by movement neurons, supporting fewer sequential samples and therefore improved speed at the expense of accuracy (*vice versa* for the accuracy condition). This scenario is equivalent to adjusting a decision bound. Consistent with this possibility, the rate of rise of movement-neuron activity was higher (lower) under the speed (accuracy) condition in the study by Heitz and Schall (2012). In the next section, we provide another, compatible explanation of these movement-neuron data.

4.2. MODULATION OF THE INTEGRATION OF ENCODED EVIDENCE

Mechanistic hypotheses on the trading of speed and accuracy by modulation of the integration of evidence can be grouped into three mutually compatible categories: modulation of the rate of integration (**Figure 5B**), modulation of the onset of integration (**Figure 5C**) and modulation of the sensitivity to the encoding of evidence (**Figure 5D**). As noted above, the first hypothesis does not refer to changes in the rate of integration resulting solely from changes in the evidence or its encoding. Rather, we refer to mechanisms hypothesized to actively target integrator circuitry in this section, regardless of upstream or downstream modulation.

4.2.1. Modulation of the rate of integration of evidence

The study by Heitz and Schall (2012) not only provides evidence for the differential modulation of sensory encoding with speed and accuracy conditions, but also for the modulation of the rate of integration of evidence (**Figure 5B**). In their study, the slope of pre-saccadic activity by movement neurons in FEF was shown to increase and decrease under speed and accuracy conditions respectively (**Figure 6C**). As noted in Section 4.1, these changes could simply result from the increase (decrease) in gain of visual neurons under speed (accuracy) conditions; however, they can be explained by the modulation of local-circuit (recurrent) dynamics (**Figure 3**), independent of upstream changes. Increasing the strength of recurrent dynamics shortens
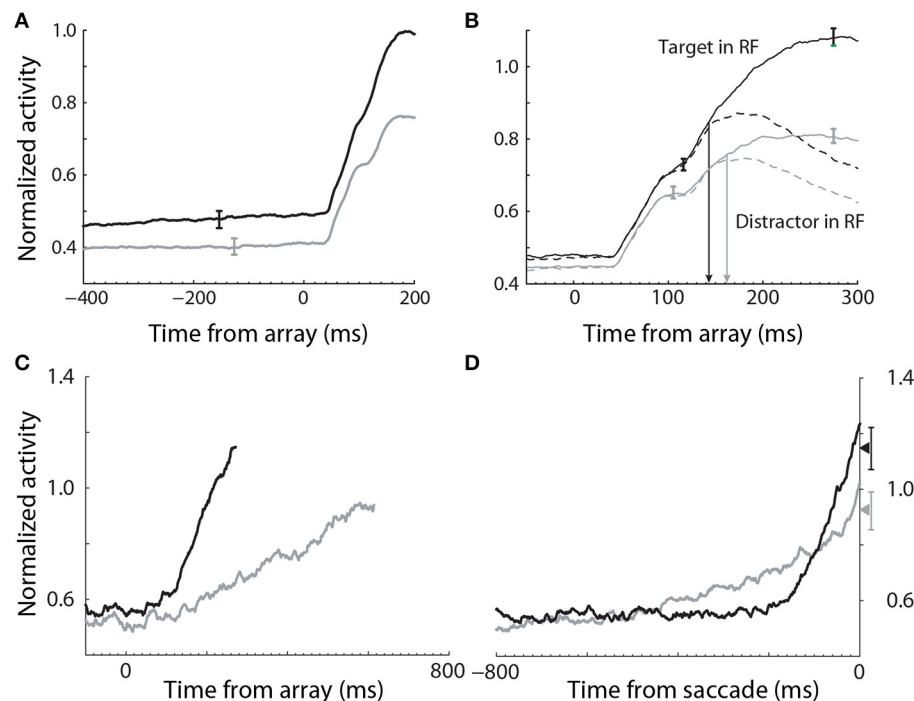
FIGURE 6 | Electrophysiological data recorded from FEF during a visual search task under speed (black) and accuracy (gray) conditions (Heitz and Schall, 2012). (A) The baseline rate of visual (evidence-encoding) neurons was higher and lower under speed and accuracy conditions respectively (stimulus onset at 0 ms). **(B)** The gain of visual neurons was higher (lower) under speed (accuracy) conditions. **(C)** The slope of movement neuron activity was higher (lower) under speed (accuracy) conditions. **(D)** The peak of movement neuron activity was higher (lower) under speed (accuracy) conditions. Data aligned to saccade initiation. Figure adapted from Heitz and Schall (2012) with permission of © Elsevier.

the effective time constant of local-circuit models (Wong and Wang, 2006; Standage et al., 2011), so the decision variable builds up more quickly, limiting the amount of integrated evidence. Decisions are consequently faster and less accurate. Conversely, decreasing the strength of recurrent dynamics lengthens the effective time constant, so the decision variable builds up more slowly and decisions are slower and more accurate. Here, it is worth noting that the computational role of the effective time constant is identical to that of the bound, operating at a different level of abstraction; it controls the duration of the integration of evidence. Thus, while it is intuitive to interpret the bound in terms of the firing rates of integrator neurons, the bound may be implemented by any mechanism that controls integration time.

Lengthening and shortening the effective time constant of a decision circuit offers a sound principle for trading speed and accuracy, but it requires a mechanism (or mechanisms) to increase and decrease the strength of recurrent dynamics under speed and accuracy conditions respectively. There are several possibilities, such as spatially non-selective excitation of excitatory neurons (Furman and Wang, 2008; Standage et al., 2013) or the conductance strength of excitatory recurrent synapses (Wong and Wang, 2006; Standage and Pare, 2011). Furman and Wang (2008) used the first of these mechanisms in simulations of an RDM task with a biophysically-based local-circuit model. They simulated the experiments by Churchland et al. (2008), who recorded from LIP neurons while monkeys chose between two or four possible

directions of motion. Not only did Furman and Wang (2008) qualitatively reproduce neural and behavioral data from the task, but they further considered the effects of speed and accuracy emphasis that were not tested experimentally. They hypothesized that the SAT is controlled by a stationary "top-down" signal, testing their hypothesis by providing non-selective spike trains to all pyramidal neurons in the network, in addition to the selective spike trains simulating motion evidence from area MT. Stronger non-selective input produced faster, less accurate decisions in the model. Furman and Wang (2008) did not show network activity under the different non-selective input rates, but it is clear from other modeling studies that the slope of network activity is higher (lower) with stronger (weaker) recurrent dynamics, corresponding to speed (accuracy) emphasis (e.g., Wong and Wang, 2006; Standage and Pare, 2011). Notably, the baseline rates of target-in *and* target-out movement neurons in the electrophysiological study by Heitz and Schall (2012) were higher (lower) under speed (accuracy) conditions, consistent with the modulation of local-circuit dynamics by a spatially non-selective signal. Note that such a signal is consistent with the use of the term "urgency" in some studies, i.e., speed (accuracy) conditions entail a higher (lower) urgency to respond (Reddi and Carpenter, 2000), though we restrict our usage of this term to time-dependent signals below, i.e., the urgency to respond increases with the duration of a decision (Churchland et al., 2008; Cisek et al., 2009; Standage et al., 2011).

Where Furman and Wang (2008) used a stationary signal to differentially modulate decision dynamics under speed and accuracy conditions, Standage et al. (2011) used a timing (urgency) signal, hypothesizing that an estimate of one's temporal constraints is sufficient to trade speed and accuracy with a fixed level of integrator activity at decision time. They used a model from the same family as that of Furman and Wang (2008), but they took a more abstract *population rate* approach, where a "transfer function" determines the proportion of an idealized neural population activated by its input (Wilson and Cowan, 1972; Gerstner, 2000). The timing signal was an increasing function of time, building up more quickly with tighter temporal constraints, but reaching a fixed maximum (see Durstewitz, 2004). The signal scaled the slope parameter of the transfer function, which in turn controlled the dynamics of the network (the higher the slope parameter, the stronger the dynamics). As such, network dynamics were weak at the start of each trial, but were strengthened with elapsed time. This progression lengthened the time constant of the network prior to entry into the decision regime, and then shortened it (**Figure 7B**). Decision-selective firing rates were fixed at decision time because the network always progressed through the same dynamic regimes, but slower buildup of the timing signal allowed the network to spend more time in regimes with a longer time constant. Thus, the slope of integrator activity was lower (higher) with longer (shorter) temporal constraints, and decisions were slower (faster) and more (less) accurate (**Figures 7C,D**). Standage et al. (2011) compared this approach to the modulation of the network by a stationary signal, showing that time-dependent modulation systematically earned more reward per unit time. In effect, time-dependent modulation of attractor dynamics makes a better use of time than stationary modulation, but human and non-human animals do not necessarily make decisions this way. The model makes testable predictions for experiments, which are an important next step for this hypothesis (see the Discussion).

What neural mechanisms could implement stationary (Furman and Wang, 2008; Roxin and Ledberg, 2008) and time-dependent (Standage et al., 2011, 2013) top-down signals for controlling the speed and accuracy of decisions? A stationary signal could be provided by persistent, goal-directed activity, for which there is abundant evidence in prefrontal and parietal cortical areas (see Wang, 2001). This mechanism would require an additional means to control the rate of persistent activity. Like integration time, the rate of persistent activity in local-circuit cortical models can be controlled by the strength of recurrent dynamics (Brunel and Wang, 2001). Thus, any mechanism that modulates recurrent dynamics in the circuitry mediating the control signal would in turn control the strength of non-selective input to downstream integrator circuitry, and thereby the SAT. To switch between speed and accuracy response modes from trial to trial (e.g., Forstmann et al., 2008; Heitz and Schall, 2012), higher and lower rates of persistent activity would need to be associated with the cues for speed and accuracy conditions respectively.

There is also abundant evidence for the encoding of elapsed time by "climbing activity," i.e., activity that peaks at the time of an anticipated event, such as a deadline (see Durstewitz, 2004). Such *prospective coding* (Rainer et al., 1999; Komura et al., 2001)

has been recorded during tasks with a timing requirement in a number of cortical areas (Niki and Watanabe, 1979; Rainer et al., 1999; Maimon and Assad, 2006; Shuler and Bear, 2006). Standage et al. (2013) built on their earlier population rate model (Standage et al., 2011) with a biophysically-based, coupled-circuit cortical model, offering a neural implementation of the timing signal, and demonstrating its modulation of downstream decision dynamics by spatially non-selective excitation. To switch between speed and accuracy response modes from trial to trial, the shorter and longer timing signals would need to be associated with the cues for speed and accuracy conditions respectively.

It is worth noting that time-dependent attractor models Standage et al. (2011, 2013) are conceptually similar to bounded integrator models in which the bound is lowered over the course of each trial (Ditterich, 2006b; Drugowitsch et al., 2012), but the former cannot be considered a neural implementation of the latter. The underlying premise of the latter is that longer processing time implies a more difficult decision and therefore a lower probability of a correct response. Lowering the bound reduces time-wasting because it speeds up decisions that are more likely to be wrong, increasing reward rate. This approach is functionally equivalent to the time-dependent multiplication of incoming evidence (Ditterich, 2006b). Expressed as a bounded integrator model, the time-dependent attractor models by Standage et al. (2011, 2013) implement the time-dependent multiplication of evidence *and* the evolving decision variable, making different predictions about the sensitivity of decisions to the timing of evidence than other bounded integrator models (see Section 5.1).

4.2.2. Modulation of the onset of integration

It is possible that speed and accuracy conditions modulate the onset of evidence integration (**Figure 5C**), as opposed to (or in addition to) the rate of integration. Purcell et al. (2012) tested this hypothesis with a leaky competing accumulator model, in which the accumulators received the activity of visually-responsive neurons in FEF, recorded during a visual search task. The accumulator corresponding to the target received the activity of target-in neurons, while the other accumulators received the activity of target-out neurons. Each accumulator received a fixed inhibitory signal serving as a gate, preventing the accumulation of activity prior to the search array, that is, the gate dictated that evidence was only accumulated if it exceeded a minimum rate. The model was fit to behavioral data from monkeys performing the search task and to electrophysiological recordings from FEF movement neurons. In simulations of an SAT experiment, adjustments to the inhibitory gate were compared to adjustments to the bound. Both parameters accounted for the SAT and maximized reward rate, but they made different predictions about the activity of movement neurons. As expected, adjustments to the bound predicted a higher (lower) rate of activity at the time of commitment to a choice under accuracy (speed) conditions, but did not impact baseline activity or the onset of integration. Adjustments to the inhibitory gate predicted higher (lower) baseline activity and earlier (later) onset of integration under speed (accuracy) conditions. To the best of our knowledge, the activity of FEF movement neurons in the study by Heitz and Schall (2012) provide the only available single-cell data to test these predictions. These data do
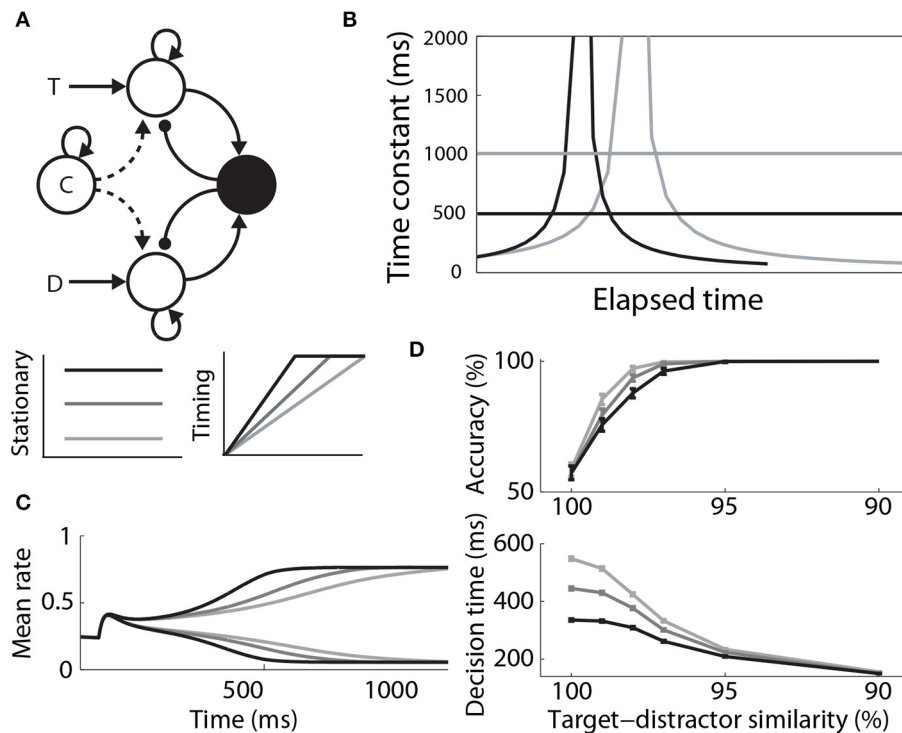
FIGURE 7 | (A) A control signal (c) modulating local-circuit decision processing. The signal could be implemented by persistent, goal directed activity (Stationary) or by climbing activity, encoding elapsed time relative to a deadline (Timing). The model schematic is the same as in **Figure 3A**, with the addition of the control signal. **(B)** Under stationary modulation, the decision network can support a single time constant of integration for a given trial, depicted by the horizontal lines. A stronger control signal furnishes a shorter time constant. Under time-dependent modulation, the time constant of integration increases in the leakage regime, before contracting in the decision regime. This progression occurs more quickly with faster buildup of the signal. Black and gray curves correspond to speed and accuracy conditions. **(C)** Target and distractor-selective firing rates in a simulated decision circuit for each timing signal. The slope of decision-selective activity is higher for shorter timing signals. **(D)** Psychometric and chronometric curves corresponding to each timing signal. **(B–D)** are adapted from Standage et al. (2011).

not support the predictions of the bound parameter. Not only do they show differential baseline activity under speed and accuracy conditions, but they also show a higher rate of activity at choice time under speed conditions (**Figure 6D**), i.e., opposite to the predicted activity. These data support the predictions for baseline activity by the gate parameter, i.e., higher baseline under speed conditions, but they do not support the prediction of differential onset of integration. Several fMRI studies with human subjects also show differential baseline activity under speed and accuracy conditions in pre-motor cortical areas (Forstmann et al., 2008; Ivanoff et al., 2008; van Maanen et al., 2011) (Section 4.3.2).

4.2.3. Modulation of the sensitivity to encoded evidence
Support for the hypothesis that integrator circuitry is more (less) sensitive to the encoding of evidence under accuracy (speed) conditions (**Figure 5D**) has been provided by a visual discrimination task, in which human subjects decided whether flashing stimuli were of the same or slightly different orientation (Ho et al., 2012). As expected, decisions were slower and more accurate under the accuracy condition (*vice versa* for speed). Because the neural mechanisms underlying fine discrimination of orientation are well-studied, these authors focused on trials on which the stimuli differed (mismatch trials). In particular, off-target

neurons (tuned away from the stimulus) are hypothesized to be more informative for fine discrimination than on-target neurons (tuned toward the stimulus), due to the steeper slope of their tuning curves at off-target orientations (see Scolari and Serences, 2012). This computational principle is depicted in **Figure 8**. In the study by Ho et al. (2012), there was no difference between blood oxygenation level dependent (BOLD) based orientation tuning curves in primary visual cortex (V1) under speed and accuracy conditions, suggesting that these conditions did not modulate the encoding of evidence on mismatch trials. However, off-target activation (tuned away from the target orientation) was higher on correct trials than error trials under the accuracy condition, that is, subjects were more accurate when off-target activation was higher. This finding suggests that subjects were more accurate when the gain of off-target neurons was higher, which further suggests that accuracy was higher because integrator populations detected this higher gain. Conversely, BOLD-based tuning curves did not differ on correct and error trials under the speed condition, suggesting that integrator populations did not detect fluctuations in the gain of off-target neurons (or on-target neurons). Taken together, the speed and accuracy data suggest that integrator populations are more sensitive to (more informative) off-target activity under
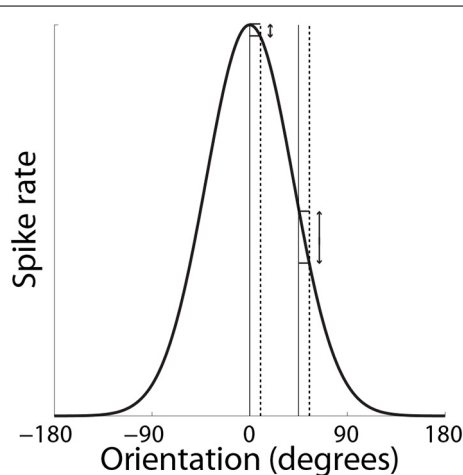
FIGURE 8 | SAT mechanism hypothesized by Ho et al. (2012). Small changes to a stimulus feature do not elicit much change in the response by neurons that are highly selective for the feature (on-target neurons). Here, the feature is orientation. The solid and dashed vertical lines on the left correspond to feature values of 0° and slightly greater than 0° respectively. The change in response by a neuron maximally responsive to 0° is shown by the corresponding horizontal lines abutting the black curve. The solid and dashed vertical lines on the right correspond to feature values of 45° and a slight increase from 45° respectively. The change in response by the same neuron (maximally responsive to 0°) is shown by the corresponding horizontal lines abutting the black curve. For a given change in feature value, the difference in the off-target response is greater than the difference in the on-target response.

accuracy conditions, resulting in higher accuracy at a cost in terms of speed. Under speed conditions, lower sensitivity to off-target activity would appear to support faster decisions, at a cost in terms of accuracy.

Ho et al. (2012) did not speculate on the mechanism by which speed (accuracy) conditions may engender lower (higher) sensitivity to more informative neurons, but it is plausible that speed conditions lower the SNR of the activity projecting to integrator circuitry, such that the fine discrimination provided by off-target activity is swallowed by noise. The lower firing rate of off-target activity (see **Figure 8**) is consistent with this possibility. Another possibility is that integrator circuitry is not differentially sensitive to off-target activity *per se*, but is preferentially *selective* for on-target and off-target neurons under speed and accuracy conditions respectively. If so, lower-rate, more informative off-target activity would take longer to accumulate to a given firing rate than higher-rate, less informative on-target activity, accounting for the SAT. Our description of this possibility does not explain how preferential selectivity would arise, but is consistent with the higher (lower) rate of rise of movement-neuron activity under speed (accuracy) conditions shown by Heitz and Schall (2012).

4.3. MODULATION OF THE AMOUNT OF INTEGRATED EVIDENCE SUFFICIENT TO MAKE A CHOICE

The hypothesis that speed and accuracy are traded by the modulation of the amount of integrated evidence has received the

lion's share of attention in mechanistic studies of the SAT, presumably because bounded integrator models are readily fit to behavioral data by adjusting the bound (see Bogacz et al., 2010b). Under the assumption of linear integration, changing the starting point is algorithmically equivalent to changing the bound. Under a neural instantiation of these terms, changes to the starting point would be manifest in changes to the baseline activity of integrator neurons, while changes to the bound would be manifest in the firing rate of integrator neurons at the time of commitment to a choice. Here, it is important to distinguish between the amount of integrated *evidence* and a neural decision variable. A decision variable may have sources of input other than the evidence (Kable and Glimcher, 2009; Doya and Shadlen, 2012), e.g., the encoding of the prior probabilities of the alternatives. Under this approach, mechanistic hypotheses on the modulation of the amount of integrated evidence sufficient to make a choice can immediately be grouped into two categories: changes to non-evidence inputs to integrator circuitry (**Figure 5E**), and changes to non-integrator inputs to thresholding circuitry (**Figure 5F**). The former tend to be limited to cortical circuitry, whereas the latter often involve cortex and the basal ganglia (BG). We also consider a third category in this section: changes to the connectivity mediating integrator inputs to thresholding circuitry (**Figure 5G**). This category is distinct from the modulation of integrated evidence described above (Section 4.2), since no mechanistic change to the integration process is entailed by changes to downstream connectivity. Note that these three general, mechanistic categories share the assumption that a fixed net input current to thresholding circuitry is required to elicit choice behavior.

4.3.1. Adjustments to non-evidence inputs to integrator circuitry

Several theoretical studies have proposed neural mechanisms for the SAT that involve differential levels of non-evidence inputs to integrator circuitry under speed and accuracy conditions (Furman and Wang, 2008; Roxin and Ledberg, 2008; Standage et al., 2013) (**Figure 5E**). A large body of electrophysiological data provides evidence for integrator activity in frontal (Kim and Shadlen, 1999; Schall et al., 2011; Ding and Gold, 2012) and parietal (Roitman and Shadlen, 2002; Thomas and Pare, 2007; Bollimunta and Ditterich, 2011) cortical areas during decision tasks (Section 3.1), so these theoretical studies have typically focused on cortical circuitry. Furman and Wang (2008) controlled the SAT by providing input spike trains to all pyramidal neurons in their biophysically-based cortical model, in addition to the selective spike trains for each of the decision alternatives. We presented this model in Section 4.2.1 because spatially non-selective input modulates the dynamics of local-circuit decision models, changing the rate of integration. However, the model does implement an adjustment to the amount of non-evidence input to integrator circuitry, albeit a small one.

The hypothesis that persistent activity controls the SAT by projecting non-selectively to integrator populations (Furman and Wang, 2008; Roxin and Ledberg, 2008) is consistent with fMRI data from a Simon task (van Veen et al., 2008), in which human subjects responded to the color of a stimulus to the left or right of fixation, while ignoring its location. This study showed an
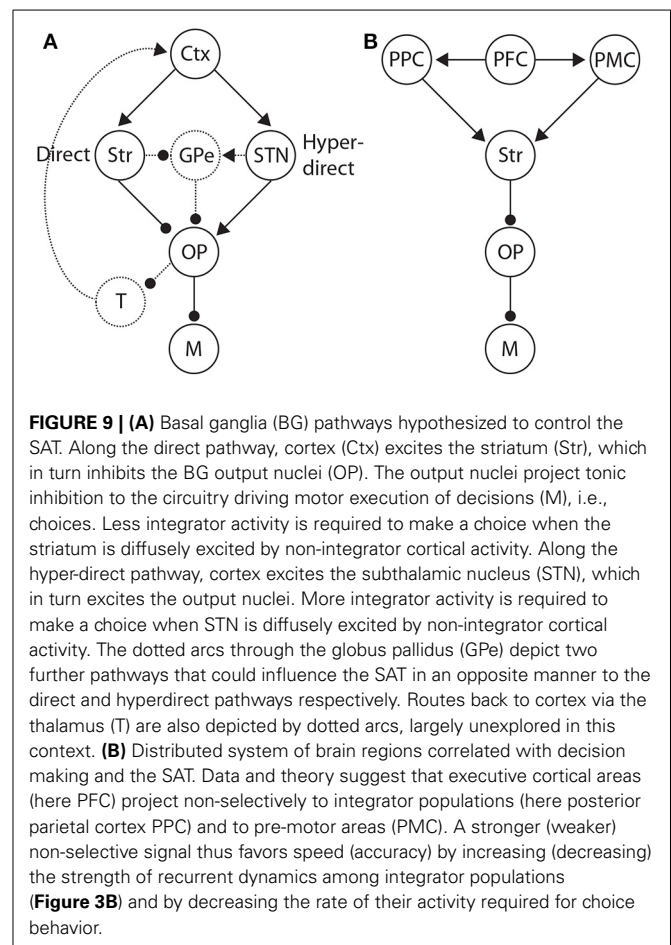
increased baseline (sustained) BOLD response in dlPFC under speed conditions relative to accuracy conditions, and an increased transient (associated with the decision process) BOLD response in the intraparietal lobule, a parietal area that may correspond to LIP in monkeys. As noted above, persistent activity has been recorded from dlPFC in studies of working memory (Fuster, 1973; Funahashi et al., 1989) and decision-correlated activity has been recorded from LIP in decision tasks (Roitman and Shadlen, 2002; Thomas and Pare, 2007), so it is plausible that dlPFC projects a stronger (weaker) control signal to integrator neurons in the intraparietal lobule under speed (accuracy) conditions, controlling the speed and accuracy of decisions. This possibility is consistent with increased (decreased) baseline activity by putative integrator neurons under speed (accuracy) conditions in the study by Heitz and Schall (2012) (Section 4.2.1), as well as with the modulation of the rate of integration by a stationary, non-selective signal (Furman and Wang, 2008).

The SAT is also controlled by non-selective excitation of integrator circuitry in the model by Standage et al. (2013). As described in Section 4.2.1, the major difference between this neural model and the one by Furman and Wang (2008) is the information content of the non-evidence input. In the model by Standage et al. (2013), the non-evidence input is an estimate of elapsed time relative to a deadline, implemented by the destabilization of background activity by strong recurrent dynamics. Like the model by Furman and Wang (2008), this model controls the SAT by modulation of the rate of integration (Section 4.2.1), but nonetheless, it does implement a time-dependent, uniform input to integrators. This input builds up more (less) rapidly under speed (accuracy) conditions.

4.3.2. Adjustments to non-integrator inputs to thresholding circuitry

A number of mechanistic hypotheses on the SAT are based on the premise that the amount of integrated evidence sufficient to make a choice is controlled by spatially non-selective input to thresholding circuitry (Frank, 2006; Simen et al., 2006; Forstmann et al., 2010; Green et al., 2012) (**Figure 5F**). According to this premise, stronger non-selective input allows lower levels of integrator activity to elicit a choice. The hypotheses differ according the processing pathways providing the non-selective inputs, and in the corresponding information content provided by these signals. Many of these hypotheses involve BG, owing to its well-established role in movement initiation (choice behavior in the present context). Excitatory input to BG arrives at the striatum, which inhibits the output nuclei along the so-called direct pathway. The output nuclei inhibit motor circuitry in their tonic (background, default) state, so excitation of the striatum releases motor circuitry from inhibition, enabling choice behavior (See **Figure 9A**).

It has been proposed that an estimate of reward rate could provide spatially non-selective input to thresholding circuitry, computed by leaky integration of reward signals (Simen et al., 2006). Such a mechanism could approximate the optimal trade-off between speed and accuracy in terms of reward-rate maximization, without speed or accuracy instructions (Simen et al., 2006). In effect, the strength of non-selective input tracks reward

FIGURE 9 | (A) Basal ganglia (BG) pathways hypothesized to control the SAT. Along the direct pathway, cortex (Ctx) excites the striatum (Str), which in turn inhibits the BG output nuclei (OP). The output nuclei project tonic inhibition to the circuitry driving motor execution of decisions (M), i.e., choices. Less integrator activity is required to make a choice when the striatum is diffusely excited by non-integrator cortical activity. Along the hyper-direct pathway, cortex excites the subthalamic nucleus (STN), which in turn excites the output nuclei. More integrator activity is required to make a choice when STN is diffusely excited by non-integrator cortical activity. The dotted arcs through the globus pallidus (GPe) depict two further pathways that could influence the SAT in an opposite manner to the direct and hyperdirect pathways respectively. Routes back to cortex via the thalamus (T) are also depicted by dotted arcs, largely unexplored in this context. **(B)** Distributed system of brain regions correlated with decision making and the SAT. Data and theory suggest that executive cortical areas (here PFC) project non-selectively to integrator populations (here posterior parietal cortex PPC) and to pre-motor areas (PMC). A stronger (weaker) non-selective signal thus favors speed (accuracy) by increasing (decreasing) the strength of recurrent dynamics among integrator populations (**Figure 3B**) and by decreasing the rate of their activity required for choice behavior.

rate under this mechanism. It is plausible that such a non-selective signal could be implemented in PFC by the increased occupancy of D1 dopamine receptors, due to slow extrasynaptic uptake (Grace, 1991; Dreyer et al., 2010). The activity of dopamine (DA) neurons in BG is extensively correlated with reward and these neurons project diffusely to PFC (and other association cortical areas), where D1 receptors are hypothesized to control attractor dynamics in support of persistent, goal-directed activity (see Durstewitz and Seamans, 2006). It is therefore possible that the rate of persistent activity in PFC could provide a reward estimate to BG, which gates choice behavior. It is not clear how such a reward-rate signal would adapt to the imposition of speed or accuracy conditions on cue, i.e., the proposed mechanism extracts an appropriate strength of signal for a given condition, but would presumably require an additional mechanism to switch between speed and accuracy modes from trial to trial.

Timing signals are another potential source of non-selective input to thresholding circuitry. Under this hypothesis, the SAT is controlled by the balance between selective input from integrator populations and non-selective input from neural populations encoding elapsed time. In the study by Green et al. (2012), human subjects performed an RDM task under reward schedules corresponding to speed and accuracy conditions. Subjects' behavior was fit by a bounded integrator model, where adjustments

to the bound were correlated with reward rate on an individual subject basis, i.e., subjects whose behavior was captured by larger adjustments to the bound earned more reward. Because a higher (lower) bound supports more (less) integration, this correlation suggests that subjects traded speed accuracy by controlling the amount of integrated evidence sufficient to make a choice. Using fMRI, these authors showed higher activation in dlPFC under the accuracy condition, and higher activation in the cerebellum under the speed condition. They further considered correlations between activation in each of these regions and that in the striatum (the effective connectivity). Note that the striatum is hypothesized to control response thresholds and thus choice behavior (see below). The effective connectivity between dlPFC and the striatum was higher under the accuracy condition and was positively correlated with the difference (high-low) between the value of the bound parameter under the two conditions. The effective connectivity between the cerebellum and the striatum was higher under the speed condition and was negatively correlated with this difference. Striatal activation did not differ between conditions, consistent with a fixed threshold. Because earlier studies have provided evidence for integrator activity in dlPFC during decisions (Kim and Shadlen, 1999; Heekeren et al., 2004; Philiastides et al., 2011) and for sub-second timing in the cerebellum (see Lewis and Miall, 2003; Ivry and Spencer, 2004), it was hypothesized that persistent changes in connectivity mediate response modes for the purpose of maximizing reward. Thus, the balance between cortico-striatal and cerebellar-striatal processing could control the SAT. This study switched speed and accuracy conditions between blocks, but each block contained very few trials (approximately 10). Subjects therefore adapted quickly to task conditions, suggesting that the underlying mechanism may be capable of switching from trial to trial on cue.

The study by Green et al. (2012) is not the only MRI study to implicate the striatum in the SAT. In the study by Forstmann et al. (2008), the BOLD signal in the pre-supplementary motor area (pre-SMA) and the striatum was stronger in response to a pre-trial cue indicating speed conditions in an RDM task, compared to accuracy or neutral conditions. When individual subject's behavioral data were fit by a bounded accumulator model, the magnitude of adjustments to the bound were positively correlated with the BOLD signal in these areas, i.e., subjects whose behavior was captured by larger adjustments showed greater activation in pre-SMA and striatum. The strength of connectivity between pre-SMA and striatum has also been correlated with individual subjects' adjustments to the bound in an RDM task, i.e., subjects whose behavior was captured by larger adjustments showed greater connectivity between these areas, as determined by structural MRI (sMRI) (Forstmann et al., 2010).

In the study by Ivanoff et al. (2008), human subjects performed an RDM task with growing motion coherence under speed and accuracy conditions. These authors classified their results according to "baseline trials" and "coherence trials," where the coherence of moving dots was 0% (over a full trial) and greater than 0% respectively. The underlying premise of this classification is that baseline trials did not provide evidence for integration, but rather,

provided only noise; whereas coherence trials provided evidence *and* noise. The BOLD signal in pre-SMA and posterior lateral prefrontal cortex (plPFC) was higher on baseline trials under the speed condition, and was higher on coherence trials under the accuracy condition. Furthermore, the difference in activation under speed and accuracy conditions on baseline trials was equal and opposite to that on coherence trials across subjects, i.e., the speed-minus-accuracy difference on baseline trials equaled the accuracy-minus-speed difference on coherence trials. These data suggest that baseline activity in these cortical regions determines the amount of integrated evidence sufficient to make a choice. In other words, the integrated evidence on coherence trials may account for the difference in activation between speed and accuracy conditions. If so, this equal, opposite difference should be found on a within-subject basis. It was found in pre-SMA, but not in plPFC.

Ivanoff et al. (2008) further showed that on coherence trials, a measure of subjects' decision criteria [the criterion metric of signal detection theory (Macmillan and Creelman, 1991)] was correlated with the BOLD signal in plPFC, but not in pre-SMA. This finding suggests that speed and accuracy conditions modulate the amount of evidence integrated by plPFC. These authors sub-classified their coherence trials according to the level of coherence at the time of subjects' decisions, defining "hits" and "false alarms" as trials on which coherence was positive and 0% at decision time respectively. The BOLD signal in pre-SMA was equal in both classes of trial. Under the assumption that brain regions supporting the integration of evidence should show greater activity on hits than false alarms (because there is evidence to integrate), these data support the hypothesis that evidence is not integrated in pre-SMA. Conversely, activation in plPFC was greater on hits than false alarms, suggesting that plPFC supports integration in the task. Overall, the study by Ivanoff et al. (2008) supports the hypothesis that pre-SMA plays an "adaptive baseline" role in the SAT, determining the amount of evidence integrated in cortical areas such as plPFC. Taken together, the studies by Forstmann et al. (2008), Forstmann et al. (2010), and Ivanoff et al. (2008) suggest that pre-SMA projects non-selectively to the striatum, where this activity is added to selective inputs from cortical integrator populations.

The above studies were extended by van Maanen et al. (2011), who considered the mechanisms by which subjects switch between response modes for speed and accuracy. Under speed conditions, trial-to-trial changes in the BOLD signal in pre-SMA were positively correlated with estimates of the starting point of accumulation in a single-trial version of a bounded accumulator model, in which the bound was fixed. In this case, a higher starting point has the same effect as a lower bound, i.e., faster, less accurate decisions. These data further support the hypothesis that pre-SMA provides a non-selective control signal to the striatum, governing the SAT. On trials that imposed a switch between speed and accuracy conditions (in either direction), a positive correlation was also found between BOLD changes in the anterior cingulate cortex (ACC) and the starting point. Interestingly, only switches from accuracy to speed were correlated with activation of the striatum, suggesting that switching between response modes may be asymmetric, i.e., different mechanisms may

mediate switching from a speed mode to an accuracy mode than *vice versa*.

The study by van Maanen et al. (2011) further showed that under accuracy conditions, BOLD changes in ACC were positively correlated with changes in the starting point in their model, but only on trials following an error. These data suggest that ACC may contribute to an emphasis on accuracy, consistent with a neural model of cortico-BG circuitry in which cortical conflict detection excites the subthalamic nucleus (STN) (Frank, 2006; Frank et al., 2007). Note that ACC is believed to play a role in conflict monitoring (Yeung et al., 2004). The model is based on earlier neural models of action selection (Gurney et al., 2001), in which rewards are associated with salient stimuli. In the model by Frank (2006), conflict arises when multiple rewarding (or unrewarding) stimuli occur simultaneously. Cortex detects this "conflict" and projects to STN, which in turn prevents action selection by inhibiting motor circuitry. The model thereby implements dynamic threshold adaptation, increasing the amount of evidence sufficient to make a choice during difficult decisions.

The underlying premises of this "STN hypothesis" (Bogacz et al., 2010b) are further supported by studies of response inhibition in "stop-signal" tasks, in which subjects are cued to withhold planned responses on a proportion of trials (Stop trials). The "direct" and "hyperdirect" pathways have been correlated with Go trials (without the stopping cue) and Stop trials respectively (Aron and Poldrack, 2006), suggesting that activation of the striatum speeds up responding and activation of STN slows it down. These data therefore suggest that speed and accuracy conditions may preferentially activate the direct and hyperdirect pathways respectively (**Figure 9A**). As described above, speed conditions have been correlated with fronto-striatal circuitry in a number of neuroimaging studies of the SAT (Forstmann et al., 2008; Ivanoff et al., 2008; Forstmann et al., 2010; van Maanen et al., 2011). However, we are unaware of any study to show a positive correlation between STN (activity or connectivity) and accuracy conditions, or a negative correlation between STN and speed conditions. The small size of STN may be a factor in this regard. The present neuroimaging data can therefore be considered to support the notion that accuracy conditions correspond to a "default" mode of decision making, modulated by speed conditions (van Veen et al., 2008; van Maanen et al., 2011). If so, switching between speed and accuracy response modes from trial to trial would only need involve fronto-striatal circuitry, as described above (Forstmann et al., 2008, 2010; Ivanoff et al., 2008). The fMRI data by van Maanen et al. (2011) suggest a more complex state of affairs, but it seems plausible that under this "striatal hypothesis" (Bogacz et al., 2010b), some baseline level of fronto-striatal activation corresponds to a default mode, where speed and accuracy conditions increase and decrease activation respectively.

4.3.3. Adjustments to the connectivity between integrators and thresholding circuitry

The hypothesis that the SAT is supported by adjustments to the connectivity between integrators and thresholding circuitry (**Figure 5G**) has been implemented in a biophysically-based, coupled-circuit model of eye-movement decisions (Lo and Wang,

2006). In the model, the integration of evidence occurs in cortex and projects directly to the superior colliculus (SC) by excitatory synaptic connectivity, and indirectly via the striatum and substantia nigra pars reticulata (SNr). Note that SC is extensively correlated with eye-movement decisions (e.g., Dorris and Munoz, 1998; Thevarajah et al., 2009). SC is tonically inhibited by SNr, so the latter pathway is disinhibitory. These authors assumed that the pre-saccadic reduction in tonic SNr activity occurs abruptly, rather than smoothly, so SC burst neurons were inactive in the model until SNr was sufficiently inhibited by the striatum. As such, burst neurons detected threshold-crossing by cortical integrator neurons, and consequently, burst firing was much more sensitive to changes in the conductance strength of cortico-striatal synapses than cortico-SC synapses. By tuning the conductance strength of cortico-striatal synapses between blocks of trials, the model traded speed for accuracy. Stronger (weaker) conductance entailed lower (higher) integrator rates under speed (accuracy) conditions, but for a given conductance strength (a given speed/accuracy condition), integrator rates were fixed across task difficulty (Roitman and Shadlen, 2002; Churchland et al., 2008). Note that the model does not appear suited to the trial-to-trial switching of speed and accuracy modes on cue, owing to the timescales of synaptic plasticity.

5. DISCUSSION AND CONCLUSIONS

Under the framework of bounded integration, there are three general classes of hypothesis on the neural implementation of the SAT: differential modulation of the encoding of evidence under speed and accuracy conditions (**Figure 5A**), differential modulation of the integration of encoded evidence (**Figures 5B–D**), and differential modulation of the amount of integrated evidence sufficient to make a choice (**Figures 5E–G**). The first category has received the least attention, but the recent study by Heitz and Schall (2012) provides strong evidence for the modulation of sensory encoding (Section 4.1).

Hypotheses on the differential modulation of integration under speed and accuracy conditions can be sub-classified according to the rate (Section 4.2.1) and onset (Section 4.2.2) of integration, and the sensitivity of integrator circuitry to the encoding of evidence (Section 4.2.3). There is considerable evidence for the first of these hypotheses. The rate of rise of putative integrator activity has been shown to increase and decrease under speed and accuracy conditions respectively (Heitz and Schall, 2012). This activity can be explained by attractor models (**Figures 3, 7**), in which speed (accuracy) conditions increase (decrease) the rate of the evolution of competitive dynamics. At least three neural models have demonstrated that a cognitive signal could control the SAT in this manner by projecting non-selectively to integrator circuitry, either by persistent mnemonic activity (Furman and Wang, 2008; Roxin and Ledberg, 2008) or by climbing activity encoding elapsed time relative to a deadline (Standage et al., 2013).

Hypotheses on the amount of integrated evidence sufficient to make a choice can be sub-classified according to adjustments to non-evidence inputs to integrator circuitry (Section 4.3.1), adjustments of non-integrator inputs to thresholding circuitry (Section 4.3.2) and adjustments to the connectivity from

integrator circuitry to thresholding circuitry (Section 4.3.3). According to the first of these hypotheses, if choice behavior requires a fixed level of activity by integrator neurons, then more (less) evidence will be required to reach this fixed level if less (more) common input is provided to all integrators. Attractor models suggest that this mechanism may be impossible to disentangle from the modulation of the rate of integration, since an increase in spatially non-selective excitation decreases their effective time constants, i.e., it increases the rate of integration. Spatially non-selective excitation, however, is not necessarily synonymous with a common input to integrators. The former entails a common input to integrator neurons *and* other neurons in the local circuitry not receiving evidence. The latter does not necessarily include these other neurons. We are unaware of any studies to systematically consider the modulation of recurrent dynamics according to this difference, but the dynamics of attractor networks are known to be influenced by the size of integrator populations relative to the number of neurons in these networks (Albantakis and Deco, 2009).

Our description of the role of BG in the adjustment of non-integrator inputs to thresholding circuitry has not considered bidirectional connectivity between cortex and BG via the thalamus, which complicates the interpretation of information flow during decisions (**Figure 9A**). The different spatial profiles of cortico-BG-thalamo-cortical loops further complicate things, since information from different cortical areas may be processed discretely within BG and returned to the areas of origin, may be integrated within BG and returned to all regions of origin, or may be partially integrated (see Nambu, 2011). Further to these complications, there are multiple processing pathways though BG. The direct and hyperdirect pathways are described above, but there is also an "indirect" pathway to the output nuclei, via the external segment of the globus pallidus (GPe, **Figure 9A**). GPe receives inhibitory projections from the striatum and makes inhibitory projections to the output nuclei. The indirect pathway thus "counteracts" the direct pathway, i.e., excitation of the striatum disinhibits motor circuitry along the direct pathway, while effectively inhibiting it via the indirect pathway (dis-disinhibition). Interestingly, STN makes excitatory projections to GPe, so the hyperdirect pathway also has a counteracting pathway, i.e., excitation of STN inhibits motor circuitry, but also disinhibits it via GPe (see Nambu, 2011). Thus, interpreting correlations between SAT behavior and activation of BG input and output nuclei is complicated by the paths this activity may follow, with each path supporting different computations. Extensive discussion of these possibilities is beyond the scope of this review, but assumptions about these and other anatomical factors influence the interpretation of the experimental data presented here.

The possibility of "self-modulation" of decision dynamics (Section 4.3.2) also warrants further comment. The cortico-BG model by Frank (2006) includes a cortical conflict detection area (potentially ACC) that raises the threshold for choice behavior by projecting to STN. Thus, more difficult tasks more strongly activate this area during decisions, raising the threshold. At first glance, this possibility appears to conflict with bounded integrator models in which reward rate is maximized by lowering the

bound during decisions (Ditterich, 2006a; Drugowitsch et al., 2012). As noted in Section 4.2.1, lowering the decision criterion reduces time-wasting because it speeds up decisions that are more likely to be wrong, but this approach may not be ideal under stringent accuracy conditions, e.g., when errors are punished by long timeouts. In this case, raising the criterion could be the better strategy. This discrepancy highlights the potential utility of separate mechanisms for speed and accuracy emphasis: it is not immediately clear how a single neural mechanism could implement the within-trial increase in the bound under accuracy conditions and decrease in the bound under speed conditions.

5.1. PREDICTIONS FOR FUTURE EXPERIMENTS

Different classes of hypothesis on the SAT make different predictions for experimental testing, as do different models within these classes. For instance, the hypothesis that the SAT is controlled by adjustments to non-evidence input to integrator circuitry (Section 4.3.1) makes a different prediction about the rate of integrator activity at the time of commitment to a choice than the hypothesis that the SAT is controlled by adjustments to non-integrator inputs to thresholding circuitry (Section 4.3.2) or adjustments to the connectivity between integrator circuitry and thresholding circuitry. Assuming a fixed current is required for choice selection, adjustments to non-evidence input to integrator circuitry imply the same rate of integrator activity at choice time across task conditions, whereas an increase (decrease) in non-integrator input to thresholding circuitry under speed (accuracy) conditions implies a lower (higher) rate of integrator activity at choice time, as does stronger (weaker) connectivity between these circuits. The only available single-cell data conflict with the latter mechanisms, showing a higher rate of putative integrator activity under speed conditions (Heitz and Schall, 2012). These authors showed that leakage by the circuitry enacting the choice could account for the difference in rate, an explanation that supports the former mechanism.

The conflict between the prediction of lower (higher) integrator rates under speed (accuracy) conditions and electrophysiological data (Heitz and Schall, 2012) raises several points of caution. Firstly, the experimental studies providing evidence for the adjustment of non-integrator inputs to thresholding circuitry employed perceptual tasks in which humans made their choices by manually pressing a button (Forstmann et al., 2008, 2010; Ivanoff et al., 2008; van Veen et al., 2008; Green et al., 2012), whereas the electrophysiological data were recorded during a task in which non-human primates made their choices with an eye-movement. We are comfortable ignoring inter-species differences at this stage of the game, but it is plausible that the pathways from frontal regions to primary motor cortex are qualitatively different in relation to the SAT than those from FEF to eye-movement circuitry (as in Heitz and Schall, 2012). On the other hand, the striatal hypothesis (Section 4.3.2) does not require that non-selective excitation of the striatum be provided by the same cortical area across response modalities. Here, it is worth noting that FEF projects directly to the circuitry mediating eye movements, but also projects to this circuitry along a pathway through the striatum, substantia nigra pars reticulata (SNr)

and SC. Because SNr tonically inhibits SC, the latter pathway potentially provides an eye-movement "version" of the striatal hypothesis described above in the context of manual movements. Suffice to say, it would be informative to run the RDM task used by Forstmann et al. (2008) in an eye-movement paradigm.

Different models that account for the SAT by the modulation of the rate of integration (Section 4.2.1) make different predictions about the weighting of evidence during decisions. Stationary attractor models (Furman and Wang, 2008; Roxin and Ledberg, 2008) predict a *primacy effect* (Wong et al., 2007), i.e., earlier evidence is weighted more heavily than later evidence. In effect, attractor dynamics amplify a decision variable, so earlier evidence is subject to amplification for longer. This prediction by stationary attractor models contrasts with that of bounded integrator models dominated by leakage, which show a *recency effect* because earlier evidence is subject to leakage for longer (see e.g., Usher and McClelland, 2001). In time-dependent attractor models (Standage et al., 2011, 2013), if the dynamics are weak at the start of a trial, then a decision variable is dominated by early leakage and late amplification. As such, the evidence will be most heavily weighted somewhere in the middle (see Standage et al., 2011). The respective predictions of these models could be tested by changing the strength of evidence at different times during decision trials. At least one study has conducted such an experiment, using an RDM task in which the coherence of the dots changed during a brief window at different times (Kiani et al., 2008). These authors found a primacy effect, but they used a fixed-duration task with a flat hazard rate, i.e., subjects responded on cue, but it was impossible to determine when the cue would arrive. It would therefore have been impossible to encode elapsed time relative to the cue. Running the same task with a fully predictable duration would be highly informative.

5.2. EVIDENCE FOR LIMITED INTEGRATION

It is not universally assumed that the neural mechanisms underlying decisions implement the principles of bounded integration as described above (Section 3). In the model by Cisek et al. (2009), momentary evidence is multiplied by elapsed time and a decision is made when the resulting quantity exceeds a decision bound. Because decisions would be susceptible to noise without temporal integration, the authors proposed that noisy evidence is low-passed filtered before being multiplied. A low-pass filter can be thought of as a leaky integrator with a short time constant, so the main difference between this "urgency-gating" model and a leaky integrator with a decreasing bound (see Ditterich, 2006b) is the length of the time constant of integration, i.e., how rapidly the evidence is leaking. Cisek et al. (2009) argued that perceptual decisions in real-world environments are likely to depend on fluctuating evidence, but integrators with long time constants are not well-suited to these conditions. Consistent with these principles, they showed that the urgency-gating model could account for behavioral data from a task with changing evidence, whereas bounded integrator models could not. In effect, the bounded integrators were not leaky enough.

Thura et al. (2012) extended this work by proposing that optimal decisions are supported by the integration of novel information only, where optimality was defined in terms of reward rate.

Formally, their model specifies the perfect integration of *differentiated* evidence, where a decision is made when the running total exceeds a decreasing bound. They showed that this procedure is optimal under the assumption of *non*-independence between sequential samples of evidence, which is likely to be the case in most natural conditions, and they proposed that this optimal procedure can be approximated by the multiplication of low-pass filtered evidence by a growing urgency signal. As such, the model is equivalent to their earlier model (Cisek et al., 2009). Their models explain the SAT in tasks with changing evidence because longer (shorter) intervals provide more (less) opportunity to integrate changes in the evidence (novelty). Under the assumption that response-time variability is primarily the result of between-trial variability in attention, arousal and related factors (Carpenter and Williams, 1995), their models further account for behavioral data from traditional tasks with fixed (within-trial) mean evidence, and they account for decision-correlated buildup activity (Section 3.1) under the assumption that this activity mainly reflects the urgency to respond.

In proposing a neural approximation of the urgency-gating model, these authors suggested that the timescale of (leaky) integration is on the order of 100 ms, consistent with evidence that perceptual decisions are based on information from a time window on this order (see Thura et al., 2012), but difficult to reconcile with the SAT on timescales of many hundreds of milliseconds. For example, in the random dot motion task by Palmer et al. (2005), accuracy was lower (higher) and decision times were shorter (longer) under a speed (accuracy) condition, where response times were as long as around 500 ms (2 s). Since the only novel evidence was provided by stimulus onset, the urgency-gating model would appear to predict shorter (longer) decision times under speed (accuracy) conditions, with no change in accuracy, i.e., the integral would have reached its asymptote before 500 ms in either condition, so additional processing time would not improve accuracy. Under the framework of attractor dynamics, however, there is no discrepancy: local-circuit dynamics are subject to modulation, where weak dynamics support a leakage regime and stronger dynamics support a decision regime (Section 3.1.1). As such, modulation of network dynamics by a cognitive signal (Section 4.2.1) can support a range of time constants in the leakage or decision regimes (see **Figure 7B**). From this viewpoint, cognitive signals projecting to integrator circuitry (or evidence-encoding circuitry) are capable of supporting the effective time constant required by a given context, from around 100 ms (Cisek et al., 2009) to several seconds (Palmer et al., 2005). Under this framework, weak dynamics may be a default mode for decision circuitry under natural conditions (changing evidence), but cognitively demanding tasks may recruit dynamics supporting longer time constants.

The framework of attractor dynamics sheds further light on the possible neural implementation of urgency-gating. A leaky integrator with a short time constant could be implemented by weak local-circuit dynamics, per the first processing stage of **Figure 4**. In this regard, Thura et al. (2012) noted that the effect of the urgency signal on the decision variables could be additive, not necessarily multiplicative. In the study by Standage et al. (2011), a network model with weak dynamics was subject to gain

modulation by a growing urgency signal, i.e., the urgency signal had a multiplicative effect on the decision process. Long time constants were an emergent property of the network, suggesting that a neural implementation of urgency-gating might require additive input. The biophysically-based model by Standage et al. (2013) suggests that this input would need to be spatially selective (targeting each decision variable, but not other local-circuit neurons), since attractor dynamics (with long effective time constants) emerged in their model with a non-selective signal (Section 4.3.1). In principle, the urgency gating model could also be implemented by the projection of the urgency signal to thresholding circuitry, implementing a time-dependent version of the striatal hypothesis (Section 4.3.2) with weak decision dynamics. These and other possibilities require further investigation. Note that there is ample evidence for urgency signals, i.e., climbing activity encoding elapsed time (see Section 4.2.1). The ways in which this activity may modulate decision processing are receiving considerable attention (Ditterich, 2006b; Churchland et al., 2008; Cisek et al., 2009; Hanks et al., 2011; Standage et al., 2011; Drugowitsch et al., 2012; Standage et al., 2013). Changing-evidence tasks represent an important direction in the study of the SAT and decision making more generally.

5.3. DISTRIBUTED INTEGRATION OF EVIDENCE AND THE SAT

The distributed nature of decision processing is an important consideration for all three general classes of hypothesis. For the most part, we have described putative integrator activity one cortical area at a time [e.g., dlPFC (Kim and Shadlen, 1999), LIP (Roitman and Shadlen, 2002; Thomas and Pare, 2007) and FEF (Ding and Gold, 2012; Heitz and Schall, 2012)], highlighting the rate of this activity at choice time in a given electrophysiological experiment. It is likely that different decision-correlated cortical areas encode different dimensions of a given task. Changes to the profile of activity in these areas may therefore differ with task conditions. For example, a higher (lower) rate of FEF movement neurons under speed (accuracy) conditions (Heitz and Schall, 2012) may be accompanied by a lower (higher) rate of activity in dlPFC and/or LIP. All three areas project (at least indirectly) to the circuitry driving eye-movements. The distributed nature of decision processing is well-appreciated by researchers in decision neuroscience, but it is often implicit in electrophysiological studies (and studies based on electrophysiological data) that the relevant decision variable is the one being recorded. There are good reasons for choice thresholds to be fixed (see Marshall et al., 2012), but a fixed choice threshold need not imply a fixed rate of decision-selective activity in each of the brain regions projecting to the relevant motor circuitry. Rather, the aggregate input to the motor circuitry may be fixed, with varying contributions from upstream areas in different conditions.

To further complicate matters, decision-correlated brain areas are often bidirectionally coupled (e.g., FEF, LIP, and SC), so these areas presumably modulate each other during decisions. It is therefore plausible that in a given area, spike rates may indeed be fixed at the time of commitment to a choice, but that peak rates reflect the post-decision dynamics of choice behavior (see Simen, 2012). In light of these considerations, there is a need for electrophysiological recordings from multiple decision-correlated areas

under speed and accuracy conditions, e.g., dlPFC, LIP, and/or FEF during eye-movement tasks. The different ways in which decision variables in these areas are modulated by speed and accuracy conditions will not only be informative about the contributions of these areas to the SAT, but also about the roles they play in decision making more generally.

Similarly, the decision dynamics described above (Section 3.1.1) are based on single-circuit models, i.e., local-circuit integration of inputs from upstream, evidence-encoding neurons. We are unaware of any neural modeling studies to systematically consider the dynamics of bidirectionally-coupled decision circuits. It is clear that single-circuit attractor models cannot provide a full account of decision making. For example, these models necessarily produce longer error trials than correct trials (see Wong and Wang, 2006; Standage et al., 2011), but correct trials are longer under some task paradigms (see Ratcliff and Smith, 2004).

5.4. A UNIFYING PERSPECTIVE

We have described the above hypotheses one at a time, largely in isolation from one another, but as indicated in Section 4, these hypotheses should not be considered mutually exclusive. The electrophysiological data by Heitz and Schall (2012) are revealing in this regard, providing evidence for the modulation of sensory encoding, the rate of integration and the strength of non-evidence inputs to integrator circuitry. These data are consistent with the hypothesis that a cognitive signal projects non-selectively to sensory-encoding populations *and* integrator populations, controlling the SAT by gain modulation. Such a signal could be implemented by dlPFC (van Veen et al., 2008; Wenzlaff et al., 2011). It is possible that such a signal also projects to thresholding circuitry. Unlike non-selective input to integrator circuitry, which controls integration times in attractor models (Furman and Wang, 2008; Roxin and Ledberg, 2008; Standage et al., 2013), non-selective input to thresholding circuitry may have a negligible effect on local-circuit dynamics, given the already-strong dynamics hypothesized to support the implementation of thresholds (Simen, 2012). Such a cognitive signal could also project to pre-motor regions [e.g., pre-SMA (Forstmann et al., 2008, 2010; Ivanoff et al., 2008)], raising their baseline rates, in turn lowering motor thresholds for choice behavior. This description of the SAT assumes that the cognitive signal is present in neutral conditions, where its rate increases and decreases under speed and accuracy conditions respectively. This hypothesis unifies much of the data presented above and is depicted in **Figure 9B**.

Despite the long history of behavioral data describing the SAT, these are early days in its mechanistic study (Bogacz et al., 2010b). Recent electrophysiological (Heitz and Schall, 2012), neuroimaging (Forstmann et al., 2008, 2010; Ivanoff et al., 2008; van Veen et al., 2008; Wenzlaff et al., 2011; Green et al., 2012; Ho et al., 2012) and biophysically-based modeling (Lo and Wang, 2006; Furman and Wang, 2008; Standage et al., 2013) studies are exemplary of the promising methods being brought to bear on this fundamental cognitive phenomenon.

REFERENCES

Albantakis, L., and Deco, G. (2009). The encoding of alternatives in multiple-choice decision making. *Proc. Natl. Acad. Sci. U.S.A.* 106, 10308–10313. doi: 10.1073/pnas.0901621106

Aron, A. R., and Poldrack, R. A. (2006). Cortical and subcortical contributions to stop signal response inhibition: role of the subthalamic nucleus. *J. Neurosci.* 26, 2424–2433. doi: 10.1523/JNEUROSCI.4682-05.2006

Bogacz, R. (2007). Optimal decision-making theories: linking neurobiology with behaviour. *Trends Cogn. Sci.* 11, 118–125. doi: 10.1016/j.tics.2006.12.006

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., and Cohen, J. D. (2006). The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol. Rev.* 113, 700–765. doi: 10.1037/0033-295X.113.4.700

Bogacz, R., Hu, P. T., Holmes, P. J., and Cohen, J. D. (2010a). Do humans produce the speed-accuracy tradeoff that maximizes reward rate? *Q. J. Exp. Psychol.* 63, 863–891. doi: 10.1080/17470210903091643

Bogacz, R., Usher, M., Zhang, J., and McClelland, J. L. (2007). Extending a biologically inspired model of choice: multi-alternatives, nonlinearity and value-based multidimensional choice. *Philos. Trans. R. Soc.* 362, 1655–1670. doi: 10.1098/rstb.2007.2059

Bogacz, R., Wagenmakers, E.-J., Forstmann, B. U., and Nieuwenhuis, S. (2010b). The neural basis of the speed-accuracy tradeoff. *Trends Neurosci.* 33, 10–16. doi: 10.1016/j.tins.2009.09.002

Bollimunta, A., and Ditterich, J. (2011). Local computation of decision-relevant net sensory evidence in parietal cortex. *Cereb. Cortex* 22, 903–917. doi: 10.1093/cercor/bhr165

Britten, K. H., Shadlen, M. N., Newsome, W. T., and Movshon, J. A. (1992). The analysis of visual motion: a comparison of neuronal and psychophysical performance. *J. Neurosci.* 12, 4745–4765.

Britten, K. H., Shadlen, M. N., Newsome, W. T., and Movshon, J. A. (1993). Responses of neurons in macaque mt to stochastic motion signals. *Vis. Neurosci.* 10, 1157–1169. doi: 10.1017/S0952523800010269

Brunel, N., and Wang, X.-J. (2001). Effects of neuromodulation in a cortical network model of object working memory dominated by recurrent inhibition. *J. Comput. Neurosci.* 11, 63–85. doi: 10.1023/A:1011204814320

Carpenter, R. H. S., and Williams, M. L. L. (1995). Neural computation of log likelihood in control of eye movements. *Nature* 377, 59–62. doi: 10.1038/377059a0

Chittka, L., Skorupski, P., and Raine, N. E. (2009). Speed-accuracy trade-offs in animal decision making. *Trends Ecol. Evol.* 24, 400–407. doi: 10.1016/j.tree.2009.02.010

Churchland, A. K., Kiani, R., and Shadlen, M. N. (2008). Decision-making with multiple alternatives. *Nat. Neurosci.* 11, 693–702. doi: 10.1038/nn.2123

Cisek, P., Puskas, G. A., and El-Murr, S. (2009). Decisions in changing conditions: the urgency-gating model. *J. Neurosci.* 29, 11560–11571. doi: 10.1523/JNEUROSCI.1844-09.2009

Cohen, J. Y., Crowder, E. A., Heitz, R. P., Subraveti, C. R., Thompson, K. G., Woodman, G. F., et al. (2010). Cooperation and competition among frontal eye field neurons during visual target selection. *J. Neurosci.* 30, 3227–3238. doi: 10.1523/JNEUROSCI.4600-09.2010

Ding, L., and Gold, J. I. (2012). Neural correlates of perceptual decision making before, during, and after decision commitment in monkey frontal eye field. *Cereb. Cortex* 22, 1052–1067. doi: 10.1093/cercor/bhr178

Ditterich, J. (2006a). Evidence for time-variant decision making. *Eur. J. Neurosci.* 24, 3628–3641. doi: 10.1111/j.1460-9568.2006.05221.x

Ditterich, J. (2006b). Stochastic models of decisions about motion direction: behavior and physiology. *Neural Netw.* 19, 981–1012. doi: 10.1016/j.neunet.2006.05.042

Dorris, M. C., and Munoz, D. P. (1998). Saccadic probability influences motor preparation signals and time to saccadic initiation. *J. Neurosci.* 18, 7015–7026.

Doya, K., and Shadlen, M. N. (2012). Decision making. *Curr. Opin. Neurobiol.* 22, 911–913. doi: 10.1016/j.conb.2012.10.003

Dreyer, J. K., Herrik, K. F., Berg, R. W., and Hounsgaard, J. D. (2010). Influence of phasic and tonic dopamine release on receptor activation. *J. Neurosci.* 30, 14273–14283. doi: 10.1523/JNEUROSCI.1894-10.2010

Drugowitsch, J., Moreno-Bote, R., Churchland, A. K., Shadlen, M. N., and Pouget, A. (2012). The cost of accumulating evidence in perceptual decision making. *J. Neurosci.* 32, 3612–3628. doi: 10.1523/JNEUROSCI.4010-11.2012

Durstewitz, D. (2004). Neural representation of interval time. *Neuroreport* 15, 745–749. doi: 10.1097/00001756-200404090-00001

Durstewitz, D., and Seamans, J. K. (2006). Beyond bistability: biophysics and temporal dynamics of working memory. *Neuroscience* 139, 119–133. doi: 10.1016/j.neuroscience.2005.06.094

Fitts, P. M. (1966). Cognitive aspects of information processing: III. set for speed versus accuracy. *J. Exp. Psychol.* 71, 849–857. doi: 10.1037/h0023232

Forstmann, B. U., Anwander, A., Schafer, A., Neumann, J., Brown, S., Wagenmakers, E.-J., et al. (2010). Cortico-striatal connections predict control over speed and accuracy in perceptual decision making. *Proc. Natl. Acad. Sci. U.S.A.* 107, 15916–15920. doi: 10.1073/pnas.1004932107

Forstmann, B. U., Dutilh, G., Brown, S., Neumann, J., von Cramon, D. Y., Ridderinkhof, K. R., et al. (2008). Striatum and pre-sma facilitate decision-making under time pressure. *Proc. Natl. Acad. Sci. U.S.A.* 105, 17538–17542. doi: 10.1073/pnas.0805903105

Frank, M. J. (2006). Hold your horses: a dynamic computational role for the subthalamic nucleus in decision making. *Neural Netw.* 19, 1120–1136. doi: 10.1016/j.neunet.2006.03.006

Frank, M. J., Scheres, A., and Sherman, S. J. (2007). Understanding decision-making deficits in neurological conditions: insights from models of natural action selection. *Philos. Trans. R. Soc. Biol. Sci.* 362, 1641–1654. doi: 10.1098/rstb.2007.2058

Funahashi, S., Bruce, C. J., and Goldman-Rakic, P. S. (1989). Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *J. Neurophysiol.* 61, 331–349.

Furman, M., and Wang, X.-J. (2008). Similarity effect and optimal control of multiple-choice decision making. *Neuron* 60, 1153–1168. doi: 10.1016/j.neuron.2008.12.003

Fuster, J. (1973). Unit activity in prefrontal cortex during delayed-response performance: neuronal correlates of transient memory. *J. Neurophysiol.* 36, 61–78.

Gerstner, W. (2000). Population dynamics of spiking neurons: fast transients, asynchronous states, and locking. *Neural Comput.* 12, 43–89. doi: 10.1162/089976600300015899

Gold, J. I., and Shadlen, M. N. (2001). Neural computations that underlie decisions about sensory stimuli. *Trends Cogn. Neurosci.* 5, 10–16. doi: 10.1016/S1364-6613(00)01567-9

Gold, J. I., and Shadlen, M. N. (2002). Banburismus and the brain: decoding the relationship between sensory stimuli, decisions, and reward. *Neuron* 36, 299–308. doi: 10.1016/S0896-6273(02)00971-6

Gold, J. I., and Shadlen, M. N. (2007). The neural basis of decision making. *Ann. Rev. Neurosci.* 30, 535–574. doi: 10.1146/annurev.neuro.29.051605.113038

Grace, A. A. (1991). Phasic versus tonic dopamine release and the modulation of dopamine system responsivity: a hypothesis for the etiology of schizophrenia. *Neuroscience* 41, 1–24. doi: 10.1016/0306-4522(91)90196-U

Green, N., Biele, G. P., and Heekeren, H. R. (2012). Changes in neural connectivity underlie decision threshold modulation for reward maximization. *J. Neurosci.* 32, 14942–14950. doi: 10.1523/JNEUROSCI.0573-12.2012

Gurney, K., Prescott, T. J., and Redgrave, P. (2001). A computational model of action selection in the basal ganglia ii. analysis and simulation of behaviour. *Biol. Cybern.* 84, 411–423. doi: 10.1007/PL00007985

Hanks, T. D., Mazurek, M. E., Kiani, R., Hopp, E., and Shadlen, M. N. (2011). Elapsed decision time affects the weighting of prior probability in a perceptual decision task. *J. Neurosci.* 31, 6339–6352. doi: 10.1523/JNEUROSCI.5613-10.2011

Heekeren, H. R., Marrett, S., Bandettini, P. A., and Ungerleider, L. G. (2004). A general mechanism for perceptual decision-making in the human brain. *Nature* 43, 859–862. doi: 10.1038/nature02966

Heitz, R. P., and Schall, J. D. (2012). Neural mechanisms of speed-accuracy tradeoff. *Neuron* 76, 616–628. doi: 10.1016/j.neuron.2012.08.030

Ho, T., Brown, S., van Maanen, L., Forstmann, B. U., Wagenmakers, E.-J., and Serences, J. T. (2012). The optimality of sensory processing during the speed-accuracy tradeoff. *J. Neurosci.* 32, 7992–8003. doi: 10.1523/JNEUROSCI.0340-12.2012

Ivanoff, J., Branning, P., and Marois, R. (2008). fmri evidence for a dual process account of the speed-accuracy tradeoff in decision-making. *PLoS ONE* 3:e2635. doi: 10.1371/journal.pone.0002635

Ivry, R. B., and Spencer, R. M. C. (2004). The neural representation of time. *Curr. Opin. Neurobiol.* 14, 225–232. doi: 10.1016/j.conb.2004.03.013

Kable, J. W., and Glimcher, P. W. (2009). The neurobiology of decision: consensus and controversy. *Neuron* 63, 733–745. doi: 10.1016/j.neuron.2009.09.003

Kiani, R., Hanks, T. D., and Shadlen, M. N. (2008). Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. *J. Neurosci.* 28, 3017–3029. doi: 10.1523/JNEUROSCI.4761-07.2008

Kim, J.-N., and Shadlen, M. N. (1999). Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nat. Neurosci.* 2, 176–185. doi: 10.1038/5739

Komura, Y., Tamura, R., Uwano, T., Nishijo, H., Kaga, K., and Ono, T. (2001). Retrospective and prospective coding for predicted reward in the sensory thalamus. *Nature* 412, 546–549. doi: 10.1038/35087595

Lewis, P. A., and Miall, R. C. (2003). Distinct systems for automatic and cognitively controlled time measurement: evidence from neuroimaging. *Curr. Opin. Neurobiol.* 13, 250–255. doi: 10.1016/S0959-4388(03)00036-9

Lo, C.-C., and Wang, X.-J. (2006). Cortico-basal ganglia circuit mechanism for a decision threshold in reaction time tasks. *Nat. Neurosci.* 9, 956–963. doi: 10.1038/nn1722

Macmillan, N. A., and Creelman, C. D. (1991). *Detection Theory: A User's Guide.* New York, NY: Cambridge University Press.

Maimon, G., and Assad, J. A. (2006). A cognitive signal for the proactive timing of action in macaque lip. *Nat. Neurosci.* 9, 948–955. doi: 10.1038/nn1716

Marr, D. (1982). *Vision.* New York, NY: W. H. Freeman and Company.

Marshall, J. A. R., Bogacz, R., and Gilchrist, I. D. (2012). Consistent implementation of decisions in the brain. *PLoS ONE* 7:e43443. doi: 10.1371/journal.pone.0043443

Nambu, A. (2011). Somatotopic organization of the primate basal ganglia. *Front. Neuroanat.* 5:26. doi: 10.3389/fnana.2011.00026

Niki, H., and Watanabe, M. (1979). Prefrontal and cingulate unit activity during timing behavior in the monkey. *Brain Res.* 171, 213–224. doi: 10.1016/0006-8993(79)90328-7

Palmer, J., Huk, A. C., and Shadlen, M. N. (2005). The effect of stimulus strength on the speed and accuracy of a perceptual decision. *J. Vis.* 5, 376–404. doi: 10.1167/5.5.1

Philiastides, M. G., Auksztulewicz, R., Heekeren, H. R., and Blankenburg, F. (2011). Causal role of dorsolateral prefrontal cortex in human perceptual decision making. *Curr. Biol.* 21, 980–983. doi: 10.1016/j.cub.2011.04.034

Purcell, B. A., Heitz, R. P., Cohen, J. Y., Schall, J. D., Logan, G. D., and Palmeri, T. J. (2010). Neurally constrained modeling of perceptual decision making. *Psychol. Rev.* 117, 1113–1143. doi: 10.1037/a0020311

Purcell, B. A., Schall, J. D., Logan, G. D., and Palmeri, T. J. (2012). From salience to saccades: multiple-alternative gated stochastic accumulator model of visual search. *J. Neurosci.* 32, 3433–3446. doi: 10.1523/JNEUROSCI.4622-11.2012

Rainer, G., Rao, S. C., and Miller, E. K. (1999). Prospective coding for objects in primate prefrontal cortex. *J. Neurosci.* 19, 5493–5505.

Ratcliff, R., and McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. *Neural Computat.* 20, 873–922. doi: 10.1162/neco.2008.12-06-420

Ratcliff, R., and Smith, P. L. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychol. Rev.* 111, 333–367. doi: 10.1037/0033-295X.111.2.333

Reddi, B. A. J., and Carpenter, R. H. S. (2000). The influence of urgency on decision time. *Nat. Neurosci.* 3, 827–830. doi: 10.1038/77739

Roitman, J. D., and Shadlen, M. N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J. Neurosci.* 22, 9475–9489.

Roxin, A., and Ledberg, A. (2008). Neurobiological models of two-choice decision making can be reduced to a one-dimensional nonlinear diffusion equation. *PLoS. Comput. Biol.* 4:e1000046. doi: 10.1371/journal.pcbi.1000046

Salinas, E., and Abbott, L. F. (1996). A model of multiplicative neural responses in parietal cortex. *Proc. Natl. Acad. Sci. U.S.A.* 93, 11956–11961. doi: 10.1073/pnas.93.21.11956

Schall, J. D. (2001). Neural basis of deciding, choosing and acting. *Nat. Rev. Neurosci.* 2, 33–42. doi: 10.1038/35049054

Schall, J. D., Purcell, B. A., Heitz, R. P., Logan, G. D., and Palmeri, T. J. (2011). Neural mechanisms of saccade target selection: gated accumulator model of the visual-motor cascade. *Eur. J. Neurosci.* 33, 1991–2002. doi: 10.1111/j.1460-9568.2011.07715.x

Scolari, M., and Serences, J. T. (2012). Adaptive allocation of attentional gain. *J. Neurosci.* 29, 11933–11942. doi: 10.1523/JNEUROSCI.5642-08.2009

Shadlen, M. N., and Kiani, R. (2013). Decision making as a window on cognition. *Neuron* 80, 791–806. doi: 10.1016/j.neuron.2013.10.047

Shuler, M. G., and Bear, M. F. (2006). Reward timing in the primary visual cortex. *Science* 311, 1606–1609. doi: 10.1126/science.1123513

Simen, P. (2012). Evidence accumulator or decision threshold - which cortical mechanism are we observing? *Front. Psychol.* 3:183. doi: 10.3389/fpsyg.2012.00183

Simen, P., Cohen, J. D., and Holmes, P. (2006). Rapid decision threshold modulation by reward rate in a neural network. *Neural Netw.* 19, 1013–1026. doi: 10.1016/j.neunet.2006.05.038

Smith, P. L., and Ratcliff, R. (2004). Psychology and neurobiology of simple decisions. *Trends Neurosci.* 27, 161–168. doi: 10.1016/j.tins.2004.01.006

Standage, D., and Pare, M. (2011). Persistent storage capability impairs decision making in a biophysical network model. *Neural Netw.* 24, 1062–1073. doi: 10.1016/j.neunet.2011.05.004

Standage, D., You, H., Wang, D.-H., and 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. *Front. Comput. Neurosci.* 5:7. doi: 10.3389/fncom.2011.00007

Standage, D., You, H., Wang, D.-H., and Dorris, M. C. (2013). Trading speed and accuracy by coding time: a coupled-circuit cortical model. *PLoS Comput. Biol.* 9:e100321. doi: 10.1371/journal.pcbi.1003021

Stone, J. V. (2014). Using reaction times and binary responses to estimate psychophysical parameters: an information-theoretic analysis. *Front. Neurosci.* 8:35. doi: 10.3389/fnins.2014.00035

Thevarajah, D., Mikulic, A., and Dorris, M. C. (2009). Role of the superior colliculus in choosing mixed-strategy saccades. *J. Neurosci.* 29, 1998–2008. doi: 10.1523/JNEUROSCI.4764-08.2009

Thomas, N. W. D., and Pare, M. (2007). Temporal processing of saccade targets in parietal cortex area LIP during visual search. *J. Neurophysiol.* 97, 942–947. doi: 10.1152/jn.00413.2006

Thura, D., Beauregard-Racine, J., Fradet, C.-W., and Cisek, P. (2012). Decision making by urgency gating: theory and experimental support. *J. Neurophysiol.* 108, 2912–2930. doi: 10.1152/jn.01071.2011

Usher, M., and McClelland, J. L. (2001). On the time course of perceptual choice: the leaky competing accumulator model. *Psychol. Rev.* 108, 550–592. doi: 10.1037/0033-295X.108.3.550

van Maanen, L., Brown, S. D., Eichele, T., Wagenmakers, E.-J., Ho, T., Serences, J., et al. (2011). Neural correlates of trial-to-trial fluctuations in response caution. *J. Neurosci.* 31, 17488–17495. doi: 10.1523/JNEUROSCI.2924-11.2011

van Veen, V., Krug, M. K., and Carter, C. S. (2008). The neural and computational basis of controlled speed-accuracy tradeoff during task performance. *J. Cogn. Neurosci.* 20, 1952–1965. doi: 10.1162/jocn.2008.20146

Wang, X.-J. (2001). Synaptic reverberation underlying mnemonic persistent activity. *Trends Neurosci.* 24, 455–463. doi: 10.1016/S0166-2236(00)01868-3

Wang, X.-J. (2002). Probabilistic decision making by slow reverberation in cortical circuits. *Neuron* 36, 955–968. doi: 10.1016/S0896-6273(02)01092-9

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

Wang, X.-J. (2012). Neural dynamics and circuit mechanisms for decision-making. *Curr. Opin. Neurobiol.* 22, 1039–1046. doi: 10.1016/j.conb.2012.08.006

Wenzlaff, H., Bauer, M., Maess, B., and Heekeren, H. R. (2011). Neural characterization of the speed-accuracy tradeoff in a perceptual decision-making task. *J. Neurosci.* 31, 1256–1266. doi: 10.1523/JNEUROSCI.4000-10.2011

Wickelgren, W. A. (1977). Speed-accuracy tradeoff and information processing dynamics. *Acta Psychol.* 41, 67–85. doi: 10.1016/0001-6918(77)90012-9

Wilson, H. R., and Cowan, J. D. (1972). Excitatory and inhibitory interactions in localized populations of model neurons. *Biophys. J.* 12, 1–24. doi: 10.1016/S0006-3495(72)86068-5

Wong, K.-F., Huk, A. C., Shadlen, M., and Wang, X.-J. (2007). Neural circuit dynamics underlying accumulation of time-varying evidence during perceptual decision making. *Front. Comput. Neurosci.* 1:6. doi: 10.3389/neuro.10.006.2007

Wong, K.-F., and Wang, X.-J. (2006). A recurrent network mechanism of time integration in perceptual decisions. *J. Neurosci.* 26, 1314–1328. doi: 10.1523/JNEUROSCI.3733-05.2006

Yeung, N., Botvinick, M. M., and Cohen, J. D. (2004). The neural basis of error detection: conflict monitoring and the error-related negativity. *Psychol. Rev.* 111, 931–959. doi: 10.1037/0033-295X.111.4.931

You, H., and Wang, D.-H. (2013). Dynamics of multiple-choice decision making. *Neural Comput.* 25, 2108–2145. doi: 10.1162/NECO-a-00473

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

ADVANTAGES OF PUBLISHING IN FRONTIERS

FAST PUBLICATION
Average 90 days
from submission
to publication

COLLABORATIVE
PEER-REVIEW
Designed to be rigorous –
yet also collaborative, fair and
constructive

RESEARCH NETWORK
Our network
increases readership
for your article

OPEN ACCESS
Articles are free to read,
for greatest visibility

TRANSPARENT
Editors and reviewers
acknowledged by name
on published articles

GLOBAL SPREAD
Six million monthly
page views worldwide

COPYRIGHT TO AUTHORS
No limit to
article distribution
and re-use

IMPACT METRICS
Advanced metrics
track your
article's impact

SUPPORT
By our Swiss-based
editorial team