



# Testing Optimal Timing in Value-Linked Decision Making

Rahul Bhui<sup>1</sup>

© Society for Mathematical Psychology 2019

## Abstract

Influential theories of the decision making process hold that a choice is made once the cumulative weight of noisily sampled information reaches a desired level. While these theories were originally motivated as optimal solutions to statistical problems, the extent to which people optimally spend time deliberating is less well explored. I conduct an experimental test of optimality in a setting where the speed of information processing reflects the difference in value between options. In this case, spending a long time without having arrived at a conclusion signals both that the problem is hard and that the options are similar in value, so the confidence level required to trigger a decision should decline over time. I find that a recently developed theory of the optimal time-varying threshold improves model fit by accurately predicting observed truncation of response time tails. Principles of optimality may thus help account for patterns of choice and response time that characterize the process of deliberation.

**Keywords** Decision making · Response time · Sequential sampling models · Optimality

When making a decision, we typically face a tradeoff between the speed of our choice and its resulting quality. An increasingly popular class of theories, known as sequential sampling models, precisely characterizes the nature of this tradeoff by detailing the process of deliberation in a psychologically and neurally plausible fashion (for review, see Forstmann et al. 2016; Ratcliff et al. 2016). Such theories describe decision making as the product of two key elements: the stochastic accumulation of evidence over time in favor of each option and the decision criterion that determines when sufficient evidence has been amassed to commit to a choice. Sequential sampling models were originally derived from efficient statistical algorithms, in which the decision criterion reflects the optimal balance between speed and accuracy (Wald 1947; Stone 1960; Bogacz et al. 2006); however, the extent to which people optimally spend time

deliberating has remained underexplored. This paper tests a new theoretical prediction made by Fudenberg et al. (2018) about the economically optimal criterion in a setting where the accumulation rate reflects the value difference between options.

Sequential sampling models have recently been applied to choice among economic goods assuming that the accumulation rate is proportional to the difference between the subjective values of items (Krajbich et al. 2010; Krajbich et al. 2015; Krajbich et al. 2012; Krajbich et al. 2014; Milosavljevic et al. 2010; Krajbich and Rangel 2011).<sup>1</sup> In this *value-based* setting, not only is the speed of processing uncertain (because values are not known in advance), but it is also directly linked to potential rewards. Hence, two things are signaled when time passes but the decision maker has not yet selected an option: the task is difficult (Drugowitsch et al. 2012; Moran 2015), and the decision maker is close to indifference between the goods. Both of these imply that continued deliberation is less profitable, and consequently, the standard of confidence required to commit to a choice should decline over time. Fudenberg et al. (2018) precisely characterize the optimal threshold in this framework

---

**Electronic supplementary material** The online version of this article (<https://doi.org/10.1007/s42113-019-0025-9>) contains supplementary material, which is available to authorized users.

---

✉ Rahul Bhui  
rbhui@g.harvard.edu

<sup>1</sup> Departments of Psychology and Economics & Center for Brain Science, Harvard University, 52 Oxford St, Cambridge, MA 02138, USA

<sup>1</sup> This body of work builds on related applications which do not explicitly tie the accumulation rate to value (e.g., Hawkins et al. 2014; Otter et al. 2008; Trueblood et al. 2014).

and show that it declines at an approximately hyperbolic rate when values are normally distributed (see also Tajima et al. 2016 who compute numerical solutions). In an experiment designed to adhere closely to their theoretical assumptions, I test their predictions against comparably parametrized versions of the drift diffusion model (DDM) with constant boundaries.

My experiment is based on a modification of the traditional random dot motion task in which participants must ascertain the direction of consistent dot motion among noise dots that distractingly move in random directions (Britten et al. 1992; Newsome et al. 1989). I use a bi-directional variant in which groups of dots move consistently in two directions rather than one (in addition to noise dots). Crucially, I set the payoff from choosing a direction to be proportional to the number of dots moving in that direction. Therefore, the ease of processing, ultimately reflected by the diffusion model drift rate, contains information about the relative value of options as the theory requires. This broad type of setting—used also in studies such as Oud et al. (2016) and Pirrone et al. (2018b)—might be called *value-linked* decision making in contrast to value-based, as the accumulation process is merely correlated with differences in value rather than being driven by them. Such an approach permits precise control over the distribution of values, which is important for the theory to be applied appropriately, but which would be challenging to enforce with consumer goods due to subjectivity of preference. This paradigm thus forms a bridge between perceptual and value-based decision making that allows the theory to be tested under suitable circumstances.

I find that collapsing boundaries provide a better fit to behavior than flat boundaries for most participants. Although even the basic fixed threshold DDM captures mean response time and accuracy data well, only collapsing boundaries accurately predict observed truncation in the tails of the response time distribution, in line with their adaptive purpose of curtailing wasteful deliberation. Fudenberg et al. (2018) derive a simple approximation to the optimal threshold that parametrizes a hyperbolic function using basic diffusion model parameters, yielding a two-parameter collapsing boundary DDM (plus non-decision time). This approximation is strongly favored over flat and generic hyperbolic thresholds by an aggregate Bayesian model selection procedure. Overall, participant behavior appears consistent with predictions of optimality in the present scenario. It must be noted that the analysis is confined to models with a relatively small number of parameters. This restriction is necessary because the identification of collapsing threshold models is known to be impaired by extra parameters used to extend the DDM (Voskuilen et al. 2016). Nonetheless, this evidence indicates at least that principled models of collapsing boundaries are

useful in that they improve predictive power while economizing on parameters.

## Uncertain-Difference Drift Diffusion Model

Many situations involve a binary choice where the difference in value between options is variable and uncertain. For instance, an employer deciding between two potential employees is unsure not only about which of them is better, but also by how much. Both candidates might be about equally effective, in which case lengthy deliberation will not produce much additional value—or one candidate might substantially increase profits while the other ruins a project, in which case lengthy deliberation is essential. In value-based DDMs, the accumulation rate is based on the difference in values between options (e.g., Krajbich et al. 2014). If one option is much better, then confidence tends to rise quickly. Thus, the confidence trajectory carries information about the value gap in addition to option rank. Spending a long time deliberating without coming to a conclusion accordingly implies that the options are similar in value. The agent should then curtail their deliberation time, contributing to collapse of the decision threshold.

Fudenberg et al. (2018) develop analytical results characterizing the optimal rate of collapse. In their setup, the agent is faced with two options,  $i \in \{l, r\}$ , which have unknown values,  $(v^l, v^r) \in \mathbb{R}^2$ . The agent holds a prior belief about these values,  $\mu_0 \in \Delta(\mathbb{R}^2)$ , and observes a two-dimensional signal  $(Z_t^i)_{t \in \mathbb{R}_+}$  which, as in the DDM, evolves according to a Wiener process with drift:

$$dZ_t^i = v^i dt + \alpha dB_t^i \tag{1}$$

where  $\alpha$  is the noisiness of the signal and  $\{dB_t^i\}$  are independent Brownian motions. They continuously update their belief about the values, holding a posterior mean for  $v^j$  denoted  $X_t^i = E[v^j | \{Z_s^i\}_{0 \leq s < t}]$  that is conditioned on the signal trajectory up to each point in time.<sup>2</sup>

The agent must determine both which option to select and when to stop deliberating. A flow cost  $c > 0$  is assumed to be incurred for each moment of time spent before a selection is made, which is interpreted as the opportunity cost of time—that is, the value of the best alternative activity (outside of the task) that the agent could be engaging in. They must choose a stopping time  $\tau$  from the set of all stopping times  $T$ , which yields a cost of  $c\tau$ . Observe that when the agent stops at time  $\tau$ , they optimally choose the option  $i$  with the highest posterior

<sup>2</sup> For more on Bayesian formulations of diffusion models, see Bitzer et al. (2014) and Fard et al. (2017).

expected value  $X_t^i$ , and so the reward from stopping is  $\max_{i=l,r} X_t^i$ . Thus, they confront the Wald optimality problem,<sup>3</sup>

$$\max_{\tau \in T} E \left[ \max_{i=l,r} X_\tau^i - c\tau \right]. \tag{2}$$

The first term represents the expected reward from the selected option, and the second term represents the cost of the chosen deliberation time. A sufficient statistic for this decision turns out to be the difference in signal values,

$$Z_t \equiv Z_t^l - Z_t^r = (v^l - v^r)t + \alpha\sqrt{2}B_t, \tag{3}$$

where  $B_t = \frac{1}{\sqrt{2}}(B_t^l - B_t^r)$  is a Brownian motion.<sup>4</sup>

The optimal policy depends critically on the prior distribution of values. The flat threshold of the classic DDM can be subsumed within the above framework; it is optimal when there is a fixed certain difference between the values of the two options—that is, when the prior is confined to two possibilities, in both of which one option is better than the other by a fixed positive amount. Formally, this means the prior  $\mu_0$  distributes the probability mass among two states,  $(v_L, v_L)$  and  $(v_L, v_H)$  where  $v_H > v_L$ , and the decision maker earns the high payoff  $v_H$  when making the correct choice ( $l$  in the former state,  $r$  in the latter state). Hence, they know the magnitude of the value difference but not its sign. Under this assumption, they know they are not indifferent even before coming to a conclusion. If the decision maker has spent a long time and  $Z_t$  is still close to 0, they essentially face the same problem they started with, and are thus as willing to continue with deliberation as they were at the outset. The current value of the process  $Z_t$  is a sufficient statistic for stopping, meaning that the past trajectory carries no additional useful information, and so the stopping threshold does not change over time.

The uncertain-difference model instead assumes a Gaussian prior on option value,  $v^i \sim N(X_0, \sigma_0^2)$ . In this case, spending a long time without coming to a conclusion carries information about the value difference. It implies that the agent is probably nearly indifferent between options and

should therefore cut short their deliberation and decide quickly. Thus, the stopping threshold decreases over time, and Fudenberg et al. (2018) characterize the optimal rule under this Gaussian assumption. The implications of their results are most clearly depicted by a function that asymptotically approximates the optimal threshold:

$$\text{boundary } b(t) = \frac{1}{2c(\sigma_0^{-2} + \alpha^{-2}t)} \tag{4}$$

This boundary function has a number of meaningful properties. First, it is hyperbolic in time (i.e., declines at rate  $1/t$ ) and declines asymptotically to zero, meaning that eventually the agent chooses almost at random. Second, it is pointwise decreasing in the cost of time  $c$  which captures the underlying reluctance to deliberate. Third, it is pointwise increasing in the prior variance  $\sigma_0$  which reflects the possible benefit of making a good choice. For example, an agent should require more evidence and spend more time deliberating when comparing houses as opposed to lunch items, because the variance in values is higher for houses. Finally, the rate of collapse is increasing in signal noise  $\alpha$ . When noise is low, spending more time is an especially strong sign that the true value difference is small and hence that continued deliberation has little benefit and should be curtailed. Note that this model does not nest constant thresholds (except in degenerate cases), in contrast to more flexible collapsing boundary models that are sometimes applied (e.g., Hawkins et al. 2015). This property is not an assumption but a theoretical result that can be tested empirically, as will be done in this paper.

## Value-Linked Random Dot Motion Task

### Participants

Twenty-four participants were recruited through the Caltech Social Science Experimental Laboratory, all of whom were college or graduate students at Caltech. They were paid a \$10 show-up fee in addition to payment for performance as described below. One participant is excluded from the following analyses due to a computer error which prevented data collection.

### Procedure

Participants engaged in two blocks of a motion discrimination task, each containing 100 trials. To adhere closely to the structural assumptions made by Fudenberg et al. (2018), a modified version of the random dot motion task (e.g., Newsome et al. 1989; Britten et al. 1992) was used. As usual, participants had to discern the motion direction of consistently

<sup>3</sup> While this is the most widely used optimality criterion across fields of study, it is not the only one that has been proposed (for discussion, see Bogacz et al. 2006; Pirrone et al. 2014; and Bhui 2019). Tajima et al. (2016) calculate numerical solutions for collapsing boundaries under alternative criteria such as reward rate maximization, though Pirrone et al. (2018b) fail to find empirical support for the resulting predictions.

<sup>4</sup> Although this setup bears some resemblance to the standard race model due to the specification of two independent accumulators, the two models should not be confused. The race model assumes that a response is triggered when either accumulator crosses a given threshold. The uncertain-difference DDM instead treats the accumulators as two sources of information which are used to inform the optimal balance between reward and time expenditure. Hence, the decision criterion can be defined based on any combination of accumulator values. The setup entails that the difference in accumulators is a sufficient statistic for solving the optimization problem in Expression (2), and thus the uncertain-difference model boils down to a version of the DDM. See also Bogacz et al. (2006) for further explication of the technical connections between various sequential sampling models.

moving dots against overlapping randomly moving noise dots. However, there were two points of departure from the most common paradigm.

First, rather than the classical stimulus consisting of one group of signal dots moving consistently either left or right and one group of noise dots moving randomly, the modified stimulus comprised two groups of signal dots (one moving consistently left and one moving consistently right) in addition to noise dots. Lam and Kalaska (2014) investigated this task and found that the difference between left and right coherences is indeed what drives the accumulation process.<sup>5</sup> Niwa and Ditterich (2008) used a related stimulus to study multialternative decision making. Their task involved three groups of possible signal dots which moved consistently at 120° to each other. Single- and multi-unit recordings in area LIP of monkeys performing this task revealed that neuronal activity reflects the net evidence in favor of choices (Bollimunta and Ditterich 2012; Bollimunta et al. 2012).

The second difference is that the reward earned in a trial was proportional to the number of dots moving in the chosen direction (the coherence). In this way, the drift rate was tied to the difference in values as encoded in value-based DDMs. Formally, this setup which I refer to as value-linked expands upon Expression (3) so that  $Z_t = (x^j - x^i)t + \alpha\sqrt{2B_t}$  for some stimulus strength  $x^j$ , and  $v^j \propto x^j$ . This is similar to the incentive structure in study 2 of Oud et al. (2016) and Pirrone et al. (2018b). For reasons unrelated to this project, the payment rate for each 10 dots was \$0.02 in the first block and \$0.04 in the second block. That is, if 30 dots were moving left and 20 dots were moving right, picking left would earn \$0.06 (first block) or \$0.12 (second block), while picking right would earn \$0.04 (first block) or \$0.08 (second block).<sup>6</sup>

In line with the theoretical assumptions used to derive collapsing thresholds, the numbers of signal dots moving in each direction were drawn independently from a discretized normal distribution. Participants were explicitly informed of this distribution which had mean 25 and standard deviation 7.<sup>7</sup> Each trial contained 100 dots in total, the remainder of which were noise dots. Two practice trials with thorough explanation were shown at the beginning of the experiment to convey the task in detail. Following all regular trials, participants received feedback displayed for two seconds about the number of dots that were moving in each direction. The average accuracy (defined as the proportion of responses picking the higher value option) was 63.6% in block 1 and 68.2% in block 2. The average

response time was 4.27 seconds in block 1 and 3.69 seconds in block 2.<sup>8</sup>

The task was programmed and displayed using the Psychophysics Toolbox in MATLAB, with 5-pixel-width circular dots moving in a 960 × 960-pixel square aperture. Dots moved at a speed of 1 pixel per frame and had a lifetime of 20 frames (at approximately 60 fps). Each trial was preceded by a fixation cross displayed for 1.5 seconds.

## Model Fitting

I set the (conditional) drift rate as the value differential, that is, the absolute value of the difference between the numbers of dots moving left and moving right.<sup>9</sup> Three models based on the DDM were fitted separately to each individual in each block:

1. Fixed boundary,  $b(t) = b$  (3 parameters: boundary  $b$ , accumulation noise  $\alpha$ , and non-decision time  $T_{nd}$ )
2. Approximately optimal boundary,  $b(t) = \frac{1}{2c(\sigma_0^2 + \alpha^2 t)}$  (3 parameters:<sup>10</sup> cost of time  $c$ , accumulation noise  $\alpha$ , and non-decision time  $T_{nd}$ )
3. Generic hyperbolic boundary,  $b(t) = \frac{1}{g+ht}$  (4 parameters: inverse initial boundary  $g$ , boundary collapse  $h$ , accumulation noise  $\alpha$ , and non-decision time  $T_{nd}$ )

The generic hyperbolic model can be considered a version of the approximately optimal model with certain assumptions relaxed. The approximately optimal model implies that the strength of boundary collapse is tied to the accumulation noise, and in order for all of its parameters to be identifiable, we must suppose that the subjective assessment of prior variance is equal to the true variance. Written in terms of the hyperbolic model,  $g = 2c\sigma_0^{-2}$  and  $h = 2c\alpha^{-2} = g(\frac{\sigma_0}{\alpha})^2$  for a known  $\sigma_0$ . Observe that in the approximately optimal model, the  $c$  parameter is completely free while the  $\alpha$  parameter is informed partly by the accumulation noise, which reduces the total number of parameters by one. If  $\sigma_0$  were allowed to vary on top of this, an extra degree of freedom would be added which severs the link between  $h$  and  $\alpha$ . This relaxation effectively renders  $g$ ,  $h$ , and  $\alpha$  (or equivalent functions thereof) the only distinct identifiable parameters, and is thus equivalent to estimation of a generic hyperbolic threshold. Loosening the restrictions imposed by the approximately optimal model in this way may or may not yield a better fit to the data, and this can be tested via model comparison.

<sup>5</sup> In the present data, an analysis in the Appendix reveals no effect of absolute value on response time, contrary to observations in Teodorescu et al. (2016) and Pirrone et al. (2018a).

<sup>6</sup> Distinctions between blocks will not be explored in the present analysis since the effects of experience, fatigue, and incentives are confounded.

<sup>7</sup> It must be noted that due to Caltech's particular nature, all students have strong quantitative backgrounds and are familiar with the normal distribution.

<sup>8</sup> Relatively long response times were also observed in Lam and Kalaska (2014).

<sup>9</sup> Since the properties of the DDM depend only on the ratios between the drift rate, decision threshold, and accumulation noise, one parameter is routinely fixed at some arbitrary level. Typically, this is the noise parameter, but for consistency with the notation of Fudenberg et al. (2018), instead I fix the (conditional) drift rate and allow the accumulation noise to be a free parameter.

<sup>10</sup> In accordance with the experimental design,  $\sigma_0$  was fixed at 7.



All models were fitted using quantile maximum probability estimation, with response time bins defined by the 10th, 30th, 50th, 70th, and 90th percentiles of the individual’s distribution (Brown and Heathcote 2003; Heathcote et al. 2004; Heathcote et al. 2002). Model predictions were simulated using a random walk approximation with time step 50 ms and roughly 400,000 replicates<sup>11</sup> (Tuerlinckx et al. 2001), and optimization was carried out via differential evolution (Mullen et al. 2011). This was computationally intensive, taking on the order of many thousands of core hours to fit the models for all subjects.

I restrict analysis to models with few parameters because identification in collapsing bounds models is impeded by degrees of freedom in the DDM (Voskuilen et al. 2016), and becomes increasingly computationally intensive. Thus, I follow past studies on collapsing bounds which tend to be similarly sparing in their parameter allowance (e.g., Palmer et al. 2005; Ditterich 2006a, b; Bowman et al. 2012), recognizing that this limits the strength of the conclusions that can be drawn. In any event, the parameter-limited regime remains of interest for several reasons. First, simpler models are easier to handle and estimate when computational power is constrained or analysts are non-experts. Second, when data is limited as may be the case in practical applications, more complex models will suffer from overfitting. Third, even when data is relatively abundant, the relevant identification problems do not seem to be eliminated fully (Voskuilen et al. 2016). In all of these cases, the question is raised of which parameters to keep and how to use them most efficiently.

## Results

To formally compare models in a way that is congruent with previous studies (e.g., Hawkins et al. 2015), I calculate the Bayesian information criterion (BIC) according to each model for every participant:  $BIC = m \log n - 2 \log L$ , where  $L$  is the maximum log likelihood,  $m$  is the number of free parameters in the model, and  $n$  is the number of data points. Here, I pool across the two blocks, so  $m = 6$  (for the fixed and approximately optimal threshold models) or 8 (for the generic hyperbolic threshold model) and  $n = 200$ . The results split by block are similar, and shown in the Appendix. I use the BIC values to compute posterior model probabilities accounting for uncertainty in model selection (Wasserman 2000). Supposing a uniform prior over the three competing models, the posterior probability of model  $M_j$  is

$$P(M_j|D) \approx \frac{\exp\left(-\frac{1}{2} BIC_j\right)}{\sum_k \exp\left(-\frac{1}{2} BIC_k\right)} \quad (5)$$

<sup>11</sup> There were 10,000 replicates per level of value difference, of which there were approximately 40.


The results of this model comparison analysis are displayed in Fig. 1, and the median thresholds are portrayed in Fig. 2 split based on which model fits each participant the best. Overall, collapsing bounds are favored for 21 of the 23 participants, 16 of which are best fit by the approximately optimal threshold.

For an aggregate comparison metric, I also enter the model evidence based on BIC into a random effects model selection procedure which estimates the frequency of each model in the population (Rigoux et al. 2014; Stephan et al. 2009). Each participant’s behavior is allowed to be generated by a different model out of those considered, according to some unknown probability distribution. This calculation yields a protected exceedance probability (*PXP*) for each model, which is the posterior probability that it has the highest frequency in the population, accounting for the possibility that the differences between models are due to chance. This aggregate comparison strongly favors the approximately optimal model ( $PXP = 0.9965$ ).

The descriptive implications of the uncertain-difference DDM can be illustrated from multiple angles. A basic prediction of the uncertain-difference model is a negative relationship between response time and accuracy; the simplest flat threshold DDM, by contrast, implies independence between the two (Ratcliff and McKoon 2008).<sup>12</sup> A logistic regression predicting accuracy from response time (controlling for value difference and trial number) reveals such a negative relationship ( $p = .003$ ), reported in Table 1. The magnitude of the estimated coefficient here implies that a 1-second increase in response time is associated with a roughly  $-0.038/4 \approx 1$  percentage point decrease in accuracy. In the current task, observed response times ranged up to 28 seconds with a 95th percentile of over 10 seconds; a 1 percentage point drop in accuracy per second is thus sizable over this range.

To clearly explain why collapsing boundaries provide a better fit to the data, the distinctive predictions of each model are visualized in Fig. 3. This depicts response time and accuracy conditional on the value difference for a subject who is best fit by the approximately optimal model. Each vertical slice can be thought to represent a condition defined by value difference, and the figure shows how the properties of the empirical and theoretical distributions vary with respect to this element. This figure reveals that all three models capture the mean response time and accuracy well (shown as solid lines) and are scarcely distinguishable in these regards. The signatures of collapsing boundaries rest instead in the tails of the response time distribution (shown as dashed lines). The optimal boundary collapse truncates the long tail of response times, especially when the signal strength is low—exactly as its adaptive function dictates.

<sup>12</sup> A negative association can also be driven by between-trial parameter variability (e.g., Laming 1968; Ratcliff 1978).



**Fig. 1** BIC-based approximation to posterior model probabilities in favor of the fixed, approximately optimal, and generic hyperbolic threshold models. Color represents the three models and columns represent individual subjects. Parameters for each subject are fit separately across blocks

Figure 4 accentuates the contrast, showing a heatmap of the difference in predicted response time frequencies between the approximately optimal model and the fixed threshold model. The abbreviated response time tail is again visible in the heatmap. In addition, the large initial magnitudes of some collapsing thresholds can produce a reluctance to conclude deliberation extremely quickly, lengthening the shortest response times. This leads to a banded heatmap pattern as with the example participant. The corresponding plots for all subjects are provided in the Appendix, and generally recapitulate these observations. The exact response times which are more indicative of the approximately optimal model depend on the specific configuration of parameters, but a striped heatmap pattern consistently appears. Collapsing boundaries are thus typically distinguishable by the truncated tails they predict in the response time distribution, visible in the heatmap and in the data.

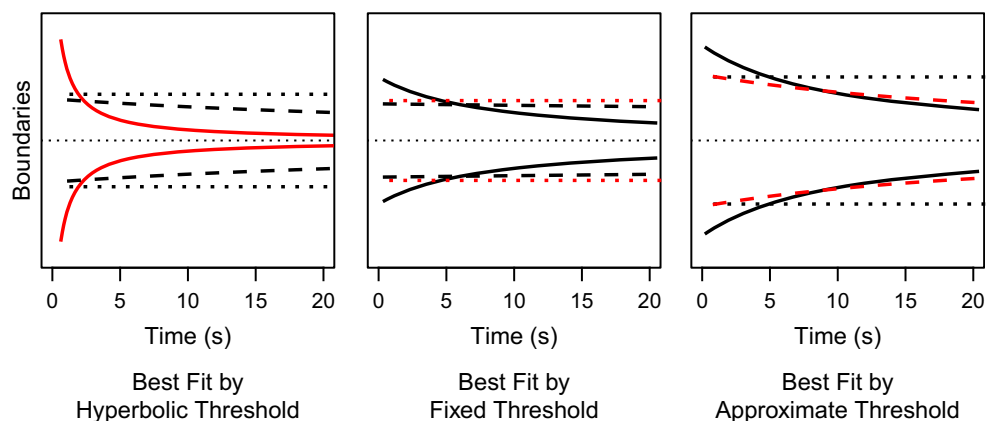
**Discussion**

I study whether people optimally adjust their decision making criteria over the course of problems in which the speed of information processing is tied to the value of alternatives. Recent theoretical results imply that the optimal confidence threshold that triggers choice declines at a particular rate inversely proportional to time already spent. To adhere as closely as possible to the central theoretical assumptions, I use a hybrid value-linked perceptual paradigm—a bi-directional

random dot motion task in which the payoff from a choice is proportional to the number of dots moving in the corresponding direction. The optimal rule fits behavior better than a fixed threshold model, indicating that principled theories of collapsing bounds may be empirically useful. While all models were able to adequately describe mean response time and accuracy, only collapsing bounds captured the truncated tails of the response time distributions, in accordance with their primary purpose. These results are broadly consistent with the hypothesis that people calibrate their standards of confidence during the course of a decision problem to optimally balance the costs and benefits of spending time. Fudenberg et al. (2018) also show that their theory helps predict behavior in the value-based experiment of Krajbich et al. (2010) with choice among snack foods.

Other studies have come to different conclusions. The closest experiment is study 2 of Oud et al. (2016), in which participants had to choose the larger of two sets of twinkling dots, and the payoff earned was based on the number of dots in the chosen set (according to two possible difficulty levels). Their clearest evidence of suboptimality was based on a time limit intervention which imposed a personally tailored deadline on each subject. Even though these deadlines provided no new information to subjects, Oud et al. (2016) found that they improved earnings. However, in absolute terms, the intervention only increased total subject earnings by an average of 15 cents (= 76.28 points × \$1/1000 points × 2 blocks). This is a miniscule fraction of the \$18.29 average subject payment. Thus, the degree of suboptimality revealed by the intervention, while statistically

**Fig. 2** Decision thresholds from median parameter values in groups of participants best fit by different models (posterior model probability at least 60%). Solid lines reflect hyperbolic thresholds, dashed lines reflect approximately optimal thresholds, and dotted lines reflect fixed thresholds which start at the non-decision time. Red lines indicate best fitting model



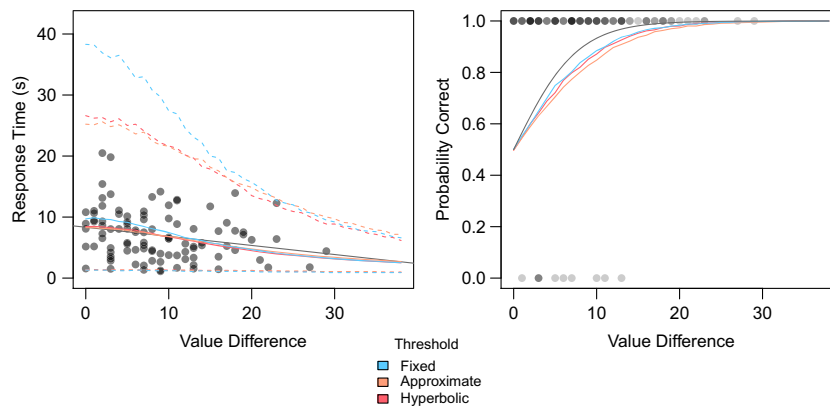
**Table 1** Accuracy logistic regression results

	Dependent variable: Accuracy
Response time	-0.038** (0.013)
Value difference	0.073*** (0.006)
Trial ( $\times 100$ )	0.014 (0.114)
Individual $\times$ block-specific intercepts	Yes


Standard errors in parentheses. \* $p < 0.05$ ; \*\* $p < 0.01$ ; \*\*\* $p < 0.001$

significant, is substantively tiny. This leaves plenty of room for theories based on optimality to improve predictive power, even if they are not perfect.

More broadly, although many have satisfactorily used fixed threshold models to assess perceptual decision making (e.g., Bode et al. 2012; Brown et al. 2008; Ding and Gold 2010, 2012; Forstmann et al. 2010; Forstmann et al. 2008; O’Connell et al. 2012; Ramakrishnan and Murthy 2013; Ramakrishnan et al. 2012; Ratcliff et al. 2009; Salinas and Stanford 2013; Schall 2003; Schurger et al. 2012; Smith and McKenzie 2011; Usher and McClelland 2001; Wang 2002; Wong and Wang 2006), many others have argued that collapsing thresholds (or other modifications with similar effect) fit the data more closely (e.g., Sanders and Linden 1967; Viviani 1979b, a; Viviani and Terzuolo 1972; Ditterich 2006a, b; Churchland et al. 2008; Cisek et al. 2009; Rao 2010; Bowman et al. 2012; Hanks et al. 2011; Thura et al. 2012; Thura and Cisek 2014; Zhang et al. 2014). The most comprehensive meta-analysis carried out to date found evidence primarily in favor of the fixed threshold DDM as compared to a generic collapsing threshold DDM or the urgency gating model (Hawkins et al. 2015), and following research has



**Fig. 3** Trial-level response time and accuracy data along with model predictions for a subject who is best fit by the approximately optimal model. Color represents the three models. (left) Response times. For all models, predicted mean response times conditional on value difference are shown as colored solid lines, and the 1st and 99th percentiles are



**Fig. 4** Heatmap depicting the difference in predicted response time distributions (conditional on value difference) between the approximately optimal and fixed threshold models for the same subject as in Fig. 3. Heatmap color reflects sign of difference (orange when approximately optimal model has higher predicted probability, blue in the opposite case), and heatmap transparency reflects magnitude of difference (darker means larger difference). Predicted mean response times conditional on value difference are shown as colored solid lines, and the 1st and 99th response time percentiles are shown as dashed lines

concluded (Voskuilen et al. 2016). However, there was substantial heterogeneity in the best-fitting model across the included studies, and the generic collapsing threshold did yield an improvement in a number of cases. This raises questions as to the possible sources of heterogeneity.

While the present experiment does not make direct contact with the above due to its unique setup, it brings an underexplored factor into the foreground. Time-varying thresholds are often defended on the basis of optimality, driven by elements such as between-trial variation in difficulty (Drugowitsch et al. 2012; Moran 2015). When incentives are tied to difficulty level, the rationale for collapse is amplified. This is clearly illustrated by the results of Bather (1962) in a similar theoretical setup to Fudenberg et al. (2018) but with a

shown as dashed lines. The ordinary least squares regression line is shown in gray. (right) Accuracy. For all models, the predicted accuracies conditional on value difference are shown as colored solid lines. The logistic regression curve with intercept fixed at 0.5 is shown in gray

fixed value difference. In this case, the optimal stopping threshold declines asymptotically at rate  $1/\sqrt{t}$ , which is compounded by the inclusion of incentives (tied to signal strength) to yield the  $1/t$  rate of decline studied in this paper. The extent to which incentives aid collapse empirically is an important topic for future research. By combining the control enabled by perceptual tasks with the ability to construct flexible reward schemes, the value-linked approach taken in this paper and elsewhere (e.g., Oud et al. 2016; Pirrone et al. 2018b) provides a useful platform for testing the role of reward in the deliberative process.

The sharpness of the uncertain-difference DDM's predictions appears to be a meaningful advantage that enhances model fit at minimal expense in terms of degrees of freedom. The development of powerful new theoretical models based on optimality demands a great deal of further empirical study testing their implications. Recent work has demonstrated that the skeletal structure of sequential sampling models alone is insufficient for making crisp and distinctive predictions. Zhang et al. (2014) and Khodadadi and Townsend (2015) have shown how any basic diffusion model with symmetric time-varying boundaries can be perfectly mimicked by a counterpart with independent accumulators. Such equivalencies strongly underscore the need for further theoretically motivated constraints to discipline these models, such as those implied by the uncertain-difference DDM. Sequential sampling models were originally inspired by optimal statistical algorithms for hypothesis testing. Particularly when more sophisticated economic incentive schemes are involved, this principle may yet offer us more insights.

**Acknowledgements** Thanks to Colin Camerer, Jaron Colas, Taisuke Imai, Ian Krajbich, and Tomasz Strzalecki for helpful comments and discussions. An earlier version of this paper was circulated under the title "Evidence on Optimally Collapsing Thresholds in Value-Linked Decision Making".

**Funding Information** Funding from the Social Sciences and Humanities Research Council of Canada and the Harvard Mind Brain Behavior Interfaculty Initiative is gratefully acknowledged.

## Compliance with Ethical Standards

The experiment was approved by the Caltech Committee for the Protection of Human Subjects.

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

## References

Bather, J. A. (1962). Bayes procedures for deciding the sign of a normal mean. *Mathematical Proceedings of the Cambridge Philosophical Society*, 58(4), 599–620.

- Bhui, R. (2019). A statistical test for the optimality of deliberative time allocation. *Psychonomic Bulletin & Review*. <https://doi.org/10.3758/s13423-018-1555-1>.
- Bitzer, S., Park, H., Blankenburg, F., & Kiebel, S. J. (2014). Perceptual decision making: drift-diffusion model is equivalent to a Bayesian model. *Frontiers in Human Neuroscience*, 8, 102.
- Bode, S., Sewell, D. K., Lilburn, S., Forte, J. D., Smith, P. L., & Stahl, J. (2012). Predicting perceptual decision biases from early brain activity. *Journal of Neuroscience*, 32(36), 12488–12498.
- 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(4), 700–765.
- Bollimunta, A., & Ditterich, J. (2012). Local computation of decision-relevant net sensory evidence in parietal cortex. *Cerebral Cortex*, 22(4), 903–917.
- Bollimunta, A., Totten, D., & Ditterich, J. (2012). Neural dynamics of choice: single trial analysis of decision-related activity in parietal cortex. *Journal of Neuroscience*, 32(37), 12684–12701.
- Bowman, N. E., Kording, K. P., & Gottfried, J. A. (2012). Temporal integration of olfactory perceptual evidence in human orbitofrontal cortex. *Neuron*, 75(5), 916–927.
- Britten, K. H., Shadlen, M. N., Newsome, W. T., & Movshon, J. A. (1992). The analysis of visual motion: a comparison of neuronal and psychophysical performance. *Journal of Neuroscience*, 12(12), 4745–4765.
- Brown, S., & Heathcote, A. (2003). QMLE: fast, robust, and efficient estimation of distribution functions based on quantiles. *Behavior Research Methods, Instruments, & Computers*, 35(4), 485–492.
- Brown, J. W., Hanes, D. P., Schall, J. D., & Stuphorn, V. (2008). Relation of frontal eye field activity to saccade initiation during a countermanding task. *Experimental Brain Research*, 190(2), 135–151.
- Churchland, A. K., Kiani, R., & Shadlen, M. N. (2008). Decision-making with multiple alternatives. *Nature Neuroscience*, 11(6), 693–702.
- Cisek, P., Puskas, G. A., & El-Murr, S. (2009). Decisions in changing conditions: the urgency-gating model. *Journal of Neuroscience*, 29(37), 11560–11571.
- Ding, L., & Gold, J. I. (2010). Caudate encodes multiple computations for perceptual decisions. *Journal of Neuroscience*, 30(47), 15747–15759.
- Ding, L., & Gold, J. I. (2012). Separate, causal roles of the caudate in saccadic choice and execution in a perceptual decision task. *Neuron*, 75(5), 865–874.
- Ditterich, J. (2006a). Evidence for time-variant decision making. *European Journal of Neuroscience*, 24(12), 3628–3641.
- Ditterich, J. (2006b). Stochastic models of decisions about motion direction: behavior and physiology. *Neural Networks*, 19(8), 981–1012.
- Drugowitsch, J., Moreno-Bote, R., Churchland, A. K., Shadlen, M. N., & Pouget, A. (2012). The cost of accumulating evidence in perceptual decision making. *Journal of Neuroscience*, 32(11), 3612–3628.
- Fard, P. R., Park, H., Warkentin, A., Kiebel, S. J., & Bitzer, S. (2017). A Bayesian reformulation of the extended drift-diffusion model in perceptual decision making. *Frontiers in Computational Neuroscience*, 11, 29.
- Forstmann, B. U., Dutilh, G., Brown, S., Neumann, J., Von Cramon, D. Y., Ridderinkhof, K. R., & Wagenmakers, E.-J. (2008). Striatum and pre-SMA facilitate decision-making under time pressure. *Proceedings of the National Academy of Sciences*, 105(45), 17538–17542.
- Forstmann, B. U., Anwander, A., Schäfer, A., Neumann, J., Brown, S., Wagenmakers, E. J., Bogacz, R., & Turner, R. (2010). Corticostriatal connections predict control over speed and accuracy in perceptual decision making. *Proceedings of the National Academy of Sciences*, 107(36), 15916–15920.



- Forstmann, B. U., Ratcliff, R., & Wagenmakers, E.-J. (2016). Sequential sampling models in cognitive neuroscience: advantages, applications, and extensions. *Annual Review of Psychology*, *67*, 641–666.
- Fudenberg, D., Strack, P., & Strzalecki, T. (2018). Speed, accuracy, and the optimal timing of choices. *American Economic Review*, *108*(12), 3651–3684.
- Hanks, T. D., Mazurek, M. E., Kiani, R., Hopp, E., & Shadlen, M. N. (2011). Elapsed decision time affects the weighting of prior probability in a perceptual decision task. *Journal of Neuroscience*, *31*(17), 6339–6352.
- Hawkins, G. E., Marley, A., Heathcote, A., Flynn, T. N., Louviere, J. J., & Brown, S. D. (2014). Integrating cognitive process and descriptive models of attitudes and preferences. *Cognitive Science*, *38*(4), 701–735.
- Hawkins, G. E., Forstmann, B. U., Wagenmakers, E.-J., Ratcliff, R., & Brown, S. D. (2015). Revisiting the evidence for collapsing boundaries and urgency signals in perceptual decision-making. *Journal of Neuroscience*, *35*(6), 2476–2484.
- Heathcote, A., Brown, S., & Mewhort, D. J. (2002). Quantile maximum likelihood estimation of response time distributions. *Psychonomic Bulletin & Review*, *9*(2), 394–401.
- Heathcote, A., Brown, S., & Cousineau, D. (2004). QMPE: estimating Lognormal, Wald, and Weibull RT distributions with a parameter-dependent lower bound. *Behavior Research Methods, Instruments, & Computers*, *36*(2), 277–290.
- Khodadadi, A., & Townsend, J. T. (2015). On mimicry among sequential sampling models. *Journal of Mathematical Psychology*, *68*, 37–48.
- Krajbich, I., & Rangel, A. (2011). Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. *Proceedings of the National Academy of Sciences*, *108*(33), 13852–13857.
- Krajbich, I., Armel, C., & Rangel, A. (2010). Visual fixations and the computation and comparison of value in simple choice. *Nature Neuroscience*, *13*(10), 1292–1298.
- Krajbich, I., Lu, D., Camerer, C., & Rangel, A. (2012). The attentional drift-diffusion model extends to simple purchasing decisions. *Frontiers in Psychology*, *3*, 193.
- Krajbich, I., Oud, B., & Fehr, E. (2014). Benefits of neuroeconomic modeling: new policy interventions and predictors of preference. *American Economic Review*, *104*(5), 501–506.
- Krajbich, I., Hare, T., Bartling, B., Morishima, Y., & Fehr, E. (2015). A common mechanism underlying food choice and social decisions. *PLoS Computational Biology*, *11*(10), e1004371.
- Lam, E. and Kalaska, J. F. (2014). *Choosing sides: the psychophysics of target choices using random dot kinematograms with mutually contradictory evidence*. Unpublished manuscript.
- Laming, D. R. J. (1968). *Information theory of choice-reaction times*. Cambridge: Academic Press.
- Milosavljevic, M., Malmaud, J., Huth, A., Koch, C., & Rangel, A. (2010). The drift diffusion model can account for the accuracy and reaction time of value-based choices under high and low time pressure. *Judgment and Decision making*, *5*(6), 437–449.
- Moran, R. (2015). Optimal decision making in heterogeneous and biased environments. *Psychonomic Bulletin & Review*, *22*(1), 38–53.
- Mullen, K. M., Ardia, D., Gil, D. L., Windover, D., & Cline, J. (2011). DEoptim: an R package for global optimization by differential evolution. *Journal of Statistical Software*, *40*(6), 1–26.
- Newsome, W. T., Britten, K. H., & Movshon, J. A. (1989). Neuronal correlates of a perceptual decision. *Nature*, *341*(6237), 52–54.
- Niwa, M., & Ditterich, J. (2008). Perceptual decisions between multiple directions of visual motion. *Journal of Neuroscience*, *28*(17), 4435–4445.
- O’Connell, R. G., Dockree, P. M., & Kelly, S. P. (2012). A supramodal accumulation-to-bound signal that determines perceptual decisions in humans. *Nature Neuroscience*, *15*(12), 1729–1735.
- Otter, T., Allenby, G. M., & van Zandt, T. (2008). An integrated model of discrete choice and response time. *Journal of Marketing Research*, *45*(5), 593–607.
- Oud, B., Krajbich, I., Miller, K., Cheong, J., Botvinick, M., & Fehr, E. (2016). Irrational time allocation in decision-making. *Proceedings of the Royal Society B: Biological Sciences*, *283*(1822), 20151439.
- Palmer, J., Huk, A. C., & Shadlen, M. N. (2005). The effect of stimulus strength on the speed and accuracy of a perceptual decision. *Journal of Vision*, *5*(5), 376–404.
- Pirrone, A., Stafford, T., & Marshall, J. A. (2014). When natural selection should optimize speed-accuracy trade-offs. *Frontiers in Neuroscience*, *8*, 73.
- Pirrone, A., Azab, H., Hayden, B. Y., Stafford, T., & Marshall, J. A. (2018a). Evidence for the speed–value trade-off: human and monkey decision making is magnitude sensitive. *Decision*, *5*(2), 129–142.
- Pirrone, A., Wen, W., & Li, S. (2018b). Single-trial dynamics explain magnitude sensitive decision making. *BMC Neuroscience*, *19*(54), 54.
- Ramakrishnan, A., & Murthy, A. (2013). Brain mechanisms controlling decision making and motor planning. *Progress in Brain Research*, *202*, 321–345.
- Ramakrishnan, A., Sureshbabu, R., & Murthy, A. (2012). Understanding how the brain changes its mind: microstimulation in the macaque frontal eye field reveals how saccade plans are changed. *Journal of Neuroscience*, *32*(13), 4457–4472.
- Rao, R. P. (2010). Decision making under uncertainty: a neural model based on partially observable markov decision processes. *Frontiers in Computational Neuroscience*, *4*, 146.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*(2), 59–108.
- Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. *Neural Computation*, *20*(4), 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*(16), 6539–6544.
- Ratcliff, R., Smith, P. L., Brown, S. D., & McKoon, G. (2016). Diffusion decision model: current issues and history. *Trends in Cognitive Sciences*, *20*(4), 260–281.
- Rigoux, L., Stephan, K. E., Friston, K. J., & Daunizeau, J. (2014). Bayesian model selection for group studies—revisited. *NeuroImage*, *84*, 971–985.
- Salinas, E., & Stanford, T. R. (2013). The countermanding task revisited: fast stimulus detection is a key determinant of psychophysical performance. *Journal of Neuroscience*, *33*(13), 5668–5685.
- Sanders, A. and Ter Linden, W. (1967). Decision making during paced arrival of probabilistic information. *Acta Psychologica*, *27*, 170–177.
- Schall, J. D. (2003). Neural correlates of decision processes: neural and mental chronometry. *Current Opinion in Neurobiology*, *13*(2), 182–186.
- Schurger, A., Sitt, J. D., & Dehaene, S. (2012). An accumulator model for spontaneous neural activity prior to self-initiated movement. *Proceedings of the National Academy of Sciences*, *109*(42), E2904–E2913.
- Smith, P. L., & McKenzie, C. R. (2011). Diffusive information accumulation by minimal recurrent neural models of decision making. *Neural Computation*, *23*(8), 2000–2031.
- Stephan, K. E., Penny, W. D., Daunizeau, J., Moran, R. J., & Friston, K. J. (2009). Bayesian model selection for group studies. *NeuroImage*, *46*(4), 1004–1017.
- Stone, M. (1960). Models for choice-reaction time. *Psychometrika*, *25*(3), 251–260.

- Tajima, S., Drugowitsch, J., & Pouget, A. (2016). Optimal policy for value-based decisionmaking. *Nature Communications*, 7(12400), 1–12.
- Teodorescu, A. R., Moran, R., & Usher, M. (2016). Absolutely relative or relatively absolute: violations of value invariance in human decision making. *Psychonomic Bulletin & Review*, 23(1), 22–38.
- Thura, D., & Cisek, P. (2014). Deliberation and commitment in the premotor and primary motor cortex during dynamic decision making. *Neuron*, 81(6), 1401–1416.
- Thura, D., Beauregard-Racine, J., Fradet, C.-W., & Cisek, P. (2012). Decision making by urgency gating: theory and experimental support. *Journal of Neurophysiology*, 108(11), 2912–2930.
- Trueblood, J. S., Brown, S. D., & Heathcote, A. (2014). The multiattribute linear ballistic accumulator model of context effects in multialternative choice. *Psychological Review*, 121(2), 179–205.
- 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(4), 443–456.
- Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: the leaky, competing accumulator model. *Psychological Review*, 108(3), 550–592.
- Viviani, P. (1979a). Choice reaction times for temporal numerosity. *Journal of Experimental Psychology: Human Perception and Performance*, 5(1), 157–167.
- Viviani, P. (1979b). A diffusion model for discrimination of temporal numerosity. *Journal of Mathematical Psychology*, 19(2), 108–136.
- Viviani, P., & Terzuolo, C. (1972). On the modeling of the performances of the human brain in a two-choice task involving decoding and memorization of simple visual patterns. *Kybernetik*, 10(3), 121–137.
- Voskuilen, C., Ratcliff, R., & Smith, P. L. (2016). Comparing fixed and collapsing boundary versions of the diffusion model. *Journal of Mathematical Psychology*, 73, 59–79.
- Wald, A. (1947). *Sequential analysis*. New York: Wiley.
- Wang, X.-J. (2002). Probabilistic decision making by slow reverberation in cortical circuits. *Neuron*, 36(5), 955–968.
- Wasserman, L. (2000). Bayesian model selection and model averaging. *Journal of Mathematical Psychology*, 44(1), 92–107.
- Wong, K.-F., & Wang, X.-J. (2006). A recurrent network mechanism of time integration in perceptual decisions. *Journal of Neuroscience*, 26(4), 1314–1328.
- Zhang, S., Lee, M. D., Vandekerckhove, J., Maris, G., & Wagenmakers, E.-J. (2014). Time-varying boundaries for diffusion models of decision making and response time. *Frontiers in Psychology*, 5, 1364.