

A Two-Stage Process Model of Sensory Discrimination: An Alternative to Drift-Diffusion

Peng Sun and  Michael S. Landy

Department of Psychology and Center for Neural Science, New York University, New York, New York 10003

Discrimination of the direction of motion of a noisy stimulus is an example of sensory discrimination under uncertainty. For stimuli that are extended in time, reaction time is quicker for larger signal values (e.g., discrimination of opposite directions of motion compared with neighboring orientations) and larger signal strength (e.g., stimuli with higher contrast or motion coherence, that is, lower noise). The standard model of neural responses (e.g., in lateral intraparietal cortex) and reaction time for discrimination is drift-diffusion. This model makes two clear predictions. (1) The effects of signal strength and value on reaction time should interact multiplicatively because the diffusion process depends on the signal-to-noise ratio. (2) If the diffusion process is interrupted, as in a cued-response task, the time to decision after the cue should be independent of the strength of accumulated sensory evidence. In two experiments with human participants, we show that neither prediction holds. A simple alternative model is developed that is consistent with the results. In this estimate-then-decide model, evidence is accumulated until estimation precision reaches a threshold value. Then, a decision is made with duration that depends on the signal-to-noise ratio achieved by the first stage.

Key words: discrimination; drift-diffusion; motion; orientation; psychophysics; reaction time

Significance Statement

Sensory decision-making under uncertainty is usually modeled as the slow accumulation of noisy sensory evidence until a threshold amount of evidence supporting one of the possible decision outcomes is reached. Furthermore, it has been suggested that this accumulation process is reflected in neural responses, e.g., in lateral intraparietal cortex. We derive two behavioral predictions of this model and show that neither prediction holds. We introduce a simple alternative model in which evidence is accumulated until a sufficiently precise estimate of the stimulus is achieved, and then that estimate is used to guide the discrimination decision. This model is consistent with the behavioral data.

Introduction

Sensory discrimination under uncertainty is an everyday task that can be visual (“Is this light yellowish or bluish?”), auditory (“Is this pitch higher than that?”), or multisensory (“How long did the thunder occur after the lightning?”). Consider, for example, the oft-studied motion direction discrimination task. A set of randomly positioned dots is displayed in which a subset moves either to the left or right and the rest move in random directions. The task is to indicate motion direction. For fixed-duration stim-

uli, signal-detection theory (SDT; Green and Swets, 1988) makes a distinction between signal and noise. For discrimination, these correspond to stimulus “value” (e.g., $\pm 5^\circ$ rather than $\pm 90^\circ$ from vertical) and stimulus “strength” (motion coherence). (Note that we do not use the term “value” here in its economic sense.) Performance is a function of the signal-to-noise ratio: identical performance can result from a large-signal/high-noise (discriminating $\pm 90^\circ$ motion directions, low coherence) or weak-signal/low-noise ($\pm 5^\circ$, high coherence) stimulus. This characterization applies to a wide variety of tasks. For example, for orientation discrimination, identical performance can result from discriminating $\pm 45^\circ$ orientations at low contrast or $\pm 5^\circ$ at high contrast.

The difficulty of a discrimination task is reflected in response accuracy as well as speed [reaction time (RT)]. SDT links response accuracy to stimulus discriminability (signal-to-noise ratio) but provides no prediction of RT. To model performance in a reaction time setting for stimuli that are extended in time, the drift-diffusion model (DDM) has been influential (Ratcliff and Rouder, 1998; Ratcliff, 2001; Palmer et al., 2005; Ratcliff and McKoon, 2008; Forstmann et al., 2016; Ratcliff et al., 2016). In

Received April 26, 2016; revised Sept. 7, 2016; accepted Sept. 9, 2016.

Author contributions: P.S. and M.S.L. designed research; P.S. performed research; P.S. and M.S.L. analyzed data; P.S. and M.S.L. wrote the paper.

This work was supported by NIH Grant EY08266 and National Science Foundation—Collaborative Research in Computational Neuroscience Grant 1420262. We thank Stephanie Badde, Paul Glimcher, and Roozbeh Kiani for helpful comments.

The authors declare no competing financial interests.

Correspondence should be addressed to Peng Sun, Department of Psychology, New York University, 6 Washington Place, Room 955, New York, NY 10003. E-mail: peng.sun@nyu.edu.

DOI:10.1523/JNEUROSCI.1367-16.2016

Copyright © 2016 the authors 0270-6474/16/3611259-16\$15.00/0

this model, each instant yields a piece of evidence about the stimulus in the form of a log likelihood ratio of the two possible decision outcomes (e.g., leftward or rightward motion direction). The observer accumulates this evidence (forming an overall log likelihood ratio: the sum of the momentary fragments of evidence) and responds when it is deemed sufficiently strong (possibly also incorporating a log prior). Typically, experiments used to study evidence accumulation use fixed, large stimulus values (such as $\pm 90^\circ$ direction discrimination), and stimulus strength (e.g., motion coherence) is varied (Shadlen and Newsome, 2001; Roitman and Shadlen, 2002; Huk and Shadlen, 2005; Palmer et al., 2005; Kiani et al., 2008; Drugowitsch et al., 2012; Bitzer et al., 2014; de Lafuente et al., 2015). Since SDT and the drift-diffusion model base performance entirely on the signal-to-noise ratio, this approach is intuitively reasonable. This account of perceptual decision-making makes full use of sensory information at any point in time (Bogacz et al., 2006; Bitzer et al., 2014).

The drift-diffusion model has had a strong impact on the field. Neural responses in several brain areas appear to implement the accumulation of sensory evidence, including lateral intraparietal cortex (Gold and Shadlen, 2001; Shadlen and Newsome, 2001; Roitman and Shadlen, 2002; Yang and Shadlen, 2007; Kiani et al., 2008; Meister et al., 2013; Park et al., 2014; de Lafuente et al., 2015; Shadlen et al., 2016) and the frontal eye fields (Gold and Shadlen, 2000, 2003; Ding and Gold, 2012). This is viewed as confirmatory evidence of the drift-diffusion model (Gold and Shadlen, 2007).

We question the assumption of the drift-diffusion model that RT is solely a function of the signal-to-noise ratio. In Experiment 1, we measured RT in three visual discrimination tasks while varying both stimulus value and strength. The effect of these two variables on RT was additive, contradicting any model such as drift-diffusion that depends solely on the signal-to-noise ratio. We developed a simple alternative model that splits the decision process into two stages. The first stage depends only on stimulus strength, whereas the second stage depends primarily on stimulus value. This model is consistent with the results of Experiment 1. The two-stage model further predicts that RT depends on accumulated evidence strength even when the response is cued, preventing further evidence accumulation. This was verified in Experiment 2 using a cued-response discrimination task.

Materials and Methods

Experiment 1a

Subjects. Three subjects (age, 32–37; two males and one female) participated in Experiment 1a. S2 is the first author. S1 and S3 are researchers in the psychology department at New York University. Both were naive to the purpose of the experiment. S1 was amblyopic and had not had much experience running visual psychophysical experiments. S2 and S3 have normal or corrected-to-normal vision.

Apparatus. The experiments were controlled by a Macintosh Intel computer running Matlab with the Psychophysics Toolbox package (Brainard, 1997). Stimuli were displayed on a 15-inch Dell Trinitron VGA monitor with 1344×1280 resolution running at a 60 Hz refresh rate. Stimuli were presented using a linearized 10-bit lookup table. The mean luminance of all the stimuli and of the background was 54 cd/m^2 . The viewing distance was fixed at 56 cm maintained by use of a chinrest (resulting in 41 pixel°).

Stimuli. (1) Orientation-discrimination task: the stimulus was a Gabor pattern superimposed on (summed with) a dynamic binary texture patch. The texture had 20% Michelson contrast. The modulating Gaussian envelope's width (1 SD) was 0.45° . The Gabor was surrounded by a 1.5° diameter black circle that helped subjects maintain fixation. The spatial frequency of the Gabor's sinewave component was 1.3 cycle° with

randomized phase. The independent variables were Gabor contrast C and orientation θ_g (relative to vertical).

(2) Location-discrimination task: the stimulus was a 2° -wide (1 SD) Gaussian blob superimposed on (added to) a dynamic binary texture with 20% Michelson contrast. The Gaussian blob was displaced horizontally by a distance θ_s from the fixation point (the fixation point overlapped and occluded the Gaussian). The contrast C and spatial offset θ_s were the two independent variables.

(3) Direction-discrimination task: the stimulus was a two-frame life-time random-dot kinematogram (RDK) stimulus containing 80 0.05° black dots. In every pair of two frames, a subset of the dots (signal dots) moved coherently in one direction whereas the other dots (noise dots) randomly changed their locations. After every frame, the previous signal dots were now designated as noise dots, and a new set of signal dots was drawn from the noise dots of the previous frame. Therefore a signal dot never moved for more than two frames. The speed of the signal dots was $3^\circ/\text{s}$. The independent variables were the fraction of the signal dots (coherence C) and the motion direction θ_m (relative to vertical) of the signal dots.

Procedure. (1) Orientation-discrimination task: subjects indicated by keypress whether the Gabor was oriented counterclockwise or clockwise relative to vertical. Subjects were required to respond as quickly as possible while maintaining a 95% accuracy rate. Feedback was provided after each trial. In addition, information on overall accuracy was provided every 50 trials. The intertrial interval was random and uniform from 0.3 to 1 s. Five different values of contrast C and six different values of orientation θ_g ($\pm 12^\circ$, $\pm 15^\circ$, and $\pm 45^\circ$; eight for S1: $\pm 12^\circ$, $\pm 15^\circ$, $\pm 22^\circ$, and $\pm 45^\circ$) were tested. Each of the 30 (40 for S1) conditions was repeated 40 times in random order, resulting in a total of 1200 (1600 for S1) trials that were completed in three (four for S1) separate blocks. In an initial control experiment, the contrast threshold for the orientation-discrimination task for $\theta_g = \pm 45^\circ$ was estimated using a three-down, one-up staircase procedure. The lowest test contrast was set to double the subject's threshold contrast. The test contrast levels were as follows: 15, 25, 35, 45, and 55% for S1; 8, 16, 24, 32, and 48% for S2 and S3.

(2) Location-discrimination task: subjects indicated by keypress whether the Gaussian blob was located left or right of the central fixation. Five different values of contrast C and six different values of θ_s ($\pm 0.25^\circ$, $\pm 0.5^\circ$, and $\pm 1^\circ$) were tested. Each of the 30 conditions was repeated 40 times in random order. The procedure was otherwise similar to the orientation-discrimination task. The lowest test contrast was again set to double each subject's threshold, but in this case, it was the threshold for the location-discrimination task for $\theta_s = \pm 0.5^\circ$. The test contrast levels were as follows: 10, 20, 30, 40, and 50% for S1 and S3; 8, 12, 20, 32, and 64% for S2.

(3) Direction-discrimination task: subjects indicated by keypress whether the motion direction was predominantly leftward or rightward. Five different coherence levels C and six different motion directions θ_m ($\pm 8^\circ$, $\pm 22^\circ$, and $\pm 90^\circ$) were tested. For the diagonal conditions ($\pm 8^\circ$ and $\pm 22^\circ$), the motion could be rotated away from upward or downward, chosen randomly. Each of the 30 conditions was repeated 40 times for a total of 1200 trials. The lowest test motion coherence was set to double of each subject's threshold coherence estimated using a three-down, one-up staircase procedure in which subjects performed the same motion-discrimination tasks for $\theta_m = \pm 90^\circ$. The test coherence levels were as follows: 12.5, 15, 20, 27.5, and 40% for S1 and S3; 7.5, 10, 15, 22.5, and 35% for S2. Other details were the same as in the other two discrimination tasks.

Experiment 1b

Experiment 1b was identical to Experiment 1a except for the following minor differences. Subject S2 and an inexperienced, naive subject, S4 (age 26, male), also participated. S4 was paid \$10 per hour. Six contrast and coherence levels were tested. These stimulus values were more evenly spaced than in Experiment 1a. The Gabor orientations θ_g tested in the orientation-discrimination task were $\pm 12^\circ$, $\pm 22^\circ$, and $\pm 45^\circ$. The test contrasts were 8, 10, 14, 22, 38, and 60% for S2 and 4, 6, 10, 18, 34, and 64% for S4. The test contrast levels of the location-discrimination task were 15, 17, 21, 28, 40, and 60% for S2 and 10, 12, 16, 22, 30, and 50% for S4. Motion directions θ_m tested in the direction-discrimination task were

± 8 , ± 22 , and $\pm 45^\circ$. The test coherences were 7.5, 9, 12.5, 17.5, 25, and 40% for S2 and 7.5, 10, 15, 22.5, 32.5, and 45% for S4.

The motivation of Experiment 1b was to use a wider range of contrast/coherence levels and to test more evenly spaced stimulus values. In some conditions in Experiment 1a, most of the stimulus values were near the reference. Such a prior distribution of values could affect participants' response strategy.

Experiment 2

Subjects. Subject S2 and a newly recruited subject, S5 (age 25, female), who was paid \$10 per hour, participated in Experiment 2. The inexperienced subject S5 was not associated with New York University and had never been tested with a RDK stimulus. Before conducting the main experiment, S5's threshold for identifying left versus right direction was obtained by a similar adaptive procedure to the one used in Experiment 1. Test coherence levels for S5 were chosen accordingly.

Procedure. In each trial, an RDK stimulus was displayed for a variable time period, followed by a cue sound that lasted 100 ms. Signal dots moved either horizontally to the left or to the right (i.e., only two possible directions). The stimulus disappeared after the onset of the cue. The time period between the onset of the stimulus and the onset of the cue is called the stimulus onset asynchrony (SOA). Subjects were required to indicate the direction (left or right) of the signal dots immediately after hearing the sound. The deadline to make a valid response was 350 ms after the onset of the cue, creating pressure for subjects to respond as quickly as possible. Key presses made outside this window (either sooner or later) would generate a message, warning that the previous response was invalid. Invalid trials were repeated at the end of each block. If responses still missed the deadline for a repeated trial, the RT of the second attempt was recorded and used in the analysis. Extensive training was provided. Initially, subjects missed the response deadline about half of the time. After extensive practice, responses were on time for all trials for S2 and 98% of trials for S5. Removing those 2% late trials for S5 would not affect any of our conclusions. In the main experiment, all SOAs were equally likely. All SOA and coherence conditions were interleaved, and each condition was repeated 50 times. In a control experiment, the occurrence of different SOAs followed a truncated exponential distribution to minimize the potential effect of varying hazard rates (Kiani et al., 2008). The mean of the exponential distribution was 420 ms. Only SOAs within the range 100–800 ms were used. Other than the specific coherence values and stimulus durations, the motion stimuli were identical to those used in the direction-discrimination task in Experiment 1.

Results

We begin by developing a version of the drift-diffusion model based on a neural population code. For a given stimulus, the stimulus strength (e.g., contrast) determines the noise level and the stimulus value (e.g., angular difference) determines the signal level. These two factors interact in the model in determining mean reaction time (RT is a function of the signal-to-noise ratio). We test this prediction in three different tasks and find no such interaction. Then, we develop a simple alternative two-stage model that is consistent with the results. Finally, we discuss a second behavioral experiment that is again inconsistent with the drift-diffusion model but, rather, provides further evidence for our two-stage model of perceptual decision-making.

Single-stage drift-diffusion model

The drift diffusion model is perhaps the most popular model to account for the perceptual decision-making process involved in many cognitive and perceptual tasks (Ratcliff, 1978; Ratcliff and Rouder, 1998; Palmer et al., 2005; Ratcliff and McKoon, 2008; Simen and Contreras, 2009; Purcell et al., 2010; Krajbich and Rangel, 2011; Krajbich et al., 2011; Drugowitsch et al., 2012; Mulder et al., 2012; Guest et al., 2015). The central machinery of a drift-diffusion model is a drifting particle whose momentary movement is subject to stochastic variation (Fig. 1). In the con-

text of decision-making, the diffusion process is thought to reflect the accumulation of noisy evidence. The decision outcome is determined by the terminating position of the particle relative to a reference position. In a two-alternative, decision-making process, the sign of the particle's relative position indicates which alternative to choose. The distance indicates the degree of belief about choosing this alternative. Practically (and also supported by evidence from neurophysiology), the momentary sensory evidence is often computed as the log likelihood ratio of the two alternatives given the stimulus (Bogacz et al., 2006; Kira et al., 2015). Here we introduce a neural population code as the source of the momentary sensory evidence in the standard drift-diffusion model. We will show that this model makes interesting predictions about the reaction time of correct responses.

Momentary evidence for discrimination

Consider, for example, a task in which the subject is to judge whether the orientation of a Gabor pattern is clockwise or anticlockwise of vertical. Because the Gabor pattern is embedded in a noisy background and because the contrast and orientation of the Gabor are not known ahead of time, this perceptual task requires the subject to make the decision based on sensory information accumulated over time. Obviously, the higher the contrast or the more oblique the orientation, the easier is the task, leading to a more rapid response.

Suppose, at any short time interval $[\tau_{t-1}, \tau_t]$, a Gabor pattern $G(C, \theta_0)$ causes a population of orientation-tuned neurons to fire spikes count $\mathbf{n}_t = (n_{1,t}, \dots, n_{M,t})$, where C and θ_0 are the contrast and orientation of the Gabor pattern, respectively, and $n_{i,t}$ is the number of spikes generated during the interval $[\tau_{t-1}, \tau_t]$ by the neuron whose tuning curve peaks at θ_i . Thus, the neural activity \mathbf{n}_t provides the momentary sensory information on which the brain may base its decision.

Assuming a Poisson process for neural activity, conditional independence between the neurons, and observer knowledge of the stimulus contrast, then the posterior distribution over orientation is as follows (Ma et al., 2006):

$$p_i(\theta | \mathbf{n}_t, C) = \frac{p(\mathbf{n}_t | \theta, C) p(\theta | C)}{p(\mathbf{n}_t | C)} \propto \frac{\prod_i e^{-g(C) f_i(\theta)} g(C)^{n_{i,t}} f_i(\theta)^{n_{i,t}} / n_{i,t}!}{\int p(\mathbf{n}_t | \theta', C) d\theta'} \quad (1)$$

where the firing rate of neuron i is assumed to be a separable function of stimulus orientation and contrast, $f_i(\theta)g(C)$; f_i is the "shape" of the tuning curve for neuron i ; and g is a common contrast gain. We further assume that the shapes of the tuning curves are identical, differing only in preferred orientation. Contrast gain g depends only on stimulus contrast C and the length of the time interval $[\tau_{t-1}, \tau_t]$. The prior $p(\theta | C)$ is assumed to be flat. Rearranging the numerator and expanding the denominator, Equation 1 becomes the following:

$$p_i(\theta | \mathbf{n}_t, C) \propto \frac{e^{-g(C) \sum_i f_i(\theta)} \left(\prod_i \frac{g(C)^{n_{i,t}}}{n_{i,t}!} \right) e^{\sum_i n_{i,t} \log f_i(\theta)}}{e^{-g(C) \sum_i f_i(\theta)} \left(\prod_i \frac{g(C)^{n_{i,t}}}{n_{i,t}!} \right) \int e^{\sum_i n_{i,t} \log f_i(\theta')} d\theta'} \quad (2)$$

where the exponential term was taken out of the integral in the denominator by assuming that $\sum_i f_i(\theta)$ is a constant independent

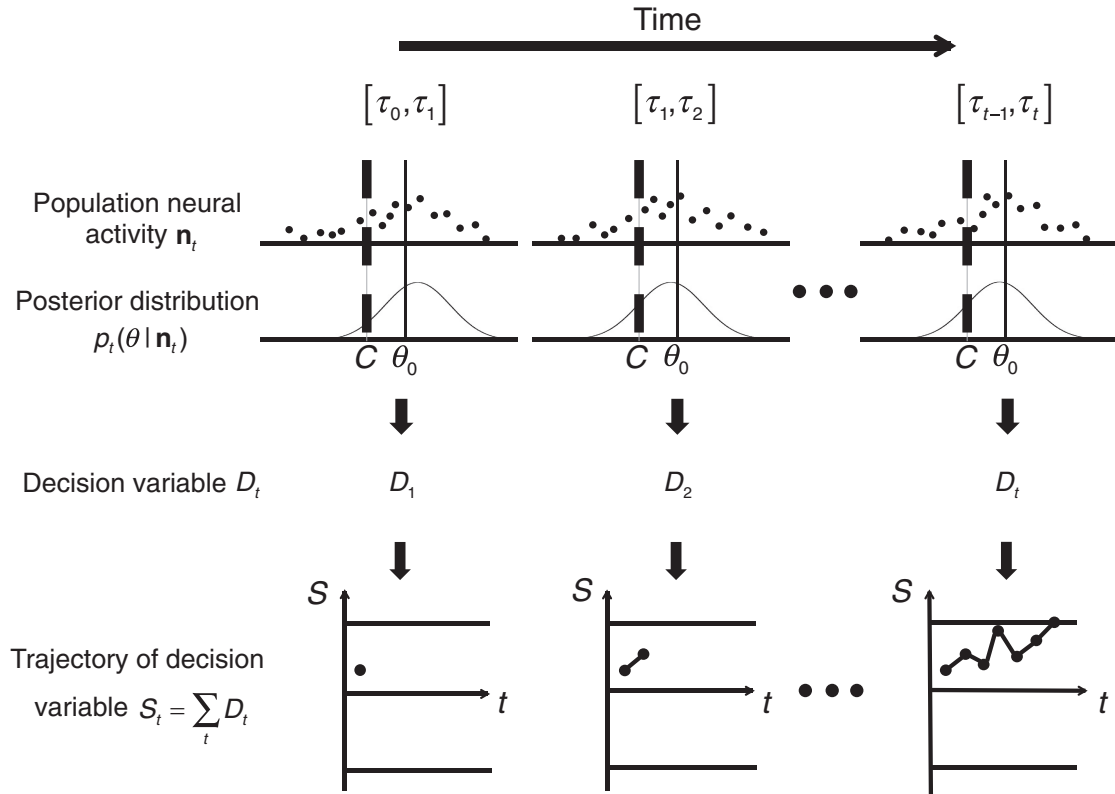


Figure 1. Schematic illustration of the single-stage model. Noisy discrimination evidence D_t is obtained at every time interval $[\tau_{t-1}, \tau_t]$ by inferring from momentary population neural activity $D_t = \log \frac{p_t(\theta > 0 | \mathbf{n}_t)}{p_t(\theta < 0 | \mathbf{n}_t)}$. Discrimination evidence is accumulated over time until it hits a boundary.

of θ . Canceling common terms, Equation 2 becomes the following:

$$p_t(\theta | \mathbf{n}_t, C) \propto \frac{e^{\sum_i n_{i+} \log f_i(\theta)}}{\int e^{\sum_i n_{i+} \log f_i(\theta')} d\theta'} = K_t e^{\sum_i n_{i+} \log f_i(\theta)}, \quad (3)$$

where K_t is a constant that depends only on momentary population activity \mathbf{n}_t but not on C or θ . Since Equation 3 does not depend on contrast, we can drop C in subsequent expressions for p_t . This also means that the subject does not need to know or estimate C to perform the task (Ma et al., 2006). The visual discrimination task can be solved by computing the following:

$$p_t(\theta > 0 | \mathbf{n}_t) = \int_{\theta > 0} p_t(\theta | \mathbf{n}_t) d\theta \quad (4)$$

$$p_t(\theta < 0 | \mathbf{n}_t) = 1 - p_t(\theta > 0 | \mathbf{n}_t),$$

where $p_t(\theta | \mathbf{n}_t)$ is given in Equation 3. The log likelihood ratio

$$D_t = \log \frac{p_t(\theta > 0 | \mathbf{n}_t)}{p_t(\theta < 0 | \mathbf{n}_t)} \quad (5)$$

serves as the momentary evidence for the orientation-discrimination task. For each time interval $[\tau_{t-1}, \tau_t]$, D_t is stochastic and may be insufficiently reliable to support a decision. But if the brain has access to this momentary evidence and accumulates it over time, then the accumulated evidence will eventually lead to a

more reliable decision. An example of such a diffusion process is illustrated in Figure 1.

Effect of C and θ_0 on the reaction time of correct responses

The single-stage drift-diffusion model predicts there is an interaction between the effects of C and θ_0 on the reaction time of correct responses. Intuitively, since firing rates of neurons are a separable product $g(C)f_i(\theta_0)$ and probabilistically pooled firing rates (Eqs. 3–4) determine the quality of the evidence, the contributions of C and θ_0 multiply so that an increase in one leads to an amplification of the effects of the other. This multiplicative effect, filtered through the diffusion process, leads to an interaction in reaction time: a reduction of C leads to a greater increase in reaction time for small (difficult) θ_0 than for large (easy) θ_0 .

To verify this intuition, it helps to consider an alternative measurement of momentary evidence $D'_t = p_t(\theta > 0 | \mathbf{n}_t) - p_t(\theta < 0 | \mathbf{n}_t)$ for which the mean stopping time of the diffusion process has a simple analytic form. Note that $p_t(\theta | \mathbf{n}_t)$ is a stochastic function of θ given that \mathbf{n}_t is a random vector. However, $E(p_t(\theta | \mathbf{n}_t))$ is a deterministic function of θ . Suppose $E(p_t(\theta | \mathbf{n}_t))$ is of approximately Gaussian shape with unbiased mean θ_0 and variance σ^2 . Then the expected value of D'_t is as follows:

$$E(D'_t) = \int_{\theta > 0} E(p_t(\theta | \mathbf{n}_t)) d\theta - \int_{\theta < 0} E(p_t(\theta | \mathbf{n}_t)) d\theta$$

$$= \left(\Phi\left(\frac{\theta_0}{\sigma}\right) \right) - \left(1 - \Phi\left(\frac{\theta_0}{\sigma}\right) \right)$$

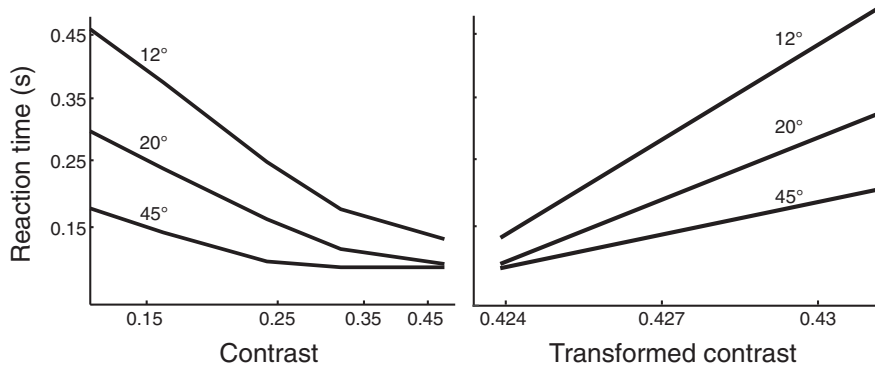


Figure 2. Simulated mean reaction time as a function of Gabor contrast and orientation. Note that nondetection time for stimulus encoding and motor delay was not included. The two figures are two different representations of the same simulated mean RTs. On the right, simulated mean RTs are plotted against transformed contrast, resulting in nearly linear curves. (See Fig. 3 and text for further details.)

$$= 2\Phi\left(\frac{\theta_0}{\sigma}\right) - 1 = 2\Phi(\theta_0 f_c(C)) - 1, \quad (6)$$

where $f_c(C) = 1/\sigma$ is a monotonically increasing function of the contrast C , because estimation uncertainty σ decreases with increasing contrast. The variance of D'_i has a more complex form. It is a function of C but is independent of θ_0 . The mean stopping time for a diffusion process using D'_i is largely inversely related to its expected value (Shadlen et al., 2006):

$$E(\text{RT}) \approx \frac{B}{E(D'_i)} = \frac{B}{2\Phi(\theta_0 f_c(C)) - 1}, \quad (7)$$

where B is a constant primarily determined by the terminating boundary. The interaction between C and θ_0 is attributable to the product in the denominator.

Because of its neural plausibility, we choose to use D_i (Eq. 5) in our drift-diffusion model, instead of D'_i . But unlike $E(D'_i)$, $E(D_i)$ is difficult to work with analytically. To verify the interaction of C and θ_0 with D_i , we simulated the model. A spike train was generated by simulating Poisson responses of 36 orientation-tuned neurons with orientation preferences evenly spaced between $-\pi/2$ and $\pi/2$. Tuning curves were circular Gaussians with SDs equal to 1 radian. A diffusion step was taken every 50 ms. The gain term in Equation 1 was arbitrarily defined as $g(C) = K \log(1 + C/C_0)$, where C_0 represents, for example, the contrast of background noise in the stimulus. K was set so that the maximal firing rate was 80 spikes/s. This arbitrary function captures the accelerating and decelerating sections of simple cells' contrast transfer function. Since the function depends on C but not on θ_0 , it does not create or remove the interactive effect between C and θ_0 . The exact shape is not critical for our purposes. The stopping time at the correct boundary was taken as the reaction time for a correct trial. The boundary height was set so that the diffusion process ended up at the correct side 95% of the time. Simulated mean reaction times are shown in Figure 2. The interaction between contrast and orientation is obvious, consistent with the intuition. Drift-diffusion models usually include a parameter for nondetection time (i.e., neural delays and time after the decision to initiate and perform the motor response); we do not include nondetection time as it is irrelevant to the interaction we are trying to demonstrate.

Experiment 1: reaction time in three visual discrimination tasks

Reaction times in three visual discrimination tasks (orientation, location, and direction discrimination) were measured to determine whether stimulus strength (e.g., contrast or motion coherence) and value (e.g., spatial or angular offset) interact in their effects on reaction time. According to the drift-diffusion model, they should.

Accuracy for all subjects in all tasks was above 92% in Experiment 1a and 88% in Experiment 1b. Accuracy and RT values across conditions were consistent with speed-accuracy trade-off and changing difficulty across conditions, i.e., lower reaction times were associated with higher accuracy. But because the ranges of accuracy were small (because of the instructions to subjects), the accuracy data were less informative than RT. Therefore, our analysis is primarily built around RT.

Figure 3A shows mean reaction time for correct responses in the orientation-discrimination task in Experiment 1a for naive subject S1 who achieved 98% accuracy across all conditions and above 95% for almost every individual condition. Data for different signs of the same stimulus value were pooled. RT decreases with increasing stimulus contrast C and stimulus orientation θ_g . (Note that the data for $\pm 22^\circ$ and $\pm 45^\circ$ were very similar, so we omit the data for $\pm 22^\circ$.) The relationship between mean RT and contrast appears to be similar across orientation conditions. Solid lines are fits of a model of the form $\text{RT} = K/\log(1 + C^\alpha/C_0) + \beta_i$. C_0 was fixed at 0.01 for all subjects, tasks, and stimulus values. α was varied across subjects and tasks to fit the data but was independent of stimulus value and strength for a given subject and task. β_i was allowed to vary across stimulus strength condition i (e.g., the orientation condition in Fig. 3). Figure 3B shows the same data as a function of $\tilde{C} = 1/\log(1 + C^\alpha/C_0)$. Similar transformations have been used to linearize mean reaction time plotted against stimulus strength for saccade data (Reddi et al., 2003; Carpenter, 2004). In this plot, K represents the slopes of the fit lines and β_i is the y -intercept. Using a single value of K for the three orientation conditions in Figure 3, this model accounts for 99.6% of the total variance. We test for interaction by comparing this model to one in which K is allowed to depend on orientation condition. The accounted variance improves to 99.8%, but at the cost of three more parameters (for the four orientation conditions, only three of which are shown in Fig. 3).

Mean RTs of correct responses for all subjects and tasks of Experiments 1a and 1b are shown in Figures 4 and 5, respectively. For easier visual examination of interaction effects, data are plotted against transformed contrast or coherence \tilde{C} . The linearization of RT curves is determined by α (see Table 1 for values for each subject and task). α carries information as to how RT varies with stimulus strength and is unimportant for our purposes here. What we are interested in is K . Visually, curves of mean RT as a function of \tilde{C} are by and large parallel for different orientations, locations, and directions (Table 2). Generally, a common value of K for all orientations explains most of the variance. Allowing K to vary for different orientations does not improve the fit by much and, in most cases, not significantly. These separately fit values of K do not differ (Fig. 6) for all cases except one subject (S4) in the




Figure 3. Mean reaction time for correct responses for subject S1 in the orientation-discrimination task of Experiment 1a. Error bars represent 95% confidence intervals. **A**, The abscissa represents stimulus contrast C . Solid lines are fits of a model that includes no interaction term (see text for details). **B**, The same data plotted as a function of transformed contrast $\tilde{C} = 1/\log(1 + C^\alpha/C_0)$, linearizing the relationship between transformed contrast and reaction time.

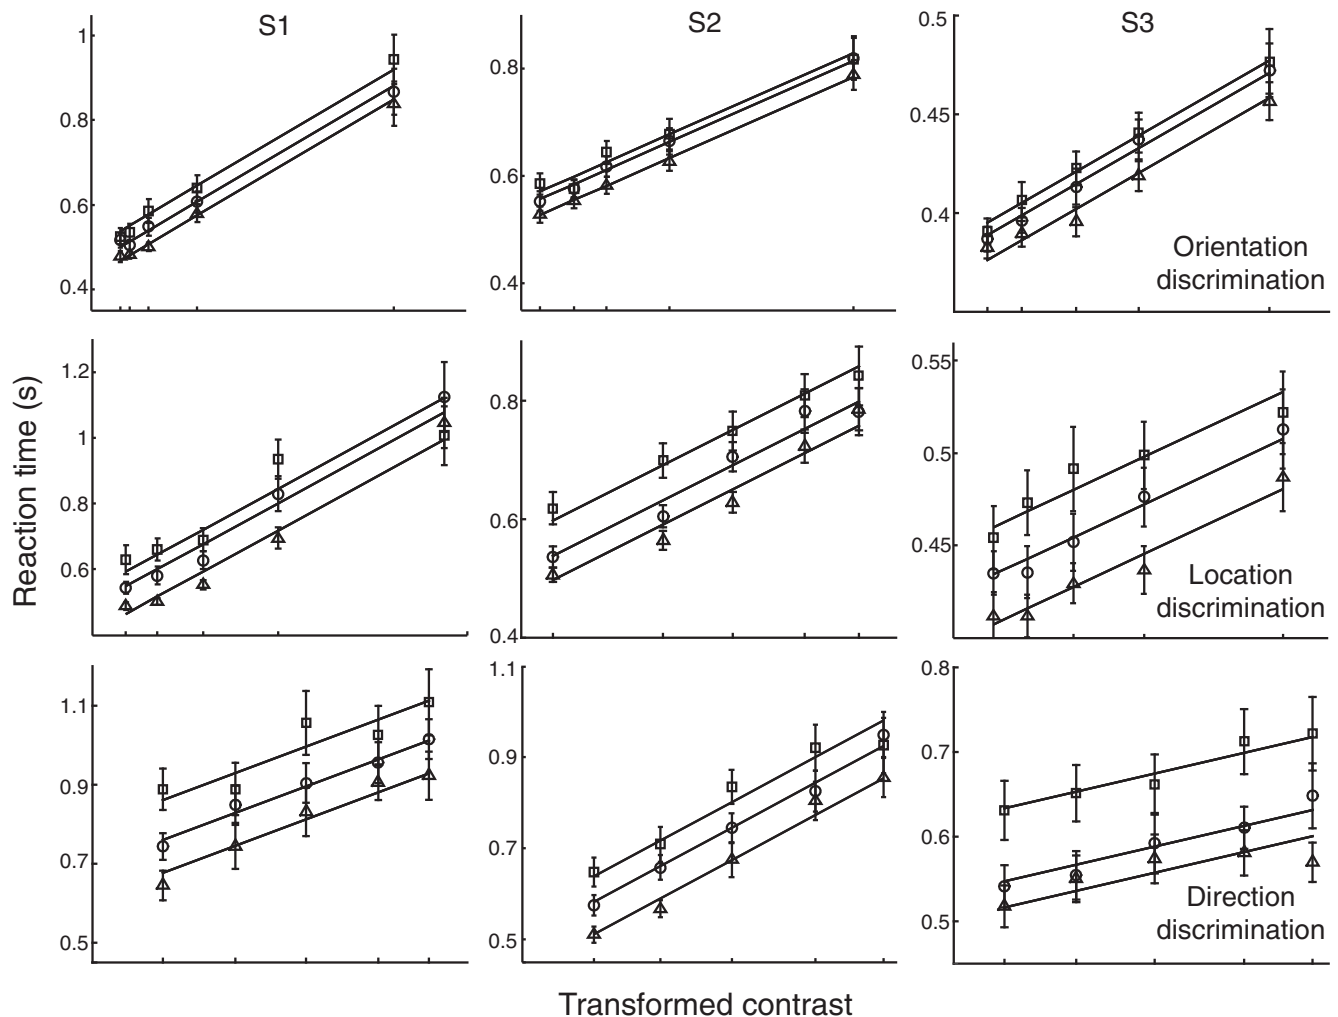


Figure 4. Experiment 1a. Mean reaction time of correct responses in three discrimination tasks (top, orientation; middle, location; bottom, direction) for three subjects. Different symbols represent different Gabor orientations ($\theta_g = \pm 12, \pm 15, \text{ and } \pm 45^\circ$), spatial displacements ($\theta_s = \pm 0.25, \pm 0.5, \text{ and } \pm 1^\circ$), or motion directions ($\theta_m = \pm 8, \pm 22, \text{ and } \pm 90^\circ$). Error bars represent 95% confidence intervals.

orientation-discrimination task (Fig. 5 (top right), Fig. 6B (first column); Table 2, top right cell).

Note that subjects were asked to maintain 95% accuracy. In fact, all subjects either surpassed this criterion or at least came

close to achieving it in all three tasks. This requirement was imposed to prevent subjects from responding early on difficult trials if they did not care about accuracy. Previous studies have shown that when a task is difficult, subjects tend to respond faster than

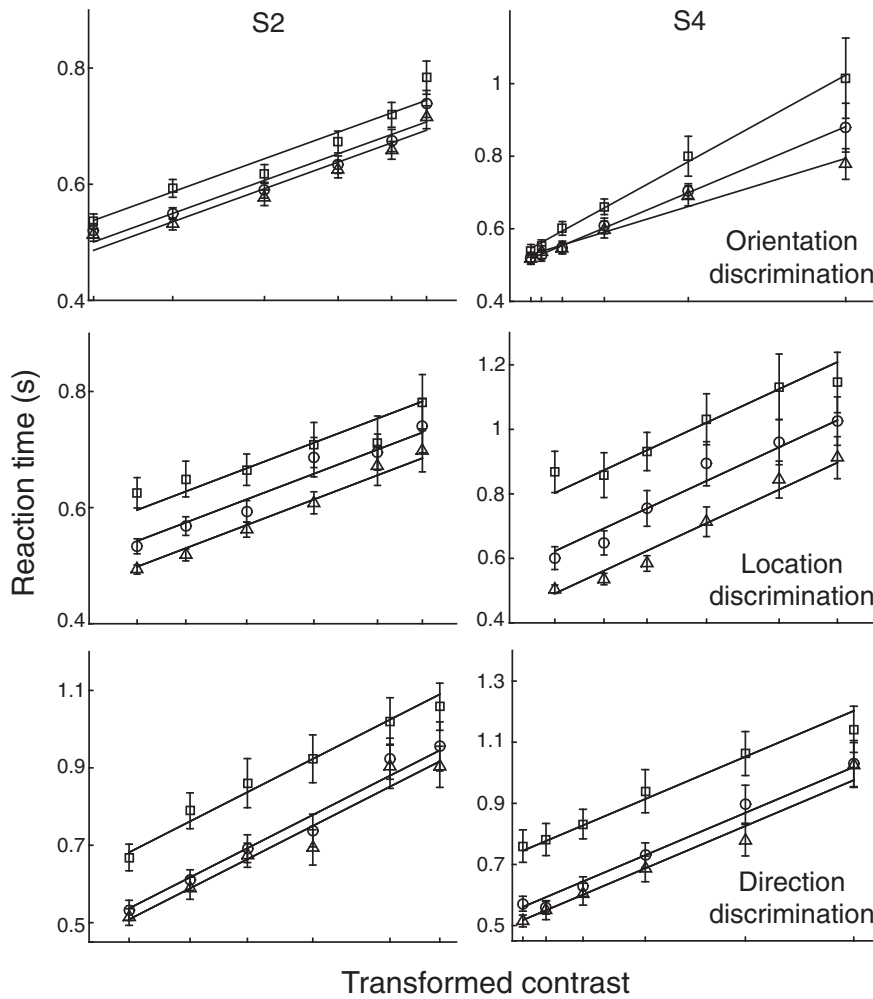


Figure 5. Experiment 1b. Mean reaction time of correct responses in three discrimination tasks (top, orientation; middle, location; bottom, direction) for two subjects. Different symbols represent different Gabor orientations ($\theta_g = \pm 12, \pm 22.5, \text{ and } \pm 45^\circ$), spatial displacements ($\theta_s = \pm 0.25, \pm 0.5, \text{ and } \pm 1^\circ$), or motion directions ($\theta_m = \pm 8, \pm 22, \text{ and } \pm 45^\circ$). Error bars represent 95% confidence intervals. Top right, These data were fit significantly and substantially better with separate slopes for each orientation condition. For all other panels, parallel fits are shown.

Table 1. Fitted α values for each subject and in each discrimination task

	Experiment 1a			Experiment 1b	
	S1	S2	S3	S2	S4
Orientation discrimination	2.23	1.03	0.46	1.5	0.91
Location discrimination	0.58	0.82	0.3	1.46	0.43
Direction discrimination	0.59	0.51	0.97	0.71	1.13

they should (Drugowitsch et al., 2012; Moran, 2015). This may reflect a strategy in which overall reward is maximized by ending difficult trials early so as to increase the total number of trials, spending more time on the easier trials that more often lead to reward (thus maximizing the overall reward rate). If this strategy were adopted consciously in our task, subjects would rush on trials with low stimulus strength or low stimulus value, possibly masking any interaction effect that would otherwise be present in the RT data. Fortunately, we found no evidence of an interaction across a wide range of performance levels, and observers maintained high performance levels across conditions, so we are confident in our interpretation of the data.

As for the exceptional case of S4 in the orientation-discrimination task, the effects of contrast and orientation

reaction time clearly interact (Fig. 5, top right). Allowing K to vary for different orientation conditions improves the fit significantly and substantially. K decreases as orientation increases, whereas the fit values of β (the y -intercepts) do not change substantially. The accuracy for this particular subject and task was 97%, in line with the other subjects for the same task in Experiments 1a and 1b. The possible causes of this interaction will be discussed in the next section.

In summary, no obvious interaction for reaction time can be found in any of the three visual discrimination tasks of Experiment 1a, and an interaction was found in only one case in Experiment 1b. In most cases, there is a strong indication that mean RT decreases with increasing stimulus strength by the same amount independent of the stimulus value. Although an interaction between stimulus strength and value may have many causes, the lack of interaction places a strong constraint on models of the decision processes involved in these tasks. This led us to consider an alternative model, as described in the next section.

The two-stage estimate-then-decide model

The drift-diffusion model consists of a single stage in which evidence is accumulated until sufficient evidence is available for a decision. The accumulating evidence combines and conflates stimulus strength and value through the signal-to-noise ratio. As a result, drift-diffusion predicts an interaction of these two stimulus variables (Fig. 2). We now develop what we consider the simplest model consistent with

the results of Experiment 1. This two-stage estimate-then-decide (ETD) model offers a simple explanation for these results: the decision process is composed of two separate subprocesses that operate sequentially. The additive relationship between the RT curves results from the cascade nature of the ETD model (McClelland, 1979); the duration of one subprocess is a function only of stimulus strength, and the other’s duration is a function of stimulus value. The processes are performed sequentially, so their durations sum, leading to no interaction between stimulus value and strength.

In the first stage of the ETD model, evidence is accumulated until a sufficiently reliable posterior distribution of θ_0 is produced. The second stage makes use of the estimated posterior distribution of θ_0 to decide whether $\theta_0 > 0$, i.e., it makes the discrimination judgment. The most notable difference between the ETD and the drift-diffusion model is that discrimination evidence is not available until a reliable posterior distribution is obtained. By separating the process into two substages, the effects of stimulus strength and value are decoupled.

Stage one: the evolution of the posterior distribution

The ETD model begins with a framework similar to drift-diffusion: a population of orientation-selective neurons with Poisson

Table 2. *F*-test for goodness of fit for the parallel case (common *K*) versus the nonparallel case (variable *K* across stimulus values)

	Experiment 1a			Experiment 1b	
	S1	S2	S3	S2	S4
Orientation discrimination	$F_{(3,11)} = 2.3$ $p = 0.13$ $v_{1k} = 99.6\%$ $v_{3k} = 99.8\%$	$F_{(2,8)} = 0.83$ $p = 0.47$ $v_{1k} = 98.9\%$ $v_{3k} = 99.1\%$	$F_{(2,8)} = 1.7$ $p = 0.23$ $v_{1k} = 98.8\%$ $v_{3k} = 99.2\%$	$F_{(2,11)} = 0.33$ $p = 0.72$ $v_{1k} = 94.3\%$ $v_{3k} = 94.6\%$	$F_{(2,11)} = .99$ $p = .000$ $v_{1k} = 94\%$ $v_{3k} = 99.9\%$
Location discrimination	$F_{(2,8)} = 3.3$ $p = 0.08$ $v_{1k} = 98.1\%$ $v_{3k} = 98.9\%$	$F_{(2,8)} = 1.5$ $p = 0.26$ $v_{1k} = 96.8\%$ $v_{3k} = 97.5\%$	$F_{(2,8)} = 1.6$ $p = 0.25$ $v_{1k} = 99.2\%$ $v_{3k} = 99.9\%$	$F_{(2,11)} = 3.8$ $p = 0.05$ $v_{1k} = 96\%$ $v_{3k} = 97.6\%$	$F_{(2,11)} = 3.2$ $p = 0.08$ $v_{1k} = 97.3\%$ $v_{3k} = 98.3\%$
Direction discrimination	$F_{(2,8)} = 0.63$ $p = 0.45$ $v_{1k} = 95.3\%$ $v_{3k} = 95.8\%$	$F_{(2,8)} = 1.4$ $p = 0.28$ $v_{1k} = 97.4\%$ $v_{3k} = 98.0\%$	$F_{(2,8)} = 4.1$ $p = 0.05$ $v_{1k} = 96.8\%$ $v_{3k} = 97.6\%$	$F_{(2,11)} = 0.63$ $p = 0.55$ $v_{1k} = 97.4\%$ $v_{3k} = 97.7\%$	$F_{(2,11)} = 2.5$ $p = 0.13$ $v_{1k} = 98.2\%$ $v_{3k} = 98.7\%$

v_{1k} and v_{3k} represent the variance accounted by the common *K* and variable *K* models, respectively.

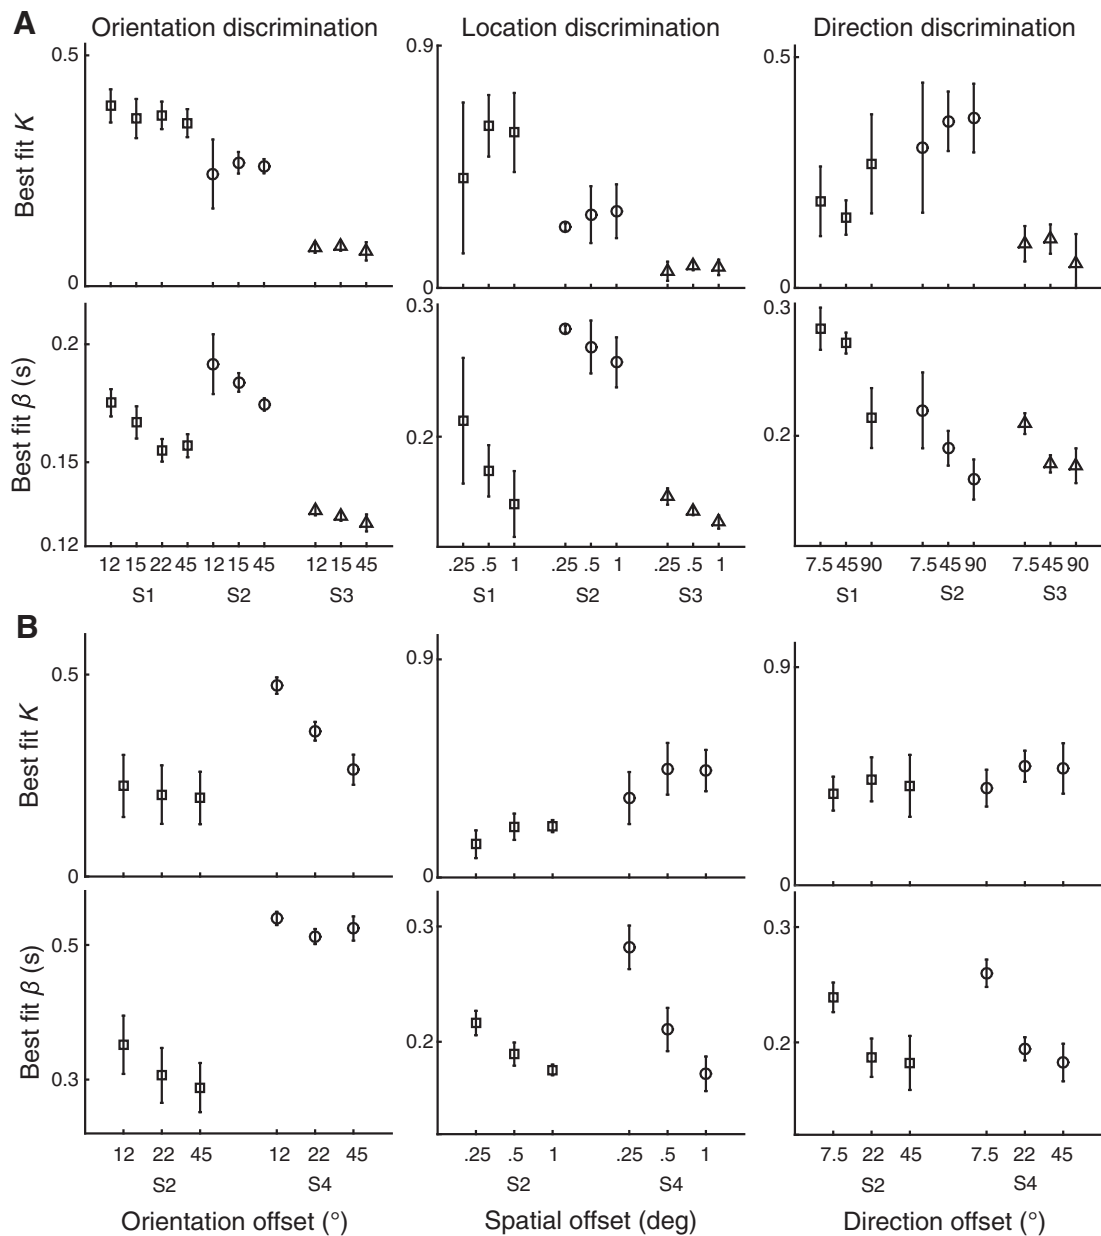


Figure 6. Best-fit values of *K* as a free parameter and best-fit values of β for different values in Experiments 1a (A) and 1b (B). Error bars represent 95% confidence intervals, assuming normality of error residuals after fitting mean reaction time data to linear models.




Figure 7. Illustration of the two-stage ETD model. In the estimation stage, neural spikes are accumulated over time until the width (SD) of the posterior distribution drops below a threshold, yielding a relatively accurate estimate of θ_0 . In the decision stage, the decision variable D is computed as $D = P(\theta > 0 | \mathbf{N}_{crit}) - P(\theta < 0 | \mathbf{N}_{crit})$, the difference between the probability of $\theta_0 > 0$ versus $\theta_0 < 0$. The decision process is then modeled as a hypothetical particle linearly rising to reach a second boundary, starting from a position that is proportional to D . D fluctuates across trials (shaded red) producing additional variation in the reaction time.

spiking statistics. In the first stage, neural activities are not pooled across the population immediately, but rather are accumulated for each individual neuron over time (Beck et al., 2008). More specifically, for a stimulus (C, θ_0) , the first stage obtains a vector of summed population neural activities $\mathbf{N}_t = (N_{1,t}, \dots, N_{M,t})$:

$$N_{i,t} = \sum_{j=1}^t n_{i,j}. \tag{8}$$

The accumulation terminates when the width (SD) of the posterior distribution $p(\theta | \mathbf{N}_t)$ (computed using Eq. 3) shrinks below a fixed threshold. The top portion of Figure 7 illustrates this process.

Stage two: judging whether $\theta_0 > 0$ based on the posterior distribution

The aim of the second-stage process is to introduce a secondary RT component that depends on the significance of the evidence given the estimated posterior distribution $p(\theta | \mathbf{N}_t)$. A simple process that can generate such an evidence-dependent RT component is a particle linearly rising to a boundary. The boundary height is fixed, and the starting position of the particle is proportional to the magnitude of the evidence. That is, given the estimated posterior distribution, to judge whether $\theta_0 > 0$, the brain pools neural activities across the population to compute the decision variable $D = p(\theta_0 > 0 | \mathbf{N}_t) - p(\theta_0 < 0 | \mathbf{N}_t)$. The travel time of the hypothetical particle in stage two is then given by the following:

$$RT \propto B - k \cdot |D|, \tag{9}$$

where B is the fixed boundary height in the second stage and k is a constant independent of D . We use the bounded quantity D rather than the log likelihood ratio, which is unbounded. In a single trial, the discrimination evidence D is constant. However, the estimated posterior distribution varies across trials and thus D also varies across trials, introducing further variation of overall reaction time. Note that the second stage does not produce response errors. All response errors are attributed to the first stage, i.e., if the posterior distribution happens to center on the wrong side. The analogy of a linearly rising particle for stage two follows the convention of a similar decision model (Carpenter and Williams, 1995; Reddi and Carpenter, 2000; Reddi et al., 2003; Carpenter, 2004; see Discussion).

The composition of two sequential processing stages is reminiscent of Carpenter’s Linear Approach to Threshold with Ergodic Rate (LATER) model of saccadic latency (Carpenter, 2004).

In the LATER model, a decision signal rises linearly from threshold for initiating the response and with a rate of rise on different trials that follows a normal distribution. We compare the two models in the Discussion.

Effect of C and θ_0 on the reaction time of correct responses

The overall reaction time is the sum of the time taken to complete each stage. If the time to complete one stage depends primarily on C and the other on θ_0 , then the effects of these two variables will not interact. Stage one terminates based on a threshold on the width of $p(\theta | \mathbf{N}_t)$. The time taken to reach that width is a function of C and is independent of stimulus value θ_0 for an isotropic neural population (Eq. 3). The time taken by stage two is a function of the likelihood ratio D . D is effectively a signal-to-noise ratio; the signal depends on θ_0 , and the noise is fixed by the threshold on the width of the posterior distribution and thus is independent of C . Thus, additivity (Figs. 3–5, parallel curves, absence of interaction) is consistent with the ETD model: stimulus strength primarily affects the duration of the first (estimation) stage whereas stimulus value primarily affects the duration of the second (decision) stage.

But, if ETD predicts absence of interaction, how does it explain the obvious interaction observed in S4’s data in the orientation-discrimination task of Experiment 1b? In fact, an interaction can happen if the subject is able to flexibly adjust the precision criteria σ according to the estimated angle $\hat{\theta}$ given $p(\theta | \mathbf{N}_t)$. The boundary might change, following the changing mode of the posterior. More specifically, for a $\hat{\theta}$ that is close to vertical, a higher precision (smaller σ) is adopted. For a $\hat{\theta}$ that is far from vertical, a low precision (large σ) is adopted. With this additional flexibility, ETD can produce an interaction as found in Figure 5 (top right).

It is important to consider whether flexible control of the boundary height in the drift-diffusion model will allow it to be consistent with the parallel RT functions that dominate Figures 3–5. This requires the DDM to raise the boundary for values of $\hat{\theta}$ that are far from the reference relative to $\hat{\theta}$ closer to the reference, which is counterintuitive. Second, to do so, it must first estimate the presented orientation rather than merely computing the log likelihood ratio, complicating the model. Finally, for the DDM to produce parallel RT curves requires specific values of boundary placement for each estimated orientation, whereas ETD produces parallel curves by default. Our data suggest that parallel curves occur most of the time. For these reasons, we find the flexible DDM a less plausible explanation.

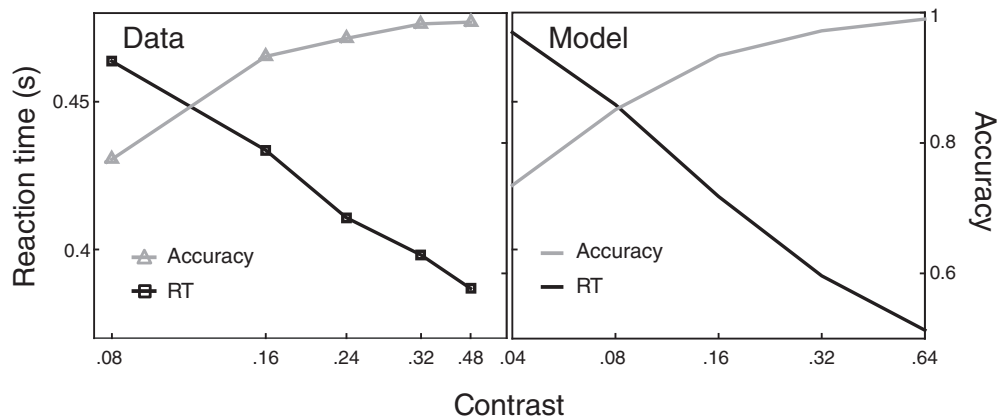


Figure 8. The ETD model and speed versus accuracy. *A*, Subject S3's data for orientation-discrimination for $\theta = \pm 12^\circ$. *B*, Simulation results for the ETD model.




Figure 9. The ETD model and RT distributions. The three left columns show histograms of three subjects' RT in the orientation-discrimination task of Experiment 1a. The top and bottom rows correspond to the lowest and highest contrasts used in the experiment for each subject. The right column shows the histogram of simulated RT for the ETD model based on parameters that put RT roughly in the same range as S3's data.

Speed versus accuracy, RT distribution and incorrect-trial RT

The ETD model is broadly consistent with many well known characteristics of human RT data reported in the literature. Here we focus on three common characteristics.

Speed and accuracy are typically linked in reaction time experiments. Difficult trials lead to slow reaction times and less accurate responses, whereas easy trials lead to fast reaction times and more accurate responses. The ETD model predicts this behavior. For stimuli of low contrast or coherence, the evolution of posterior distribution takes longer to achieve the target precision. Despite the integration of more stimulus information, the posterior distribution is still more likely to appear on the wrong side of the reference for more difficult trials. Figure 8 shows an example of this phenomenon for S3 in the orientation-discrimination task for $\theta = \pm 12^\circ$ and corresponding simulation results for the ETD model. This condition is chosen for illustration because the accuracy range is the largest. The accuracy ranges for other subjects and conditions are too small to clearly show this effect. The parameters for the ETD model were chosen so that mean RT fell roughly in the same range as the data, and ETD clearly shows the same pattern of RT and response accuracy.

In RT experiments, the distribution of RT is asymmetric with a long tail for longer RTs (Luce, 1986). More importantly, the variance of the RT distribution tends to increase with the mean. These characteristics are true of our experimental data and consistent with simulations of the ETD model (Fig. 9).

Finally, we focus on the difference between RTs for trials with correct and incorrect responses. Incorrect RTs are typically

slower than correct RTs, although not always (Luce, 1986). This pattern is also observed in our experiments (Fig. 10). However, the standard DDM predicts identical RTs for correct and incorrect responses (Ratcliff, 1978; Shadlen et al., 2006; Ratcliff and McKoon, 2008). Likewise, the population-code, drift-diffusion model developed above predicts equal correct and incorrect RT distributions (Fig. 10).

We are aware of three possible remedies for the standard DDM to generate slower RTs for incorrect responses: (1) the fixed drift rate in the standard DDM varies randomly across trials (Ratcliff and McKoon, 2008); (2) the noise term of the diffusion process is non-Gaussian (Shadlen et al., 2006); or (3) the decision variable is modulated by a gain function that increases monotonically with time (Ditterich, 2006). However, there is no obvious modification of the population-code, drift-diffusion model equivalent to solution (1) or (2) to generate slower incorrect RTs. A direct application of solution (3) to our population-code, drift-diffusion model produces slower error-trial RTs. However, in our hands it also produces predictions that are inconsistent with other aspects of human behavior such as error-trial RTs that are independent of stimulus strength (simulation data not shown). Therefore, in addition to parallel RT curves, slower incorrect RTs is another feature of the human data that the population-code, drift-diffusion model cannot explain.

On the other hand, the ETD model predicts slower incorrect RTs (Fig. 10). The intuition is as follows. The ETD produces error responses when the posterior distribution happens to fall on the wrong side of the reference. On average, the distance between the




Figure 10. Mean RT for correct and incorrect trials. The first three panels show data for the orientation-discrimination task of Experiment 1a. Only RTs for the lowest contrast and the smallest orientation offset are plotted. Right panels show mean RTs predicted by the DDM and ETD models, based on parameters that put the RTs roughly in the same range as S3's data.

posterior distribution and the reference is smaller when it is on the wrong side than when it is on the correct side. A smaller distance results in a longer duration of the second stage of ETD (smaller D in Eq. 9). Therefore, on average, incorrect responses have larger RTs.

Leaky accumulation in the estimation stage

The current implementation of the ETD model assumes perfect accumulation of spike counts in the estimation stage. However, leaky accumulation is probably more realistic. This could be achieved by convolving the accumulated spike counts with an exponential function that decays with time. In so far as the estimation and decision processes are separate and take place sequentially, an ETD model with leaky accumulation in the estimation stage will still produce parallel RT curves.

Specification of the decision stage

The decision stage represents a process that operates based on accumulated sensory information. Its operation involves deriving a decision plan and implementing the derived plan. Specifying the exact computation of the decision stage is beyond the scope of the current study. What is critical is that its duration depends on sensory information accumulated during the estimation stage. Here, this dependency is modeled as a linear function (Eq. 9). Other forms of this relationship are possible (e.g., division rather than subtraction), and we do not claim superiority of Eq. 9 over other alternatives.

Experiment 2: reaction time in a cued-response, direction-discrimination task

The results for the direction-discrimination task in Experiment 1 are particularly interesting. Our interpretation goes against a rich body of literature that models the perceptual decision processes in motion-discrimination tasks as a diffusion process similar to the one that is rejected here (Ratcliff and Rouder, 1998; Palmer et al., 2005; Kiani et al., 2008; Drugowitsch et al., 2012; Mulder et al., 2012). The motion-discrimination task typically used in those studies is a special case of the one used in the current experiment, so it seems reasonable to expect that the underlying perceptual decision process in our task is the same as in those studies. However, one might argue that stimulus conditions used in those experiments differ from ours. In a typical experimental setting, only two opposite motion directions (e.g., left vs right) are tested, whereas 10 directions were interleaved in our direction-discrimination task. Perhaps the inclusion of so many possible motion directions triggered a different strategy than when only two directions are used. Although we find it difficult to articulate a reason for such a strategy shift, we next test the model using a two-alternative, cued-response task.

In a cued-response task, subjects do not respond until a “go” signal is given. The cue time is usually defined as the time between the onset of the stimulus and the onset of the go signal, the stimulus onset asynchrony. In this task, the estimation stage is terminated by the cue. SOA determines the length of time during which sensory evidence is accumulated. When the first-stage estimation process is terminated by a response cue, the evidence (effectively a signal-to-noise ratio) is passed on to the second-stage decision process. RT is entirely a function of the decision stage. Shorter SOAs lead to reduced evidence and hence a flatter posterior distribution over, for example, motion direction. This leads to a weaker decision signal D for the decision stage and consequently a longer reaction time after the cue. Weaker stimulus strength or smaller stimulus value also result in a weaker decision signal D and therefore longer RT. In summary, in a cued-response task, the ETD model predicts that RT should decrease with increasing stimulus strength, with increasing stimulus value, and with increasing SOA.

On the other hand, the single-stage drift-diffusion model predicts constant reaction times for different SOAs and stimulus strengths and values. A typical diffusion model is composed of a diffusion process and a nondecision component that contributes an additive term to the overall reaction time. Importantly, this additive term is independent of the diffusion process and of stimulus properties (Palmer et al., 2005). The cued-response task terminates the diffusion process so that RT is determined only by nondecision time and hence should be independent of SOA and stimulus value. The decision maker only needs to respond with the sign of the current position of the particle.

Note that for the above logic to work, the response cue should not be too early. For direction discrimination, nondecision time is typically ~ 300 ms, which includes stimulus encoding time and motor latency (Palmer et al., 2005). Monkey physiological data suggest that monkey MT activity peaks and stabilizes 100–200 ms after motion onset, independent of stimulus strength (Gold and Shadlen, 2007). MT neurons are thought to provide the encoded sensory input to the brain region that performs the decision process. At the same time, there is a latency to process the response cue itself. It seems reasonable to avoid response-cue SOAs shorter than 100 ms to ensure that the response cue, once processed, occurs after the latency to process the initial stimulus information so that evidence accumulation has already begun.

In Experiment 2, RT was measured as a function of SOA and coherence in a direction-discrimination task in which response time was cued. Tested SOAs range between 120 and 800 ms, presumably covering a large section of the decision process according to fit drift-diffusion model parameters reported in the literature. In Figure 11, the top row shows mean reaction time as

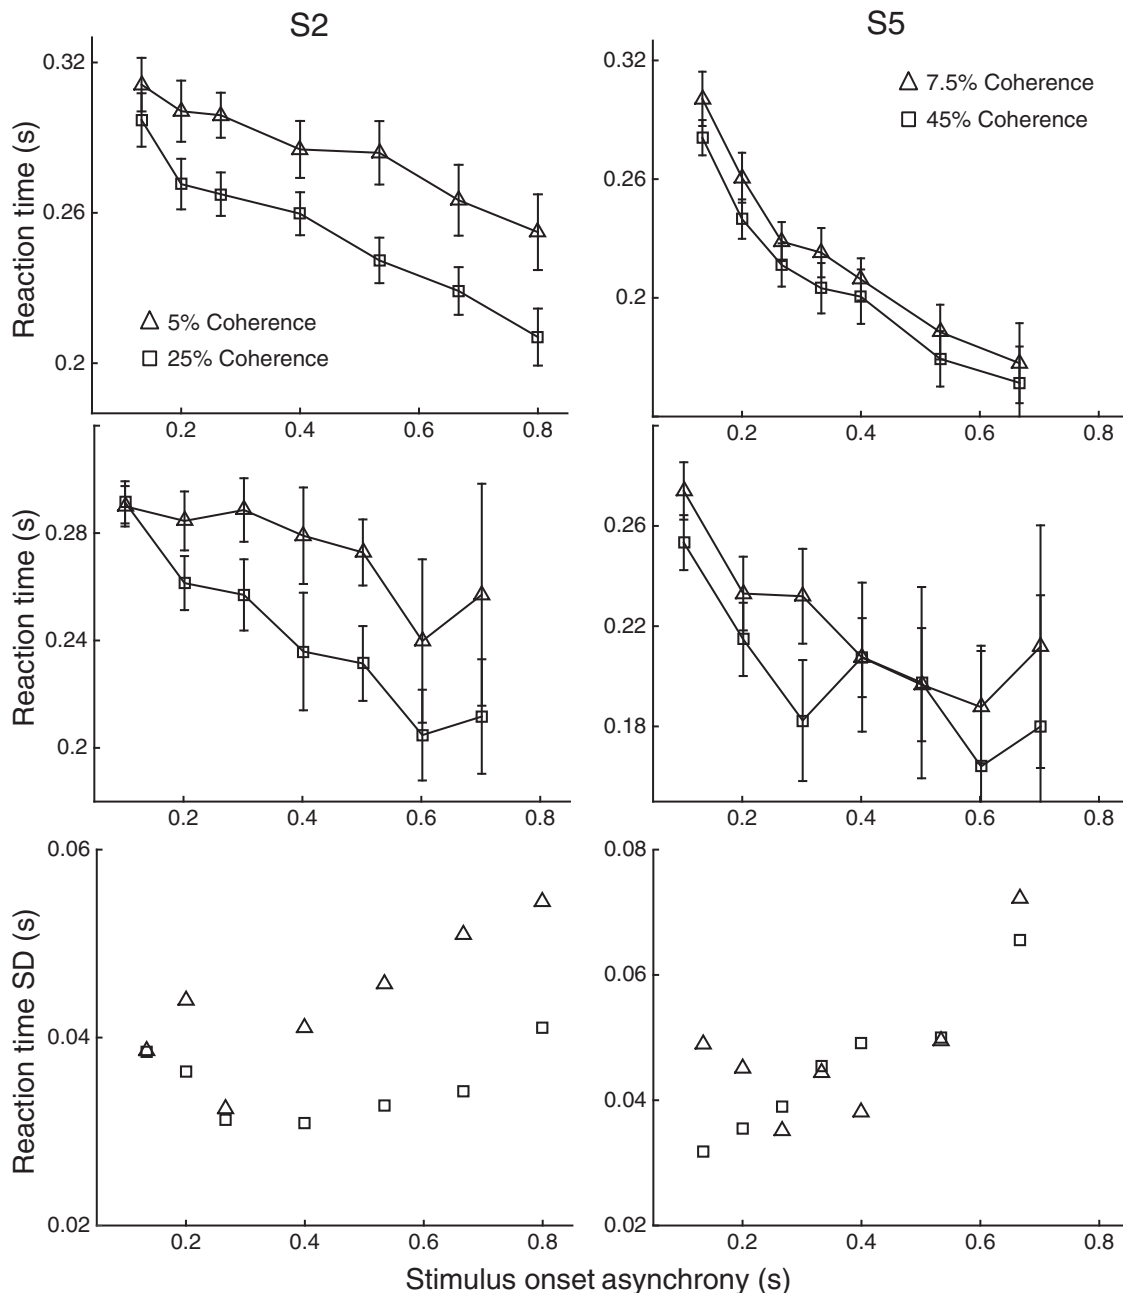


Figure 11. Top row, Mean reaction time in the cued-response, direction-discrimination task for two subjects. Different SOAs occurred equally often. The difference in RT between the shortest and the longest SOA is significant for both subjects ($p < 0.0001$) as is the difference between the high and low coherences for both subjects ($p < 0.01$). Middle row, The mean of binned RT data points (every 100 ms) in the control experiment where occurrence of different SOAs followed a truncated exponential distribution. The difference in RT between the shortest and the longest SOA is significant for both subjects ($p < 0.0001$) as is the difference between the high and low coherences for both subjects ($p < 0.01$ and $p = 0.03$). Bottom row, SD of RT.

a function of SOA for two coherence levels. Consistent with the two-stage model, RT decreases with increasing SOA and increasing coherence. This pattern of relationships remained intact even after the hazard rate was approximately equated across SOA in the control experiment (Fig. 11, middle row).

RT varies systematically with SOA and motion coherence as predicted by the ETD model and is inconsistent with the single-stage drift-diffusion model. However, there are possible modifications of the single-stage model that would be consistent with these results, which we discuss in turn.

Postcue evidence accumulation

The cued-response experiment is intended to tap into the instantaneous state of the decision process at the time of the cue. How-

ever, although the stimulus has disappeared after cue onset, residual neural signals in the “pipeline” may continue to be integrated until a boundary is reached (Kiani et al., 2008). Although a very narrow RT window was imposed to prevent the postcue evidence accumulation, it is hard to rule out this possibility completely. For short SOA or weak stimuli, the particle in the DDM is more likely to be far from the boundary at the time of the cue. If evidence continues to be accumulated after the cue, it will take longer to reach the desired boundary for particles that are further away. That is, RT should decrease with increasing coherence and SOAs, exactly as observed here.

However, this account of postcue evidence accumulation for the DDM makes a different prediction for the variance of RT than

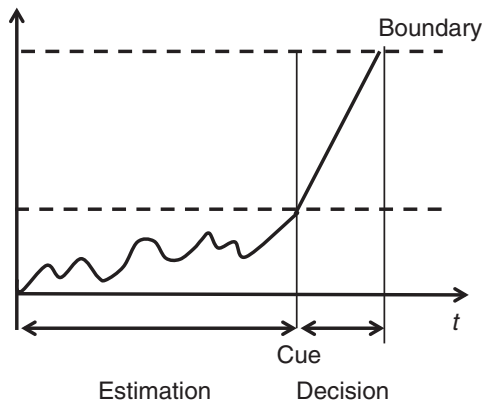


Figure 12. Simplified ETD model for a two-alternative decision process in a cued-response experiment. The estimation stage is abstracted as a diffusion process for which the particle's height represents the strength of accumulated evidence (the precision of the current estimate). In the decision stage, the particle continues to rise linearly after the cue to a fixed boundary. This structure is consistent with Equation 9.

the ETD model. The cue essentially divides the process into two drift-diffusion processes. RT is the duration of the second post-cue drift-diffusion process. Therefore, like standard DDM, its variance should increase with its mean.

On the other hand, the ETD model predicts that, in a cued-response experiment, the variance of reaction time should increase with SOA, not with mean RT. To provide an intuition for this, consider a simplified ETD model that consists of a standard diffusion process followed by a ballistic process in which the particle continues to linearly rise to a fixed boundary (Fig. 12). This simplified ETD bears some resemblance to the LATER model (Carpenter et al., 2009). It is also consistent with the fact that, during a perceptual decision-making process, average LIP neural activities initially rise at different rates for stimuli of different coherences, but their rates converge toward the end of the process (Gold and Shadlen, 2007). Recall that for ETD, RT in the cued-response task depends on the evidence strength at the time of the cue. Therefore, the variance of RT increases with the variance of the evidence strength, i.e., the variance of the magnitude of D in Equation 9. Since D is the sum of a series of independent random variables over time, its variance increases with SOA. Therefore, the RT in a cued-response task will increase with SOA. This is what we find in the data. Note that RT variance appears to decrease slightly with increasing SOA for the shortest SOAs (near 200 ms). This initial dip may correspond to the incomplete encoding process before the decision process. For monkeys, the decision process has been found to occur at least 200 ms after stimulus onset (Gold and Shadlen, 2007; Kiani et al., 2008).

Note that the simplified ETD model predicts equal variance for strong- and weak-coherence stimuli. S5's data appear to be consistent with this prediction, but S1's data do not: they show a larger variance for weaker stimuli. However, it is easy to make the simplified ETD model produce a larger RT variance for weaker stimuli in a cued-response task by increasing the noise SD for weaker stimuli.

Divided attention

The time required to respond to the auditory response cue may depend on the amount of attention allocated to monitoring auditory input. Accumulating visual sensory evidence could limit the resources available for monitoring the auditory cue. For long SOAs, the accumulation process may terminate before the cue, allowing the subject to allocate additional resources to monitor

and respond to the upcoming cue. Such a model would make two predictions. First, it predicts a constant RT for short SOAs (shorter than any possible termination of the diffusion process) and a sudden change to decreasing RT once the SOA is long enough to allow the diffusion process to reach the boundary for some trials. The sudden change in mean RT should be more obvious for stimuli of strong coherence than for stimuli of weak coherence. This is because the boundary crossing time for strong stimuli will be less variable; therefore, the transition should be sharper. Second, the RT distribution in general should be bimodal, with one mode corresponding to the case where evidence accumulation has not ended before the cue and RTs are long, and the other when it has ended and RTs are short. For very short SOAs, the RT distribution is dominated by the slow mode. For very long SOAs, the RT distribution is dominated by the fast mode. As the SOA increases from short to long, the RT distribution should transition from the slow to the fast mode. As a result, the variance of RT should rise and then fall. In examining the data, however, we find none of these predictions hold. The improvement in RT with SOA is gradual, taking place as early as 200 ms and continuing all the way to SOAs of 800 ms (Fig. 11). The shape of this improvement does not differ between weak and strong stimuli. The RT distributions do not appear bimodal, and RT variance by and large increases with increasing SOA. These observations argue against this account of shared processing resources.

A motor component sensitive to the strength of accumulated evidence

The nondecision component after the diffusion process may incorporate a mechanism that depends on the degree of belief resulting from the preceding diffusion process, as implied by Joo et al. (2016). This way RT will depend on SOA and coherence in the same way as observed here. However, with this modification, the single-stage diffusion model essentially becomes a two-stage model. In fact, to produce the pattern of RTs observed here, the decision process must comprise two separate stages: one deals with sensory information accumulation and the other deals with decision-making based on the information that has been accumulated.

To conclude, the drift-diffusion model and ETD model make distinct predictions about RT in a cued-response task. Our results favor the account of the ETD model for the decision-making process underlying typical visual direction-discrimination tasks where only two motion directions are to be discriminated.

Discussion

We have provided evidence against the drift-diffusion model of decision-making. First, in three visual discrimination tasks, we find that the influences of stimulus value and strength on RT are mostly additive, whereas drift-diffusion predicts an interaction (because RT should only depend on the signal-to-noise ratio). Second, we find that RT in a cued-response task varies with both SOA and stimulus strength. For drift-diffusion, the decision required at the time of the cue is merely to report the sign of the current accumulated evidence; RT need not be affected by stimulus value or SOA. We introduce a two-stage ETD model consistent with both findings.

The drift-diffusion model

The drift-diffusion model has been suggested widely as a model for decision-making in the brain. For example, in a direction-discrimination task, neural responses in cortical area LIP typi-

cally increase throughout the trial until reaching a common firing rate when the animal initiates the saccade that indicates its decision. Thus, the accumulation process might be identified with the pooled firing of neurons in LIP (Platt and Glimcher, 1999; Gold and Shadlen, 2001; Shadlen and Newsome, 2001; Roitman and Shadlen, 2002; Yang and Shadlen, 2007; Kiani et al., 2008; de Lafuente et al., 2015; Shadlen et al., 2016). The instantaneous firing rate in LIP should reflect the decision outcome. However, there are alternative accounts of the role of LIP in decision-making (Meister et al., 2013; Park et al., 2014; Latimer et al., 2015). In particular, Park et al. (2014) found that individual LIP spike trains encoded multiple factors of a given trial (fixation point, appearance of saccadic targets, noisy motion stimulus, saccadic response) along with interspike correlations. Although accumulated evidence toward a decision could be decoded from LIP activity, that activity reflected all of these components combined.

Other problems with the drift-diffusion model have been noted previously. The model seems “optimal” in the sense that it accumulates sensory evidence in a manner (as summed log likelihood ratios) that is consistent with an ideal Bayesian computation. Coupled with a starting point for the particle corresponding to prior log-odds, the particle’s position may be interpreted as the current log-posterior odds. Comparing log-posterior odds with a fixed threshold is appropriate if by optimal the intention is to maximize the expected gain from the current trial. However, animals do not always appear to optimize the reward in each trial. In some situations, they spend less time on difficult trials as if they are trying to optimize rewards per minute rather than rewards per trial (Drugowitsch et al., 2012, 2014, 2015). This problem has been addressed by suggesting that there is a cost of time (Hanks et al., 2011; Thura et al., 2012) or by imposing bounds that collapse as the trial progresses (Hanks et al., 2011). In effect, the diffusion process rapidly “forgets” the initial prior, and yet if the diffusion takes too long, it is preferable to invoke the prior and move on to the next trial. To complicate matters further, a recent study shows that subjects may not even optimize rewards per minute; they spend far too much time on difficult trials for which reward was deliberately made small (Oud et al., 2016).

The drift-diffusion model is also tied to tasks in which there are only two opposing options. Extending the two-alternative drift-diffusion model to explain the multialternative decision process remains a challenge. An obvious solution is to introduce a diffusion process for each possible choice and appropriate interactions between them (Niwa and Ditterich, 2008; Ratcliff and Starns, 2013). However, the ETD model begins with stimulus estimation, making it a simple matter to extend the model to more complex tasks by using those estimates to make the subsequent decision. In real-life situations, estimation of the visual world occurs generically and ubiquitously; we are always estimating world properties. However, decision-making is task specific. It makes sense to have a general-use estimation stage and a task-specific decision stage that follows it.

Implications for perceptual decision-making

The most provocative claim in this study is our rejection of the drift-diffusion model of perceptual decision-making compared with the proposed ETD model. This calls into question what exactly is accumulated during a simple perceptual decision process such as motion direction discrimination. The drift-diffusion model accumulates evidence for “which” action should be taken; the ETD model accumulates evidence for “what” stimulus was presented. However, the success of the drift-diffusion model

should not be overlooked. For the ETD model to provide a serious alternative to drift-diffusion, one must show that the ETD model can account for all the known characteristics of perceptual decision-making consistent with the predictions of drift-diffusion. In this study, we tested three major characteristics of perceptual decision-making and showed that the ETD model is consistent with all three without introducing any further complications to the model. Exhausting the entire list, however, is beyond the scope of the current study. However, we are optimistic based on the success of a different but related model, LATER (Carpenter and Williams, 1995; Reddi and Carpenter, 2000; Reddi et al., 2003; Carpenter, 2004).

The LATER model

In LATER, on each trial a particle rises at fixed speed to a boundary; the travel time constitutes the model’s prediction of RT. The speed of particle travel varies randomly from trial to trial after a normal distribution. The second stage of our two-stage model is very similar to the LATER model. The LATER model was originally proposed to characterize saccade latency after target onset. For medium to high target contrasts, LATER predicts the saccade latency distribution well (Carpenter and Williams, 1995). Studies and commentaries have suggested that the LATER model predicts the RT distribution under a wide range of testing conditions (Reddi et al., 2003) and that the LATER model and the drift-diffusion model “show signs of convergence” (Ratcliff, 2001).

The ETD model versus LATER

The main criticism of LATER is that although it has a stochastic latency component (predicting RT distributions), it does not have a stochastic response component (predicting response errors and their covariation with RT). To explain response errors, the LATER model has been elaborated to include a diffusion process preceding the linear rise stage (Carpenter et al., 2009). This diffusion process represents a detection stage that integrates noisy signals over time. The detection stage may falsely detect a signal (false alarm) and trigger an incorrect response.

The ETD model is similar in structure to the elaborated LATER model: its first stage is a stochastic process, and its second stage is a constant-speed rise to threshold. However, the ETD model differs from the elaborated LATER model in some notable ways. First, the ETD model explicitly spells out the mechanism of the preceding stochastic process. In the LATER model, the first-stage diffusion process is meant to represent the process of detecting a signal, but its physiological interpretation is unclear. Second, the LATER model does not explain the origin of the trial-by-trial variability of the second stage. In the ETD model, the variability results from the variability in the standard deviation and location of the posterior distribution over stimulus value estimated by the first-stage process. Third, in the elaborated LATER model, the mean particle speed is fixed across stimulus strength, therefore the mean duration of the second stage is independent of stimulus properties (Carpenter et al., 2009). In the ETD model, the speed is fixed but the starting location is proportional to the accumulated sensory evidence, thus the duration depends on a range of factors, including stimulus strength, stimulus value, and duration of accumulation.

Bayesian inference

In deriving predictions of a population-code drift-diffusion model, we assumed that the brain extracted all of the information contained in the population response via Bayesian inference. Would a suboptimal drift-diffusion model yield a qualitatively

different prediction? We expect not, but obviously there is a huge class of possible suboptimal drift-diffusion models. We have implemented one such model in which the pooled activity of all neurons tuned clockwise of the reference angle was compared with that of neurons tuned counterclockwise. Again, an interaction of stimulus value and strength on RT was found (as in Fig. 2). In fact, it is the pooling of neural activities across the population that effectively combined the impact of stimulus strength (contrast) and value (stimulus orientation) into a single factor contributing to reaction time in the drift-diffusion model. In so far as momentary evidence pools neurons across population, an interaction of stimulus value and strength appears likely in the drift-diffusion model.

Conclusion

To conclude, we have provided evidence against the drift-diffusion model for visual discrimination. We proposed a two-stage estimate-then-decide model as an alternative. In stage one, the brain accumulates sensory evidence until a reliable estimate of the stimulus value is obtained. Based on this estimate, a decision is made in stage two. Stage two takes a variable length of time, depending on the estimated value and the uncertainty of the estimate at the time of the decision. We offer this two-stage model as an alternative account of perceptual decision-making.

References

- Beck JM, Ma WJ, Kiani R, Hanks TD, Churchland AK, Roitman JD, Shadlen MN, Latham PE, Pouget A (2008) Probabilistic population codes for Bayesian decision making. *Neuron* 60:1142–1152. [CrossRef Medline](#)
- Bitzer S, Park H, Blankenburg F, Kiebel JS (2014) Perceptual decision making: drift-diffusion model is equivalent to a Bayesian model. *Front Hum Neurosci* 8:102. [CrossRef Medline](#)
- Bogacz R, Brown E, Moehlis J, Holmes P, Cohen JD (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. [CrossRef Medline](#)
- Brainard DH (1997) The Psychophysics Toolbox. *Spat Vis* 10:433–436. [Medline](#)
- Carpenter RH (2004) Contrast, probability, and saccadic latency: evidence for independence of detection and decision. *Curr Biol* 14:1576–1580. [Medline](#)
- Carpenter RH, Williams ML (1995) Neural computation of log likelihood in control of saccadic eye movements. *Nature* 377:59–62. [CrossRef Medline](#)
- Carpenter RH, Reddi BA, Anderson AJ (2009) A simple two-stage model predicts response time distributions. *J Neurophysiol* 587:4051–4062. [CrossRef Medline](#)
- de Lafuente V, Jazayeri M, Shadlen MN (2015) Representation of accumulating evidence for a decision in two parietal areas. *J Neurosci* 35:4306–4318. [CrossRef Medline](#)
- Ding L, Gold JI (2012) Neural correlates of perceptual decision making before, during, and after decision commitment in monkey frontal eye field. *Cereb Cortex* 22:1052–1067. [CrossRef Medline](#)
- Ditterich J (2006) Evidence for time-variant decision making. *Eur J Neurosci* 24:3628–3641. [Medline](#)
- Drugowitsch J, Moreno-Bote R, Churchland AK, Shadlen MN, Pouget A (2012) The cost of accumulating evidence in perceptual decision making. *J Neurosci* 32:3612–3628. [CrossRef Medline](#)
- Drugowitsch J, DeAngelis GC, Klier EM, Angelaki DE, Pouget A (2014) Optimal multisensory decision-making in a reaction-time task. *eLife* 3. [CrossRef Medline](#)
- Drugowitsch J, DeAngelis GC, Angelaki DE, Pouget A (2015) Tuning the speed-accuracy trade-off maximize reward rate in multisensory decision-making. *eLife* 4:e06678. [CrossRef Medline](#)
- Forstmann BU, Ratcliff R, Wagenmakers EJ (2016) Sequential sampling models in cognitive neuroscience: advantages, applications and extensions. *Annu Rev Psychol* 67:641–666. [CrossRef Medline](#)
- Gold JI, Shadlen MN (2000) Representation of a perceptual decision in developing oculomotor commands. *Nature* 404:390–394. [Medline](#)
- Gold JI, Shadlen MN (2001) Neural computations that underlie decisions about sensory stimuli. *Trends Cogn Sci* 5:10–16. [Medline](#)
- Gold JI, Shadlen MN (2003) The influence of behavioral context on the representation of a perceptual decision in developing oculomotor commands. *J Neurosci* 23:632–651. [Medline](#)
- Gold JI, Shadlen MN (2007) The neural basis of decision making. *Annu Rev Neurosci* 30:535–574. [Medline](#)
- Green DM, Swets JA (1988) Signal detection theory and psychophysics. Los Altos, CA: Peninsula.
- Guest D, Howard CJ, Brown LA, Gleeson H (2015) Aging and the rate of visual information processing. *J Vis* 15:10. [CrossRef Medline](#)
- Hanks TD, Mazurek ME, Kiani R, Hopp E, Shadlen MN (2011) Elapsed decision time affects the weighting of prior probability in a perceptual decision task. *J Neurosci* 31:6339–6352. [CrossRef Medline](#)
- Huk AC, Shadlen MN (2005) Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. *J Neurosci* 25:10420–10436. [Medline](#)
- Joo SJ, Katz LN, Huk AC (2016) Decision-related perturbations of decision-irrelevant eye movements. *Proc Natl Acad Sci U S A* 113:1925–1930. [CrossRef Medline](#)
- Kiani R, Hanks TD, Shadlen MN (2008) Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. *J Neurosci* 28:3017–3029. [CrossRef Medline](#)
- Kira S, Yang T, Shadlen MN (2015) A neural implementation of Wald's sequential probability ratio test. *Neuron* 85:861–873. [CrossRef Medline](#)
- Krajbich I, 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. [CrossRef Medline](#)
- Krajbich I, Armel C, Rangel A (2011) Visual fixations and the computation and comparison of value in simple choice. *Nat Neurosci* 13:1292–1298. [CrossRef Medline](#)
- Latimer KW, Yates JL, Meister MLR, Huk AC, Pillow JW (2015) Single-trial spike trains in parietal cortex reveal discrete steps during decision-making. *Science* 349:184–187. [CrossRef Medline](#)
- Luce RD (1986) Response times: their role in inferring elementary mental organization. New York: Oxford UP.
- Ma WJ, Beck JM, Latham PE, Pouget A (2006) Bayesian inference with probabilistic population codes. *Nat Neurosci* 9:1432–1438. [Medline](#)
- McClelland JL (1979) On the time relations of mental processes: an examination of systems of processes in cascade. *Psychol Rev* 86:287–330. [CrossRef](#)
- Meister MLR, Henning JA, Huk AC (2013) Singal multiplexing and single-neuron computations in lateral intraparietal area during decision-making. *J Neurosci* 33:2254–2267. [CrossRef](#)
- Moran R (2015) Optimal decision making in heterogeneous and biased environments. *Psychon Bull Rev* 22:38–53. [CrossRef](#)
- Mulder MJ, Wagenmakers EJ, Ratcliff R, Boekel W, Forstmann BU (2012) Bias in the brain: a diffusion model analysis of prior probability and potential payoff. *J Neurosci* 32:2335–2343. [CrossRef Medline](#)
- Niwa M, Ditterich J (2008) Perceptual decisions between multiple directions of visual motion. *J Neurosci* 28:4435–4445. [CrossRef Medline](#)
- Oud B, Krajbich I, Miller K, Cheong JH, Botvinick M, Fehr E (2016) Irrational time allocation in decision-making. *Proc R Soc Lond B Biol Sci* 283:pii: 20151439. [CrossRef Medline](#)
- Palmer J, Huk AC, Shadlen MN (2005) The effect of stimulus strength on the speed and accuracy of a perceptual decision. *J Vis* 5:376–404. [CrossRef Medline](#)
- Park IM, Meister MLR, Huk AC, Pillow JW (2014) Encoding and decoding in parietal cortex during sensorimotor decision-making. *Nat Neurosci* 17:1395–1403. [CrossRef Medline](#)
- Platt ML, Glimcher PW (1999) Neural correlates of decision variables in parietal cortex. *Nature* 400:233–238. [Medline](#)
- Purcell BA, Heitz RP, Cohen H, Schall JD, Logan GD, Palmeri TJ (2010) Neurally constrained modeling of perceptual decision making. *Psychol Rev* 117:1113–1143. [CrossRef Medline](#)
- Ratcliff R (1978) A theory of memory retrieval. *Psychol Rev* 85:59–108.
- Ratcliff R (2001) Putting noise into neurophysiological models of simple decision making. *Nat Neurosci* 4:336–337. [CrossRef Medline](#)
- Ratcliff R, McKoon G (2008) The diffusion decision model: theory and data for two-choice decision tasks. *Neural Comput* 20:873–922. [CrossRef](#)
- Ratcliff R, Rouder JN (1998) Modeling response times for two-choice decisions. *Psychol Sci* 9:347–356.

- Ratcliff R, Starns JJ (2013) Modeling confidence judgments, response times and multiple choices in decision making: recognition memory and motion discrimination. *Psychol Rev* 120:697–719. [CrossRef Medline](#)
- Ratcliff R, Smith PL, Brown SD, McKoon G (2016) Diffusion decision model: current issues and history. *Trends Cogn Sci* 20:260–281. [CrossRef Medline](#)
- Reddi BA, Carpenter RH (2000) The influence of urgency on decision time. *Nat Neurosci* 3:827–830. [Medline](#)
- Reddi BA, Asrress KN, Carpenter RH (2003) Accuracy, information, and response time in a saccadic decision task. *J Neurophysiol* 90:3538–3546. [CrossRef Medline](#)
- Roitman JD, Shadlen MN (2002) Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J Neurosci* 22:9475–9489. [Medline](#)
- Shadlen MN, Newsome WT (2001) Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J Neurophysiol* 86:1916–1936. [Medline](#)
- Shadlen MN, Hanks TD, Churchland AK, Kiani R, Yang T (2006) The speed and accuracy of a simple perceptual decision: a mathematical primer. In: *Probabilistic approaches to neural coding* (Doya K, Ishii S, Pouget A, Rao RPN, eds), pp 207–236. Cambridge, MA: MIT.
- Shadlen MN, Kiani R, Newsome WT, Gold JL, Wolpert DM, Zylberberg A, Ditterich J, de Lafuente V, Yang T, Roitman JD (2016) Comment on “Single-trial spike trains in parietal cortex reveal discrete steps during decision-making.” *Science* 351:1406. [CrossRef Medline](#)
- Simen P, Contreras D (2009) Reward rate optimization in two-alternative decision making: empirical tests of theoretical predictions. *J Exp Psychol Hum Percept Perform* 35:1865–1897. [CrossRef Medline](#)
- Thura D, Beauregard-Racine J, Fradet CW, Cisek P (2012) Decision making by urgency gating: theory and experimental support. *J Neurophysiol* 108:2912–2930. [CrossRef Medline](#)
- Yang T, Shadlen MN (2007) Probabilistic reasoning by neurons. *Nature* 447:1075–1080. [Medline](#)