

Decisions in Changing Conditions: The Urgency-Gating Model

Paul Cisek, Geneviève Aude Puskas, and Stephany El-Murr

Groupe de Recherche sur le Système Nerveux Central, Département de Physiologie, Université de Montréal, Montréal, Québec H3C 3J7, Canada

Several widely accepted models of decision making suggest that, during simple decision tasks, neural activity builds up until a threshold is reached and a decision is made. These models explain error rates and reaction time distributions in a variety of tasks and are supported by neurophysiological studies showing that neural activity in several cortical and subcortical regions gradually builds up at a rate related to task difficulty and reaches a relatively constant level of discharge at a time that predicts movement initiation. The mechanism responsible for this buildup is believed to be related to the temporal integration of sequential samples of sensory information. However, an alternative mechanism that may explain the neural and behavioral data is one in which the buildup of activity is instead attributable to a growing signal related to the urgency to respond, which multiplicatively modulates updated estimates of sensory evidence. These models are difficult to distinguish when, as in previous studies, subjects are presented with constant sensory evidence throughout each trial. To distinguish the models, we presented human subjects with a task in which evidence changed over the course of each trial. Our results are more consistent with “urgency gating” than with temporal integration of sensory samples and suggest a simple mechanism for implementing trade-offs between the speed and accuracy of decisions.

Introduction

Research into the temporal aspects of decision making has provided support for a class of theories that may be called “sequential sampling,” or “bounded integrator” models (Stone, 1960; Laming, 1968; Ratcliff, 1978; Carpenter and Williams, 1995; Usher and McClelland, 2001; Wang, 2002; Mazurek et al., 2003; Reddi et al., 2003; Smith and Ratcliff, 2004; Wong and Wang, 2006; Bogacz and Gurney, 2007; Grossberg and Pilly, 2008). According to these models, simple decisions involve the sequential sampling of the sensory stimulus and temporally integrating the information present in that stimulus until a threshold (or “bound”) is reached, at which point the decision is taken (see Fig. 1A). The rate of this integration is related to the quality of information present in the stimulus in favor of a given choice, and the threshold is related to motivational factors such as payoff, risk, and urgency (Reddi and Carpenter, 2000). Several variations of bounded integrator models exist, and they offer a remarkably successful account of behavior in simple decision tasks. In particular, assuming that the rate of integration is subject to variability, these models can explain error rates and distributions of reaction times (RTs) in a wide variety of tasks (Ratcliff, 1978; Carpenter and Williams, 1995; Reddi and Carpenter, 2000; Reddi et al., 2003; Smith and Ratcliff, 2004). Furthermore, recent neurophys-

iological data on decision tasks have provided evidence for accumulation processes in the superior colliculus (Munoz and Wurtz, 1995; Munoz et al., 2000; Ratcliff et al., 2003, 2007; Shen and Paré, 2007), the lateral intraparietal area (LIP) (Roitman and Shadlen, 2002; Leon and Shadlen, 2003), the frontal eye fields (Gold and Shadlen, 2000, 2003), and the prefrontal cortex (Kim and Shadlen, 1999). Finally, integration of samples toward a bound is reminiscent of the “sequential probability ratio test” (SPRT) (Wald, 1945; Bogacz et al., 2006), an optimal procedure for making decisions on the basis of information that arrives over time (Wald and Wolfowitz, 1948). Consequently, integrator models are now widely accepted as an explanatory mechanism for decision making in many simple tasks.

However, in most of the studies supporting integrator models, the information presented to subjects varied across trials but was constant during the course of each trial. As described below, in such conditions one cannot distinguish whether the buildup observed in neural activity (and inferred from behavioral data) is caused by temporal integration of sensory information or by a growing signal related to elapsed time. It is therefore conceivable that the neural and behavioral data from constant-information tasks can be explained by an alternative model, in which neural activity is the product of current stimulus information and a growing signal related to the urgency for making a choice (see Fig. 1B). Although this model, which we call “urgency gating,” behaves very similarly to integrator models in constant-information tasks, it proposes a very different mechanism to explain neural activity and behavior. To distinguish the models, we presented human subjects with a decision task in which the information favoring one choice over another changed during each trial. Some of these results have previously been presented in abstract form (Puskas et al., 2007; Cisek et al., 2008).

Received April 15, 2009; revised June 30, 2009; accepted July 20, 2009.

This study was supported by the Natural Sciences and Engineering Research Council of Canada, the EJLB Foundation, the Faculty of Medicine summer student program (Université de Montréal), and an infrastructure grant from Fonds de la Recherche en Santé du Québec. We thank Drs. Yoshua Bengio, Andrea Green, John Kalaska, Galen Pickard, and David Thura, and two anonymous reviewers for helpful comments on this work.

Correspondence should be addressed to Dr. Paul Cisek, Département de Physiologie, Université de Montréal, C.P. 6128 Succursale Centre-ville, Montréal, QC H3C 3J7, Canada. E-mail: paul.cisek@umontreal.ca.

DOI:10.1523/JNEUROSCI.1844-09.2009

Copyright © 2009 Society for Neuroscience 0270-6474/09/2911560-12\$15.00/0

Materials and Methods

Modeling formalism

Many variations of integrator models exist, but all may be formalized as follows:

$$x_i(t) = g \int_0^t E_i(\tau) d\tau, \quad (1)$$

where $x_i(t)$ is a putative neural variable corresponding to choice i , g is a scalar gain term, and $E_i(\tau)$ is some internal estimate of the sensory information, present at time τ , that constitutes evidence in favor of choice i . The variable x_i starts at some initial level $x_i(0)$ related to the previous probability that choice i is correct and grows over time at a rate related to the sensory evidence until it hits a threshold T , at which time the subject commits to choice i . All integrator models share this basic set of assumptions. Where they differ is on how they define the evidence function $E_i(t)$ and on how they determine the time at which the decision is taken (the “stopping rule”). For example, “independent race” models (Vickers, 1970) suggest that there exist separate variables $x_i(t)$ for each option, each of which independently accumulates an estimate $E_i(t)$ that is computed solely on the basis of sensory information in favor of that option. Whenever any of these independent processes reaches its threshold, the decision is made in favor of the corresponding option. In contrast, the “diffusion” model (Stone, 1960; Laming, 1968; Ratcliff, 1978; Ratcliff et al., 2003; Smith and Ratcliff, 2004) suggests that there is only a single variable $x(t)$, which accumulates $E(t)$, which is defined as the difference between sensory evidence for option A over option B. In other words, $E(t) = E_A(t) - E_B(t)$. The decision is taken when $x(t)$ either grows above a positive threshold $+T$, selecting option A, or falls below a negative threshold $-T$, selecting option B. Another variant (Shadlen et al., 1996) defines separate variables $x_A(t)$ and $x_B(t)$, each of which integrates independent evidence, and uses a stopping rule based on the difference of accumulated totals. Still another variant, called the “leaky competing accumulator” model (Usher and McClelland, 2001), proposes separate accumulators that mutually inhibit each other. Bogacz et al. (2006) showed that, under reasonable parameter choices, most of these models (with the exception of the independent race model) are functionally equivalent to the diffusion model. Furthermore, the diffusion model is formally equivalent to the SPRT and is very effective at reproducing human behavior. Below, we consider four variations of integrator models, all based on diffusion, using different definitions of what is meant by “sensory evidence.”

In addition to integrator models, we consider a different class of models, which are similar in that there is also a buildup of neural activity toward a threshold, but the buildup is caused by a time-varying gain (Ditterich, 2006), not by temporal integration. We call these urgency-gating models, which can be expressed as follows:

$$x_i(t) = g \cdot E_i(t) \cdot u(t), \quad (2)$$

where $u(t)$ is some function of time that is not related to evidence for any particular choice. This model (Fig. 1B), suggests that neural activity $x_i(t)$ is the product of the momentary evidence $E_i(t)$ and a signal $u(t)$, which reflects the growing urgency to make a response (Ditterich, 2006; Churchland et al., 2008). Although this model is very different from the integrator models (1), the two can be shown to be mathematically equivalent in the case of constant evidence tasks, as follows.

First, note that if $E_i(t)$ is not a function of time, then it can be replaced by a constant E_i within each trial. Because a constant can be moved outside of an integral, Equation 1 can now be rewritten as follows:

$$x_i(t) = g \int_0^t E_i d\tau = g \cdot E_i \int_0^t d\tau = g \cdot E_i \cdot t. \quad (3)$$

Similarly, for the urgency-gating model, we can set $E_i(t)$ to a constant and simply set urgency proportional to elapsed time, $u(t) = t$, and so Equation 2 becomes the following:

$$x_i(t) = g \cdot E_i(t) \cdot u(t) = g \cdot E_i \cdot t. \quad (4)$$

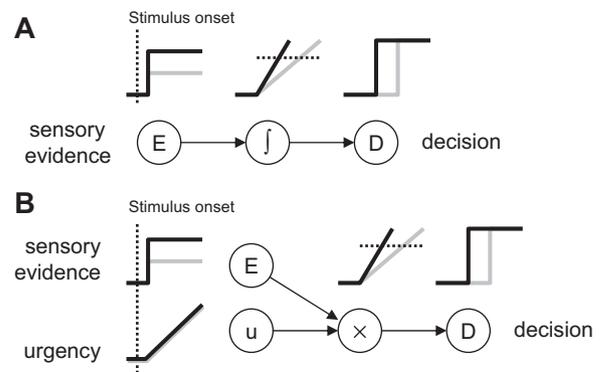


Figure 1. Two alternative models of simple decisions. **A**, The bounded integrator model, in which sensory information (E) is integrated over time (J) and compared with a threshold, resulting in a decision (D). When evidence is strong (black lines), the buildup of activity in the neural integrator is faster and produces a decision earlier in time than when evidence is weaker (gray lines). **B**, The urgency-gating model, in which sensory information is multiplied by a growing signal related to the urgency to respond (u). The product of these two signals causes buildup of activity that, as in **A**, is faster for strong evidence (black) than for weak evidence (gray).

To summarize, with the assumptions of constant evidence and linear urgency growth, the integrator and urgency-gating models are equivalent. Therefore, for any task that presents subjects with constant evidence during each trial, the models make very similar predictions about neural activity and behavior. In both, the neural variable grows at a rate proportional to the subject’s estimate of the strength of evidence, and reaches some threshold level of activity at the time the decision is made. In Equation 1, the growth is attributable to integration of sensory evidence over time (Fig. 1A). In Equation 2, it is attributable to multiplication of momentary evidence by a growing urgency to respond (Fig. 1B). Because both models propose a growth of activity at a rate proportional to evidence, both make similar predictions about how sensory evidence influences neural activity, error rates, and reaction time distributions.

The above derivation of Equation 3 does not address the issue of noise. However, it is clear that noise exists in the sensory signal as well as in the internal processes of sensory transduction and computation of decision variables. The presence of noise is one reason why temporal integration is seen as essential for decision making. However, it is not the only option. A low-pass filter can also deal with noise, without necessarily retaining all properties of pure integration such as a long-lasting memory of past states. Therefore, we propose that the estimate of momentary evidence [$E_i(t)$] in the urgency-gating model is low-pass filtered before it is multiplied by the urgency signal.

In the context of the comparison between temporal integration and urgency gating, it is useful to distinguish two kinds of noise in the decision process: (1) intratrial variability in $E_i(t)$, which is attributable to moment-to-moment neural activity fluctuations; and (2) intertrial variability in E_p , attributable to variations in levels of arousal, attention, etc., which are different from trial to trial but are relatively constant over the course of each trial. Analyses of reaction time distributions using “reciprobit plots” (Carpenter and Williams, 1995) suggest that the primary cause of variability in reaction times is intertrial variability (differences in arousal/attention) and that intratrial noise does not have a major impact at the behavioral level. This makes sense because during each trial, the brain can average across the uncorrelated activity fluctuations of many thousands of neurons (Shadlen et al., 1996), but it cannot average over changes in underlying baselines that vary between trials. Clearly, the presence of intertrial noise in E_i does not affect the derivation of Equation 3 and affects both kinds of models identically. That is not the case for intratrial noise, which results in a time-dependent noise distribution in Equation 3 but a time-independent distribution in Equation 4. Nevertheless, we conjecture that, if intratrial noise is relatively weak (as suggested by reciprobit plot analyses), and reaction times are relatively short, then this difference will be too subtle to distinguish in data from experiments that used constant-evidence tasks.

How then can we more effectively distinguish between the integrator and urgency-gating models? The key to doing so is to use experimental tasks in which the evidence for or against a given choice is changing over the course of an individual trial (Huk and Shadlen, 2005; Kiani et al., 2008). If evidence is changing, then $E_i(t)$ is no longer a constant, cannot be taken outside of the integration, and therefore Equations 1 and 2 are no longer equivalent and now make distinct predictions about both behavioral and neural phenomena. The present experiment is aimed at testing some of these predictions using behavioral data from human subjects.

Experimental design

Twenty-two human subjects (ages, 18–60; 7 males, 15 females; 2 left-handed) performed a reach decision task shown in Figure 2A. All gave informed consent before the experiment, and the procedure was approved by the University of Montréal ethics committee.

Each trial began with a central circle (2.5 cm diameter) and two target circles (2.5 cm diameter) placed 180° apart at a distance of 5 cm from the center. Within the central circle, 15 small circular tokens were randomly arranged. Subjects began each trial by moving the cursor into the central circle. At this point, the tokens began to jump, one-by-one every 200 ms (“predecision interval”), from the central circle to one or the other of the targets (50/50 chance). The subject’s task was to move the cursor to the target that he/she believed would ultimately receive the majority of the tokens.

Importantly, the subject was allowed to make the decision as soon as he/she felt sufficiently confident. Once the choice was reported by moving the cursor into one of the targets, the remaining tokens jumped more quickly to their final targets. In separate blocks, this “postdecision interval” was set to either 20 ms (in “fast” blocks) or to 170 ms (in “slow” blocks). Subjects were asked to continue each block until they made a total of 70 correct choices, indirectly motivating them to optimize successes per unit time. Thus, the subjects were presented with a trade-off: either be conservative and wait until all tokens have moved, when the decision can be made with confidence, or guess ahead of time, which is not as reliable but yields potential successes more quickly. For example, if the first five tokens all move into the same target, then the probability that that target is correct is 94.5%, so it is a good idea to take a guess at that point instead of waiting until the remaining tokens have moved. In fast blocks, such a guess would save the subject a total of 1800 ms over the strategy of waiting until the end, and would have only a 5% chance of failing.

Each subject completed four to six blocks, alternating between fast blocks, which encourage hasty behavior, and slow blocks, which encourage more conservative behavior. Subjects were told about the timing parameters for both blocks ahead of time and allowed to establish their own policies for trading off speed versus accuracy. On each trial, subjects were provided with feedback indicating correct or incorrect decisions, but there was no feedback or instruction regarding the timing of their choices.

The design of the task allowed us to calculate, at each moment in time, the “success probability” $p_i(t)$ associated with choosing each target i (Fig. 2B). If at a particular moment in time the right target contains N_R tokens, whereas the left contains N_L tokens, and there are N_C tokens remaining in the center, then the probability that the target on the right will ultimately

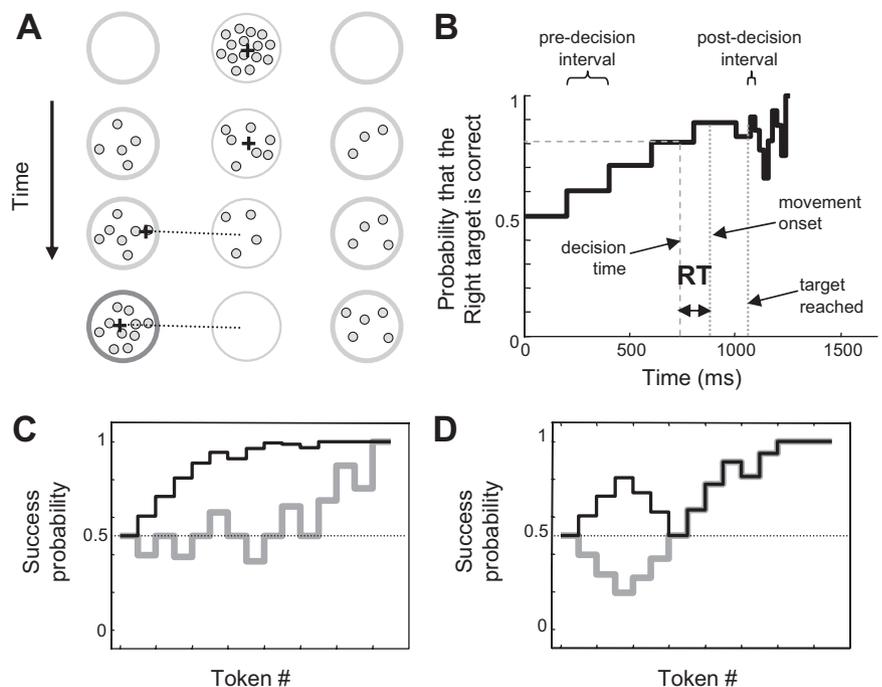


Figure 2. Experimental design. **A**, Behavioral task. Top row, The subject begins each trial by placing the cursor (plus sign) within the central circle. Second row, Next, the tokens begin to move from the center to one of the two targets at a slow speed (predecision interval: 200 ms between token movements). Third row, The subject makes a choice by moving the cursor. Bottom row, The remaining tokens move more quickly to the targets (postdecision interval: 20 ms in fast blocks and 170 ms in slow blocks) and feedback is given to the subject. **B**, The temporal profile (thick black line) of the probability of a given target is computed using Equation 5. The time of the decision (vertical dashed line) is computed by subtracting the subject’s mean reaction time from movement onset time, allowing computation of the success probability at that moment (horizontal dashed line). **C**, Profiles of success probability for easy trials (black line) and ambiguous trials (gray line). **D**, Profiles of success probability for bias-for (black) and bias-against trials (gray).

be the correct one (i.e., the success probability of guessing right) is as follows:

$$p(R|N_R, N_L, N_C) = \frac{N_C!}{2^{N_C}} \sum_{k=0}^{\min(N_C, 7-N_L)} \frac{1}{k!(N_C - k)!} \quad (5)$$

As far as the subjects knew, the correct target and the individual token movements were completely random. However, to test specific hypotheses about the dynamics of decision making, we interspersed among the fully random trials four specific classes of trials characterized by particular temporal profiles of success probability. Subjects were not told about the existence of these trials. For example, 15% of trials were so-called “easy” trials, in which tokens tended to move consistently toward one of the targets, quickly driving the success probability $p_i(t)$ for each toward either 0 or 1. There were several variations of easy trials, and an example is shown in Figure 2C, black line. Another 15% of trials were “ambiguous” (Fig. 2C, gray line), in which the initial token movements were balanced, making the $p_i(t)$ function hover near 0.5 until late in the trial. Another 10% of trials were called “bias-for” trials (Fig. 2D, black line) in which the first three tokens moved to the correct target, the next three toward the opposite target, and the remaining ones resembled an easy trial. Another 10% were called “bias-against” trials (Fig. 2D, gray line), which were identical to bias-for trials except the first six token movements were reversed (i.e., during the early part of bias-against trials, the bias was toward the wrong target). The remaining 50% of trials were fully randomized. Thus, the final distribution of trials was as follows: 50% random, 15% easy, 15% ambiguous, 10% bias-for, and 10% bias-against. In all cases, even when the temporal profile of success probability of a trial was predesigned, the actual correct target was randomly selected on each trial. Over the four to six blocks of trials, each subject performed an average of 554 trials. Unless stated otherwise, analyses included both

correct and error trials, and success probability was computed with respect to the target chosen by the subject.

Before the 15 token task described above, each subject also performed 20–40 trials of a simple choice reaction time task. This task was identical except there was only one token that moved from the center to one of the targets and the subject was instructed to respond as quickly as possible. We detected the time of movement onset and used that to determine each subject's mean RT. This provided us with an estimate of the sum of the delays attributable to sensory processing of the stimulus display as well as to response initiation, muscle contraction, etc. Then, in the 15 token task, we detected the time of movement onset and subtracted each subject's mean RT (from the choice reaction time task) to estimate the time actually used to make the decision—the “decision time”—as shown in Figure 2*B*. We then used Equation 5 to compute the success probability at the time of the decision.

Of course, critical to interpretation of the data from this task is a precise definition of what is meant by sensory evidence. One reasonable definition is the following: “the information currently present within the visual stimulus that indicates which choice is more likely to be correct.” In our task, this evidence is related to the distribution of the tokens at a given moment in time, which determines the probability that one or the other target will ultimately be the correct one. This probability can be explicitly computed for a given choice i using Equation 5 to get $p_i(t)$. Although we do not expect that subjects were able to explicitly compute Equation 5, it is conceivable that they were able to construct an approximation (see Results). An alternative definition of sensory evidence in our task is the new information provided by each token movement, in other words, “the change in the stimulus that favors one choice over another.” This can be explicitly computed as $dp_i(t)/dt$, and again we expect that subjects can construct a reasonable approximation.

With a given definition of sensory evidence, we can discuss several variations of integrator and urgency-gating models that use sensory information to arrive at a decision.

Simulations

We compared human behavior to six putative models of decision making: four variations of integrator models and two variations of the urgency-gating model. All of these were presented with the same kinds of trials that were presented to the human subjects (focusing on slow blocks only), and the same analyses were applied to their performance.

Model 1: pure diffusion, integration of currently available sensory information.

$$x_i(t) = g \int_0^t E_i(\tau) d\tau \text{ or, alternatively, } \frac{dx_i}{dt} = g \cdot E_i(t) \quad (6)$$

$$E_i(t) = p_i(t) - 0.5 + N(t). \quad (7)$$

In this model, sensory evidence for a choice i is related to the success probability of that choice given the current distribution of tokens, and this quantity is integrated over time without any additional leak. The gain g is set to 1.5 and the term $N(t)$ represents Gaussian noise with mean zero and SD of 3. The threshold was set at ± 500 . This model is equivalent to several previously described diffusion models (Stone, 1960; Laming, 1968; Ratcliff, 1978; Mazurek et al., 2003), which integrate the information that is present in the stimulus at each moment in time. It is also equivalent to the leaky competing accumulator model of Usher and McClelland (2001), in which the net leakage parameter k is equal to the competition strength parameter β (Bogacz et al., 2006).

Model 2: diffusion with leak, integration of currently available sensory information.

$$\frac{dx_i}{dt} = g \cdot E_i(t) - L \cdot x_i(t) \quad (8)$$

$$E_i(t) = p_i(t) - 0.5 + N(t). \quad (9)$$

In this model, sensory evidence is defined as above, but there is a leak term (with parameter $L = 0.0005$, producing a strong leak) in the dynamic equation for $x_i(t)$. This model is equivalent to the leaky competing accumulator model in which the leak k is stronger than the competition strength β .

Model 3: diffusion without leak, integration of novel sensory information.

$$x_i(t) = g \int_0^t E_i(\tau) d\tau \text{ or } \frac{dx_i}{dt} = g \cdot E_i(t) \quad (10)$$

$$E_i(t) = \frac{dp_i(t)}{dt} + N(t). \quad (11)$$

In this model, sensory evidence is defined as the change in the success probability for a given choice. In other words, $E_i(t)$ is nonzero only at the moment when a token moves and is zero in-between token movements regardless of the current distribution of tokens. The gain g is set to 1.

Model 4: diffusion with leak, integration of novel sensory information.

$$\frac{dx_i}{dt} = g \cdot E_i(t) - L \cdot x_i(t) \quad (12)$$

$$E_i(t) = \frac{dp_i(t)}{dt} + N(t). \quad (13)$$

Like model 3, above, this model also defines evidence as the change in probability, but includes an additional leak term as in model 2. However, the value of L is reduced to 0.0003, since any larger values make it nearly impossible for neural activity to ever reach the threshold.

Model 5: urgency-gating model without filtering.

$$x_i(t) = g \cdot E_i(t) \cdot u(t) \quad (14)$$

$$E_i(t) = p_i(t) - 0.5 + N(t) \quad (15)$$

$$u(t) = t. \quad (16)$$

In this model, unlike models 1–4, the growth of activity is attributable entirely to the urgency term $u(t)$, which for simplicity is here defined simply as elapsed time. The gain g is set to 0.4 and the noise SD is 0.2.

Model 6: urgency-gating model with a low-pass filter. Clearly, model 5 is very susceptible to noise, especially late in each trial. Thus, we propose a final model, which is similar except that sensory evidence is low-pass filtered before gating by urgency, as follows:

$$x_i(t) = g \cdot E_i(t) \cdot u(t) \quad (17)$$

$$\tau \frac{dE_i(t)}{dt} = -E_i(t) + (p_i(t) - 0.5 + N(t)) \quad (18)$$

$$u(t) = t. \quad (19)$$

This model is similar to model 5, except that sensory evidence is low-pass filtered using a linear differential equation with a time constant of $\tau = 200$ ms. The gain g is set to 3 and the noise SD is 0.7.

Results

Behavioral results

In the choice reaction time task, the mean reaction times of subjects ranged from 214 to 416 ms (mean, 279 ms; SD, 45 ms). The mean RT of each individual subject was used to calculate that subject's decision times in the 15 token task. No significant decision time differences were found for choices made toward the right versus left target.

One of our first questions was to investigate whether subjects modify their decision policy as the timing parameters of the task are varied. To address this, we compared each subject's behavior in two conditions, each performed in separate blocks of ~ 100 trials. In fast blocks, the interval between token movements was 200 ms before any of the targets was reached, and accelerated to

20 ms afterward. This encouraged guessing because subjects saved a lot of time by taking a chance and could try again almost right away. In slow blocks, the token interval accelerated to 170 ms after a target was reached. This encouraged more conservative behavior because the benefit of choosing early was reduced. As expected, most subjects behaved more hastily in the fast blocks and more conservatively in the slow blocks. This is shown in Figure 3A for one subject, whose decision times in fast blocks (mean, 1105 ms) were significantly shorter (Kolmogorov–Smirnov test, $p < 10^{-25}$) than in slow blocks (mean, 1664). Furthermore, the success probability at decision time was lower in fast blocks than in slow blocks (Fig. 3B) (KS test, $p < 10^{-9}$). Most subjects (20 of 22) showed significantly faster responses in fast versus slow blocks (Fig. 3C) and nearly one-half (9 of 22) had lower success probability in fast than in slow blocks (Fig. 3D). This suggests that, in general, subjects adjusted their guessing policy to trade-off speed versus accuracy.

Next, we asked whether the specific pattern of token movements observed during a particular trial has an effect on the subject's behavior. To do this, we first focused on comparing behavior in easy versus ambiguous trials. As expected, 19 of 22 subjects made decisions significantly later in ambiguous trials than in easy ones (Fig. 4A,C) (KS test, $p < 0.05$). More interesting is that all of them also made decisions at a significantly lower level of success probability in ambiguous trials (Fig. 4B,D) (KS test, $p < 0.05$). That is, subjects appeared more willing to guess in ambiguous trials than in easy trials.

Next, we compared behavior during bias-for and bias-against trials (Fig. 2D), focusing on trials in which decisions were made after the first six token movements. This comparison is of interest because the two classes of models (integrator vs urgency gating) make distinct predictions about the timing of decisions in these trials. In particular, because integrator models retain a “memory,” they suggest that, after the first six token movements, neural activity related to the correct target will be higher in bias-for trials than in bias-against trials, and therefore closer to threshold. Therefore, these models predict faster decision times in bias-for trials than bias-against trials. In contrast, the urgency-gating models do not predict a significant difference.

Surprisingly, in agreement with urgency-gating models, there was no significant difference between decision times in the bias-for and bias-against trials (Fig. 5A,C). This was found for 21 of 22 subjects (KS test, $p > 0.05$). The success probability at decision time also was similar in the two kinds of trials (Fig. 5B,D) for all subjects (KS test, $p > 0.05$).

Figure 6 summarizes trends in comparing correct trials versus errors during the slow blocks. Across all trials, the mean decision time in error trials was longer than in correct trials for 11 of 22 subjects (KS test, $p < 0.05$), and the opposite was true for one subject (Fig. 6A). Figure 6B–E shows the same analysis restricted to the specific trial types. No consistent trends were found in

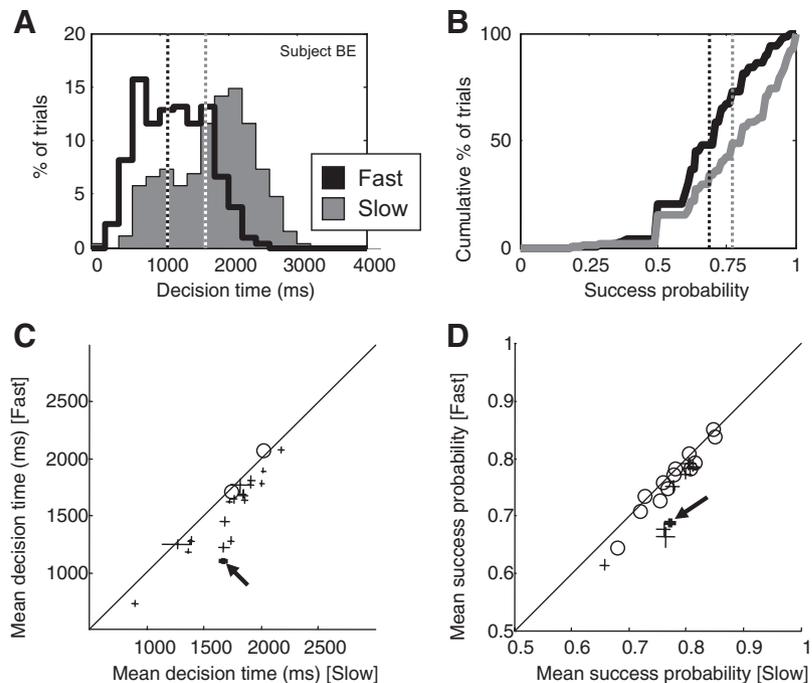


Figure 3. Comparison of behavior during fast versus slow blocks, using both correct and error trials. **A**, Distributions of the decision times of a representative subject. Dark line, Fast block ($N = 319$). Gray shaded region, Slow block ($N = 275$). The mean decision time in the fast block (dotted black line; 1105 ± 482 ms) was shorter than in the slow block (dotted gray line; 1664 ± 672 ms) (KS test, $p < 10^{-25}$). **B**, Cumulative distribution of success probability at the time of the decision during fast (black) and slow (gray) blocks, for the same subject. The mean in the fast block (0.69 ± 0.15) was lower than in the slow block (0.77 ± 0.18) (KS test, $p < 10^{-9}$). **C**, Average decision times of each subject during slow (x-axis) and fast (y-axis) blocks. The pluses indicate the mean and SE for subjects for whom the difference was significant (KS test, $p < 0.05$). The circles represent subjects for whom the difference was not significant. The arrow indicates the subject shown above. **D**, Average success probability at decision time of each subject during slow (x-axis) and fast (y-axis) blocks. The format is the same as in **C**.

these restricted analyses, except for the trivial observation that errors in bias-against trials tended to occur before the seventh token movement (Fig. 6E). More interesting was an analysis of the success rate as a function of the number of tokens that moved before a subject made their decision. Across all trials, the success rate was low for very fast decisions, increased later in the trial, and then decreased again (Fig. 6F). This was partially attributable to the fact that subjects generally waited for > 10 tokens only in trials that were more difficult (for example, the ambiguous trials) and in which success was closer to chance levels. Figure 6G–J shows the same analysis restricted to the specific trial types. As expected, in all cases the success rate is clearly dependent on the pattern of token movements. For example, in bias-for trials most errors occur between the third and seventh token, and in bias-against trials most errors occur before the seventh token.

Control analyses and simulations

Before interpreting our results with respect to particular decision models, it is important to consider whether subjects could have been using an explicit cognitive strategy to make their decisions. For example, could they have discovered the presence of the special trial types (easy, ambiguous, etc.) and then optimized their behavior to take advantage of this knowledge? We have several reasons to be confident that this was not the case. In particular, the subjects whose data are included here were never told about these trials, and during postexperiment interviews none of them reported detecting any special trials. Furthermore, even if subjects had been explicitly told to look for special trials, it would have been very hard to detect them: First, because these trials

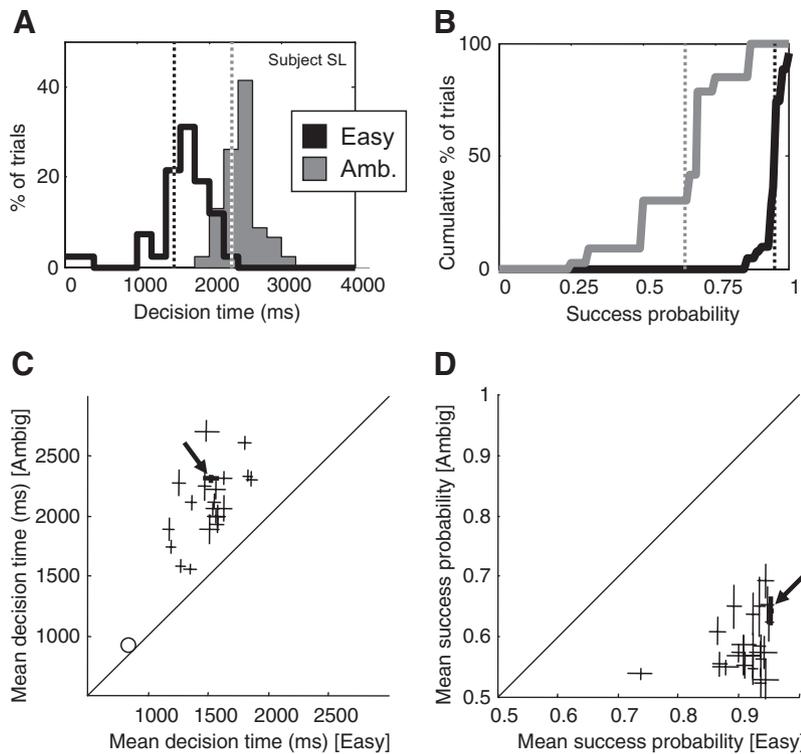


Figure 4. Comparison of behavior during correct easy versus ambiguous trials (Fig. 2C), embedded in the slow blocks. **A**, Distribution of decision times for one subject. Dark line, Easy ($N = 42$). Gray shaded region, Ambiguous ($N = 16$). The mean decision time in easy trials (1520 ± 421 ms) was shorter than in the ambiguous trials (2223 ± 269 ms) (KS test, $p < 10^{-8}$). **B**, Cumulative distribution of success probability at decision time, which was higher in the easy trials (mean, 0.95 ± 0.03) than in ambiguous trials (mean, 0.60 ± 0.11) (KS test, $p < 10^{-11}$). **C**, Average decision times for all subjects (same format as Figure 3C). **D**, Average success probabilities for all subjects (same format as Figure 3D).

were interspersed among 50% of trials that were completely random, each special trial type was relatively rare. Second, there were several variations of each special trial type, making it difficult to memorize any specific pattern. Third, even among the random trials, there are many that may partially resemble segments of the special trials, making it extremely difficult to discover any category boundaries. Fourth, the actual correct target was always randomized, and this was a much more salient piece of information for subjects to think about. Fifth, the token movements were quite fast, making it very hard to keep track of the specific pattern. Again, none of the subjects were told about these trials and none reported finding them when interviewed at the end. Finally, if against these odds a subject had indeed discovered the categories, then that subject's behavior would clearly exhibit optimal strategies. For example, if a subject identified a given trial as bias-for or bias-against, then he/she should always make a decision after exactly seven tokens (i.e., deciding between 1400 and 1600 ms), but this was not observed for any subject (Fig. 5C).

Assuming that subjects did not use any explicit strategy to make decisions, we examined how their behavior may agree or disagree with several varieties of decision models. To do so, we performed a series of simulations using the six models described in Materials and Methods. We presented each model with 100 repetitions of easy, ambiguous, bias-for, and bias-against trials, and compared their behavior (in terms of decision time and success probability at decision time) to human data.

Like all of the models tested, the diffusion model (model 1) (Fig. 7A) correctly exhibited faster decisions in easy versus ambiguous trials (KS test, $p < 0.01$). It also correctly exhibited a

lower average success probability in ambiguous versus easy trials (Fig. 7A, left, green box) (KS test, $p < 0.01$). However, it incorrectly predicted faster decisions in bias-for versus bias-against trials (Fig. 7A, right, red box) (KS test, $p < 0.01$). The reason for this is that, immediately after the first six token movements (time, 1200 ms), the neural activity was higher, and closer to threshold, in bias-for trials than in bias-against trials (Fig. 7A, shaded region in the fourth panel).

This result may appear counterintuitive. It might appear that a diffusion model (Stone, 1960; Laming, 1968; Ratcliff, 1978; Smith and Ratcliff, 2004; Ratcliff et al., 2007) will predict the same timing in bias-for and bias-against trials because it accumulates the difference in sensory evidence for the two options. However, note that, in bias-for trials, there is no moment in time at which evidence favors the wrong target (i.e., the success probability function never falls below 0.5), and so the difference in evidence is always positive (favoring the correct choice). If the sign of an integrated quantity is sometimes positive and never negative, then the final result of the integration will be greater than zero. In contrast, for the first six tokens of bias-against trials, the evidence always favors the wrong target, so the difference in evidence

is always negative, and therefore the result of integration after six tokens will be less than zero. Therefore, any model that integrates differences in currently available sensory information will predict earlier responses in bias-for than in bias-against trials. Note that all of our simulations include substantial intratrial noise (as opposed to intertrial noise) to demonstrate that our conclusions hold even if that is the only source of noise in the system. It is trivial to show that, if intertrial noise is the major source of variability, then our conclusions will only be stronger.

Next, we investigated whether the addition of a large leakage term to the diffusion model (i.e., yielding model 2) may reduce the difference in behavior between bias-for and bias-against trials. In this view, by the time the decision is made, differences between the accumulated activities in the early part of bias-for and bias-against trials would have decayed away, and behavior would be similar in both kinds of trials. However, even a very strong leak (so strong that the model had difficulty ever reaching the decision threshold) did not eliminate the difference between the trials (Fig. 7B, right, red box). Nevertheless, we examined our human data to see whether a strong leak could explain it. In particular, we looked at decision times from a subset of bias-for and bias-against trials in which a subject made the decision within 400 ms after the sixth token movement. These early decisions might still retain some bias that has not yet decayed away. However, we did not see any significant differences between decision times for bias-for versus bias-against trials in any of the 16 subjects who made enough of these fast choices to make the comparison possible (mean difference, -9 ms; SD, 68 ms). When we limited the analysis to the first 200 ms after the sixth token movement (six subjects made at least a few of these very early choices), we

again saw no significant differences (mean difference, -3 ms; SD, 23 ms). Thus, even a very strong leak term cannot explain our results.

Is it possible to explain our results if we postulate that the integrators can get “reset”? In other words, suppose that subjects can recognize the condition of complete ambiguity (when there are three tokens in each target) as a special case and reset their neural integrators back to baseline. Since beyond that point the bias-for and bias-against trials are identical, then so would be the behavior. To evaluate this possibility, we looked among the random trials for variations of bias-for and bias-against trials in which the first few token movements were not three and three. Specifically, we identified “bias-updown” trials as ones in which the first three tokens moved in the correct direction and then one moved in the opposite direction, and “bias-downup” trials as those in which the first token moved in the wrong direction and the next three in the correct direction. We also constrained the trials such that the profile of success probability was approximately the same for the remainder of the trial. Critically, as shown in Figure 5E, success probability in bias-updown trials never reached the critical value of 0.5 that could potentially trigger a reset of the integrators. Nevertheless, the reaction time distributions of bias-updown trials were not faster than those of bias-downup trials (Fig. 5F).

In summary, no model that temporally integrates currently available sensory information can explain our results. Therefore, we next considered models that do not integrate the currently available evidence, but rather, integrate new information provided by novel sensory events (i.e., token movements, or the change in sensory information). Indeed, a diffusion model in which the change in sensory information was being integrated (model 3) did correctly produce similar decision times in bias-for versus bias-against trials (Fig. 7C, right, green box) (KS test, $p > 0.1$). However, it in turn failed to reproduce the result shown on Figure 4B, because its success probability was always the same in ambiguous and easy trials (Fig. 7C, left, red box) (KS test, $p > 0.1$). The reason for this is straightforward: if the neural activity is an integral of the derivative of some quantity $\left(\int \frac{dp}{dt}\right)$, then it is simply equivalent to that quantity [i.e., to $p(t)$]. Thus, it always reaches the threshold at the same level of success probability. Adding leak (i.e., model 4) did not improve matters because now the easy trials incorrectly exhibited lower success probability than ambiguous trials (Fig. 7D, left, red box) (KS test, $p < 0.01$) and decisions were slower in bias-for than in bias-against trials (Fig. 7D, right, red box) (KS test, $p < 0.01$), again in disagreement with data. In short, no tested variation of the diffusion model could reproduce

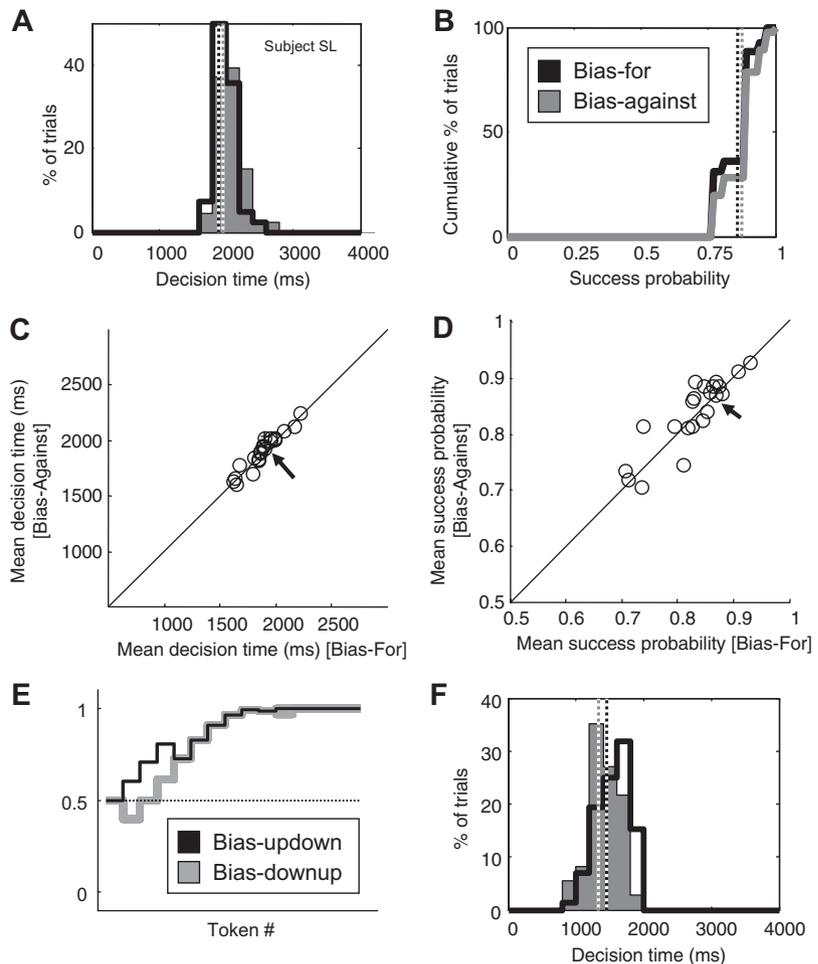


Figure 5. Comparison of behavior during correct bias-for versus bias-against trials (Fig. 2D), embedded in the slow blocks. We exclusively focus on those trials in which decisions were made after the first six token movements. **A**, Distribution of decision times for one subject. Dark line, Bias-for ($N = 42$). Gray shaded region, Bias-against ($N = 46$). The mean decision time in bias-for trials (1885 ± 145 ms) was not significantly different from bias-against trials (1957 ± 191 ms) (KS test, $p > 0.05$). **B**, Cumulative distribution of success probability at decision time, which was also not significantly different in bias-for (mean, 0.86 ± 0.06) and bias-against trials (mean, 0.88 ± 0.06) (KS test, $p > 0.05$). **C**, Average decision times for all subjects (same format as Figure 3C). **D**, Average success probabilities for all subjects (same format as Figure 3D). **E**, Profiles of success probability for bias-updown (black) and bias-downup (gray) trials. Note that because these trials were relatively rare, we pooled data across subjects to yield 72 bias-updown and 37 bias-downup trials. **F**, Distributions of decision times for bias-updown (mean, 1447 ± 233 ms) and bias-downup trials (mean, 1327 ± 223 ms). The difference was not significant.

both the finding that the success probability at decision time is higher in easy than ambiguous trials (Fig. 4B,D) as well as the finding that decision times were similar in bias-for and bias-against trials (Fig. 5A,C).

In contrast, both versions of the urgency-gating model reproduced these critical results. The model without a low-pass filter (model 5) correctly reproduced the lower success probability in ambiguous versus easy trials (Fig. 7D, left, green box) (KS test, $p < 0.01$) as well as similar decision times in bias-for and bias-against trials (Fig. 7D, right, green box) (KS test, $p > 0.1$). However, it is highly susceptible to noise as evident from plots of example neural activity patterns. As discussed above, the brain can overcome such noise on each trial by averaging over a large number of neurons with uncorrelated fluctuations of activity (Shadlen et al., 1996). Such a process may be approximated by the addition of a low-pass filter to the model, which effectively reduces the gain of intratrial noise. As shown in Figure 7F, the addition of a low-pass filter does not appreciably change the behavioral results, and the model still correctly produces lower suc-

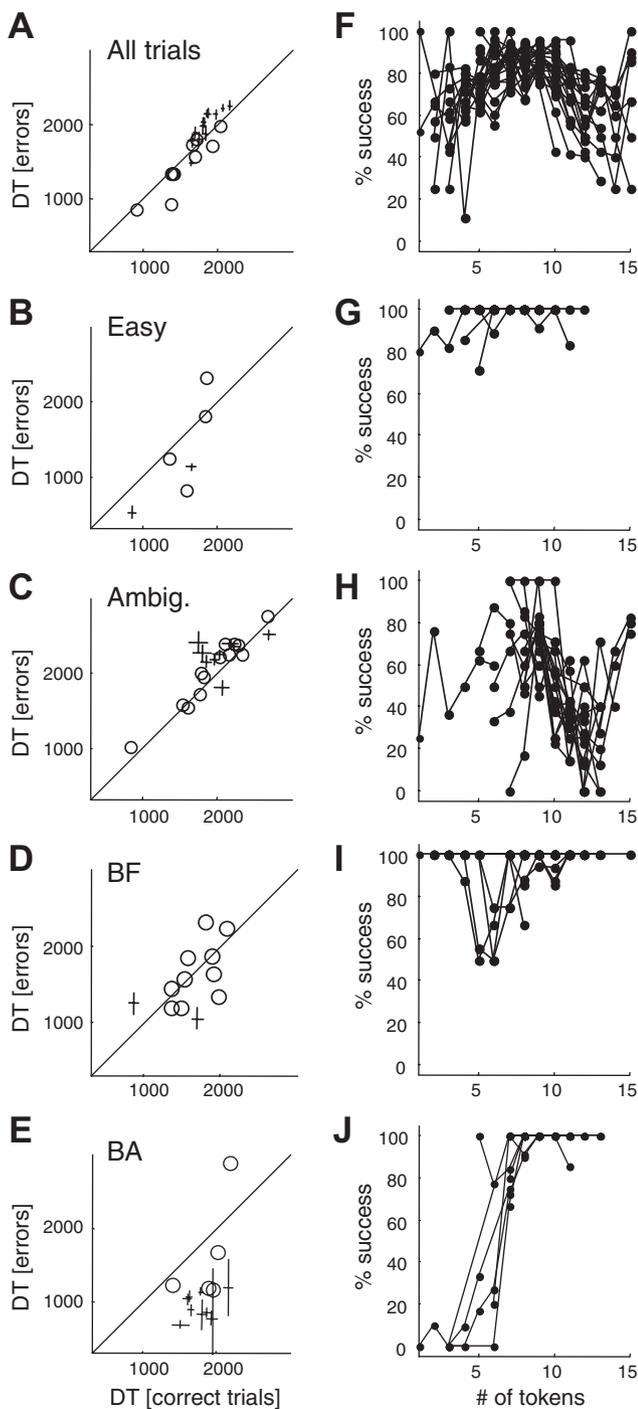


Figure 6. *A*, Comparison of mean decision times (DT) in correct (*x*-axis) versus error trials (*y*-axis) across all trials in slow blocks (same format as Figure 3C). *B*, Comparison of mean decision times in correct versus error trials using data only from easy trials. Data are shown only for those subjects who made at least one error in easy trials ($N = 6$). *C*, Comparison of correct versus error decision times in ambiguous (Ambig.) trials. *D*, Same for bias-for (BF) trials. *E*, Same for bias-against (BA) trials. *F*, For each subject, the lines show the success rate across all trials as a function of the number of tokens that moved before the decision. For clarity, points for which there were fewer than five total trials are omitted. *G*, Same as *F* except only for easy trials. *H*, Same for ambiguous trials. *I*, Same for bias-for trials (all, regardless of when decision was made). *J*, Same for bias-against trials.

cess probability in ambiguous versus easy trials as well as similar decision times in bias-for and bias-against trials.

Note that, if the model that integrates novel sensory information (model 3) is gated by a growing urgency function, then it will

effectively become a version of the urgency-gating model (model 6). This is because integrating the derivative of some signal with an integrator that has a short time constant is similar to low-pass filtering that signal. Thus, we can conclude that our results support models in which an urgency signal multiplicatively gates a filtered estimate of current evidence, and that one way to compute that estimate is through relatively fast integration of novel sensory events.

Decreasing accuracy criterion

If we suppose that the urgency-gating model accurately describes the process underlying decision making in our task, then we can make an additional prediction about the level of confidence at which subjects will be making decisions across all trials, not just the special trials emphasized in Figures 4 and 5. In particular, if we make an educated guess about the $E(t)$ function that is actually used by our subjects, then we can predict that its value at the time of the decision should decrease as a function of decision time. The reason can be seen by setting Equation 2 equal to a constant threshold T and solving for $E_i(t)$ as follows:

$$g \cdot E_i(t) \cdot u(t) = T \quad (20)$$

$$E_i(t) = \frac{T}{g \cdot u(t)} = \frac{T}{g \cdot t} \quad (21)$$

Of course, we cannot truly know the exact form of the $E(t)$ function used by our subjects. It would be difficult to believe that subjects can precisely calculate Equation 5, but we can expect that they can make a reasonable estimate. For example, a simple “first-order” estimate of sensory evidence is the sum of log-likelihood ratios (SumLogLR) of individual token movements as follows:

$$E_S(n) = \sum_{j=1}^n \log \frac{p(e_j|S)}{p(e_j|U)}, \quad (22)$$

where $p(e_j|S)$ is the likelihood of a token event e_j during trials in which the selected target is correct, and $p(e_j|U)$ is its likelihood during trials in which the unselected target is correct. Although this may at first appear complex, it simply amounts to counting the number of tokens which move in each direction. This expression for $E_S(t)$ can then be used to estimate the posterior probability of target S being correct using the following relationship:

$$p(S|n) = \frac{e^{E_S(n)}}{1 + e^{E_S(n)}} \quad (23)$$

Strictly speaking, this estimate of probability is wrong. It ignores the conditional probability between sequential token movements and the correct response: That is, the likelihood $p(e_1, e_2|S)$ is not simply the product of $p(e_1|S)$ and $p(e_2|S)$ because e_1 and e_2 are conditionally dependent. To compute an accurate estimate of $p(t)$, Equation 23 would have to take that conditional probability into account. Nevertheless, this first-order estimate actually does quite well for the first 10 token movements: For those first 10 tokens, the estimate provided by Equation 23 linearly correlates with the real success probability computed using Equation 5 with a slope of 0.82 and an R^2 of 0.99 ($p < 0.001$).

With a reasonable estimate of how subjects may compute $E(t)$, we set out to test whether the value reached by this quantity at decision time decreases as a function of time, as predicted by the urgency-gating model. To do so, we grouped trials according to the number of tokens that moved before the decision time and calculated the value of $E(t)$ for the selected target at the time of the

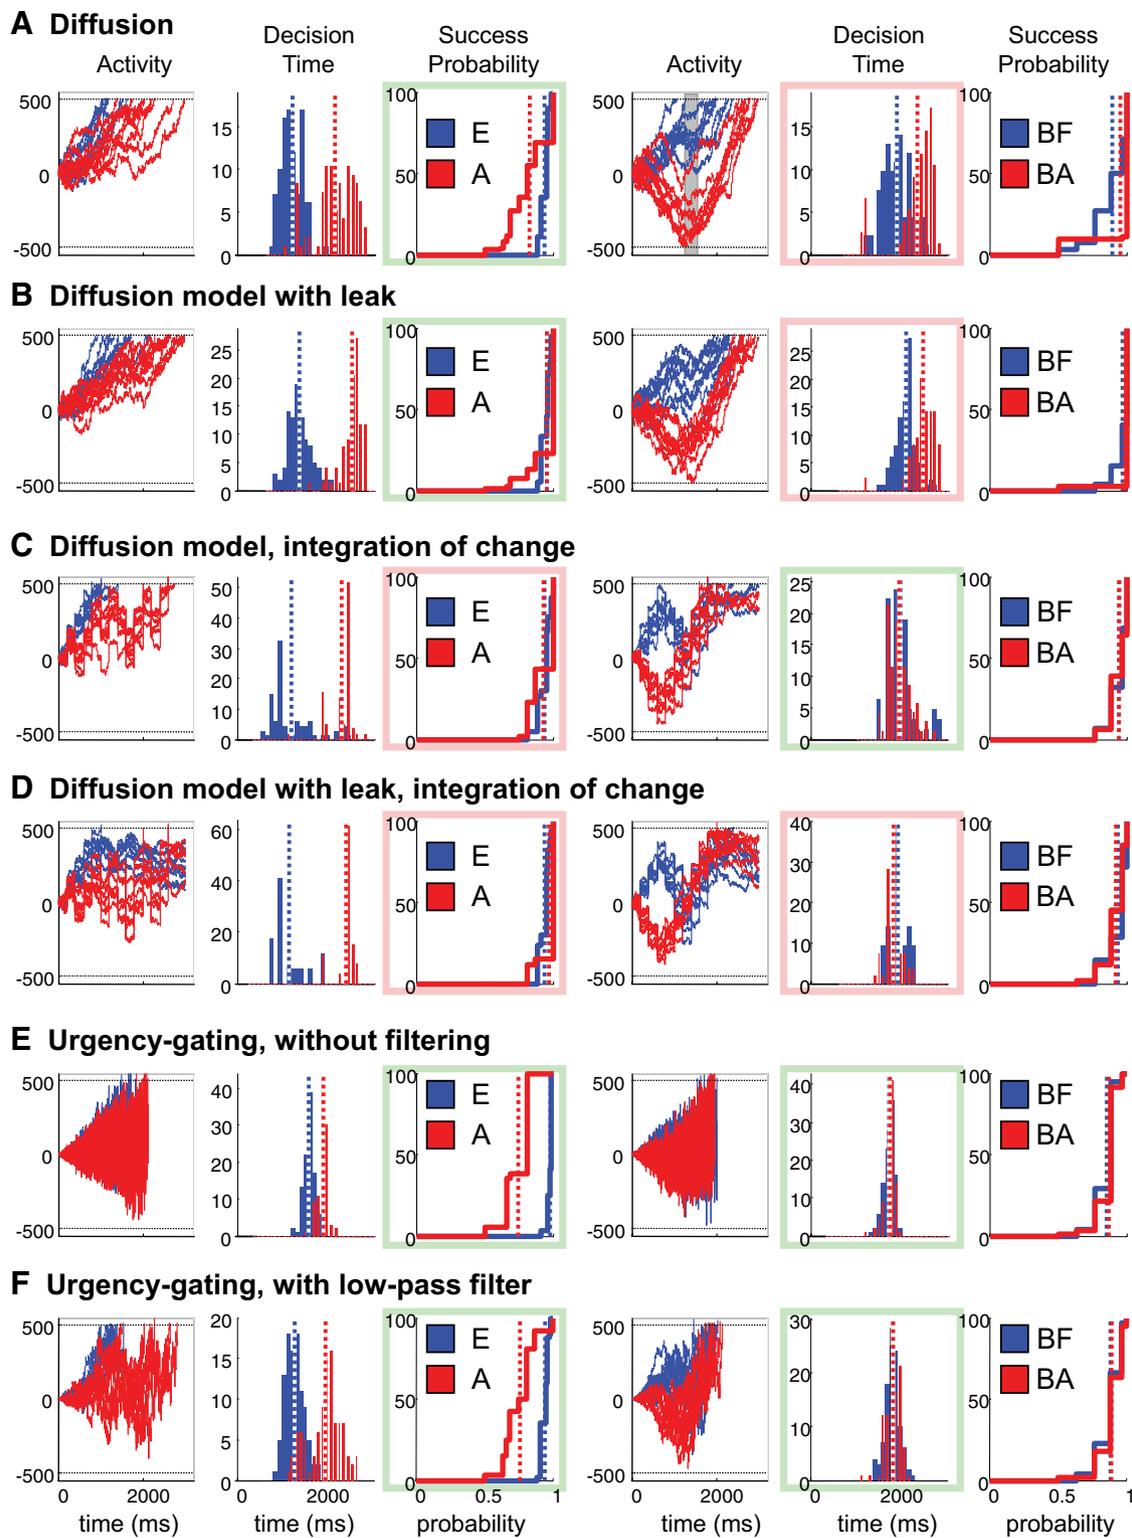


Figure 7. Simulations of the six models. For each model, the leftmost three plots compare behavior between easy (E) (blue) and ambiguous (A) (red) trials, and the rightmost three plots compare behavior between bias-for (BF) (blue) and bias-against (BA) (red) trials. The first plot of each triplet shows the simulated neural activity of 10 example trials of each type, as a function of time. The second shows histograms of the decision times as a percentage of all trials in which a choice was made. The third plot shows the success probabilities at decision time as a cumulative percentage. The red and green frames highlight the critical comparisons discussed in the text, with red representing disagreement with data and green representing agreement. The frames in the third column compare success probability in easy versus ambiguous trials (as done in Fig. 4B), whereas those in the fifth column compare decision times in bias-for and bias-against trials (as done in Fig. 5A). **A**, Diffusion model (model 1). The shaded region indicates the time period when there are three tokens in each target during bias-for and bias-against trials. **B**, Diffusion model with leak (model 2). **C**, Diffusion model that integrates the change in sensory information (model 3). **D**, Same as **C** but with a leak (model 4). **E**, Urgency-gating model without filtering (model 5). **F**, Urgency-gating model with a low-pass filter (model 6).

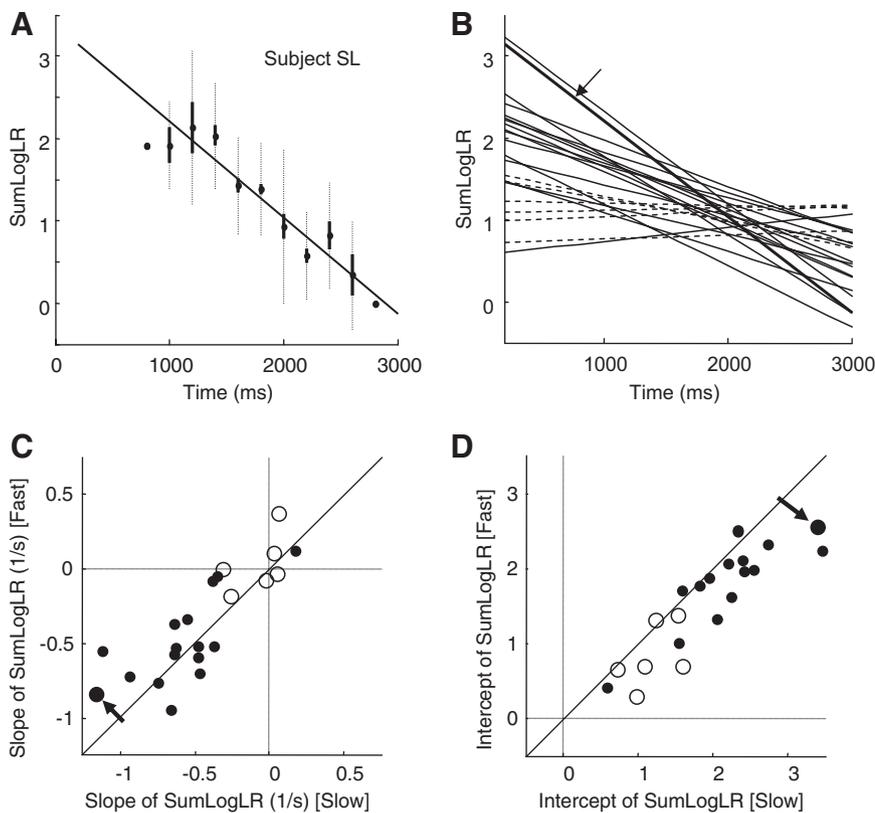


Figure 8. *A*, Analysis of SumLogLR at decision time for decisions made at different times, for an individual subject's data from slow blocks ($N = 266$). For each 200 ms bin, we show the mean (dot), SEM (thick lines), and SD (thin lines) of the SumLogLR at the time the decision was made. The oblique line shows a linear regression through the data (slope, -0.24 ; intercept, 3.41 ; $R^2 = 0.29$; $p < 10^{-10}$), suggesting a decreasing threshold. *B*, Results of linear regressions for all of the subjects (the subject in *A* is indicated by the arrow). The solid lines show regressions that were significant ($p < 0.05$), and the dashed lines show those that were not. *C*, Slopes of regression lines of individual subjects, comparing regressions from slow blocks (x -axis) versus fast blocks (y -axis). The dots indicate when the regression from the slow block was significant. *D*, Intercepts of regression lines from slow (x -axis) versus fast (y -axis) blocks.

decision. The result for one subject is shown in Figure 8*A*. A simple linear regression through the data shows a significant fit ($p < 10^{-10}$) with a negative slope. A significant regression was found for 16 of 22 subjects (Fig. 8*B*), and 15 of these had a negative slope. In summary, there was a trend for later decisions to be made at a lower level of $E(t)$ than decisions made early in the trial, consistent with the predictions of the urgency-gating model. Furthermore, the slope tended to be shallower (Fig. 8*C*) (paired t test, $p < 0.01$) and the y -intercept lower (Fig. 8*D*) (paired t test, $p < 0.01$) during the fast blocks than the slow blocks, suggesting that, on average, the urgency signal follows a different time-course during the two blocks and thus controls the trade-off between speed and accuracy of decisions. This finding further predicts that the level of confidence that subjects have in the decisions they make should decrease with longer decision times.

Discussion

Several models propose that simple decisions involve the temporal integration of sequential sensory samples until a threshold is reached (Stone, 1960; Laming, 1968; Ratcliff, 1978; Carpenter and Williams, 1995; Usher and McClelland, 2001; Wang, 2002; Mazurek et al., 2003; Grossberg and Pilly, 2008). These models explain error rates and reaction time distributions, and are supported by neurophysiological studies showing buildup activity in a number of brain structures during decision-making tasks (Munoz and Wurtz, 1995; Kim and Shadlen, 1999; Gold and Shadlen,

2000; Roitman and Shadlen, 2002; Ratcliff et al., 2003, 2007; Shen and Paré, 2007).

The present study, however, prompts us to reconsider two aspects of these models. First, we propose that evidence for a given choice should not be computed by temporally integrating the information currently present in the stimulus. Instead, it should involve either summation of only new information (provided by a change in the state of the stimulus) or be a low-pass-filtered signal related to the state of the sensory information. This is consistent with studies suggesting that the time window of integration for perceptual decision making is on the order of 100 ms (Ludwig et al., 2005; Ghose, 2006). Second, we suggest that the long buildup of neural activity in constant-evidence tasks is not caused by an integration process but is primarily attributable to a growing urgency signal that is unrelated to any particular choice.

An influential argument in favor of integrator models arises from their similarity to the SPRT (Wald, 1945), a statistical test for deciding whether current evidence for a given hypothesis is sufficient to ensure a criterion level of accuracy. Because the SPRT can be performed through summation of independent pieces of evidence and comparison to a threshold, it has been suggested that it is effectively implemented by integrator models (Bogacz et al., 2006; Bogacz and Gurney, 2007; Gold and Shadlen, 2007). However, there is a difference between how probability is calculated and how it is traded off against time. The SPRT is optimal in the sense of producing the best accuracy after a given time, or requiring the fewest number of samples to reach a given accuracy (Wald and Wolfowitz, 1948), but it does not implement a trade-off between time and accuracy. Animals cannot afford to have a fixed threshold of accuracy but must be willing to tolerate lower success rates to reduce the time spent in making a decision (Chittka et al., 2009). An accuracy criterion that decreases over time accomplishes this and can be implemented through a multiplication of evidence by a growing urgency signal and comparison with a constant neural threshold.

Furthermore, the similarity between the diffusion model and the SPRT only holds if the sequential samples of information are statistically independent (Bogacz et al., 2006), which is not the case in most tasks that have been studied. In particular, if a sample is already predicted by previous samples, then it should be ignored. In other words, if integration takes place, then it should be integration of novel sensory information, and not of the information present in the stimulus at a given time. Because in constant-evidence tasks novel information only arrives at cue presentation, the integration should result in a step function of neural activity (or a fast saturation in the case of noisy information). It should not resemble the long growth of activity observed in neural studies and inferred from reaction time distributions. That growth, we propose, is primarily attributable to an urgency signal that implements a trade-off between speed and accuracy.

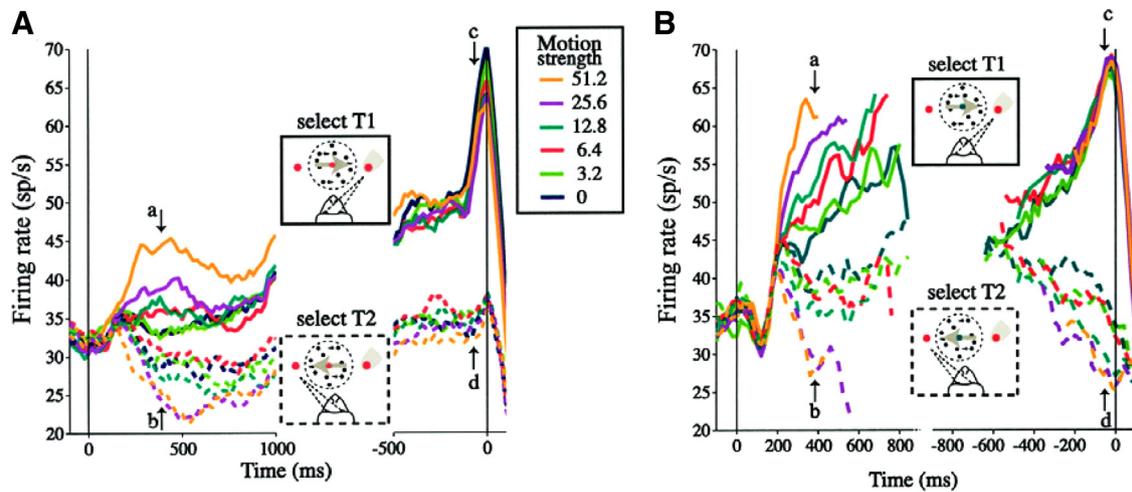


Figure 9. Neural data from the experiment of Roitman and Shadlen (2002), reprinted with permission. **A**, Data from the FD version of the motion discrimination task. The solid lines show average neural activity of 54 LIP cells during trials in which the monkey correctly selected the target in the response field of the cell, and the dotted lines show activity when the monkey correctly selected the opposite target. Different colors indicate trials with different motion coherence (see inset). The data on the left are aligned to the onset of motion, and data on the right to the onset of the saccade. **B**, Data from the RT version of the task, same format.

There is already widespread evidence for buildup signals in many brain regions and many experimental paradigms. For example, neural activity related to elapsed time has been reported in prefrontal cortex during duration reproduction (Jech et al., 2005), and in LIP during time interval (Leon and Shadlen, 2003) and motion discrimination (Churchland et al., 2008). During instructed delays with different possible durations, as each of the likely GO signal times approaches there is a buildup of neural activity in LIP during saccade tasks (Janssen and Shadlen, 2005), and in motor cortex during reaching tasks (Renoult et al., 2006), with corresponding changes of corticospinal excitability (van Elswijk et al., 2007). More generally, buildup activity has been reported in a variety of brain regions even during motor tasks that do not involve any decisions (Hanes and Schall, 1996; Munoz et al., 2000; Ivry and Spencer, 2004; Roesch and Olson, 2005; Tanaka, 2007; Thomas and Paré, 2007; Lebedev et al., 2008). It is therefore reasonable to suggest that such buildup can influence decision-making processes, which appear to involve the same structures that are involved in sensorimotor control (Glimcher, 2003; Romo et al., 2004; Cisek and Kalaska, 2005; Gold and Shadlen, 2007).

We cannot know whether our conclusions can be applied to other studies. It is likely that decision-making mechanisms are at least partially task dependent (Ghose, 2006). Nevertheless, urgency gating provides a simple potential explanation for data from the well known motion discrimination tasks. During these tasks, neurons in the medial temporal area (MT) reflect motion evidence very rapidly, equilibrating to a relatively steady coherence-related firing rate within 150–200 ms (Britten et al., 1992). This is compatible with a temporal filter model (Ludwig et al., 2005) in which MT activity is a filtered version of the noisy motion signals arriving from earlier visual areas. During a “fixed duration” (FD) version of the task (Fig. 9A), in which the monkey must report its decision only after an external GO signal, LIP activity equilibrates at a coherence-dependent firing rate in ~300 ms, suggesting that no additional integration of the MT signal takes place (Roitman and Shadlen, 2002) (see also Shadlen and Newsome, 2001). However, during a RT version (Fig. 9B), in which the monkey can report its decision at any time, LIP activity continues to grow for as long as 800 ms (Roitman and Shadlen, 2002). Furthermore, the apparent threshold of LIP activity for

saccade initiation in the RT task (60–70 Hz) is higher than the activity at which the same cells equilibrate in the FD task (<50 Hz), despite the fact that performance in both conditions is comparable. What can explain this task dependence of LIP activity?

Previous models have suggested that, during the FD task, the LIP signal saturates because there is a large leak term (Grossberg and Pilly, 2008) or because a threshold is crossed and activity stops at its current level after a delay (Mazurek et al., 2003). The urgency-gating model suggests a straightforward alternative: In the RT task, monkeys are allowed to control the trade-off between speed and accuracy, and do so through a growing urgency signal. This, multiplied by the coherence-dependent MT input, produces LIP activity that exhibits a coherence-dependent growth rate. In the FD task, monkeys must ensure that LIP activity does not reach the saccade threshold prematurely, so the urgency signal is low and constant until the GO signal. This, multiplied by the MT input, would yield coherence-dependent but relatively constant LIP activity, as observed.

Recently, Ditterich (2006) showed that the behavioral and neural data from the RT task can only be explained by models that include a gain that grows linearly with time. Importantly, he showed that, as long as there is a growing gain, then it is not strictly necessary for the model to involve any type of temporal integration. His analysis is therefore compatible with the urgency-gating model. Although Ditterich favored the presence of integration on the grounds of signal-to-noise considerations, those considerations can also be met if MT signals are low-pass filtered before arriving in LIP, as in model 6.

Nevertheless, we cannot claim that urgency gating can explain the variety of results observed during the motion discrimination tasks. For example, Kiani et al. (2008) presented brief motion pulses to monkeys performing the FD task, finding a short window for motion integration (~300–350 ms). However, Huk and Shadlen (2005) presented similar motion pulses during the RT task and obtained a much longer window. It remains unclear why such different results were obtained in these two studies.

These caveats aside, our data propose a reconsideration of the central concept of many recent models of decision making—buildup of activity through temporal integration. Summation of evidence from sequential samples is a good way to estimate the

posterior probability of success for a given choice, and the brain may indeed use something like that when it is appropriate [e.g., in the present task or the task of Yang and Shadlen (2007)]. However, such a process should only sum novel information and therefore may not be responsible for the long-duration growth of neural activity observed in many neurophysiological experiments and inferred from behavioral data in constant-evidence tasks. Instead, we propose that this growth of activity is primarily attributable to multiplication of sensory evidence (which may be computed through a low-pass filter or through integration of the change in the stimulus) with a motor initiation-related buildup signal. In this view, what determines the timing of actions is not the termination of pure decision-making processes followed by movement preparation and execution. Instead, the strength of evidence for a given choice is combined with a motor signal related to the urgency to make a choice, and it is the two together that turn decisions into action.

References

- Bogacz R, Gurney K (2007) The basal ganglia and cortex implement optimal decision making between alternative actions. *Neural Comput* 19:442–477.
- 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.
- Britten KH, Shadlen MN, Newsome WT, Movshon JA (1992) The analysis of visual motion: a comparison of neuronal and psychophysical performance. *J Neurosci* 12:4745–4765.
- Carpenter RH, Williams ML (1995) Neural computation of log likelihood in control of saccadic eye movements. *Nature* 377:59–62.
- Chittka L, Skorupski P, Raine NE (2009) Speed-accuracy tradeoffs in animal decision making. *Trends Ecol Evol* 24:400–407.
- Churchland AK, Kiani R, Shadlen MN (2008) Decision-making with multiple alternatives. *Nat Neurosci* 11:693–702.
- Cisek P, Kalaska JF (2005) Neural correlates of reaching decisions in dorsal premotor cortex: specification of multiple direction choices and final selection of action. *Neuron* 45:801–814.
- Cisek P, El-Murr S, Puskas GA (2008) Decision-making in changing conditions: evidence against temporal integration models. *Soc Neurosci Abstr* 34:715.4.
- Ditterich J (2006) Stochastic models of decisions about motion direction: behavior and physiology. *Neural Netw* 19:981–1012.
- Ghose GM (2006) Strategies optimize the detection of motion transients. *J Vis* 6:429–440.
- Glimcher PW (2003) The neurobiology of visual-saccadic decision making. *Annu Rev Neurosci* 26:133–179.
- Gold JL, Shadlen MN (2000) Representation of a perceptual decision in developing oculomotor commands. *Nature* 404:390–394.
- Gold JL, Shadlen MN (2003) The influence of behavioral context on the representation of a perceptual decision in developing oculomotor commands. *J Neurosci* 23:632–651.
- Gold JL, Shadlen MN (2007) The neural basis of decision making. *Annu Rev Neurosci* 30:535–574.
- Grossberg S, Pilly PK (2008) Temporal dynamics of decision-making during motion perception in the visual cortex. *Vision Res* 48:1345–1373.
- Hanes DP, Schall JD (1996) Neural control of voluntary movement initiation. *Science* 274:427–430.
- 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.
- Ivry RB, Spencer RM (2004) The neural representation of time. *Curr Opin Neurobiol* 14:225–232.
- Janssen P, Shadlen MN (2005) A representation of the hazard rate of elapsed time in macaque area LIP. *Nat Neurosci* 8:234–241.
- Jech R, Dusek P, Wackermann J, Vymazal J (2005) Cumulative blood oxygenation-level-dependent signal changes support the “time accumulator” hypothesis. *Neuroreport* 16:1467–1471.
- 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.
- Kim JN, Shadlen MN (1999) Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nat Neurosci* 2:176–185.
- Laming D (1968) *Information theory of choice reaction time*. New York: Wiley.
- Lebedev MA, O’Doherty JE, Nicolelis MA (2008) Decoding of temporal intervals from cortical ensemble activity. *J Neurophysiol* 99:166–186.
- Leon MI, Shadlen MN (2003) Representation of time by neurons in the posterior parietal cortex of the macaque. *Neuron* 38:317–327.
- Ludwig CJ, Gilchrist ID, McSorley E, Baddeley RJ (2005) The temporal impulse response underlying saccadic decisions. *J Neurosci* 25:9907–9912.
- Mazurek ME, Roitman JD, Ditterich J, Shadlen MN (2003) A role for neural integrators in perceptual decision making. *Cereb Cortex* 13:1257–1269.
- Munoz DP, Wurtz RH (1995) Saccade-related activity in monkey superior colliculus. I. Characteristics of burst and buildup cells. *J Neurophysiol* 73:2313–2333.
- Munoz DP, Dorris MC, Paré M, Everling S (2000) On your mark, get set: brainstem circuitry underlying saccadic initiation. *Can J Physiol Pharmacol* 78:934–944.
- Puskas GA, Thivierge JP, El-Murr S, Cisek P (2007) Making decisions as the evidence is changing. *Soc Neurosci Abstr* 33:507.12.
- Ratcliff R (1978) A theory of memory retrieval. *Psychol Rev* 83:59–108.
- Ratcliff R, Cheria A, Segraves M (2003) A comparison of macaque behavior and superior colliculus neuronal activity to predictions from models of two-choice decisions. *J Neurophysiol* 90:1392–1407.
- Ratcliff R, Hasegawa YT, Hasegawa RP, Smith PL, Segraves MA (2007) Dual diffusion model for single-cell recording data from the superior colliculus in a brightness-discrimination task. *J Neurophysiol* 97:1756–1774.
- Reddi BA, Carpenter RH (2000) The influence of urgency on decision time. *Nat Neurosci* 3:827–830.
- Reddi BA, Asrress KN, Carpenter RH (2003) Accuracy, information, and response time in a saccadic decision task. *J Neurophysiol* 90:3538–3546.
- Renoult L, Roux S, Riehle A (2006) Time is a rubberband: neuronal activity in monkey motor cortex in relation to time estimation. *Eur J Neurosci* 23:3098–3108.
- Roesch MR, Olson CR (2005) Neuronal activity dependent on anticipated and elapsed delay in macaque prefrontal cortex, frontal and supplementary eye fields, and premotor cortex. *J Neurophysiol* 94:1469–1497.
- 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.
- Romo R, Hernández A, Zainos A (2004) Neuronal correlates of a perceptual decision in ventral premotor cortex. *Neuron* 41:165–173.
- 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.
- Shadlen MN, Britten KH, Newsome WT, Movshon JA (1996) A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *J Neurosci* 16:1486–1510.
- Shen K, Paré M (2007) Neuronal activity in superior colliculus signals both stimulus identity and saccade goals during visual conjunction search. *J Vis* 7:15.1–15.13.
- Smith PL, Ratcliff R (2004) Psychology and neurobiology of simple decisions. *Trends Neurosci* 27:161–168.
- Stone M (1960) Models for choice reaction time. *Psychometrika* 25:251–260.
- Tanaka M (2007) Cognitive signals in the primate motor thalamus predict saccade timing. *J Neurosci* 27:12109–12118.
- Thomas NW, Paré M (2007) Temporal processing of saccade targets in parietal cortex area LIP during visual search. *J Neurophysiol* 97:942–947.
- Usher M, McClelland JL (2001) The time course of perceptual choice: the leaky, competing accumulator model. *Psychol Rev* 108:550–592.
- van Elswijk G, Kleine BU, Overeem S, Stegeman DF (2007) Expectancy induces dynamic modulation of corticospinal excitability. *J Cogn Neurosci* 19:121–131.
- Vickers D (1970) Evidence for an accumulator model of psychophysical discrimination. *Ergonomics* 13:37–58.
- Wald A (1945) Sequential tests of statistical hypotheses. *Ann Math Stat* 16:117–186.
- Wald A, Wolfowitz J (1948) Optimum character of the sequential probability ratio test. *Ann Math Stat* 19:326–339.
- Wang XJ (2002) Probabilistic decision making by slow reverberation in cortical circuits. *Neuron* 36:955–968.
- Wong KF, Wang XJ (2006) A recurrent network mechanism of time integration in perceptual decisions. *J Neurosci* 26:1314–1328.
- Yang T, Shadlen MN (2007) Probabilistic reasoning by neurons. *Nature* 447:1075–1080.