

The Drift Diffusion Model can account for the accuracy and reaction time of value-based choices under high and low time pressure

Milica Milosavljevic^{*1}, Jonathan Malmaud^{1,2*}, Alexander Huth^{1*},
Christof Koch^{1,2,3}, and Antonio Rangel^{1,4}

¹ Computation and Neural Systems, California Institute of Technology, Pasadena, CA

² Division of Biology, California Institute of Technology, Pasadena, CA

³ Division of Engineering and Applied Science, California Institute of Technology, Pasadena, CA

⁴ Division of Humanities and Social Sciences, California Institute of Technology, Pasadena, CA

Abstract

An important open problem is how values are compared to make simple choices. A natural hypothesis is that the brain carries out the computations associated with the value comparisons in a manner consistent with the Drift Diffusion Model (DDM), since this model has been able to account for a large amount of data in other domains. We investigated the ability of four different versions of the DDM to explain the data in a real binary food choice task under conditions of high and low time pressure. We found that a seven-parameter version of the DDM can account for the choice and reaction time data with high-accuracy, in both the high and low time pressure conditions. The changes associated with the introduction of time pressure could be traced to changes in two key model parameters: the barrier height and the noise in the slope of the drift process.

Keywords: drift-diffusion model, value-based choice, response time.

1 Introduction

The drift diffusion model (DDM) is one of the cornerstones of modern psychology (Ratcliff, 1978; Ratcliff & McKoon, 2008; Ratcliff & Smith, 2004; Smith & Ratcliff, 2004) and, increasingly, of behavioral neuroscience (Bogacz, 2007; Gold & Shadlen, 2007; Link, 1992; Smith & Ratcliff, 2004). The model has received increased attention over the last years for several reasons. First, it has provided more accurate descriptions of accuracy and reaction time data than alternative models in a wide range of psychological tasks, including perceptual discrimination and go-no-go tasks (Ratcliff & Rouder, 1998, 2000; Ratcliff & Smith, 2004; Ratcliff, Van Zandt, & McKoon, 1999). Second, the model is a special case of many of its competitors, which is a sign of its generality (Bogacz, 2007; Bogacz, Brown, Moehlis, Holmes, & Cohen, 2006; Usher & McClelland, 2001). Finally, the model has been applied to explain neurophysiological

data in various perceptual discrimination tasks (Gold & Shadlen, 2007; Heekeren, Marrett, & Ungerleider, 2008; Philiastides, Ratcliff, & Sajda, 2006; Ratcliff, Cherian, & Segraves, 2003; Ratcliff, Hasegawa, Hasegawa, Smith, & Segraves, 2007; Ratcliff, Philiastides, & Sajda, 2009) and has a compelling neuronal interpretation.

An important open problem in behavioral neuroscience is how the brain compares values to make simple choices. This problem is particularly interesting because there is ample evidence suggesting that the comparison process is not deterministic, and that it does not always choose the best option (Hare, Camerer, Knoepfle, O'Doherty, & Rangel, 2009; Hare, Camerer, & Rangel, 2009; Hare, O'Doherty, Camerer, Schultz, & Rangel, 2008; Padoa-Schioppa & Assad, 2006; Tom, Fox, Trepel, & Poldrack, 2007). The success of the DDM in the realm of perceptual decision making has lead several groups in neuroeconomics to speculate that the same computational model might be used by the brain to make simple value based choices (Gold & Shadlen, 2007; Rangel, 2008; Rangel, Camerer, & Montague, 2008; Rangel & Hare, 2010; Wallis, 2007). This type of choice refers to situations in which the brain chooses among several possible stimuli

^{*}The first three authors contributed equally. We gratefully thank the NGA, the NSF and the Mathers and Moore Foundations for funding this research. We also thank Roger Ratcliff for giving us invaluable comments during the review process.

associated with different reward values at consumption (e.g., alternative food items) by assigning a value to every item under consideration and comparing the values to select one of them.

This paper investigates the extent to which the DDM can explain the accuracy and reaction time data in real simple food choices. Subjects made real choices between pairs of appetitive snack foods and had to eat the food that they chose in a randomly selected trial. This task is conceptually similar to previous experiments on perceptual discrimination (Ratcliff & Rouder, 1998) in which human subjects had to decide which of two stimuli was brightest. The key difference with our experiment is that our subjects made choices between stimuli associated with different levels of reward at consumption, which is an instance of value-based choice. It is interesting to ask whether both types of tasks are described by the same computational model because, given the large degree of specialization in the brain, there is no a priori reason that the same algorithms would be used in the perceptual and reward domains.

We compare four different versions of the DDM that vary in the number of free parameters that they contain, and in whether the barriers are constant or decrease with time to speed up decision making (Cisek, Puskas, & El-Murr, 2009; Ditterich, 2006a). Since the best fitting model is likely to depend on the speed with which the decisions have to be made, we carry out separate model fitting comparisons for the two different time pressure experimental conditions, which allow us to identify which aspects of the DDM are responsible for any changes in performance.

We found that a popular seven-parameter version of the Ratcliff DDM model (Ratcliff & McKoon, 2008; Ratcliff & Rouder, 1998; Ratcliff, et al., 1999) can account for the data with high-accuracy in both the high and low time pressure conditions. Furthermore, we also found that the changes associated with the introduction of time pressure could be traced to changes in two key model parameters: the barrier height and the noise in the slope of the drift process.

Understanding the conditions under which the DDM can explain the behavioral data in simple value based choice is important for several reasons. First, it is a necessary first step in exploring the extent to which this model can account for the underlying neural computations. Second, since the DDM has been shown to provide an accurate description of data in other domains, the finding that the DDM also provides a good computational description of value-based choices provides insight into the nature of some basic algorithms that might be at work in many different psychological processes.

To the best of our knowledge, the performance of the DDM has not been tested before in the realm of value-

based choice, although it has been extensively tested on other realms. In fact, to the best of our knowledge, it has provided an accurate quantitative characterization of the key aspects of the data in every domain to which it has been applied. For example, the DDM has been tested in human and non-human primates using the Newsome-Shadlen random dot motion perceptual discrimination task (Ditterich, 2006b; Mazurek, Roitman, Ditterich, & Shadlen, 2003; Philiastides, et al., 2006; Ratcliff, et al., 2003; Ratcliff, et al., 2007; Ratcliff & McKoon, 2008; Shadlen & Newsome, 2001).

A few studies have explored the ability of models related to the DDM to explain various common choice patterns. For example, several studies have explored the ability of Decision Field Theory, which is a variant of the DDM, to account for patterns of choice in multi-attribute settings such as choices among lotteries, but have not provided a full experimental test against the data (Busemeyer, Jessup, Johnson, & Townsend, 2006; Busemeyer & Johnson, 2004). Other studies have investigated the ability of the Competing Accumulator Model to accomplish similar goals (Usher & McClelland, 2001, 2004). Although these previous studies have shown that relatively simple computational models of decision-making can account for some stylized facts of the behavioral literature, their full properties in the realm of simple value based choice have not been investigated.

1.1 The Drift Diffusion Model and its variants

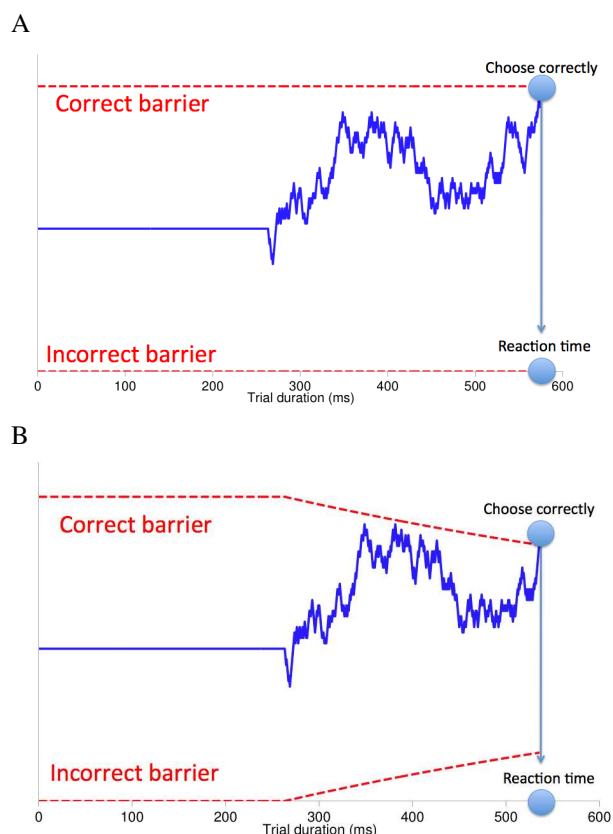
We compare the following four versions of the DDM model.

First, we consider the **simple DDM** (sDDM), which is illustrated in Figure 1A. At every instant, the model encodes a relative value signal that measures the accumulated “evidence” in favor of the hypothesis that the item presented on the left has a higher value than the item on the right (positive values indicate that the left item is better; negative values indicate the opposite). The left item is chosen when the relative value signal crosses the upper barrier; the right item is chosen when the lower barrier is crossed. The relative value signal (RVS) evolves according to the equation

$$X(t) = X(t - 1) + \mu + \epsilon(t),$$

where the drift rate μ denotes the speed at which the barriers are approached, and $\epsilon(t)$ represents white Gaussian noise centered at zero with constant variance σ . Note that $\mu + \epsilon(t)$ measures the local change in the RVS signal in favor of the left alternative at the instant t . We assume that the RVS begins the integration process with a value of zero.

Figure 1: A) Schematic representation of the simple Diffusion Decision Model (sDDM) with three parameters. B) Schematic representation the simple DDM with barriers that decay exponentially towards 0 with time.



The drift rate is a function of the value differences between the left and right items, which we denote by d . In most of the estimation exercises we assume that this function is linear, but we also show that this assumption is consistent with our data.

The sDDM is characterized by the following four parameters: (1) the symmetric location of the barriers ($\pm b$), (2) the linear slope of the drift rate d_m (so that in any trial $\mu = d_m \cdot (v_{left} - v_{right})$, where v_{left} , v_{right} denotes the value of the items), (3) the variance of the diffusion process σ , and (4) a fixed latency time T_m that measures a fixed amount of time out of every trial that takes place prior to the initiation of the comparison process (i.e., the time that passes from the appearance of the items and the beginning of the DDM computations) and after the conclusion of the process (e.g., the time that it takes to execute the motor commands necessary to implement the choice of the DDM). It can be shown that, without loss of generality, the variance parameter can be fixed, since the slope and barrier location are identified only up to ratios of these parameters (Ratcliff & McKoon, 2008; Ratcliff,

et al., 1999). For this reason in the rest of the analyses we set $\sigma = 0.1$, which is a commonly used normalization (Ratcliff & McKoon, 2008). Thus, the model has only three parameters that need to be estimated. This normalization is imposed in all of the four models.

Note several properties of the sDDM. First, although the relative value signal typically moves towards the correct barrier, the process has noise and thus mistakes are made. This is illustrated in Figure 1A. Second, the frequency of mistakes increases with the amount of noise and decreases with the value difference between the two items. Third, accuracy can be improved by increasing the amplitude of the barriers, but this comes at the cost of longer reaction times.

Second, we consider a more complicated version of the DDM with additional parameters, which we refer to as the **full DDM** (fDDM), which has been shown to provide a more accurate description of reaction time distributions than the sDDM (Ratcliff & McKoon, 2008; Ratcliff & Smith, 2004; Ratcliff, et al., 1999). The logic of the model is similar to the basic DDM except for the introduction of several additional parameters that provide the necessary flexibility to fit all of the moments of the response-time data in other experimental domains.


The fDDM is characterized by eight parameters: (1) It includes the four parameters from the sDDM, (2) a standard deviation d_{SD} characterizing the noise with which the drift rate is sampled each period (so that every trial with values v_{left} and v_{right} we have that $d \sim N(d_m \cdot (v_{left} - v_{right}), d_{SD}^2)$, (3) a range T_{range} parameter characterizing the support of a uniform distribution from which the initial latency of the process is sampled every trial (so that every trial $T \sim U(T_m - \frac{T_{range}}{2}, T_m + \frac{T_{range}}{2})$), (4) a parameter z_m that models the mean a bias in the start point of the RVS accumulation process, which allows subjects to be biased towards or away one of the two barriers, with $z_m > 0$ denoting biases towards the left barrier, and $z_m < 0$ denoting biases towards the right barrier, and (5) a range parameter z_{range} characterizing the support of a uniform distribution from which the bias parameter for the start of the RDV signal is sampled every period (so that every trial, the start point of the accumulation process is $X(T) = z$, with $z \sim U(z_m - \frac{z_{range}}{2}, z_m + \frac{z_{range}}{2})$).

Third, we consider a version of the sDDM, which we call the **simple collapsing barrier DDM** (scbDDM), which differs from the sDDM only in that the height of the upper and lower barriers decay exponentially towards 0 with time according to the equation

$$b(t) = e^{-rt}$$

where $r \geq 0$ is a rate constant. This model collapses to the sDDM when $r = 0$. It is characterized by five parameters: the same for parameters describing the sDDM and

Figure 2: Sample experimental trial.



the barrier rate of decline r . The same normalization regarding the variance described above is maintained here.

The motivation for considering time-variant decision thresholds comes from the fact that several papers have previously found that they are useful in accounting for the observation that reaction times tend to be higher in incorrect than correct trials in many psychological tasks (Churchland, Kiani, & Shadlen, 2008; Ditterich, 2006b). As shown in Figure 1B, the collapsing of the barrier can be thought of as an urgency signal that kicks in to preclude the subject from taking excessively long times to reach a decision when the two items are of similar value and thus the drift rate is close to zero. Note, however, that versions of the fDDM have been previously shown to generate slower reaction times in error trials without collapsing barriers (Ratcliff & Rouder, 1998; Ratcliff, et al., 1999), which motivates one of the key questions in the paper: how much additional explanatory power do they provide at the expense of introducing one additional parameter?

Fourth, we consider a version of the fDDM, which we call the **full collapsing barrier DDM (fcbDDM)**, which differs from the fDDM only in the addition of the parameter r measuring the rate at which the barriers collapse towards zero.

2 Methods

Subjects. Eight Caltech students with normal or corrected to normal vision participated in two, one-hour experimental sessions. They were compensated \$20 per session. Subjects were required not to eat for 3 hours before the experiment in order to increase their subjective value for the food items.

Experimental Task. On every trial subjects saw high-resolution color images of 50 different food items including an equal mix of salty and sweet foods (e.g., potato chips and candy bars; see Figure 2 for sample images). Items are widely available in local stores and were highly familiar to our subjects. Subjects were seated in a dimly lit room with their heads positioned in a forehead and chin rest. Eye-position data were acquired from the right

eye at 1000 Hz using the Eyelink 1000 infrared eyetracker (SR Research, Osgoode, Canada). The distance between the computer screen and subject was 80 cm, giving a total visual angle of $28^\circ \times 21^\circ$. The images were presented on a CRT screen (120 Hz) using MATLAB Psychophysics toolbox and Eyelink toolbox extensions (Brainard, 1997; Cornelissen, Peters, & Palmer, 2002).

Each experimental session (one for each of the two conditions described below) began with a liking-rating task. Subjects were shown one food item at a time, centered on the monitor screen and were asked to rate how much they would like to eat each food item at the end of the experiment (-2 to 2 discrete scale). The image was shown until subjects made a response. We used these ratings as independent measures of each subject's subjective value for the items.

In each condition, during the main task, subjects made 750 choices between randomly selected pairs of food items. Figure 2 depicts the timing of the trial. Each trial began by requiring subjects to fixate on a central cross for 800 ms. After the eye-tracker registered a successful fixation, the cross disappeared leaving the blank screen for 200 ms. Immediately after, two different food items were shown simultaneously for 20 ms, centered at 6.2° in the left and right hemifields. Next, two faint circular choice targets were displayed at the same location as the food items to indicate alternative saccade landing positions. Subjects made a choice by fixating on the left or the right target. A 500 ms blank screen separated the trials. At the end of the entire experiment subjects were required to eat whatever food they chose from one randomly selected trial, thus giving them an incentive to select the highest value item of the pair on each trial. The items and locations were randomly selected with the constraint that the two images should have different liking-ratings. d is the absolute value of the difference in liking rating between the two food items; in any one trial, d can be at most 4 and is, by design, never 0. The 200 ms gap between fixation and food display was added because it has been shown to accelerate saccade initiation (Fischer & Weber, 1993), which reduces the impact of motor delays on measured reaction times.

In order to minimize the visual demands of the task, subjects were encouraged, but not required to fixate exactly on the choice targets. The identity of the choice was recorded as soon as a saccade was initiated, as measured by crossing a threshold located 2.2° from the center of the screen on both sides. Reaction times were measured as the time difference between the onset of the images and the initiation of the saccade.

Subjects participated in two different task conditions: high and low time pressure treatments. In the low time pressure (LTP) condition they were asked to indicate their choice only after they were sure which item they preferred. In the high time pressure (HTP) condition they were asked to make their choices as quickly as possible. The order was counterbalanced across subjects. We selected these two treatments to compare the ability of the DDM to account for the data across the range of conditions under which consumers make decisions in the field.

Model fitting procedure. We fit the sDDM and fDDM versions of the models at the individual subject level using the DMA-Toolbox (Vandekerckhove & Tuerlinckx, 2008). We assumed that the slope of the drift process increased linearly with the value difference between the two items (left minus right), and that the upper (lower) barrier corresponded to a left (right) choice.

Note that in this literature it is common to let the bias parameter z_m to be different than zero only if there is evidence for either a response or a reaction time bias to one of the two locations. This was the case in our dataset. Five subjects exhibited either a significant response bias towards left or right (binomial t-test vs. 50% null), or significantly different reaction times for left and right responses (paired two-sided t-test).

The toolbox had to be modified to allow for the maximum likelihood estimation of the scbDDM and fcbDDM versions of the models, since the original code allowed only for constant barrier heights. Here we provide a brief description of the changes made to the DMA-Toolbox to address this problem. The modified toolbox is available from the authors upon request.

As in the basic case, we assumed that the slope of the drift process increased linearly with the value difference between the two items (left minus right), and that the upper (lower) barrier corresponded to a left (right) choice.

Since there is no analytical formula to explicitly calculate the model's predicted distribution of choice accuracy and reaction time in the presence of collapsing barriers, we simulated 1000 trials using a Monte Carlo procedure to approximate these distributions for each candidate set of model parameters to be tested during maximum-likelihood estimation procedure. The results of the simulation were used to calculate the marginal probability of making a correct choice, given the location of the correct stimulus, and to calculate the cumulative distribution

function of the reaction time distribution conditioned on making a correct or incorrect choice and the location of the correct stimulus. The likelihood of the model was computed as the product of the likelihood of the individual trials. The likelihood for each concrete set of parameters was based on the estimate of the probability of choosing the option chosen in that trial, as well as the probability of landing on the same reaction time for that trial. Reaction time bins were given by the 5th, 10th, 30th, 50th, 70th, and 90th percentiles of each subject's reaction time distribution.

To carry out each trial simulation, a random-walk approximation to the drift diffusion process was used (Tuerlinckx, et al., 2001). At the start of each simulation, a barrier height, drift diffusion rate, and non-decision reaction time were sampled according to the candidate parameter set, which contains the mean and variance or range of these quantities. A state variable representing the location of the drift diffusion process was initialized to the bias parameter, and then updated using a time step of 10ms until the state crossed one of the barriers using the rule described below. The location of the absorbing barrier (upper or lower) was used to determine which stimulus was chosen and hence whether this was a correct or incorrect trial, and the number of update steps needed to reach a barrier was scaled and added to the non-decision reaction time and recorded as the net reaction time of the trial.

In each time step of the simulation, there is a probability p that the state variable will increase by a quantity δ , and a probability $(1 - p)$ will decrease by δ . The quantities p and δ are functions of the drift diffusion rate and the time step of the simulation. They are defined so that in the limit, as the time step approaches zero, the random-walk approximation converges to the true drift diffusion process. Specifically, δ is defined as $\sigma \cdot \sqrt{\tau}$ while p is defined as $.5 \cdot (1 + v \cdot \frac{\sqrt{\tau}}{\sigma})$. Here, σ is the standard deviation of the noise process and is fixed at .1, τ is the time step and is set to 10ms, and v is the drift rate for this trial.

The implementation of these changes to the DMA toolbox required two major changes. First, the module that calculates the probability of an observed trial given a set of model parameters was completely rewritten to use the above procedure. Second, the optimization algorithm to find maximum likelihood estimates in the "genalg" module of the toolbox was changed as well. By default, this module uses a simplex optimization procedure. However, Monte Carlo error introduces many erroneous local minima into the objective function and we found that simplex optimization would often become trapped in a local minimum. Thus, in order to increase our ability to find global minima, we instead used a "multistart" procedure as implemented in the MATLAB Global Optimization toolkit. Ten candidate parameter sets were chosen approximately

Table 1: Individual performance by condition, averaged over all values of d .

Subject	N	Accuracy (%)	Mean RT (S.E.M.)		
			All Trials	Correct Trials	Error Trials
LOW TIME PRESSURE					
1	749	75.8	436 (4.77)	444 (5.73)	411 (7.83)
2	749	81.0	514 (4.03)	510 (4.16)	530 (11.46)
3	750	84.4	623 (5.44)	608 (5.77)	704 (13.37)
4	738	90.0	533 (5.89)	534 (6.02)	521 (22.78)
5	719	97.8	811 (10.26)	807 (10.22)	982 (98.29)
6	750	85.3	681 (6.11)	671 (6.34)	737 (18.60)
7	746	85.7	480 (6.23)	484 (6.37)	452 (17.53)
8	738	66.8	520 (6.73)	510 (6.67)	540 (13.10)
MEAN	742	83.2	574	578	552
HIGH TIME PRESSURE					
1	749	76.5	343 (2.88)	344 (3.24)	339 (6.26)
2	744	73.1	497 (3.13)	496 (3.64)	498 (6.14)
3	747	75.2	479 (4.13)	475 (4.40)	492 (9.94)
4	747	85.3	404 (3.64)	398 (3.76)	439 (11.26)
5	745	80.1	406 (2.44)	405 (2.57)	409 (6.57)
6	738	71.2	520 (3.70)	521 (4.45)	518 (6.60)
7	747	63.5	325 (2.61)	335 (3.35)	307 (3.92)
8	748	80.7	436 (3.26)	434 (3.49)	446 (8.49)
MEAN	745	75.7	426	426	426

uniformly in the space of plausible parameter sets, and a separate simulated annealing procedure was run using each parameter set as a starting point of the annealing process. The best solution from this procedure was then run through a simplex algorithm to procedure the final estimate of the optimal parameter set for each subject. The results of this procedure were robust to changes in the random seed used for annealing and choosing the starting candidate parameter sets, suggesting the parameter sets we found were close to the true global optimum.

Model selection. We computed the Bayes Information Criterion (BIC) of each model, which is given as $-2 \cdot \log - \text{likelihood}(data|model) + k$, where k is the number of free parameters in the model, and N is the number of trials used to fit the model. The relative fit of each pair of candidate models was tested as follows. First, we computed the difference in BIC scores for each individual subject and model. Second, we computed populations statistics by carrying out a two-sided paired t-test over the individual BIC score differences. If this difference was significant at the .05 level, then the model with the lower BIC was said to be a better fit of the data. This

method allowed us to rank the models from best to worse fitting. The ordering of models did not change when the Akaike Information Criterion (AIC) was used instead of the BIC.


3 Results

Behavioral results. Table 1 and Figure 3 depict the main behavioral results of the experiment. Subjects chose the best item (as determined by their own liking-ratings) more frequently in the low time pressure than in the high time pressure condition (83.2%, SEM=3.27 vs. 75.7%, SEM=2.37; two-sided paired t-test, $t=1.9$, $p=.09$), even after controlling for the value distance d between the two items (as measured by the liking-rating of the best minus the worse items). Furthermore, choice accuracy (defined as the probability of choosing the item with the highest liking rating) increased with the value distance in both conditions ($t=-7.21$, $p=.0004$ for a two-sided paired t-test of accuracy at $d=1$ versus $d=4$ in the speed condition; $t=-6.84$, $p=.0005$ for the same test in the accuracy condition).

Table 2: Average BIC values for each model and condition. SEMs across subjects are listed in parenthesis.

	LOW TIME PRESSURE	HIGH TIME PRESSURE	ALL TRIALS
sDDM	2765 (52.66)	2552 (196.32)	2659 (92.53)
fDDM	2676 (50.20)	2491 (199.70)	2584 (101.71)
scbDDM	2684 (30.44)	2532 (200.29)	2596 (104.68)
fcbbDDM	2679 (51.21)	2496 (197.85)	2588 (101.46)

Figure 3: A) Choice accuracies in the low time pressure (red dashed) and high-time pressure (blue solid) conditions (N=8). B) Reaction times in both conditions. Bars denote SEMs. Horizontal tics are offset for clarity.



Consistent with the experimental manipulation, subjects made choices faster in the high time pressure condition than in the low time pressure condition (426 ms, SEM=24.91 vs. 574 ms, SEM=43.53; two-sided paired t-test, $t=3.68$, $p=.007$), even after controlling for value distance. Reaction times decreased with d , in the low time pressure condition (from the mean of 616 ms for $d=1$ to 517 ms for $d=4$; $t=4.12$, $p=.006$) but only slightly so in the high time pressure condition (from the mean of 433ms for $d=1$ to 401ms for $d=4$; $t=3.56$, $p=.01$).

Model fitting in the low time pressure condition. We carried out all possible pair-wise comparisons of the BIC measures of quality of fit among the four models using only the data from the low-time pressure condition. (See methods for details.) We found that the fDDM has a significantly lower BIC score than the sDDM ($t=-3.25$, $p<.014$) and the scbDDM ($t=-2.25$, $p<.05$). While its BIC score is lower than the fcbbDDM, the difference did not reach significance ($t=-1.07$, $p<.32$). We also found that adding a collapsing barrier to the sDDM improved its fitting ($t=-.83$, $p<.43$), but that adding it to the fDDM did not ($t=1.07$, $p<.32$), although in neither case did the difference reach significance. Table 2 summarizes the BIC scores for every model and condition. Table 3 provides a full description of the estimated parameters for every subject in the fDDM in each condition.

Model fitting in the high time pressure condition. We also carried out all possible pair-wise comparisons of the BIC measures of quality of fit among the four models using only the data from the high-time pressure condition. The reason for carrying out this model comparison separately was to rule out the *a priori* intuitive possibility that collapsing barriers might play a role under high time pressure, but not under low time pressure. We found very similar results: the fDDM has a lower BIC score than the sDDM ($t=-3.71$, $p<.008$), the scbDDM ($t=-.25$, $p<.81$), and the fcbbDDM ($t=-1.89$, $p<.10$). Adding a collapsing barrier to the sDDM improved its fitting ($t=3.32$, $p<.01$), but adding it to the fDDM did not ($t=1.89$, $p<.1$).

Model fitting in the joint dataset. Finally, we carried out all possible pair-wise comparisons in the full dataset to investigate which model accounted best for the full-dataset. We found very similar results: the fDDM has a lower BIC score than the sDDM ($t=-5.12$, $p<.0015$), the scbDDM ($t=-1.74$, $p<.13$), and the fcbbDDM ($t=-2.17$, $p<.067$). Adding a collapsing barrier to the sDDM improved its fitting ($t=3.00$, $p<.02$), but adding it to the fDDM did not ($t=-2.17$, $p<.067$).

Fits of the fDDM model. The previous results show that the fDDM provides the best relative fit of the data, but is it good in absolute terms? To investigate this question

Table 3: Estimated parameters for every subject in the fDDM by time pressure condition. The p-values listed at the bottom are for a comparison of the distribution of individual parameters across the two time pressure conditions.

Subject	b (units)	d _m (unit/s)	d _{SD} (units/s)	T _m (ms)	T _{range} (ms)	Z _m (units)	Z _{range} (units)
HIGH TIME PRESSURE							
1	0.065	0.094	0.091	245	65	-0.0005	0.001
2	0.063	0.094	0.142	397	142	-0.0005	0.001
3	0.053	0.174	0.415	434	335	-0.0005	0.026
4	0.075	0.157	0.247	301	186	-0.0005	0.011
5	0.062	0.165	0.217	336	141	0	0.028
6	0.054	0.107	0.180	456	228	0	0.045
7	0.064	0.162	0.325	351	164	0.0003	0.016
8	0.044	0.065	0.129	280	151	0	0.021
mean	0.060	0.127	0.218	350	176	-0.0002	0.019
LOW TIME PRESSURE							
1	0.080	0.079	0.286	316	23	0	0.02
2	0.085	0.15	0.139	322	149	-0.0005	0.03
3	0.12	0.164	0.239	436	245	0.055	0.016
4	0.088	0.134	0.036	392	375	-0.003	0.012
5	0.255	0.164	0.084	420	334	-0.0005	0.186
6	0.093	0.123	0.058	521	312	-0.0065	0.015
7	0.090	0.083	0.177	359	186	0	0.078
8	0.115	0.089	0.012	270	15	-0.0005	0.083
mean	0.102	0.123	0.129	379	204	0.0055	0.055
p-values for two-sided paired t-test comparison across conditions							
	0.033	0.78	0.09	0.19	0.53	0.45	0.12

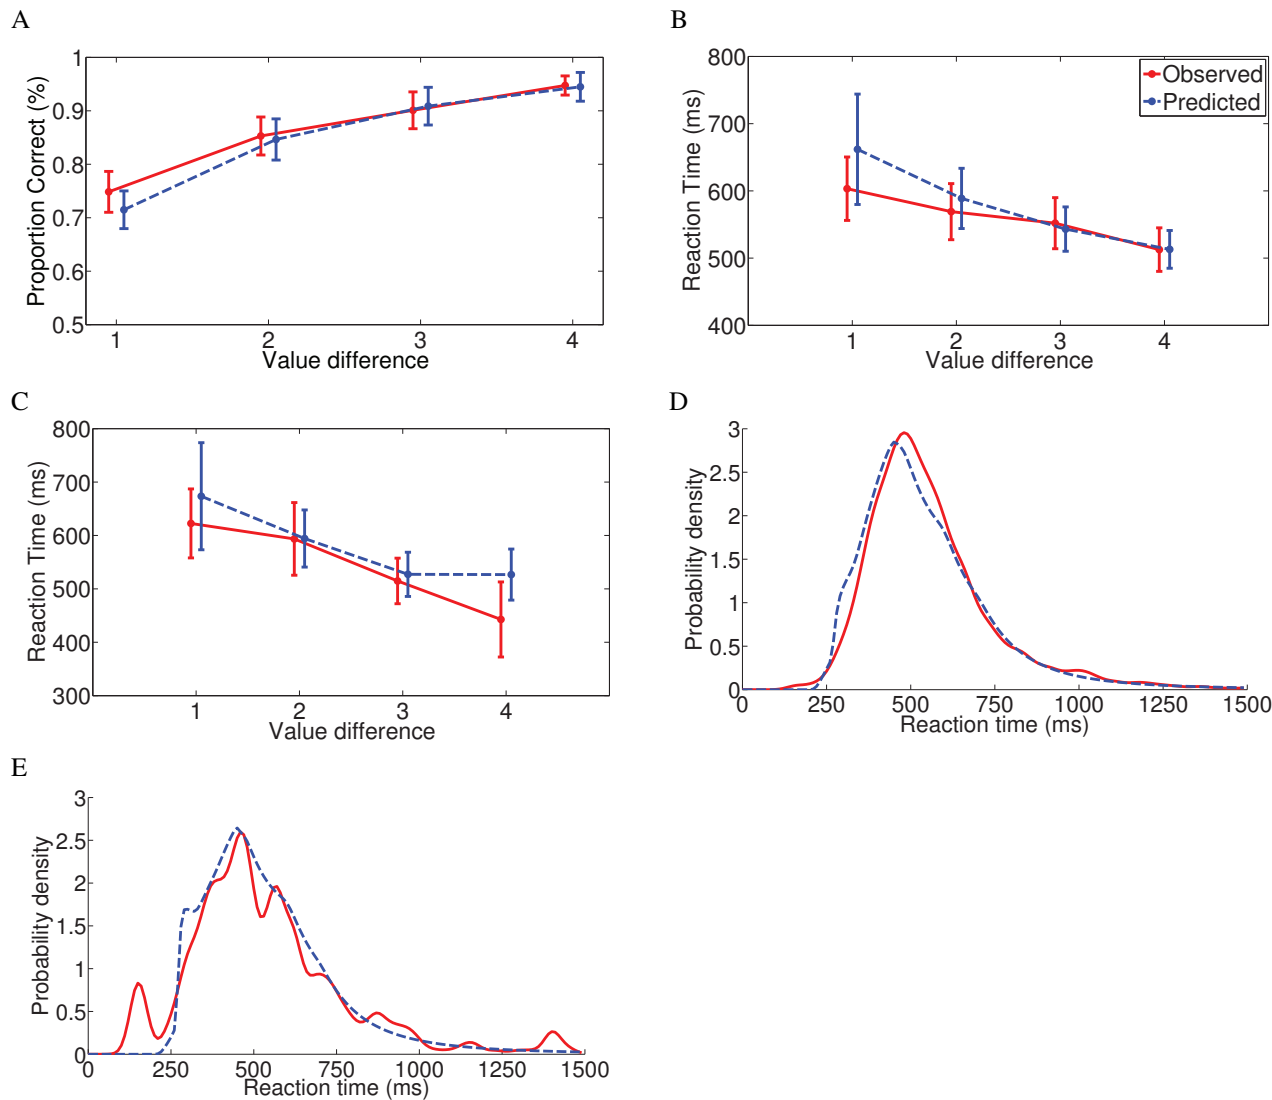
we carried out 1000 simulations of each observation in the dataset, by running the fDDM with the estimated parameters for each condition given the relative value d for those trials. Figure 4 displays the results for the low time pressure trials. Figure 5 displays them for the high time pressure trials. In Figure 4A-C and 5A-C, paired t-tests at every level of difficulty d fail to reject the null hypothesis that these data were sampled from the theoretical fDDM reaction time distribution for that level of difficulty.

Testing the linearity of the drift rate. All of the estimation and simulation exercises described above assume that the drift rate is linear in the value difference variable d . We tested the validity of this assumption by estimating a version of the sDDM in which the parameters were allowed to change by difficulty (given by value best minus value worse). The results, depicted in Figure 6, show that the linear assumption is consistent with the data in the low time pressure ($r^2 = .99$, $p < .00005$) and high time

pressure ($r^2 = .95$, $p < .0005$) conditions.

The effect of time pressure. We compared the estimated parameters of the fDDM that were fitted subject by subject in each of the two conditions separately. This allowed us to assess which parameters of the DDM are essential to explain the differences between the two conditions. The results are described in Table 3. The only two parameters that changes significantly or marginally significantly from the low to the high time pressure conditions are a decrease in the barrier height and an increase in the noise of the drift slope. These changes are highly intuitive. A decrease in the barrier height speeds up choices by making it easier to reach the barrier. A decrease in processing time might decrease the accuracy with which the underlying value of the items can be measured, which implies that the integration slope should be noisier. Note that the results are likely to be marginally significant due to the low number of subjects.

Figure 4: Fits of the estimated fDDM in the low time pressure condition. A) Probability that the best item is chosen as a function of difficulty (equal to value best — value worse) in correct trials. B) Mean reaction time as a function of difficulty in correct trials. C) Mean reaction time as a function of difficulty in incorrect trials. D) Reaction time distribution in correct trials. E) Reaction time distribution in incorrect trials. Bars denote SEMs. The horizontal ticks are offset for clarity.



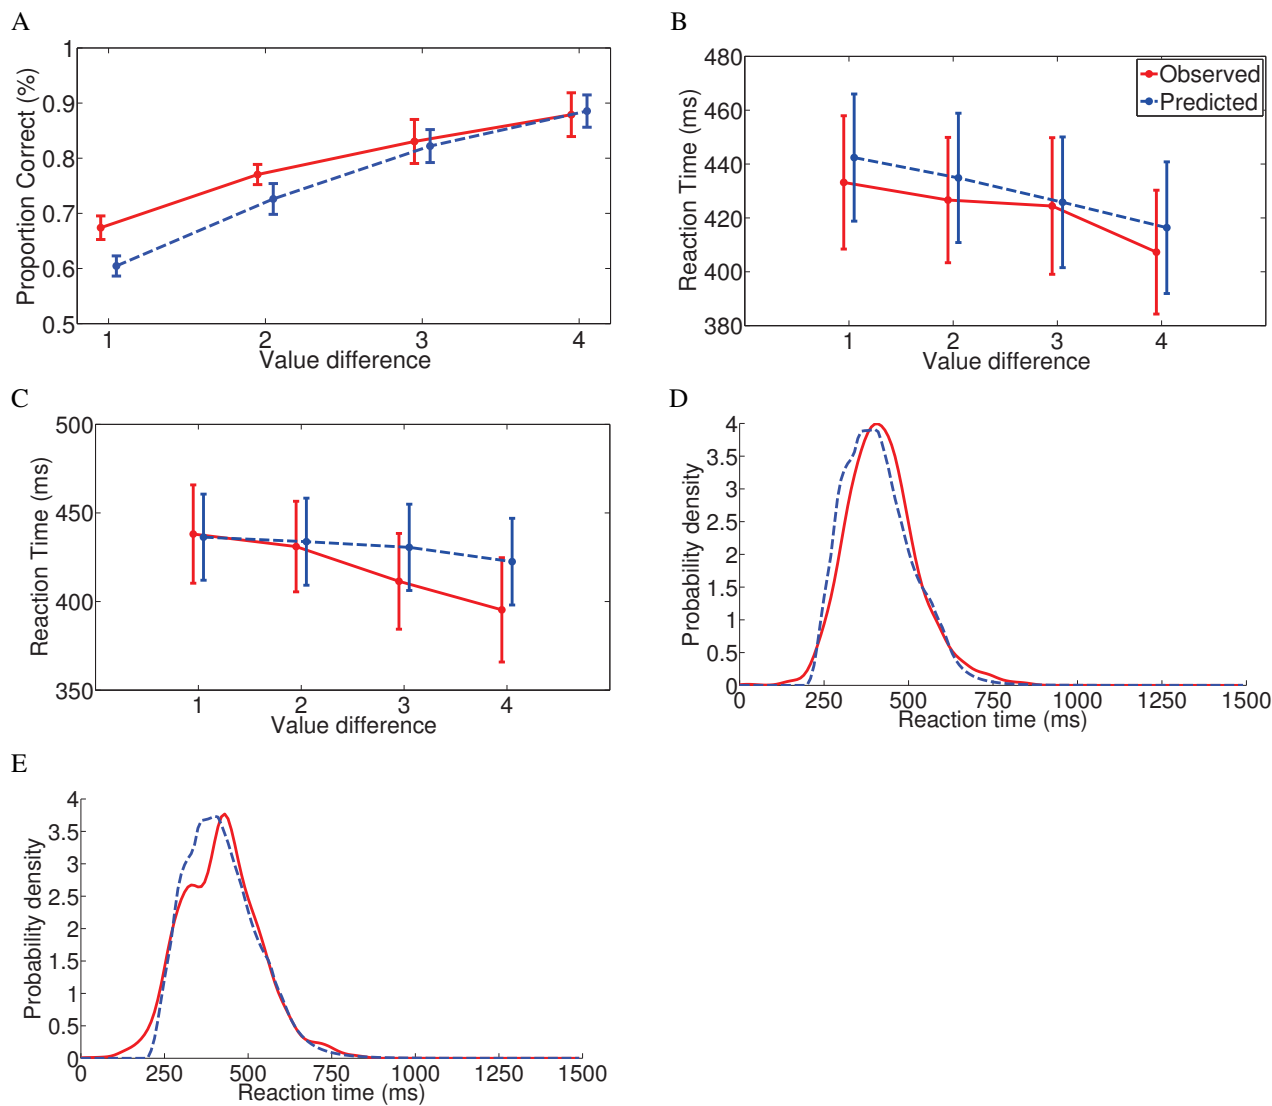
4 Discussion

Several behavioral neuroscience groups have speculated that the brain might use a version of the DDM to make simple value based choices. A first step in testing this hypothesis is to investigate the extent to which the model can explain basic data. Our results show that a popular seven-parameter version of the DDM (Ratcliff & McKoon, 2008; Ratcliff & Rouder, 1998; Ratcliff & Smith, 2004; Ratcliff, et al., 1999) can account for the data with high-accuracy in both the high and low time pressure conditions. Furthermore, we also found that the changes as-

sociated with the introduction of time pressure could be traced to changes in two key model parameters: the barrier height and the noise in the slope of the drift process.

We explored the extent to which the fits of the DDM improved with the introduction of collapsing barriers, which some groups have argued are important to account for the full distribution of reaction times in correct and incorrect trials (Churchland, et al., 2008; Ditterich, 2006b). We found that the introduction of this feature was useful when the baseline model was a simple DDM, but not when it was the popular full version of the DDM.

Figure 5: Fits of the estimated fDDM in the high time pressure condition. A) Probability that the best item is chosen as a function of difficulty (equal to value best — value worse). B) Mean reaction time as a function of difficulty in correct trials. C) Mean reaction time as a function of difficulty in incorrect trials. D) Reaction time distribution in correct trials. E) Reaction time distribution in incorrect trials. Bars denote SEMs. The horizontal ticks are offset for clarity.

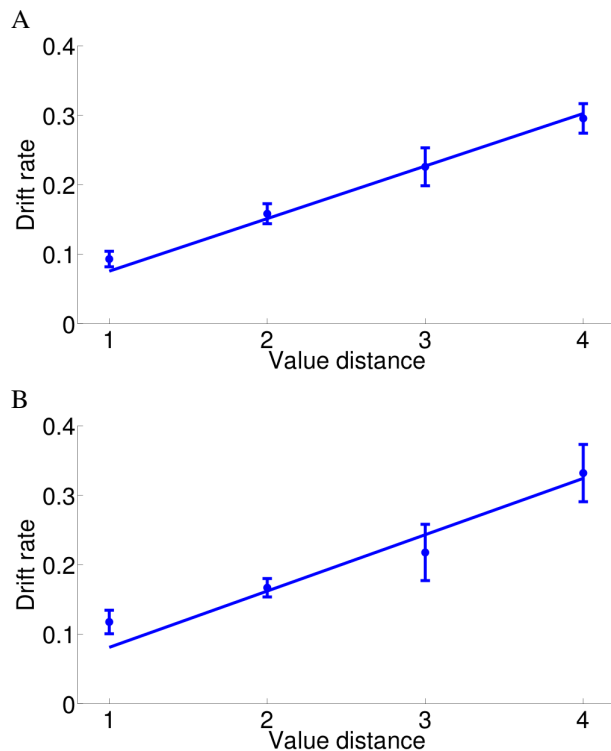


Previous studies have also explored the effect of time on the ability of the DDM to explain the data in perceptual discrimination tasks (Ratcliff, 2002; Ratcliff & Rouder, 1998; Ratcliff, Thapar, & McKoon, 2003; Thapar, Ratcliff, & McKoon, 2003). Their findings parallel ours: a change in parameters of the DDM, specially an increase in the size of the barriers as time pressure increases, is able to provide a highly accurate description of the response and reaction time data.

One of the most successful applications of the DDM is to perceptual decision making tasks (Churchland, Kiani,

& Shadlen, 2008; Ditterich, 2006b; Ditterich, Mazurek, & Shadlen, 2003; Gold & Shadlen, 2001, 2002; Gold & Shadlen, 2007; Hanks, Ditterich, & Shadlen, 2006; Huk & Shadlen, 2005; Mazurek, et al., 2003; Palmer, Huk, & Shadlen, 2005; Philiastides et al., 2006; Ratcliff et al., 2003; Ratcliff et al., 2007; Ratcliff et al., 2003; Roitman & Shadlen, 2002; Shadlen & Newsome, 1996, 2001; Smith, Lee, Wolfgang, & Ratcliff, 2009). There is an important difference between the applications to perceptual and value based decision-making. In the standard Newsome-Shadlen random dot motion task, subjects are

Figure 6: Estimated fitted drift rates in the sDDM by value distance. A) Low time pressure condition. B) High time pressure condition. Bars denote SEMs.



exposed to a stochastic stimulus that is assumed to generate perceptual noise signals in area MT. Under appropriate assumptions, it can be shown that the drift-diffusion model implements an optimal decision-making process that amounts to a sequential-likelihood ratio test (Bogacz, 2007; Bogacz et al., 2006; Gold & Shadlen, 2001, 2002; Gold & Shadlen, 2007). In our model the stimuli are non-stochastic, in the sense that the image is non-changing. However, we hypothesize that in order to construct value the brain needs to integrate a series of noisy signals about the value of the stimuli, in this case generated internally. In particular, we hypothesize that the brain assigns value to the stimuli by sequentially and stochastically extracting features of the stimuli, retrieving the learnt values for such features, and then integrating those values. Although the objective nature of the noise is quite different in both cases, the computational problem has similar properties.

An important question for future research is how does the brain implement each of the features of the full DDM during value based choices, which we have shown provides a good quantitative description of the data.

References

- Bogacz, R. (2007). Optimal decision-making theories: linking neurobiology with behaviour. *Trends in Cognitive Science, 11*, 118–125.
- Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (2006). The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced choice tasks. *Psychological Review, 113*, 700–765.
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision, 10*, 433–436.
- Busemeyer, J. R., Jessup, R. K., Johnson, J. G., & Townsend, J. T. (2006). Building bridges between neural models and complex decision making behaviour. *Neural Netw, 19*, 1047–1058.
- Busemeyer, J. R., & Johnson, J. G. (2004). *Computational models of decision making*. Cambridge, MA: Blackwell.
- Decision-making with multiple alternatives, 11 Cong. Rec. 693–702 (2008).
- Churchland, A. K., Kiani, R., & Shadlen, M. N. (2008). Decision-making with multiple alternatives. *Nature Neuroscience, 11*, 693–702.
- Cisek, P., Puskas, G. A., & El-Murr, S. (2009). Decisions in changing conditions: the urgency-gating model. *Journal of Neuroscience, 29*, 11560–11571.
- Cornelissen, F. W., Peters, E. M., & Palmer, J. (2002). The Eyelink Toolbox: eye tracking with MATLAB and the Psychophysics Toolbox. *Behavior Research Methods Instrumentats and Computers, 34*, 613–617.
- Ditterich, J. (2006a). Evidence for time-variant decision making. *European Journal of Neuroscience, 24*, 3628–3641.
- Ditterich, J. (2006b). Stochastic models of decisions about motion direction: behavior and physiology. *Neural Networks, 19*, 981–1012.
- Ditterich, J., Mazurek, M. E., & Shadlen, M. N. (2003). Microstimulation of visual cortex affects the speed of perceptual decisions. *Nature Neuroscience, 6*, 891–898.
- Fischer, B., & Weber, H. (1993). Express saccades and visual attention. *Behavioral and Brain Sciences, 16*, 553–610.
- Gold, J. I., & Shadlen, M. N. (2001). Neural computations that underlie decisions about sensory stimuli. *Trends in Cognitive Sciences, 5*, 10–16.
- Gold, J. I., & Shadlen, M. N. (2002). Banburisms and the brain: Decoding the relationship between sensory stimuli, decisions, and reward. *Neuron, 36*, 299–308.

- Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, *30*, 535–574.
- Hanks, T. D., Ditterich, J., & Shadlen, M. N. (2006). Microstimulation of macaque area LIP affects decision-making in a motion discrimination task. *Nature Neuroscience*, *9*, 682–689.
- Hare, T., Camerer, C., Knoepfle, D., O'Doherty, J., & Rangel, A. (2009). Value computations in VMPFC during charitable decision-making incorporate input from regions involved in social cognition. *Journal of Neuroscience*, *13*, 583–590.
- Hare, T., Camerer, C., & Rangel, A. (2009). Self-control in decision-making involves modulation of the vMPFC valuation system. *Science*, *324*, 646–648.
- Hare, T. A., O'Doherty, J., Camerer, C. F., Schultz, W., & Rangel, A. (2008). Dissociating the role of the orbitofrontal cortex and the striatum in the computation of goal values and prediction errors. *Journal of Neuroscience*, *28*, 5623–5630.
- Heekeren, H. R., Marrett, S., & Ungerleider, L. G. (2008). The neural systems that mediate human perceptual decision making. *Nat Rev Neurosci*.
- Huk, A. C., & Shadlen, M. N. (2005). Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. *Journal of Neuroscience*, *25*, 10420–10436.
- Link, S. W. (1992). *The wave theory of difference and similarity*. Cambridge: Psychology Press.
- Mazurek, M. E., Roitman, J. D., Ditterich, J., & Shadlen, M. N. (2003). A role for neural integrators in perceptual decision making. *Cerebral Cortex*, *13*, 1257–1269.
- Padoa-Schioppa, C., & Assad, J. A. (2006). Neurons in the orbitofrontal cortex encode economic value. *Nature*, *441*, 223–226.
- Palmer, J., Huk, A. C., & Shadlen, M. N. (2005). The effect of stimulus strength on the speed and accuracy of a perceptual decision. *Journal of Vision*, *5*, 376–404.
- Philiastides, M. G., Ratcliff, R., & Sajda, P. (2006). Neural representation of task difficulty and decision making during perceptual categorization: a timing diagram. *Journal of Neuroscience*, *26*, 8965–8975.
- Rangel, A. (2008). The computation and comparison of value in goal-directed choice. In P. W. Glimcher, C. F. Camerer, E. Fehr & R. A. Poldrack (Eds.), *Neuroeconomics: Decision Making and the Brain*. New York: Elsevier.
- Rangel, A., Camerer, C., & Montague, P. R. (2008). A framework for studying the neurobiology of value-based decision making. *Nature Review of Neuroscience*, *9*, 545–556.
- Rangel, A., & Hare, T. (2010). Neural computations associated with goal-directed choice. *Current Opinion in Neurobiology*, *20*, 262–270.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*, 59–108.
- Ratcliff, R. (2002). A diffusion model account of reaction time and accuracy in a brightness discrimination task: Fitting real data and failing to fit fake but plausible data. *Psychonomic Bulletin and Review*, *9*, 278–291.
- Ratcliff, R., Cherian, A., & Segraves, M. (2003). A comparison of macaque behavior and superior colliculus neuronal activity to predictions from models of two-choice decisions. *Journal of Neurophysiology*, *90*, 1392–1407.
- Ratcliff, R., Hasegawa, Y. T., Hasegawa, R. P., Smith, P. L., & Segraves, M. A. (2007). Dual diffusion model for single-cell recording data from the superior colliculus in a brightness-discrimination task. *Journal of Neurophysiology*, *97*, 1756–1774.
- Ratcliff, R., & Rouder, J. N. (2000). A diffusion model account of masking in two-choice letter identification. *Journal of Experimental Psychology: Human Perception and Performance*, *26*, 127–140.
- Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation*, *20*, 873–922.
- Ratcliff, R., Philiastides, M. G., & Sajda, P. (2009). Quality of evidence for perceptual decision making is indexed by trial-to-trial variability of the EEG. *Proceedings of the National Academy of Sciences*, *106*, 6539–6544.
- Ratcliff, R., & Rouder, J. (1998). Modelling response times for two-choice decisions. *Psychological Science*, *9*, 347–356.
- Ratcliff, R., & Smith, P. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychological Review*, *111*, 333–367.
- Ratcliff, R., Thapar, A., & McKoon, G. (2003). A diffusion model analysis of the effects of aging on brightness discrimination. *Perception and Psychophysics*, *65*, 523–535.
- Ratcliff, R., Van Zandt, T., & McKoon, G. (1999). Connectionist and diffusion models of reaction time. *Psychological Review*, *106*, 261–300.
- Roitman, J. D., & Shadlen, M. N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *Journal of Neuroscience*, *22*, 9475–9489.

- Shadlen, M. N., & Newsome, W. T. (1996). Motion perception: seeing and deciding. *Proceedings of the National Academy of Sciences*, *93*, 628–633.
- Shadlen, M. N., & Newsome, W. T. (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *Journal of Neurophysiology*, *86*, 1916–1936.
- Smith, P. L., Lee, Y. E., Wolfgang, B. J., & Ratcliff, R. (2009). Attention and the detection of masked radial frequency patterns: Data and model. *Vision Research*, *49*, 1363–1377.
- Smith, P. L., & Ratcliff, R. (2004). Psychology and neurobiology of simple decisions. *Trends in Neuroscience*, *27*, 161–168.
- Thapar, A., Ratcliff, R., & McKoon, G. (2003). A diffusion model analysis of the effects of aging on letter discrimination. *Psychology and Aging*, *18*, 415–429.
- Tom, S. M., Fox, C. R., Trepel, C., & Poldrack, R. A. (2007). The neural basis of loss aversion in decision-making under risk. *Science*, *315*, 515–518.
- Tuerlinckx, F., Maris, E., Ratcliff, R., & De Boeck, P. (2001). A comparison of four methods for simulating the diffusion process. *Behavior Research Methods, Instruments, & Computers*, *33*, 443–456.
- Usher, M., & McClelland, J. (2001). The time course of perceptual choice: the leaky, competing accumulator model. *Psychological Review*, *108*, 550–592.
- Usher, M., & McClelland, J. L. (2004). Loss aversion and inhibition in dynamical models of multialternative choice. *Psychological Review*, *111*, 757–769.
- Vandekerckhove, J., & Tuerlinckx, F. (2008). Diffusion model analysis with MATLAB: A DMAT primer. *Behavioral Research Methods*, *40*, 61–72.
- Wallis, J. D. (2007). Orbitofrontal cortex and its contribution to decision-making. *Annual Review of Neuroscience*, *30*, 31–56.