Estimation of time-varying decision thresholds from the choice and reaction times with no assumption on the shape

Yul Hyoung Ryul Kang

Department of Neurobiology, Columbia University

Abstract

When a decision is made based on a series of samples of evidence, the threshold for decision may change over time. Here we propose a fast heuristic algorithm that estimates the time-varying thresholds without making an assumption on their shape. The algorithm gives an approximate best estimate of the time-varying thresholds for all time considered when all other parameters are fixed. The algorithm outperforms conventional methods that gradually adjust the threshold estimates during fitting, when fitting time is limited.

Introduction

When a subject makes a binary choice based on a series of samples of scalar evidence, its speed and accuracy are typically modeled well by a diffusion process with two absorbing boundaries (Gold and Shadlen 2007; Bogacz et al. 2006). In these models, a decision is made when the cumulative sum of the evidence exceeds one of the two boundaries (or thresholds). It has been shown that when the quality of evidence is fixed (i.e., when there is only one signal to noise ratio), the number of accurate decisions per unit time can be maximized with thresholds that are constant over time (Wald 1973; Wald and Wolfowitz 1948). But the optimal form of the thresholds is different when signal to noise ratios vary between trials. There, trials with poor quality of evidence tend to take a long time to reach the threshold, and conversely, when a long time has passed and the threshold is not reached, it tends to be a trial with poor quality of evidence. Therefore, it is optimal to decrease the threshold over time, so as to move quickly onto the next decision, which is likely to be easier (Drugowitsch et al. 2012; Shadlen et al. 2006).

Although it has been reported that the normative solution for the optimal threshold over time can be found using the Bellman Equation (Drugowitsch et al. 2012), previous methods to estimate the shape of the thresholds had limitations. Many previous studies imposed a particular functional form (quadratic, Weibull, etc., see, e.g., Hawkins et al. 2015), and fit its parameters, but they may have missed the true shape of the thresholds if it cannot be expressed by the

functional form chosen. At least one study used cosine basis functions (Drugowitsch et al. 2012), which can fit thresholds of any form given a sufficient number of basis functions. However, fitting the model is inefficient because it involves fitting many parameters using a gradient descent procedure, and it requires arbitrary hyperparameters, including the width of the cosine basis function.

Here we propose an efficient heuristic algorithm that does not require any parameter dedicated to thresholds. Although we defer rigorous justification of the algorithm to future works, we show that the algorithm can fit simulated data with no assumption on the shape of the thresholds, and do so significantly faster than previous methods with parametric thresholds.

Methods

In the following text, we sometimes refer to the previous and the new approaches as parametric and nonparametric methods, following the way they specify the thresholds. Time and the level of accumulated evidence are discretized in steps of Δt and Δy. We ignore the discretization error as long as the error becomes arbitrarily small with smaller steps. We denote a random variable with an uppercase italic letter, and a scalar variable or constant with a lowercase italic letter. We denote a vector/matrix with a bold lowercase/uppercase letter, respectively (Table 1).

Symbol	Definition
$b_{z,t}$	Threshold for choice z at time t
$\mathbf{B}, \mathbf{B}_{[0,t]}$	The whole thresholds across all times and choices ($\mathbf{B}$), and the thresholds up to t ($\mathbf{B}_{[0,t]}$)
$\mathbf{D}$	The data; the set of the pair of choice and reaction times, $\{(t_{R,i}, z_i)\}$
i	Trial index
n	Number of trials
T_D	Decision time; the part of the T_R that depends on the difficulty of the decision
T_N	Nondecision time; the part of the T_R that does not depend on the difficulty of the decision
$T_\mathrm{R}, t_{\mathrm{R},i}$	Reaction time and its observed value on trial i; the time between the stimulus onset and the response; $T_\mathrm{R} = T_\mathrm{D} + T_\mathrm{N}$
t_max	Maximum reaction time observed
Y_t	Evidence accumulated up to time t
y_max	Maximum level of accumulated evidence considered
Z, z_i	Choice $\in \{\mathrm{Left}, \mathrm{Right}\}$ and its observed value on trial i
$\Delta t, \Delta y$	Sizes of the steps of the discretized time and accumulated evidence axes
$\boldsymbol{\theta}$	Parameters of the diffusion model; $\boldsymbol{\theta} \coloneqq \left({\boldsymbol{\theta}_\mathrm{D}}^\top, {\boldsymbol{\theta}_\mathrm{N}}^\top, {\boldsymbol{\theta}_\mathrm{B}}^\top\right)^\top$
$\boldsymbol{\theta}_\mathrm{N}, \boldsymbol{\theta}_\mathrm{B}, \boldsymbol{\theta}_\mathrm{D}$	Parameters for the nondecision time, thresholds, and decision time ($\boldsymbol{\theta}_\mathrm{D}$ has parameters for the decision time except those for the thresholds)
$\boldsymbol{\theta}_{\mathrm{B}-}$	Parameters except those for the thresholds; $\boldsymbol{\theta}_{\mathrm{B}-} \coloneqq \left({\boldsymbol{\theta}_\mathrm{D}}^\top, {\boldsymbol{\theta}_\mathrm{N}}^\top\right)^\top$

Table 1. Important notations

Algorithm 1. Fitting a diffusion model with parametric or nonparametric thresholds

```
// Equations are simplified to aid intuition. See text for full expressions.
Load the data D ← {(t_{R,i}, z_i)}
Initialize parameters for the diffusion and nondecision times, θ_D and θ_N // together called θ_{B−}
if method=parametric:
    Initialize parameters for the thresholds, θ_B
end
do
    Initialize P(Y_{t=0} = y, T_D > t) and P(T_D = 0, Z = z) // Eqs.10–11
    for t in Δt to t_max do
        Compute the distribution of accumulated evidence assuming no thresholds at t,
            P(Y'_t = y), using P(Y_{t−Δt} = y, T_D > t − Δt) and θ_D // Eq. 12
        if method=parametric:
            Compute the threshold b_t using θ_B // Eq. 9
        else: // method=nonparametric
            Initialize the threshold b_t
            Compute the probability of outcomes p_{t|θ_N} that does not depend on b_t,
                using the data D and θ_N // Eqs. 19–21
            do
                Compute the probability of outcomes p_{t|b_t,θ_{B−}} that does depend on b_t,
                    Using b_t, P(Y'_t = y), and θ_D // Eq. 22–24
                Compute P(b_t | θ_{B−}) ≈ P(p_{t|b_t,θ_{B−}} | n p_{t|θ_N}) // Eq. 25
                Adjust b_t to increase P(b_t | θ_{B−}) // Eqs. 25–26
            until P(b_t | θ_{B−}) converges
        end
        Compute the probability of decision P(T_D = t, Z = z)
            using P(Y'_t = y) and b_t // Eqs. 14–15
        Compute distribution of the accumulated evidence (with thresholds at t),
            P(Y_t = y, T_D > t),
            using P(Y'_t = y) and b_t // Eq. 16
    end
    Compute the nondecision time distribution P(T_N = t | Z = z) using θ_N // Eq. 5
    Compute the reaction time distribution P(T_R = t, Z = z)
        using P(T_D = t, Z = z) and P(T_N = t | Z = z) // Eq. 3
    Compute the cost, −Σ_i log P(T_R = t_{R,i}, Z = z_i | θ), using P(T_R = t, Z = z) and D // Eq. 2
    if method=parametric:
        Adjust θ_D, θ_N and θ_B to decrease the cost
    else: // method=nonparametric
        Adjust θ_D and θ_N to decrease the cost
    end
until the cost converges
return θ_D, θ_N, and B
```

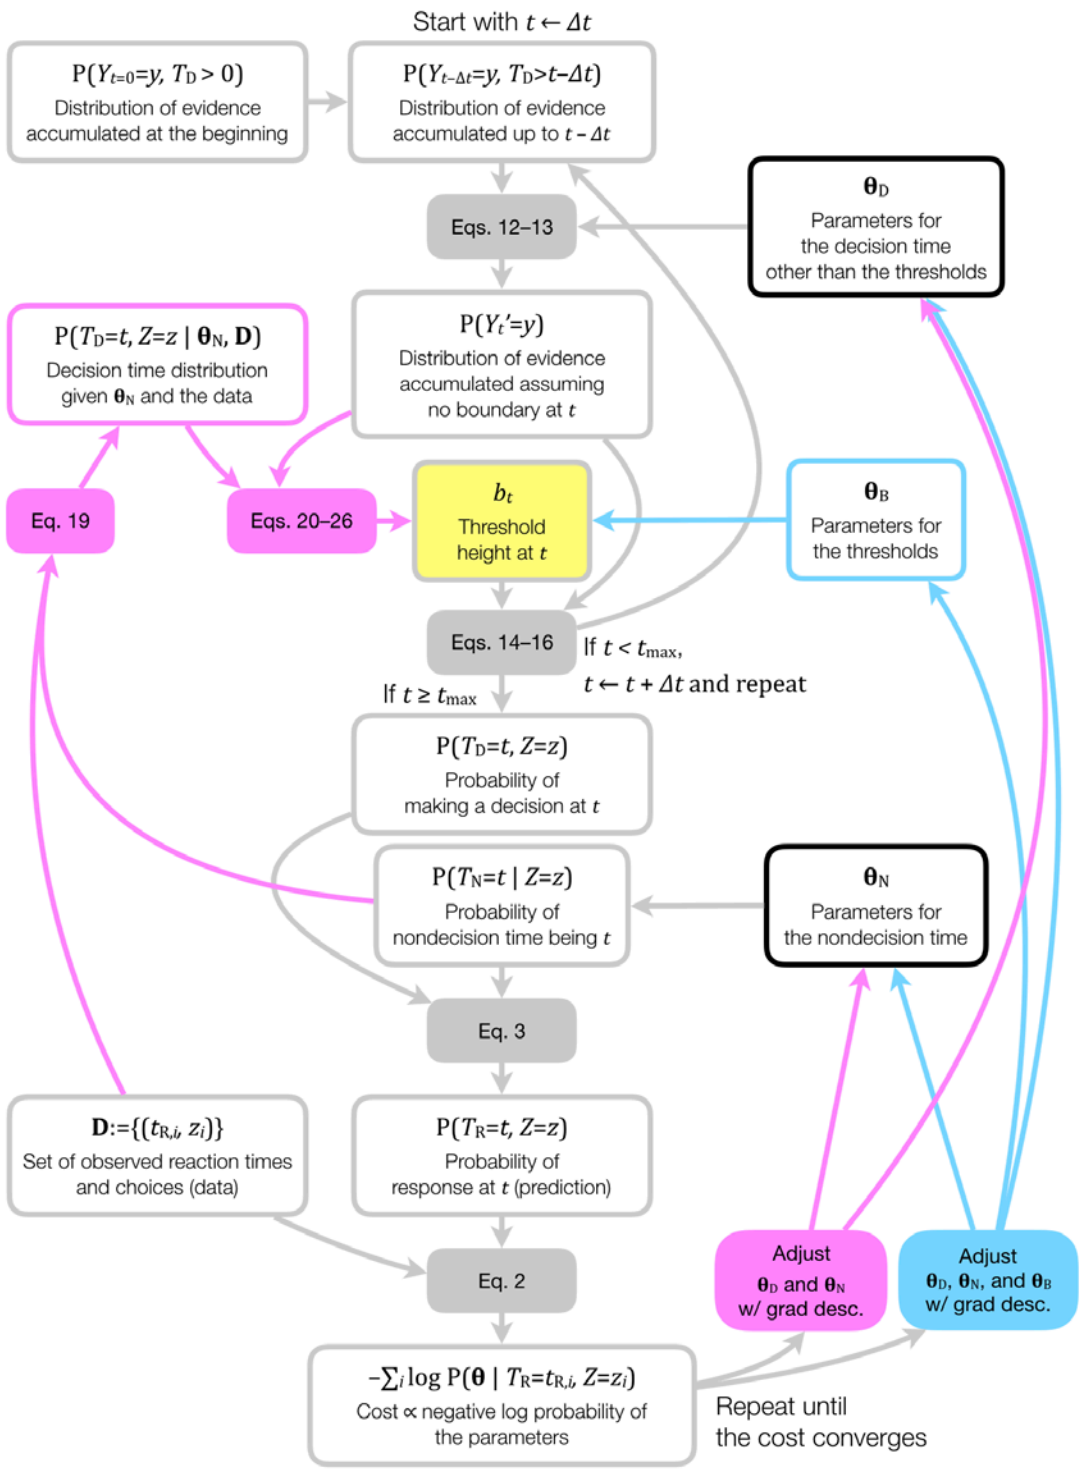

Figure 1. Flow of information

Expressions are simplified to aid intuition (see text for full expressions). Arrows denote dependence. Pink and blue colors denote quantities and dependence specific to the nonparametric and the parametric methods, respectively. Gray symbols are common to both. b_t is highlighted because it is the key quantity that the two methods calculate differently. Boxes denoting common parameters are colored black. "Grad desc." indicates gradient descent.

Previous approach

We first review a previous, parametric approach to fitting the diffusion model. The key is that here the threshold heights are specified directly by parameters $\boldsymbol{\theta}_{\mathrm{B}}$, which will be introduced later in this section. In the next section, we will introduce a new, nonparametric method that proposes the threshold heights without any devoted parameter. The commonality and the difference between the two methods are summarized in Algorithm 1 and Figure 1.

For concreteness, we use an example of a direction discrimination task where on each trial the subject judges the direction of a motion stimulus to be left or rightward, although the mathematics holds for any 2-alternative forced choice experiment (Shadlen et al. 2006).

On each trial i, we observe a sample $(t_{\mathrm{R},i}, z_i)$ of the choice $Z \in \{\mathrm{Left}, \mathrm{Right}\}$ and the reaction time $T_{\mathrm{R}} > 0$, which is the time between the stimulus onset and the response. We decompose T_{R} into two parts, "decision time" (T_{D}) that depends on the ambiguity of the stimulus (hence the difficulty of the decision) and "nondecision time" (T_{N}) that consists of motor and sensory delays that are independent of the difficulty, such that

$$T_{\mathrm{R}} \coloneqq T_{\mathrm{D}} + T_{\mathrm{N}}. \qquad \text{(Eq. 1)}$$

In previous studies, a vector of parameters $\boldsymbol{\theta}$ was optimized to maximize the likelihood of the parameters given the data. Following the Bayes rule, it is the same as maximizing the log posterior probability of the data, or minimizing the negative log posterior, or "cost":

$$\begin{aligned}\widehat{\boldsymbol{\theta}} &= \mathrm{argmax}_{\boldsymbol{\theta}} \prod_i \mathrm{P}(\boldsymbol{\theta} \mid T_{\mathrm{R}} = t_{\mathrm{R},i}, Z = z_i) \\ &= \mathrm{argmax}_{\boldsymbol{\theta}} \prod_i \mathrm{P}(T_{\mathrm{R}} = t_{\mathrm{R},i}, Z = z_i \mid \boldsymbol{\theta})\mathrm{P}(\boldsymbol{\theta}) \\ &= \mathrm{argmin}_{\boldsymbol{\theta}} \left(-\sum_i \log \mathrm{P}(T_{\mathrm{R}} = t_{\mathrm{R},i}, Z = z_i \mid \boldsymbol{\theta})\mathrm{P}(\boldsymbol{\theta}) \right)\end{aligned} \qquad \text{(Eq. 2)}$$

$\mathrm{P}(\boldsymbol{\theta})$ is the prior, for which we use a relatively uninformative one that is flat within a range in this chapter (Table 2). We omit $\mathrm{P}(\boldsymbol{\theta})$ in later equations for simplicity. $\mathrm{P}(T_{\mathrm{R}} = t_{\mathrm{R},i}, Z = z_i \mid \boldsymbol{\theta})$ is the likelihood of observing the pair $(t_{\mathrm{R},i}, z_i)$ on trial i given the proposed parameters. Note that the parameters $\boldsymbol{\theta}$ include those that specify the thresholds, $\boldsymbol{\theta}_{\mathrm{B}}$: We can write $\boldsymbol{\theta} \coloneqq \left(\boldsymbol{\theta}_{\mathrm{D}}{}^{\top}, \boldsymbol{\theta}_{\mathrm{N}}{}^{\top}, \boldsymbol{\theta}_{\mathrm{B}}{}^{\top}\right)^{\top}$, where $\boldsymbol{\theta}_{\mathrm{N}}$ specifies the distribution of the nondecision time T_{N}, and $\boldsymbol{\theta}_{\mathrm{D}}$ and

$\boldsymbol{\theta}_{\mathrm{B}}$ jointly specifies the distribution of the decision time T_{D}, where $\boldsymbol{\theta}_{\mathrm{D}}$ has the parameters that does not directly specify the thresholds.

The likelihood $\mathrm{P}(T_{\mathrm{R}} = t_{\mathrm{R},i}, Z = z_i \mid \boldsymbol{\theta})$ in the last line of Eq. 2 is computed by convolving the predicted joint distribution of the choice and decision time $\mathrm{P}(T_{\mathrm{D}} = t, Z = z \mid \boldsymbol{\theta}_{\mathrm{D}}, \boldsymbol{\theta}_{\mathrm{B}})$ with the proposed nondecision time distribution $\mathrm{P}(T_{\mathrm{N}} = t \mid Z = z, \boldsymbol{\theta}_{\mathrm{N}})$, i.e.:

$$\mathrm{P}(T_{\mathrm{R}} = t, Z = z \mid \boldsymbol{\theta}) = \sum_{\tau=0}^{t} \mathrm{P}(T_{\mathrm{D}} = t - \tau, Z = z \mid \boldsymbol{\theta}_{\mathrm{D}}, \boldsymbol{\theta}_{\mathrm{B}}) \mathrm{P}(T_{\mathrm{N}} = \tau \mid Z = z, \boldsymbol{\theta}_{\mathrm{N}}) \tag{Eq. 3}$$

Because the dependence of the decision time on $\boldsymbol{\theta}_{\mathrm{B}}$ is only through the thresholds $\mathbf{B}$ that it specifies, we can write

$$\mathrm{P}(T_{\mathrm{D}} = t, Z = z \mid \boldsymbol{\theta}_{\mathrm{D}}, \boldsymbol{\theta}_{\mathrm{B}}) = \mathrm{P}(T_{\mathrm{D}} = t, Z = z \mid \boldsymbol{\theta}_{\mathrm{D}}, \mathbf{B}). \tag{Eq. 4}$$

In this chapter, we use appropriately discretized gamma distributions for the nondecision time, parameterized in a way that $\boldsymbol{\theta}_{\mathrm{N}}$ consists of two pairs of numbers: the mean $\mu_{\mathrm{N},z}$ and the coefficient of variation $v_{\mathrm{N},z}$ (the ratio of the standard deviation and the mean), for each of the two choices z:

$$\mathrm{P}(T_{\mathrm{N}} = t \mid Z = z, \boldsymbol{\theta}_{\mathrm{N}}) \coloneqq f_{\mathrm{N}}(t, z; \boldsymbol{\theta}_{\mathrm{N}}) \propto \mathrm{Gamma}\big(t;\ \mu_{\mathrm{N},z}, v_{\mathrm{N},z}\big) \tag{Eq. 5}$$

The likelihood of T_{D}, $\mathrm{P}(T_{\mathrm{D}} = t - \tau, Z = z \mid \boldsymbol{\theta}_{\mathrm{D}}, \mathbf{B})$, is computed by modeling T_{D} as the first crossing time of a one-dimensional diffusion process:

$$dY_t = \mu dt + \sigma dW_t \tag{Eq. 6}$$

$$Y_{t=0} = 0 \tag{Eq. 7}$$

with two absorbing boundaries $b_{z=\mathrm{Left},t} \leq b_{z=\mathrm{Right},t}$ for the two choices, which in general vary over time. μ is the drift rate, σ is the diffusion coefficient, Y_t is evidence accumulated by time t, and W_t is the Wiener process. By convention, positive Y (and μ) supports the rightward choice, and negative leftward. Small absolute μ corresponds to a difficult decision. The decision time T_{D} is the first crossing time

$$T_{\mathrm{D}} \coloneqq \inf\{t\colon Y_t \leq b_{z=\mathrm{Left},t} \text{ or } Y_t \geq b_{z=\mathrm{Right},t}\} \tag{Eq. 8}$$

where $\inf\{\cdot\}$ is the infimum. In previous approaches, the parameters $\boldsymbol{\theta}_{\mathrm{B}}$ specified the boundaries directly, i.e.,

$$b_{z,t} \coloneqq f_{\mathrm{B}}(z, t;\ \boldsymbol{\theta}_{\mathrm{B}}) \qquad \text{(Eq. 9)}$$

where $f_{\mathrm{B}}(\cdot)$ was, e.g., a beta cumulative distribution function (Kang et al. 2017) or a weighted sum of cosine basis functions (Drugowitsch et al. 2012). We will collectively denote the proposed boundaries across all times and choices $\mathbf{B}$. (The key and the only difference of the new method is that $\mathbf{B}$ is not proposed using dedicated parameters $\boldsymbol{\theta}_{\mathrm{B}}$: it is instead inferred from the data using the rest of the parameters $\boldsymbol{\theta}_{\mathrm{N}}$ and $\boldsymbol{\theta}_{\mathrm{D}}$, as we will see in the next section.)

In the rest of this section, we will make concrete how $b_{z,t}$ gives rise to the likelihood of T_{D}. In this chapter, we use a rough yet simple approximation (for more accurate methods, see e.g., Smith 2000). We start with the initial condition, where at $t = 0$, no evidence is accumulated:

$$\mathrm{P}(Y_{t=0} = y, T_{\mathrm{D}} > t \mid \boldsymbol{\theta}_{\mathrm{D}}, \mathbf{B}) = 1 \text{ iff } y = 0 \qquad \text{(Eq. 10)}$$

and no decision is made yet:

$$\mathrm{P}(T_{\mathrm{D}} = t, Z = z \mid \boldsymbol{\theta}_{\mathrm{D}}, \mathbf{B}) = 0 \qquad \text{(Eq. 11)}$$

Then, for the next time step $t \geq \Delta t$, we first compute the distribution of the accumulated evidence assuming no boundary (we denote such accumulated evidence as Y' to distinguish it from the one *with* absorbing boundaries, Y):

$$\begin{aligned}
&\mathrm{P}(Y'_t = y', T_{\mathrm{D}} \geq t \mid \boldsymbol{\theta}_{\mathrm{D}}, \mathbf{B}) \\
&\quad = \sum_{y=-\infty}^{+\infty} \mathrm{P}(Y'_t = y' \mid Y_{t-\Delta t} = y, T_{\mathrm{D}} \geq t, \boldsymbol{\theta}_{\mathrm{D}}, \mathbf{B}) \mathrm{P}(Y_{t-\Delta t} = y, T_{\mathrm{D}} \geq t \mid \boldsymbol{\theta}_{\mathrm{D}}, \mathbf{B})
\end{aligned} \qquad \text{(Eq. 12)}$$

Here, note that $\mathrm{P}(Y_{t-\Delta t} = y, T_{\mathrm{D}} \geq t \mid \boldsymbol{\theta}_{\mathrm{D}}, \mathbf{B})$= $\mathrm{P}(Y_{t-\Delta t} = y, T_{\mathrm{D}} > t - \Delta t \mid \boldsymbol{\theta}_{\mathrm{D}}, \mathbf{B})$ comes from Eq. 10 in the previous time step. The transition probability is parameterized by $\boldsymbol{\theta}_{\mathrm{D}}$ as a Gaussian distribution (and is independent of $\mathbf{B}$, the latter because we are assuming no boundary for Y'):

$$\begin{aligned}
&\mathrm{P}(Y'_t = y' \mid Y_{t-\Delta t} = y, T_{\mathrm{D}} \geq t, \boldsymbol{\theta}_{\mathrm{D}}, \mathbf{B}) \\
&\quad = \mathrm{P}(Y'_t = y' \mid Y_{t-\Delta t} = y, T_{\mathrm{D}} \geq t, \boldsymbol{\theta}_{\mathrm{D}}) \\
&\quad \coloneqq \phi(y' - y;\ \mu\Delta t, \sigma^2 \Delta t)
\end{aligned} \qquad \text{(Eq. 13)}$$

where $\mu\Delta t$ is the mean, $\sigma^2 \Delta t$ is the variance, μ is the drift rate, and σ is the diffusion coefficient. We fix σ to 1 without loss of generality (Palmer, Huk, and Shadlen 2005), leaving a single parameter, μ, in $\boldsymbol{\theta}_{\mathrm{D}}$. We then find the approximate probabilities of the rightward and leftward

choice, at time t given a height of the threshold $b_{z,t}$, as the sum of the probabilities outside the thresholds (for more accurate methods, see, e.g., Smith 2000):

$$\mathrm{P}(Z = \mathrm{Right}, T_\mathrm{D} = t \mid \boldsymbol{\theta}_\mathrm{D}, \mathbf{B}) \approx \mathrm{P}(Y'_t \geq b_{z=\mathrm{Right},t}, T_\mathrm{D} \geq t \mid \boldsymbol{\theta}_\mathrm{D}, \mathbf{B}) \quad \text{(Eq. 14)}$$

$$\mathrm{P}(Z = \mathrm{Left}, T_\mathrm{D} = t \mid \boldsymbol{\theta}_\mathrm{D}, \mathbf{B}) \approx \mathrm{P}(Y'_t \leq b_{z=\mathrm{Left},t}, T_\mathrm{D} \geq t \mid \boldsymbol{\theta}_\mathrm{D}, \mathbf{B}) \quad \text{(Eq. 15)}$$

Then the probability of being undecided up to t with accumulated evidence Y_t *with* absorbing boundaries at t is set to that of the corresponding Y'_t (the one without boundaries), except when Y'_t is outside the boundaries, in which case the probability is set to 0 (again, for more accurate methods, see, e.g., Smith 2000):

$$\mathrm{P}(Y_t = y, T_\mathrm{D} > t \mid \boldsymbol{\theta}_\mathrm{D}, \mathbf{B}) \approx \begin{cases} \mathrm{P}(Y'_t = y, T_\mathrm{D} \geq t \mid \boldsymbol{\theta}_\mathrm{D}, \mathbf{B}) & \text{if } b_{z=\mathrm{Left},t} < y < b_{z=\mathrm{Right},t} \\ 0 & \text{otherwise} \end{cases} \quad \text{(Eq. 16)}$$

The steps in Eqs. 10–16 give the predicted probability of the decision time and the choice $\mathrm{P}(T_\mathrm{D} = t, Z = z \mid \boldsymbol{\theta}_\mathrm{D}, \mathbf{B})$ for all $t \geq 0$, which in turn gives $\mathrm{P}(T_\mathrm{R} = t_{\mathrm{R},i}, Z = z_i \mid \boldsymbol{\theta})$ through Eq. 3, which is optimized as in Eq. 2. Now we are ready to see whxat is different in the new approach.

Symmetric thresholds with one difficulty level

In this section, we introduce a new heuristic approach to propose the thresholds $\mathbf{B}$ without relying on $\boldsymbol{\theta}_\mathrm{B}$. We instead propose them by inferring them from the data, using the rest of the parameters $\boldsymbol{\theta}_\mathrm{N}$ and $\boldsymbol{\theta}_\mathrm{D}$, which we collectively denote as $\boldsymbol{\theta}_{\mathrm{B}-}$. Note that the proposal of the thresholds is the only difference of the new approach from the previous approaches explained in the previous section (Figure 1). In this section, we consider the case of symmetric thresholds, where $b_{z=\mathrm{Left},t} = -b_t$ and $b_{z=\mathrm{Right},t} = +b_t$, and consider the asymmetric case in a later section. We also consider only one drift rate (hence one difficulty level) for now, and generalize it to multiple difficulty levels in a later section.

To anticipate where we are getting at, we propose the threshold at each time step by maximizing the approximate probability of the threshold given the three possible outcomes (leftward, rightward, and no choice) at the time step, which we denote as $\mathrm{P}(b_t = b \mid \boldsymbol{\theta}_{\mathrm{B}-})$ below (Eq. 17; Figure 2C). The challenge is that, because the time of the decision T_D (the first crossing time of the diffusion process, Eq. 8) is observable only through the reaction time T_R after a delay of the nondecision time T_N (Eq. 1), we do not know what the outcomes were at time t. So we infer the

probabilities of the outcomes (leftward, rightward, and no choice) in two ways, one way that does not depend on b_t, and the other way that does. The first way that does not depend on b_t is to use the parameters proposed for the nondecision time, and we put the resultant probabilities in a 3-vector, $\mathbf{p}_{t|\boldsymbol{\theta}_\mathrm{N}}$ (Figure 2A; Eqs. 19–21). The second way that does depend on b_t also depends on $\boldsymbol{\theta}_\mathrm{D}$ and thresholds proposed up to $t - \Delta t$ (which we denote as $\mathbf{B}_{[0,t-\Delta t]}$), and we will put its results in another 3-vector, $\mathbf{p}_{t|b,\boldsymbol{\theta}_{\mathrm{B}-}}$ (Figure 2B; Eqs. 22–24). Then

$$\mathrm{P}(b_t = b \mid \boldsymbol{\theta}_{\mathrm{B}-}) \approx \mathrm{P}(\mathbf{p}_{t|b,\boldsymbol{\theta}_{\mathrm{B}-}} \mid n\mathbf{p}_{t|\boldsymbol{\theta}_\mathrm{N}}) \qquad \text{(Eq. 17)}$$

where n is the total number of trials (Figure 2C; see also Eq. 25). Then we set b_t to b that maximizes (the log of) this probability (Figure 2D; see also Eq. 26):

$$b_t = \mathrm{argmax}_b \log \mathrm{P}(b_t = b \mid \boldsymbol{\theta}_{\mathrm{B}-}) \qquad \text{(Eq. 18)}$$

Once we have the b_t to propose, the rest of the steps are the same as the previous approach (Figure 1, from b_t and below). To summarize, with the b_t, we compute the decision time distribution for that time step, $\mathrm{P}\left(Z = z, T_\mathrm{D} = t \,\middle|\, \boldsymbol{\theta}_\mathrm{D}, \mathbf{B}_{[0,t]}\right)$, as in Eqs. 14–15 (we use $\mathbf{B}_{[0,t]}$ in the condition, since b_t became available in addition to the thresholds for the preceding time steps, $\mathbf{B}_{[0,t-\Delta t]}$.) We can repeat this procedure up to the maximum reaction time observed (t_max) to get the entire decision time distribution $\mathrm{P}(Z = z, T_\mathrm{D} = t \mid \boldsymbol{\theta}_\mathrm{D}, \mathbf{B})$ for all $t \in [0, t_\mathrm{max}]$. We in turn use it to compute the reaction time distribution $\mathrm{P}\left(T_\mathrm{R} = t_{\mathrm{R},i}, Z = z_i \,\middle|\, \boldsymbol{\theta}_{\mathrm{B}-}\right)$ through Eq. 3, which is optimized as in Eq. 2 (with $\boldsymbol{\theta}$ replaced with $\boldsymbol{\theta}_{\mathrm{B}-}$). So as mentioned earlier, the proposal of the thresholds is the only step that differs from the previous approaches: all other parameters, in $\boldsymbol{\theta}_{\mathrm{B}-}$ ($\boldsymbol{\theta}_\mathrm{D}$ and $\boldsymbol{\theta}_\mathrm{N}$), are optimized in the same way. We defer rigorous justification and analysis of the new method to future works, but the method gives reasonable fits fast, as we can see in Results.

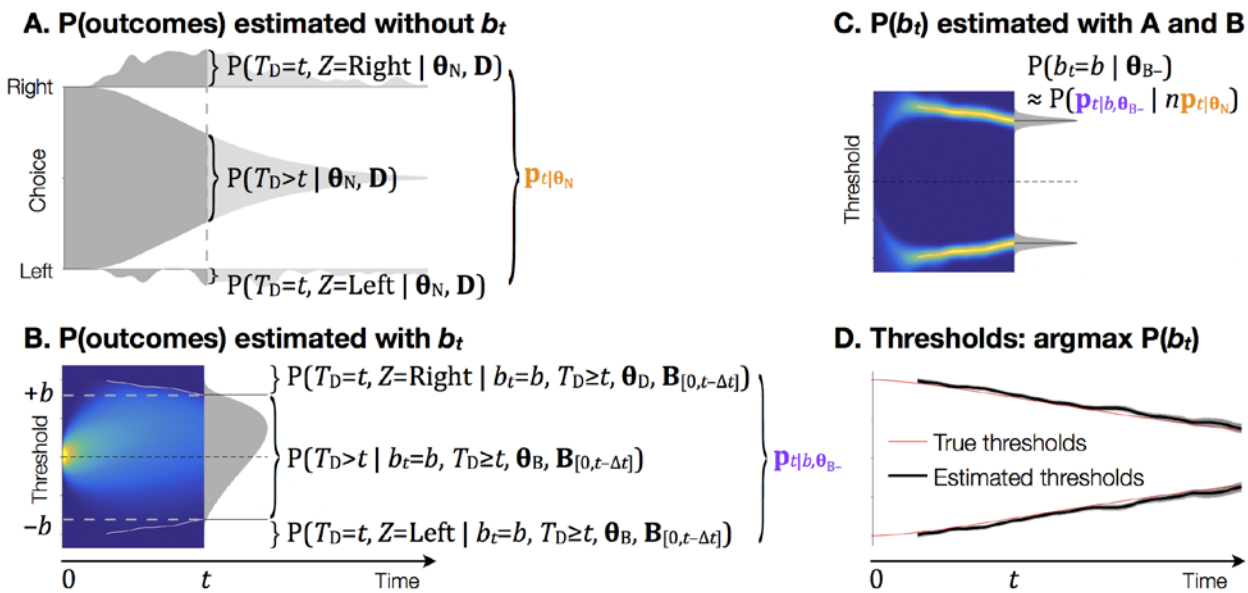

Figure 2. Estimation of the thresholds

A. Probability of the rightward, leftward, and no choice estimated *without* b_t (top, bottom, and middle shade and Figure 1, top pink frame), by convolving the observed distribution T_R with the proposed distribution of T_N flipped in time (Eqs. 19–21). Shades are lighter after t for visibility of the expressions. $\mathbf{p}_{t|\boldsymbol{\theta}_N}$ is a 3-vector of the three probabilities.

B. Probability of the rightward, leftward, and no choice estimated *with* b_t. Gray shade shows the distribution of the accumulated evidence assuming no threshold at t. If we impose the thresholds $b_t = b$, then the sums of the probabilities outside and within the thresholds approximate the probabilities of making a rightward or a leftward choice at t (top and bottom), and no choice (middle). $\mathbf{p}_{t|b,\boldsymbol{\theta}_{B-}}$ is a 3-vector that consists of the three probabilities, which sum to 1 (Eqs. 22–24). White solid lines are threshold heights estimated up to $t - \Delta t$ ($\mathbf{B}_{[0,t-\Delta t]}$; displayed from the 5th-percentile of the estimated decision time distribution). Pseudocolor shows the distribution of the accumulated evidence for $t' < t$: $P(T_D \geq t', Y_{t'} = y \mid \boldsymbol{\theta}_D)$.

C. *Top half.* The likelihood of the threshold for the rightward choice being b at each time up to t (pseudocolor) and at t (gray shade; Eq. 25); *Bottom half.* The same for the leftward choice, which is a mirror image of the top, since we are assuming symmetric thresholds in this example.

D. Thresholds estimated as the value maximizing $P(b_t = b \mid \boldsymbol{\theta}_{B-})$ for each t (Eq. 26). Black solid lines are the maximum likelihood estimates, and shades are the standard errors (likely underestimated; see text after Eq. 26). The true thresholds used in the simulation are shown in red lines. The estimated thresholds are drawn between the 5th- and 95th-percentiles of the estimated decision time distribution.

In the rest of this section, we explain details of the new method outlined above. To compute the probabilities of the outcomes $\mathbf{p}_{t|\boldsymbol{\theta}_N}$ without relying on b_t (appeared at the end of Eq. 17 and Figure 2A), we start by computing the distribution of the decision times given the data **D** (a set of observed reaction times and choices $\{(t_{R,i}, z_i)\}$) and $\boldsymbol{\theta}_N$. Because reaching the threshold for

the choice z at time t would have happened in one of the trials where the choice z was made at a later time $t_{\mathrm{R}} \geq t$, with a probability of $\mathrm{P}(T_{\mathrm{N}} = t_{\mathrm{R}} - t \mid \boldsymbol{\theta}_{\mathrm{N}})$, we can count all such trials weighted with the nondecision time probability to compute $\mathrm{P}(T_{\mathrm{D}} = t, Z = z \mid \boldsymbol{\theta}_{\mathrm{N}}, \mathbf{D})$. In practice, we convolve the observed reaction time distribution with the nondecision time distribution reversed in time (not shown in the equation; the top pink frame in Figure 1 and Figure 2A):

$$\begin{aligned}
&\mathrm{P}(T_{\mathrm{D}} = t_{\mathrm{D}}, Z = z \mid \boldsymbol{\theta}_{\mathrm{N}}, \mathbf{D}) \\
&\qquad \approx \frac{1}{n} \sum_{i \in \{i : z_i = z\}} \mathrm{P}\left(T_{\mathrm{N}} = t_{\mathrm{R},i} - t_{\mathrm{D}} \,\middle|\, T_{\mathrm{R}} = t_{\mathrm{R},i}, Z = z, \boldsymbol{\theta}_{\mathrm{N}}\right) \\
&\qquad \approx \frac{1}{n} \sum_{i \in \{i : z_i = z\}} f_{\mathrm{N}}(t_{\mathrm{R},i} - t_{\mathrm{D}}, z; \boldsymbol{\theta}_{\mathrm{N}}) \qquad \text{(Eq. 19)}
\end{aligned}$$

where n is the number of trials, $\{i : z_i = z\}$ is the set of trial indices with choice z, $\boldsymbol{\theta}_{\mathrm{N}}$ is the vector of parameters regarding the nondecision time, and $f_{\mathrm{N}}(t_{\mathrm{N}}, z; \boldsymbol{\theta}_{\mathrm{N}})$ gives the probability of the nondecision time t_{N} given the choice z and parameters $\boldsymbol{\theta}_{\mathrm{N}}$, such that $\sum_{t_N \geq 0} f_{\mathrm{N}}(t_{\mathrm{N}}, z; \boldsymbol{\theta}_{\mathrm{N}}) = 1$ (Eq. 5). Obviously, $f_{\mathrm{N}}\left(t_{\mathrm{R},i} - t_{\mathrm{D}}, z; \boldsymbol{\theta}_{\mathrm{N}}\right) = 0 \; \forall \; t_{\mathrm{R},i} < t_{\mathrm{D}}$.

Note that the decision time distribution computed this way is not guaranteed to maximize the probability of observing the data given $\boldsymbol{\theta}_{\mathrm{N}}$ (see Appendix). We still chose to use Eq. 19 because it is fast and simple, and we defer finding a better choice to future works. We confirmed that Eq. 19 nonetheless gives a better-than-chance approximation of the true simulated decision time distribution with a shuffle test (see Appendix and Figure 7).

With the decision time distribution estimated in Eq. 19, we also compute the probability that the decision is not made until t as (Figure 2A):

$$\mathrm{P}(T_{\mathrm{D}} > t \mid \boldsymbol{\theta}_{\mathrm{N}}, \mathbf{D}) = 1 - \sum_{\tau \in [0,t]} \sum_{z \in \{\mathrm{Left}, \mathrm{Right}\}} \mathrm{P}(T_{\mathrm{D}} = \tau, Z = z \mid \boldsymbol{\theta}_{\mathrm{N}}, \mathbf{D}) \qquad \text{(Eq. 20)}$$

where square brackets denote a closed interval. Then we put the probability of leftward, rightward, and no choice into a 3-vector, $\mathbf{p}_{t \mid \boldsymbol{\theta}_{\mathrm{N}}}$:

$$\mathbf{p}_{t \mid \boldsymbol{\theta}_{\mathrm{N}}} \coloneqq \begin{bmatrix} p_{\mathrm{Left}, t \mid \boldsymbol{\theta}_{\mathrm{N}}} \\ p_{\mathrm{Right}, t \mid \boldsymbol{\theta}_{\mathrm{N}}} \\ p_{\mathrm{Undecided}, t \mid \boldsymbol{\theta}_{\mathrm{N}}} \end{bmatrix}$$

$$\coloneqq \begin{bmatrix} \mathrm{P}(T_\mathrm{D} = t, Z = \mathrm{Left} \mid \boldsymbol{\theta}_\mathrm{N}, \mathbf{D}) \\ \mathrm{P}(T_\mathrm{D} = t, Z = \mathrm{Right} \mid \boldsymbol{\theta}_\mathrm{N}, \mathbf{D}) \\ \mathrm{P}(T_\mathrm{D} > t \mid \boldsymbol{\theta}_\mathrm{N}, \mathbf{D}) \end{bmatrix} \quad \text{(Eq. 21)}$$

Note that the three elements do not sum to 1: they are the proportions of the trials out of all trials, and will be used to estimate the effective number of trials with each outcome at t (Eq. 25). The following equality may help: $p_{\mathrm{Undecided},\tau|\boldsymbol{\theta}_\mathrm{N}} + \sum_{\tau=0}^{t}\left(p_{\mathrm{Left},\tau|\boldsymbol{\theta}_\mathrm{N}} + p_{\mathrm{Right},\tau|\boldsymbol{\theta}_\mathrm{N}}\right) = 1$ (see Eq. 20)

Now we compute the probabilities of outcomes that does depend on b_t, $\mathbf{p}_{t|b,\boldsymbol{\theta}_{\mathrm{B}-}}$ (appeared in Eq. 17 and Figure 2B). To do so, we first assume that we have already computed the probabilities of a decision and no decision up to $t - \Delta t$. For $t = \Delta t$, this is easy, since we know that no decision is made and no evidence is accumulated at $t - \Delta t = 0$, i.e., at the beginning (Eqs. 10–11). We can also assume $b_{t=0} > 0$, to avoid the trivial case where all decision is made immediately without any deliberation. For $t > \Delta t$, these probabilities depend on $\boldsymbol{\theta}_\mathrm{D}$, as well as the thresholds proposed up to $t - \Delta t$, $\mathbf{B}_{[0,t-\Delta t]}$ (which we assume we have, too.) Then we compute the probability of leftward and rightward choices given a certain value $b_t = b$, as we did in Eqs. 14–15 (Figure 2B): we first compute the distribution of accumulated evidence Y_t' assuming no boundaries at t (Eq. 12), then impose the boundaries at $\pm b$ and sum the probabilities outside each boundary to get the probability of each choice. A slight difference is that we conditionalize with $T_D \geq t$ to make the probability of outcomes sum to 1, to make it work with the Dirichlet probability distribution function later (Eq. 25):

$$p_{\mathrm{Right},t|b,\boldsymbol{\theta}_{\mathrm{B}-}} \coloneqq \mathrm{P}\left(Y_t' \geq +b_t \mid b_t = b, T_\mathrm{D} \geq t, \boldsymbol{\theta}_\mathrm{D}, \mathbf{B}_{[0,t-\Delta t]}\right) \quad \text{(Eq. 22)}$$

$$p_{\mathrm{Left},t|b,\boldsymbol{\theta}_{\mathrm{B}-}} \coloneqq \mathrm{P}\left(Y_t' \leq -b_t \mid b_t = b, T_\mathrm{D} \geq t, \boldsymbol{\theta}_\mathrm{D}, \mathbf{B}_{[0,t-\Delta t]}\right) \quad \text{(Eq. 23)}$$

We also compute the probability that the decision is not made until t as the sum of the probabilities inside the two boundaries:

$$p_{\mathrm{Undecided},t|b,\boldsymbol{\theta}_{\mathrm{B}-}} \coloneqq \mathrm{P}\left(-b_t < Y_t' < +b_t \,\middle|\, b_t = b, T_\mathrm{D} \geq t, \boldsymbol{\theta}_\mathrm{D}, \mathbf{B}_{[0,t-\Delta t]}\right) \quad \text{(Eq. 24)}$$

With the three probabilities in Eqs. 22–24, we construct a vector of length three, $\mathbf{p}_{t|b,\boldsymbol{\theta}_{\mathrm{B}-}}$, whose elements sum to 1 (Figure 2B).

Once we have the two ingredients, $\mathbf{p}_{t|\boldsymbol{\theta}_\mathrm{N}}$ and $\mathbf{p}_{t|b,\boldsymbol{\theta}_{\mathrm{B}-}}$ for Eq. 17, we compute the approximate likelihood of the threshold height b heuristically, using the Dirichlet distribution (Figure 2C):

$$\mathrm{P}(b_t = b \mid \boldsymbol{\theta}_{\mathrm{B}-}) \approx \mathrm{Dir}(\mathbf{p}_{t|b,\boldsymbol{\theta}_{\mathrm{B}-}} \mid n\mathbf{p}_{t|\boldsymbol{\theta}_\mathrm{N}})$$

$$\propto \prod_{z' \in \{\text{Left,Right,Undecided}\}} \left(p_{z',t|b,\boldsymbol{\theta}_{\text{B}-}}\right)^{np_{z',t|\boldsymbol{\theta}_{\text{N}}}} \qquad \text{(Eq. 25)}$$

where $\text{Dir}(\mathbf{x}|\boldsymbol{\alpha}) := \frac{1}{\text{B}(\boldsymbol{\alpha}+1)} \prod_{z'} x_{z'}^{\alpha_{z'}}$ is the probability density function of the Dirichlet distribution where z' indicates the category (in this case, $z' \in \{\text{Left}, \text{Right}, \text{Undecided}\}$), $x_{z'}$ is the probability of category z' and $\mathbf{x}$ is a vector containing $x_{z'}$ $\forall z'$, $\alpha_{z'}$ is the number of observed trials whose outcome is category z' and $\boldsymbol{\alpha}$ is a vector containing $\alpha_{z'}$ $\forall z'$, and $\text{B}(\boldsymbol{\alpha}+1)$ is the multivariate Beta function, which does not depend on b and hence is omitted in the last line of Eq. 25 (we define the elements of $\boldsymbol{\alpha}$ to be smaller than usual by 1 for simplicity of the notation). The Dirichlet distribution gives the probability that, given the number of occurrences of categories in $\boldsymbol{\alpha}$ sampled from a multinomial distribution, the probabilities of the categories are the elements of $\mathbf{x}$. To make it concrete, note that the elements of $\mathbf{x}$ sum to 1, and the elements of $\boldsymbol{\alpha}$ sum to $n' := n \sum_{z'} p_{z',t|\boldsymbol{\theta}_{\text{N}}} = n\text{P}(T_{\text{D}} \geq t \mid \boldsymbol{\theta}_{\text{N}}, \mathbf{D})$, which is effectively the estimated number of trials that crossed the thresholds at or after t.

Then, as anticipated in Eq. 18, we set b_t as the value that maximizes its probability in Eq. 25 (Figure 2C and D):

$$\begin{aligned} b_t &= \text{argmax}_b \log \text{P}(b_t = b \mid \boldsymbol{\theta}_{\text{B}-}) \\ &\approx \text{argmax}_b \log \text{Dir}(\mathbf{p}_{t|b,\boldsymbol{\theta}_{\text{B}-}} \mid n\mathbf{p}_{t|\boldsymbol{\theta}_{\text{N}}}) \\ &= \text{argmax}_b \sum_{z' \in \{\text{Left,Right,Undecided}\}} np_{z',t|\boldsymbol{\theta}_{\text{N}}} \log p_{z',t|b,\boldsymbol{\theta}_{\text{B}-}} \end{aligned} \qquad \text{(Eq. 26)}$$

To see the rationale for Eqs. 25–26, imagine an extreme case where the nondecision time is a known constant t_{N}^*, such that the threshold-crossing time on each trial is directly observable through the data, i.e., $t_{\text{D},i} = t_{\text{R},i} - t_{\text{N}}^*$ (i.e., there is no uncertainty about $n\mathbf{p}_{t|\boldsymbol{\theta}_{\text{N}}}$). Further imagine that we have correct values of $\boldsymbol{\theta}_{\text{D}}$ and of all thresholds up to $t - \Delta t$, $\mathbf{B}_{[0,t-\Delta t]}$, so that $\text{P}(Y_t' = y \mid \boldsymbol{\theta}_{\text{D}}, \mathbf{B}_{[0,t-\Delta t]})$ exactly matches the true distribution as well. Then, the probability that $\mathbf{p}_{t|b,\boldsymbol{\theta}_{\text{B}-}}$ gave to the number of "observations" $n\mathbf{p}_{t|\boldsymbol{\theta}_{\text{N}}}$ (about which we have no uncertainty in this case), is exactly given by the Dirichlet distribution (to the extent that the approximation in Eqs. 14–15 holds, and except that we are ignoring the prior over b_t). Since $\mathbf{p}_{t|b,\boldsymbol{\theta}_{\text{B}-}}$ in this case is a deterministic function of b (since we know $\boldsymbol{\theta}_{\text{D}}$ and $\mathbf{B}_{[0,t-\Delta t]}$ with certainty), the probability of

$\mathbf{p}_{t|b,\boldsymbol{\theta}_{\mathrm{B}-}}$ equals the probability of b, making the approximation in Eq. 25 exact. Since the Dirichlet distribution is unimodal, given infinite amount of data $n \to \infty$, which makes the elements of $\mathbf{p}_{t|\boldsymbol{\theta}_{\mathrm{N}}}$ exactly match the true value of $\mathrm{P}(T_{\mathrm{D}} = t, Z = z)$ and $\mathrm{P}(T_{\mathrm{D}} > t)$ under the generative model, Eq. 26 would give the true b_t that matches the generative model's.

In practice, the preconditions are not met, at least not exactly, so we should use caution in interpreting $\mathrm{P}(b_t = b \mid \boldsymbol{\theta}_{\mathrm{B}-})$ and b_t estimated from it. To begin with, neither $\boldsymbol{\theta}_{\mathrm{D}}$ nor $\mathbf{B}_{[0,t-\Delta t]}$ would match the true value initially in general, and $\mathbf{p}_{t|\boldsymbol{\theta}_{\mathrm{N}}}$ estimated using Eqs. 19–21 is a rough and biased approximate when T_{N} has nonzero variance (see Appendix). Also, when the amount of data is finite, $\mathbf{B}_{[0,t-\Delta t]}$ would have variability, which is ignored here. Therefore, if we repeat the experiment with the same generative model, the variability of b_t estimated from Eq. 26 would be bigger than what $\mathrm{P}(b_t = b \mid \boldsymbol{\theta}_{\mathrm{B}-})$ as in Eq. 25 suggests. This seem to be the case in our anecdotal observations, and we leave the full characterization of the confidence interval to future works. For now, we should keep in mind that the widths of the confidence intervals in Figures 2–4 are likely underestimated, which we noted in the legends.

Turning back to Eq. 26 with the caveats in mind, b_t for each time step can be found in $O(\log\lceil y_{\max}/dy\rceil)$ with a gradient ascent procedure with other parameters $\boldsymbol{\theta}_{\mathrm{B}-}$ fixed ($y_{\max}$ is the maximum level of accumulated evidence considered). It requires computing $\mathrm{P}(Y'_t = y, T_{\mathrm{D}} \geq t|\boldsymbol{\theta}_{\mathrm{D}}, \mathbf{B}_{[0,t-\Delta t]})$ once per time step (as in the parametric methods), and we denote the time complexity for that with $O(s)$, which is not smaller than $O(\lceil y_{\max}/dy\rceil)$. Since $O(\log\lceil y_{\max}/dy\rceil) \leq O(\lceil y_{\max}/dy\rceil)$, the new method does not increase the overall time complexity. (We measured and reported the execution time, as mentioned at the end of Section "Comparison of the performance with the parametric models" and corresponding Results). We do not provide a guarantee that the gradient ascent procedure will find the global maximum here, but in practice it worked reasonably, as we can see in Results.

As mentioned in the outset, we can repeat the above procedure, Eqs. 19–26, to obtain b_t for all time $t \in [0, t_{\max}]$, and compute the decision time distribution $\mathrm{P}(Z = z, T_{\mathrm{D}} = t \mid \boldsymbol{\theta}_{\mathrm{D}}, \mathbf{B})$ with Eqs. 10–16 and the reaction time distribution $\mathrm{P}\big(T_{\mathrm{R}} = t_{\mathrm{R},i}, Z = z_i \,\big|\, \boldsymbol{\theta}_{\mathrm{B}-}\big)$ with Eq. 3, and optimize it as in Eq. 2 (with $\boldsymbol{\theta}$ replaced with $\boldsymbol{\theta}_{\mathrm{B}-}$). See Algorithm 1 for review.

Symmetric thresholds with multiple difficulty levels

When there are multiple difficulty levels (in our example, different direction and noise levels of the motion stimulus), there are two additional steps we need to perform. First, we need to compute the drift rates of the diffusion process from the difficulty levels. For the motion direction discrimination task, the following linear transformation from the signed motion coherence c (positive when the motion is rightward and negative when it is leftward, with the absolute magnitude given by the probability of movement of each dot toward the signal direction; Shadlen and Newsome 2001) to the drift rate μ_c is known to work well (Palmer, Huk, and Shadlen 2005):

$$\mu_c = \kappa(c - c_0) \qquad \text{(Eq. 27)}$$

with two free parameters for $\boldsymbol{\theta}_{\mathrm{D}}$, the sensitivity κ and the bias c_0.

Second, we need to pool results across difficulty levels. It can be done simply by summing the log likelihoods across conditions, and substituting it into Eq. 26:

$$b_t = \mathrm{argmax}_b\left(\sum_c \log \mathrm{Dir}\left(\mathbf{p}_{t|b,c,\boldsymbol{\theta}_{\mathrm{B}-}} \mid n_c \mathbf{p}_{t|c,\boldsymbol{\theta}_{\mathrm{N}}} + 1\right)\right). \qquad \text{(Eq. 28)}$$

where n_c is the number of trials whose signed coherence is c, $\mathbf{p}_{t|b,c,\boldsymbol{\theta}_{\mathrm{B}-}}$ is computed just like $\mathbf{p}_{t|b,\boldsymbol{\theta}_{\mathrm{B}-}}$ with μ replaced with μ_c in Eq. 13, and $\mathbf{p}_{t|c,\boldsymbol{\theta}_{\mathrm{N}}}$ is computed just like $\mathbf{p}_{t|\boldsymbol{\theta}_{\mathrm{N}}}$ with trials whose signed coherence is c.

Asymmetric thresholds

The method can be easily generalized to two asymmetric thresholds. Here, we calculate the likelihoods of the leftward, rightward and no decision as a function of both thresholds (cf. Eqs. 14–15 and 22–24):

$$\begin{aligned} p_{\mathrm{Left},t|b_{\mathrm{L}},b_{\mathrm{R}},\boldsymbol{\theta}_{\mathrm{B}-}} &\coloneqq \mathrm{P}\left(Z = \mathrm{Left}, T_{\mathrm{D}} = t \;\middle|\; b_{\mathrm{Left},t} = b_{\mathrm{L}}, b_{\mathrm{Right},t} = b_{\mathrm{R}}, \boldsymbol{\theta}_{\mathrm{D}}, \mathbf{B}_{[0,t-\Delta t]}\right) \\ &\approx \mathrm{P}(Y'_t \le b_{\mathrm{L}}, T_{\mathrm{D}} \ge t \mid \boldsymbol{\theta}_{\mathrm{D}}, \mathbf{B}_{[0,t-\Delta t]}) \qquad \text{(Eq. 29)} \end{aligned}$$

$$\begin{aligned} p_{\mathrm{Right},t|b_{\mathrm{L}},b_{\mathrm{R}},\boldsymbol{\theta}_{\mathrm{B}-}} &\coloneqq \mathrm{P}\left(Z = \mathrm{Right}, T_{\mathrm{D}} = t \;\middle|\; b_{\mathrm{Left},t} = b_{\mathrm{L}}, b_{\mathrm{Right},t} = b_{\mathrm{R}}, \boldsymbol{\theta}_{\mathrm{D}}, \mathbf{B}_{[0,t-\Delta t]}\right) \\ &\approx \mathrm{P}(Y'_t \ge b_{\mathrm{R}}, T_{\mathrm{D}} \ge t \mid \boldsymbol{\theta}_{\mathrm{D}}, \mathbf{B}_{[0,t-\Delta t]}) \qquad \text{(Eq. 30)} \end{aligned}$$

$$\begin{aligned} p_{\mathrm{Undecided},t|b_{\mathrm{L}},b_{\mathrm{R}},\boldsymbol{\theta}_{\mathrm{B}-}} &\coloneqq \mathrm{P}\left(T_{\mathrm{D}} > t \;\middle|\; b_{\mathrm{Left},t} = b_{\mathrm{L}}, b_{\mathrm{Right},t} = b_{\mathrm{R}}, \boldsymbol{\theta}_{\mathrm{D}}, \mathbf{B}_{[0,t-\Delta t]}\right) \\ &\approx \mathrm{P}(b_{\mathrm{L}} < Y'_t < b_{\mathrm{R}}, T_{\mathrm{D}} \ge t \mid \boldsymbol{\theta}_{\mathrm{D}}, \mathbf{B}_{[0,t-\Delta t]}) \qquad \text{(Eq. 31)} \end{aligned}$$

and with them, we construct a vector of length three, $\mathbf{p}_{t|b_\mathrm{L},b_\mathrm{R},\boldsymbol{\theta}_{\mathrm{B}-}}$. Then, as in Eq. 25, the likelihood of the thresholds is given by the Dirichlet distribution:

$$\begin{aligned}
&\mathrm{P}\left(b_{\mathrm{Left},t} = b_\mathrm{L}, b_{\mathrm{Right},t} = b_\mathrm{R} \,\middle|\, \boldsymbol{\theta}_{\mathrm{B}-}\right) \\
&\quad = \mathrm{Dir}\left(\mathbf{p}_{t|b_\mathrm{L},b_\mathrm{R},\boldsymbol{\theta}_{\mathrm{B}-}} \mid n\mathbf{p}_{t|\boldsymbol{\theta}_\mathrm{N}}\right) \\
&\quad \propto \prod_{z' \in \{\mathrm{Left,Right,Undecided}\}} \left(p_{z',t|b_\mathrm{L},b_\mathrm{R},\boldsymbol{\theta}_\mathrm{D},\boldsymbol{\theta}_\mathrm{N}}\right)^{np_{z',t|\boldsymbol{\theta}_\mathrm{N}}}
\end{aligned} \tag{Eq. 32}$$

The threshold for t is set to the point with the maximum likelihood:

$$(b_{\mathrm{Left},t}, b_{\mathrm{Right},t}) = \mathrm{argmax}_{(b_\mathrm{L},b_\mathrm{R})}\, \mathrm{P}\left(b_{\mathrm{Left},t} = b_\mathrm{L}, b_{\mathrm{Right},t} = b_\mathrm{R} \,\middle|\, n\mathbf{p}_{t|\boldsymbol{\theta}_\mathrm{N}}\right) \tag{Eq. 33}$$

We can find the estimate using a gradient ascent procedure that optimizes $(b_{\mathrm{Left},t}, b_{\mathrm{Right},t})$ with other parameters fixed, with theoretical time complexity of $O(\lceil t_\mathrm{max}/\Delta t\rceil \log y_\mathrm{max})$, which is the same as the symmetric case. We defer characterizing the actual execution time to future works.

Example fits

To evaluate the performance of the new method, we simulated the data by sampling the reaction time T_R and choice Z from $\mathrm{P}(T_\mathrm{R}, Z \mid \boldsymbol{\theta})$ predicted by the diffusion model with known parameters (Eqs. 14–15), and fit them. First, as a sanity check, we tried giving the correct parameters $\boldsymbol{\theta}_{\mathrm{B}-}$ as the initial values (although not fixing them) and letting the new method find the thresholds (Figure 3, panels with the title "Original", and Figure 4). We also tried making one parameter of the generative model different from the initial value of the free parameters and fitting them (Figure 3, other panels). See legends for generative parameters, and Table 2 for parameter settings used in fitting.

For absolute stimulus strengths $|c|$, we used 0, 0.032, 0.064, 0.128, 0.256, and 0.512, and generated positive and negative drift rates in a balanced way. For example, when we generated 1200 total trials, there were 100 trials for $c = -0.512$, 100 for $c = +0.512$, and so on, and there were 200 trials for $c = 0$ (because $+0 = -0$).

To parametrically generate collapsing thresholds, we used the same functional form as (Kang et al. 2017):

$$b_{z,t} = sb_0(1 - \mathrm{Betacdf}(t';\, \beta_1, \beta_2)) \tag{Eq. 34}$$

where $\mathrm{Betacdf}(t;\beta_1,\beta_2) \coloneqq \mathrm{B}(\beta_1,\beta_2)\int_0^x t^{\beta_1-1}(1-t)^{\beta_2-1}dt$ is the cumulative beta distribution function (where $\mathrm{B}(\cdot)$ is the beta function, which gives a normalization factor), $t' \coloneqq t/t_{\max}$ is the relative time, and the sign s is $+1$ for $z = \mathrm{Right}$ and -1 for $z = \mathrm{Left}$ (examples in Figure 3–Figure 4). In the following, we use $t_\beta \coloneqq t_{\max}\beta_1/(\beta_1+\beta_2)$ and $b_{\log} \coloneqq \log_{10}(\beta_1\beta_2)$ instead of β_1 and β_2 to aid interpretation. Here, t_β is approximately the time of the steepest collapse (i.e., $t_\beta \approx \mathrm{argmax}_t \mid db_t/dt \mid$), and $b_{\log}$ correlates with the slope of the collapse.

In generating the datasets with asymmetric thresholds (Figure 4), we used two sets of parameters for the thresholds, mentioned in the legend. When fitting the asymmetric thresholds data, we fixed c_0 to 0, because a trend in both thresholds in the same direction (possible for asymmetric thresholds) is indistinguishable from a bias in the drift rate.

Comparison of the performance with the parametric models

To systematically compare the performance of the new method (hereafter, the "nonparametric" method) with parametric methods (Figures 5 and 6), we generated the data with random settings of the threshold shapes and other parameters. Then we compared the cost (as in Eq. 2) between the nonparametric and the parametric methods on each step of gradient descent ("iteration"). For the comparison, we only used the symmetric thresholds, for which we could fit the model with a reasonable amount of data. It would be of interest to test the nonparametric method with asymmetric thresholds after improving it with extensions such as smoothness priors in future works.

We simulated the data with two models that differ in the shape of the thresholds used, which we call "betacdf" and "bump" thresholds and datasets. For the "betacdf" datasets, we used the same shape of the thresholds as in the previous section (Eq. 34 and Figure 5, left two columns). We used the smoothly and monotonically varying thresholds because their shape was similar to that of the commonly used parametric thresholds', which we describe later in this section. We hypothesized that the parametric models would fit these datasets better than the nonparametric model. The parameters for the simulation were sampled independently and uniformly from a discrete set (denoted with curly braces) or from an open interval (denoted with parentheses) as follows: $\kappa \in \{5, 10, 20, 40\}$, $\mu_{\mathrm{N}} \in (0.25, 0.4)$ s, $\sigma_{\mathrm{N}} \in \{0.05, 0.1, 0.2\}$ s (giving corresponding $v_n = \sigma_{\mathrm{N}}/\mu_{\mathrm{N}}$), $b_0 \in (0.5, 1.5)$, $t_\beta \in (0.2, 1)$ s, $b_{\log} \in (0, 4)$.

For the "bump" datasets, we used thresholds with two step changes in the height in the opposite direction from each other (Figure 5, right two columns). We used the abruptly and non-monotonically varying thresholds because they did not match the shape of the thresholds of the common parametric models. We hypothesized that the nonparametric model would fit these datasets better than the parametric models. The initial height b_0 was sampled uniformly from an open interval $(0.5, 1.5)$. The two step changes were made at 0.5 and 1.0 s from the stimulus onset. When $b_0 < 1$, the first step was upward and the second step was downward, and when $b_0 > 1$, it was the opposite. The size of each step was sampled uniformly and independently from an open interval $(0.2, 0.8)$. The parameters for the drift and nondecision times, κ, μ_{N}, and v_{N} were sampled in the same way as the betacdf datasets.

For each of the two generative models, we generated 180 random parameter settings (i.e., made 180 betacdf datasets and 180 cosbasis datasets), and for each parameter setting (i.e., for each dataset), we simulated 2400 total trials with the same set of stimulus strengths as in the previous section.

To fit the data, we used the nonparametric method as well as two parametric methods, and compared their performance (Figure 6). The nonparametric model was optimized as described in the previous sections until convergence.

The two parametric models differed only in terms of what functional form each used for the thresholds. One model used the beta cumulative distribution function, which is the same form as the betacdf datasets', and we will call it the "betacdf" model. It had three free parameters for the thresholds, b_0, t_β, and $b_{\log}$, as described in the previous section. We will call the datasets generated with the betacdf thresholds as "betacdf datasets", and fits made with the betacdf thresholds as "betacdf fits" to disambiguate them.

The other model used cosine basis functions, and we will call it the "cosbasis" model. For this model, we used the following function for the threshold (similar to Drugowitsch et al. 2012):

$$b_t = \frac{1}{2}\sum_{j=1}^{n_{\mathrm{b}}} b_j\left[\cos\left\{2\pi(\mathrm{reltime}(t) - \phi_j)\right\} + 1\right] \qquad \text{(Eq. 35)}$$

where $\cos(x)$ gives the cosine of x for $-\pi \leq x \leq \pi$ and 0 otherwise. We used the number of the basis functions $n_{\mathrm{b}} \coloneqq 7$, and spaced them 1/4 cycles apart from each other, by using $\phi_j \coloneqq \frac{1}{4}(j - 2)$. So it offered 7 free parameters for the threshold height, $\{b_j\}$. All b_j were initially set to 0.5 (corresponding to $b_t = 1$ for all t). We also had two more parameters that define the relative time, $\mathrm{reltime}(t)$, which varies between 0 and 1 from the stimulus onset until the maximum time modeled ($t_{\max}$), following $\mathrm{reltime}(t) \coloneqq \mathrm{Betacdf}(\frac{t}{t_{\max}}; 10^{\beta_{\mathrm{b},1}}, 10^{\beta_{\mathrm{b},2}})$, where $\mathrm{Betacdf}(\cdot)$ is the cumulative beta distribution function as described in the previous section, and $\beta_{\mathrm{b},1}$ and $\beta_{\mathrm{b},2}$ are free parameters, both set to 0 at the beginning (which gives $t' = \frac{t}{t_{\max}}$). In total, the cosbasis model had 9 parameters for the thresholds.

Symbol	Definition	Initial value (min–max)	# Used in models		
			Nonpara-metric	Betacdf	Cosbasis
κ	Gain of the drift	5 (0.1–60)	1	1	1
c_0	Bias of the drift	0 (-0.256–0.256)	1	1	1
$\mu_{\mathrm{N},z}$	Mean nondecision time for choice z	0.3 s (0.15–0.75)	2	2	2
$v_{\mathrm{N},z}$	Coefficient of variation of nondecision time for choice z ($\sigma_{\mathrm{N},z}/\mu_{\mathrm{N},z}$)	0.067 (0.05–0.95)	2	2	2
b_0	Initial threshold height	1 (0.01–3)	0	1	0
t_β	Approximate time of the steepest collapse of the threshold	0.5 s (0.15–4.85)	0	1	0
$b_{\log}$	Slope of the threshold	1 (0–5)	0	1	0
b_j	Coefficient for the j-th raised cosine basis function ($j \in [1, 7]$)	0.5 (0–1.5)	0	0	7
$\beta_{\mathrm{b},1}, \beta_{\mathrm{b},1}$	Parameter of the Betacdf that modulates the relative time	0 (-0.5–0.5)	0	0	2

Table 2. Free parameters

All models, including the nonparametric model, had 6 additional free parameters regarding the drift (2) and the nondecision time (4). The starting point and the range of all free parameters are summarized in Table 2. Note that c_0 was included as a free parameter in fitting following a common procedure for fitting the real experimental data, although we fixed it to 0 in simulating the data for simplicity. We also included $\mu_{\mathrm{N},z}$ and $\sigma_{\mathrm{N},z}$ for each choice $z \in \{\mathrm{Left}, \mathrm{Right}\}$, although we set $\mu_{\mathrm{N},z=\mathrm{Left}} = \mu_{\mathrm{N},z=\mathrm{Right}}$ and $\sigma_{\mathrm{N},z=\mathrm{Left}} = \sigma_{\mathrm{N},z=\mathrm{Right}}$ in simulating the data, for the same reason.

To compare the performance as a function of the number of the iterations, we calculated the cost relative to the smallest cost at convergence across all three models within each simulated dataset (such that the minimum becomes 0 for each dataset). We then plotted the mean across the 180 datasets within each fitting method within each iteration and plotted it (Figure 6, first column). Because each method took different number of iterations until convergence, we used the cost at convergence for averages for iterations after convergence (e.g., if the cosbasis model converged on Iteration 30 with cost = 100 for Betacdf Dataset 15, we used 100 as the cosbasis model's cost on all later $[>30]$ iterations for Betacdf Dataset 15). We also plotted the percentage of simulations (among 180) that each parametric method had a lower cost than the nonparametric method on each iteration (Figure 6, second column).

While the nonparametric method offers the same time complexity as parametric methods, it requires extra time to search for the best threshold height per time step (Eq. 26). So we also computed the time per calculation of the cost for all three fitting methods, by averaging the time spent in 100 calculations after 3 warm-up calculations. Also, in addition to the number of iterations until convergence, we compared (1) the total execution time until convergence, and (2) the "pure" execution time devoted to the calculation of cost, by multiplying the number of iterations with the time spent per calculation of the cost for each fitting method.

Results

We first ran a sanity check, by setting the initial values of all but 0–1 free parameter to the generative model's and seeing if the new method can estimate the threshold. For symmetric thresholds, the method captured the shapes with as few as 600 total trials (Figure 3). In case of asymmetric thresholds, the method captured the qualitative difference between the two thresholds. Not surprisingly, the fit was worse compared to the symmetric case with the same number of trials (Figure 4, top row), but it improved with more trials (Figure 4, bottom row).

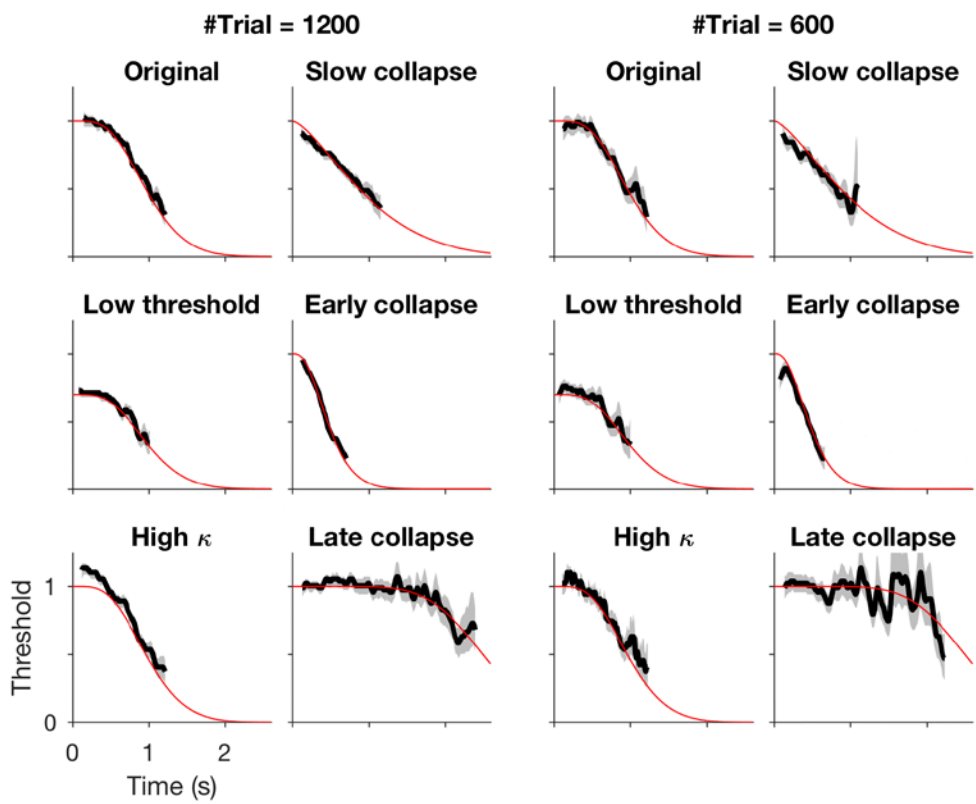

Figure 3. Comparison of the fit and the ground truth for symmetric thresholds
The black solid line is the maximum likelihood estimate, the shade is the 95% confidence interval based on $\mathrm{P}(b_t = b \mid \boldsymbol{\theta}_{\mathrm{B}-})$ at each t (likely underestimated; see text after Eq. 26), and the red line is the true threshold. Estimates are plotted between 5 and 95 percentiles of the estimated decision times. Parameters used are: $\kappa = 5$ (or 10 for the "high κ"), $b_{\log} = 1$ (or 0.5 for the "slow collapse"), $t_\beta = 0.4$ s (or 0.2 and 0.8 s for the "early collapse" and "late collapse"), and $b_0 = 1$ (or 0.7 for the "low threshold"). All other parameters had the default values listed in Table 2. Left two columns are fit with the total of 1200 trials, right two columns with 600 trials.

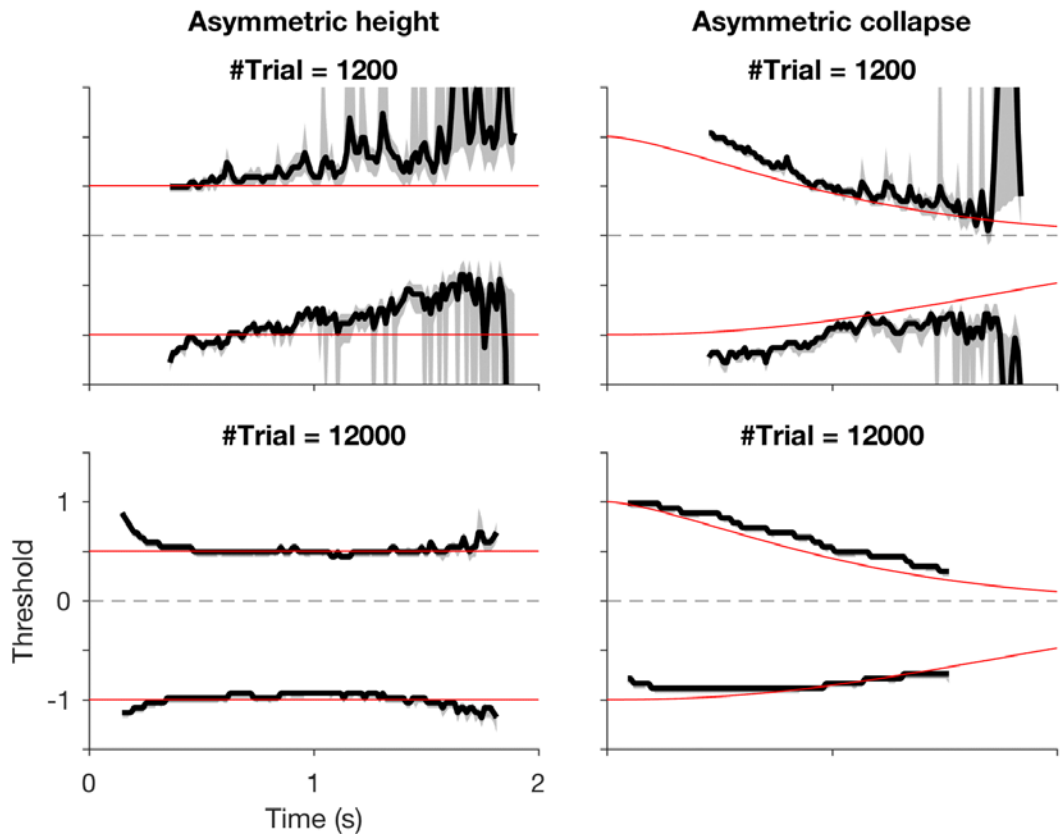

Figure 4. Comparison of the fit and the ground truth for asymmetric thresholds
The black solid line is the maximum likelihood estimate, the shade is the 95% confidence interval based on $\mathrm{P}\left(b_{z,t} = b \,\middle|\, \boldsymbol{\theta}_{\mathrm{B}-}\right)$ for each t and z (likely underestimated; see text after Eq. 26), and the red lines are the true threshold. Estimates are drawn between 5th- and 95th-percentiles of the estimated decision times. The top row is fit with 1200 trials, bottom with 12000 trials. We used $\kappa = 4$ for both examples. For the "asymmetric height" examples, we used $b_{z=\mathrm{Right},t} = 0.5$ and $b_{z=\mathrm{Left},t} = -1$ for all t. For the "asymmetric collapse" examples, we used $b_{z=\mathrm{Right},t=0} = b_{z=\mathrm{Left},t=0} = 1$ and $t_{\beta,z=+1} = 1$ s and $t_{\beta,z=-1} = 2$ s, respectively. All other parameters had the default values listed in Table 2.

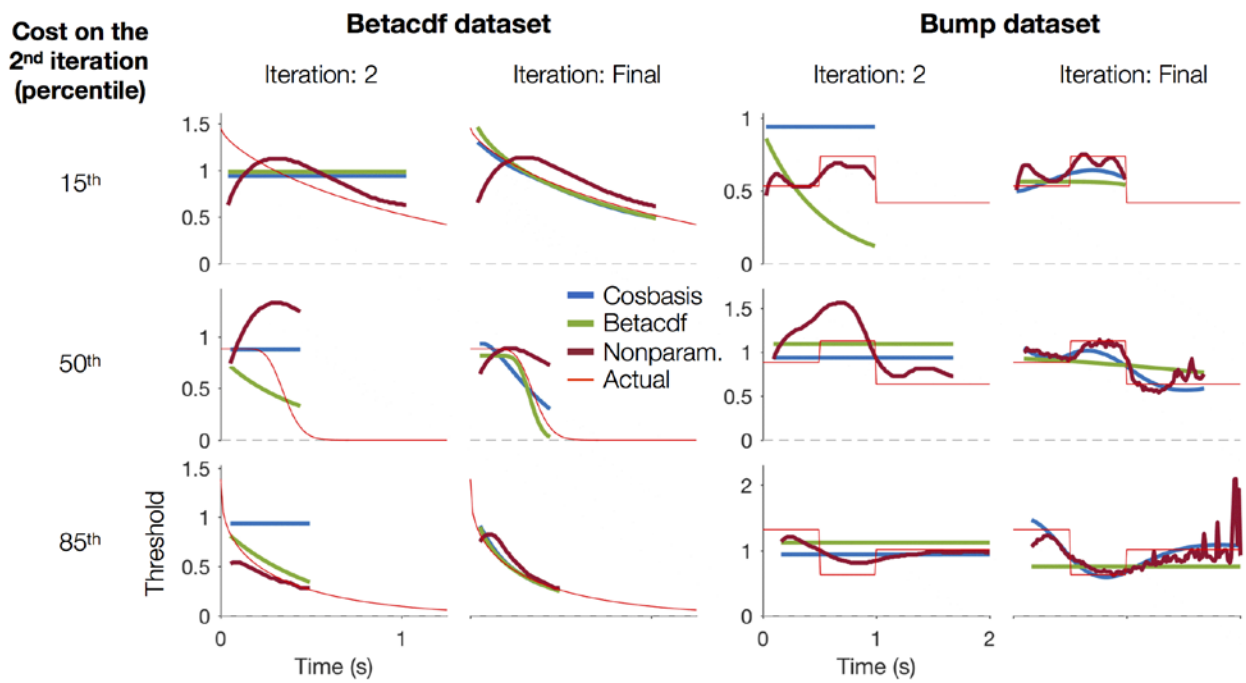

Figure 5. Example fits

First column. Example fits to the data simulated with the betacdf thresholds. The thin red line is the true threshold, and the thick blue, green, and red lines are the thresholds estimated by the cosbasis, betacdf, and nonparametric models on the 2nd iteration of the gradient descent. The fit thresholds are drawn between 5th- and 95th-percentile decision times. The top row is the simulated dataset with the 15th-percentile cost for the nonparametric model on the 2nd iteration of gradient descent, among all datasets. The second and the third rows are datasets with the 50th- and 85th-percentile costs.

Second column. The fit thresholds after convergence, for the same datasets as the first column.

Third and fourth columns. The same as the first two columns, except for the data simulated with the bump thresholds.

We then compared the performance of the nonparametric model with that of two parametric models, betacdf and cosbasis, with random parameter settings as a function of the number of iterations of the gradient descent (See Section "Comparison of the performance"). Visual examination of the fits revealed that the nonparametric fits capture the shape of the thresholds as early as the 2nd iteration, for both the betacdf and bump datasets (Figure 5).

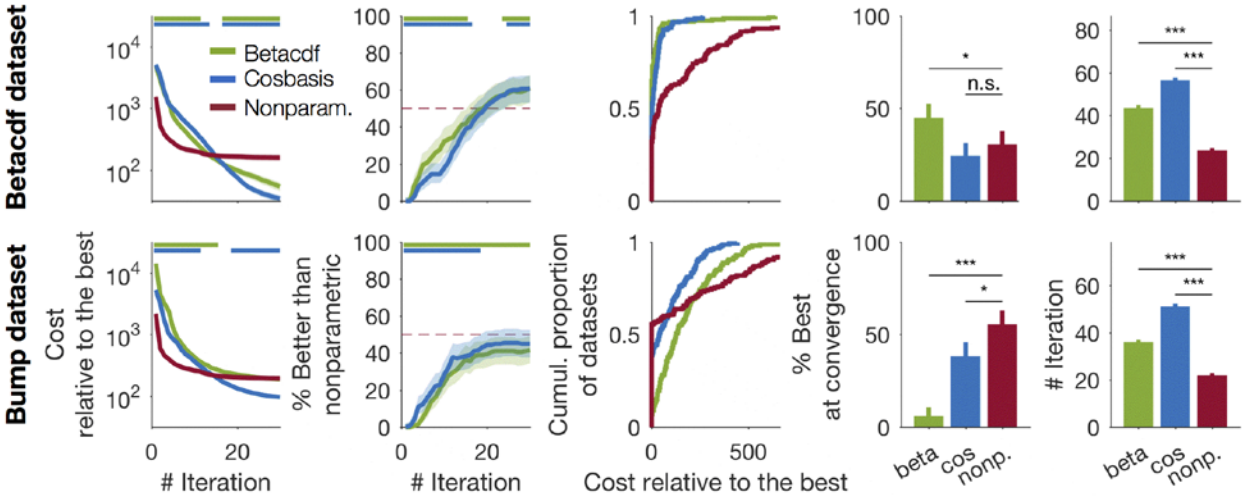

Figure 6. Comparison of performance between fitting methods

First column. Cost (relative to the smallest across all fitting methods after convergence within each simulated dataset) while fitting the betacdf and bump datasets, as a function of the number of iterations (see Methods). Lines and shades are the mean and SEM (which was often smaller than the thickness of the lines). Green, blue, and red correspond to the betacdf, cosbasis, and nonparametric fits, respectively. Horizontal lines at the top indicate the intervals where the difference is significant ($p < 0.05$, two-tailed), determined by 10,000 random pairwise shuffling of the fitting method label with the nonparametric method.

Second column. Percentage of the simulated dataset for which the given parametric method gave lower cost than the nonparametric method. Lines and shades are the mean and the 95% binomial CI. The dashed horizontal line marks the 50% level. Horizontal lines at the top indicate the intervals where the difference from 50% is significant ($p < 0.05$, two-tailed), determined as above.

Third column. Cumulative distribution of the cost relative to the best (see above) after convergence. Rightmost parts of the graphs are truncated.

Fourth column. Proportion of the datasets for which each fitting method gives the lowest cost (across the 3 methods) at convergence (%) and the 95% CI. Markers indicate the significance of pairwise comparisons (sign test; ***: $p < 0.001$; **: $p < 0.01$; *: $p < 0.05$; n.s.: $p > 0.1$).

Fifth column. Average number of iterations needed for convergence (± SEM). The meaning of the markers is the same as above.

Finally, we compared the performance of the parametric and nonparametric fitting methods quantitatively. On the first column of Figure 6, we can see that all fitting methods have a high cost at the beginning of the fitting process, which decreases in later iterations. Notably, the nonparametric fits gave a lower average cost compared to parametric fits from the first iteration of the gradient descent for both betacdf and bump datasets (shuffle test: $p \leq 0.0002$), and kept the lead for 11–15 iterations (Figure 6, first column; $p < 0.05$ up to iteration 11, 13, 15, and 11

compared to the combinations of the betacdf fit-betacdf dataset, cosbasis fit-betacdf dataset, betacdf fit-bump dataset, and cosbasis fit-bump dataset, respectively.)

We also evaluated the nonparametric fits in terms of the *proportion* of the simulated datasets (each with a random parameter setting) for which the nonparametric fits gave a higher cost than the parametric fits on each iteration (Figure 6, second column). The results favored the nonparametric fits more than the average cost did: the nonparametric fits gave a better fit in significantly more datasets than the parametric fits on the first iterations (shuffle test: $p < 0.05$ up to the first 15, 16, 133, and 18 iterations compared to the combinations of the betacdf fit-betacdf dataset, cosbasis fit-betacdf dataset, betacdf fit-bump dataset, and cosbasis fit-bump dataset, respectively.) By either metric (average cost or proportion of lower cost), the nonparametric model was a significantly better choice than the parametric models if we were to restrict the number of iterations allowed for fitting.

Not surprisingly, for the betacdf dataset, the nonparametric fits were not better than the parametric fits at convergence. Nonparametric fits gave the smallest cost less frequently than the betacdf fits (Figure 6, fourth column, top row; $p = 0.03$), although it did not differ significantly from cosbasis fits ($p = 0.31$), and its average cost was higher (paired t-test: $p < 0.001$) compared to both betacdf and cosbasis fits. This is expected since the parametric models are constrained to have smoothly-varying thresholds, as is the model that generated the dataset.

However, for the bump dataset, the nonparametric fits gave the best fit significantly more frequently than either the betacdf or the cosbasis fits (Figure 6, fourth column, bottom row; sign test: $p < 0.001$ and $p = 0.02$), although it converged significantly faster (Figure 6, rightmost column, bottom row; paired t-test: $p < 0.001$ compared to both betacdf and cosbasis fits).

Its advantage in the speed was qualitatively the same (the order among fitting methods was the same and the p-values were all < 0.001), when we tried comparing (1) the total execution time, or (2) the "pure" execution time devoted to the calculation of cost, by multiplying the number of iteration with the time per calculation of cost. Here, the time per calculation of cost were 67 ± 8, 70 ± 4, and 91 ± 5 ms [mean ± SD] for the betacdf, cosbasis, and nonparametric models, based on 100 calculations after 3 warm-up calculations. The ratio of the time compared to the nonparametric model's were 74% and 77% for the betacdf and cosbasis models.

Interestingly, the average cost of the nonparametric fits at convergence for the bump dataset was higher than that of the cosbasis fits (paired t-test: $p < 0.001$), although it was not significantly different from that of the betacdf fits ($p = 0.57$). Inspection of the cumulative distribution function of the cost at convergence showed why (Figure 6, third column, bottom row). The nonparametric fits gave the lowest cost most frequently (red), but when it was not the best, it tended to give higher cost than the parametric models. We discuss the implication of the characteristics of the nonparametric model's performance below.

Discussion

We showed that we can estimate time-varying decision thresholds without an assumption on their shape with a simple heuristic algorithm. The method does not require any hyperparameter for the threshold, yet it showed significantly better performance compared to parametric methods when a limited number of gradient descent steps are run: during the initial steps, it was even better than the parametric models of the same functional form as the one generated the data (Figure 6, first column). At convergence, the method performed better than parametric methods more frequently when the shape of the thresholds of the model that generated the data did not match the parametric models' (Figure 6, fourth column, bottom row), although it needed fewer iterations to converge (Figure 6, rightmost column, bottom row).

Given the method's initial advantage in performance from the first step of the gradient descent, it would be especially useful for situations where computing time is extremely limited, such as for online analyses of the data from high-throughput neural recording, financial markets, or large-scale internet-based services. Notably, the new method showed the initial advantage even compared to models whose functional form of the threshold matched what generated the data. If this comparative advantage holds for other functional forms as well, we may be able to choose the new method without worrying about exactly what kind of threshold generated the data, as long as we need to limit the fitting time enough.

When the shape of the true thresholds did not match the shape assumed by the parametric models, the nonparametric model performed the best on convergence more frequently than the parametric models, but when its fit was bad, the cost tended to be higher (Figure 6, third column, bottom row). Such performance may be suited for applications to a competitive situation where being the best model is the only important consideration (i.e., when it does not matter how bad the fit is

when it is not the best among all.) To improve the model's fits in the future, it may be worthwhile to examine if there is a particular regime of the parameter settings that the nonparametric model fails to fit.

In this preliminary report, we did not use cross-validated costs, and the obvious next step for us is to confirm the results with cross-validation. However, we suspect that the results shown would not change qualitatively, because our conclusions are based on simulated data where we know the true thresholds, and the match (or mismatch) of the fitted thresholds (Figure 5) with the true thresholds is unlikely to have been driven entirely by outliers.

There are a number of directions that the method can be extended. The simplest would be to introduce a smoothness prior. The appropriate smoothness may be found using cross validation. Another direction is to use a better method to estimate the decision time distribution from the reaction time distribution than the current one (Eq. 20) which is approximate and biased (see Appendix). While many other directions are possible, it would need careful consideration to retain the advantage in speed, because if the speed is not a concern we can simply resort to parametric models (e.g., the cosbasis model with many bases). In this regard, it may be worthwhile to look for a way to get the best of both worlds. For example, the method may be combined with parametric models, by quickly providing information about the functional form. We may first run the nonparametric model for a few iterations to find out around how much time after the stimulus onset the threshold changes the most, and subsequently fit the data with cosine basis functions, with more basis functions placed around the time of the most change. Such elaboration may improve the fit while preserving the speed. The method we described here may serve as a simple yet useful starting point for more elaborate methods.

Appendix. Estimating the decision time distribution from reaction times and $\boldsymbol{\theta}_{\mathbf{N}}$

In Eq. 19, we estimate the distribution of the decision time T_{D} from the observed reaction times (included in the data $\mathbf{D}$) and parameters for the nondecision time $\boldsymbol{\theta}_{\mathrm{N}}$ for the case with choice z, which is reproduced here:

$$\mathrm{P}(T_{\mathrm{D}} = t_{\mathrm{D}}, Z = z \mid \boldsymbol{\theta}_{\mathrm{N}}, \mathbf{D})$$

$$\approx \frac{1}{n} \sum_{i \in \{i: z_i = z\}} \mathrm{P}\left(T_\mathrm{N} = t_{\mathrm{R},i} - t_\mathrm{D} \,\middle|\, T_\mathrm{R} = t_{\mathrm{R},i}, Z = z, \boldsymbol{\theta}_\mathrm{N}\right)$$

$$\approx \frac{1}{n} \sum_{i \in \{i: z_i = z\}} f_\mathrm{N}(t_{\mathrm{R},i} - t_\mathrm{D}, z; \boldsymbol{\theta}_\mathrm{N}) \qquad \text{(Eq. 36)}$$

To simplify the problem without losing generality, let's condition all variables on $Z = z$:

$$\mathrm{P}(T_\mathrm{D} = t_\mathrm{D} \mid Z = z, \boldsymbol{\theta}_\mathrm{N}, \mathbf{D})$$

$$\approx \frac{1}{|i: z_i = z|} \sum_{i \in \{i: z_i = z\}} \mathrm{P}\left(T_\mathrm{N} = t_{\mathrm{R},i} - t_\mathrm{D} \,\middle|\, T_\mathrm{R} = t_{\mathrm{R},i}, Z = z, \boldsymbol{\theta}_\mathrm{N}\right)$$

$$\approx \frac{1}{|i: z_i = z|} \sum_{i \in \{i: z_i = z\}} f_\mathrm{N}(t_{\mathrm{R},i} - t_\mathrm{D}, z; \boldsymbol{\theta}_\mathrm{N}) \qquad \text{(Eq. 37)}$$

Here, $|\cdot|$ denotes the cardinality, division with which ensures that the right-hand side sums to 1. The decision time distribution computed this way is not guaranteed to maximize the probability of observing the data given $\boldsymbol{\theta}_\mathrm{N}$. To see this, consider Eq. 1, reproduced here with asterisks attached, to denote that the symbols relate to the true random variables of the generative model, in the case where the choice is z:

$$T_\mathrm{R}^* \coloneqq T_\mathrm{D}^* + T_\mathrm{N}^* \qquad \text{(Eq. 38)}$$

Let's assume that our $\boldsymbol{\theta}_\mathrm{N}$ matched the generative model's. Then, if we denote with $\tilde{T}_\mathrm{D}$ the T_D we calculated above with Eq. 37, and with $\tilde{T}_\mathrm{R}$ the empirical distribution $\{t_{R,i}\}$ in the case $i \in \{i: z_i = z\}$, then Eq. 37 can be rewritten as:

$$\tilde{T}_\mathrm{D} \coloneqq \tilde{T}_\mathrm{R} - T_\mathrm{N}^* \qquad \text{(Eq. 39)}$$

This way of calculation is nice in that the expectation of $\tilde{T}_\mathrm{D}$ matches the true value:

$$\begin{aligned} \mathrm{E}[\tilde{T}_\mathrm{D}] &= \mathrm{E}[\tilde{T}_\mathrm{R}] - \mathrm{E}[T_\mathrm{N}^*] \\ &= \mathrm{E}[T_\mathrm{R}^*] - \mathrm{E}[T_\mathrm{N}^*] \\ &= \mathrm{E}[T_\mathrm{D}^*] \end{aligned} \qquad \text{(Eq. 40)}$$

However, the expectation of the variance $\mathrm{E}\left[\mathrm{Var}[\tilde{T}_\mathrm{D}]\right]$ is bigger than the true value, since variance is added even when the random variable is subtracted, and since we are interested in the case where $\mathrm{Var}[T_\mathrm{N}^*] > 0$:

$$
\begin{aligned}
\mathrm{E}\left[\mathrm{Var}[\tilde{T}_{\mathrm{D}}]\right] &= \mathrm{E}\left[\mathrm{Var}[\tilde{T}_{\mathrm{R}}]\right] + \mathrm{E}\left[\mathrm{Var}[T_{\mathrm{N}}^{*}]\right] \\
&= \mathrm{Var}[T_{\mathrm{R}}^{*}] + \mathrm{Var}[T_{\mathrm{N}}^{*}] \\
&> \mathrm{Var}[T_{\mathrm{R}}^{*}] - \mathrm{Var}[T_{\mathrm{N}}^{*}] \\
&= \mathrm{Var}[T_{\mathrm{D}}^{*}]
\end{aligned}
\qquad \text{(Eq. 41)}
$$

The last equality comes from applying $\mathrm{Var}[\cdot]$ to both sides of Eq. 38 and rearranging. Therefore, we expect $\tilde{T}_{\mathrm{D}}$ to be biased to have too big a variance.

There may be better ways to estimate T_{D} given $\boldsymbol{\theta}_{\mathrm{N}}$ and $\mathbf{D}$. We can imagine using some kind of deconvolution of $\tilde{T}_{\mathrm{R}}$, probably after smoothing for numerical stability. We can also imagine parameterizing the distribution of T_{D} with, e.g., cosine basis functions, and optimizing it to give the best likelihood of observing the data, through steps like Eqs. 2–3. However, we chose to use Eq. 36 (i.e., Eq. 19) due to its simplicity and speed. There certainly could be a better method, and we defer finding it to future works. We nonetheless confirmed that Eq. 36 gives a better-than chance approximation of the true distribution of T_{D}^{*}, as follows.

To see if Eq. 36 (i.e., Eq. 19) gives a better-than-chance approximation of the true simulated decision time distribution, we computed the Kullback-Leibler divergence:

$$
\mathrm{D}_{\mathrm{KL,original}} \coloneqq \mathrm{D}_{\mathrm{KL}}\left(\mathrm{P}(T_{\mathrm{D}} = t, Z = z \mid \boldsymbol{\theta}_{\mathrm{D}}^{*}, \mathbf{B}^{*}) \,||\, \mathrm{P}\left(T_{\mathrm{D}} = t, Z = z \mid \widehat{\boldsymbol{\theta}}_{\mathrm{N}}, \mathbf{D}\right)\right) \qquad \text{(Eq. 42)}
$$

Here, $\mathrm{P}(T_{\mathrm{D}} = t, Z = z \mid \boldsymbol{\theta}_{\mathrm{D}}^{*}, \mathbf{B}^{*})$ is the distribution from the generative model, computed with Eqs. 14–15 with the $\boldsymbol{\theta}_{\mathrm{D}}$ and $\mathbf{B}$ used in simulating the data (denoted by asterisks), as described in Section "Comparison of the performance with the parametric models". $\mathrm{P}\left(T_{\mathrm{D}} = t, Z = z \mid \widehat{\boldsymbol{\theta}}_{\mathrm{N}}, \mathbf{D}\right)$ is computed as in Eq. 36 with the value of $\boldsymbol{\theta}_{\mathrm{N}}$ estimated by the new (nonparametric) method at convergence, $\widehat{\boldsymbol{\theta}}_{\mathrm{N}}$. We computed $\mathrm{D}_{\mathrm{KL,original}}$ for all 180 "bump" datasets, then examined their difference from $\mathrm{D}_{\mathrm{KL,shuffled}}$, which are computed after shuffling the association between $(\boldsymbol{\theta}_{\mathrm{D}}^{*}, \mathbf{B}^{*})$ and $\widehat{\boldsymbol{\theta}}_{\mathrm{N}}$ across the 180 datasets. The results clearly showed that $\mathrm{D}_{\mathrm{KL,shuffle}}$ tends to be bigger than $\mathrm{D}_{\mathrm{KL,original}}$ (sign test; $p < 10^{-15}$; Figure 7), confirming that Eq. 36 (i.e., Eq. 19) gives a better approximation of the original decision time distribution than of a randomly chosen decision time distribution. The performance is definitely not perfect (only 89% of $\mathrm{D}_{\mathrm{KL,original}}$ was better than $\mathrm{D}_{\mathrm{KL,shuffle}}$): it may be due to similarity between the decision time distributions,

bias in the algorithm as described above, or likely, both. We defer it to future works to improve the estimation of the decision time distribution.

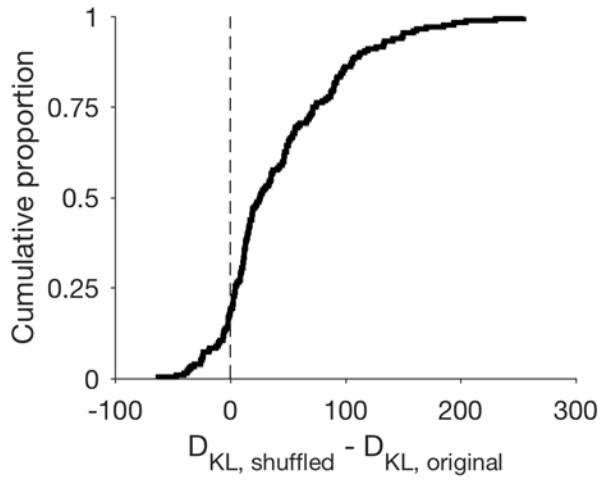

Figure 7. Cumulative distribution of the difference in D_{KL}
See above text and Eq. 42. Positive value on the abscissa indicates that the original is better.

References

Bogacz, Rafal, Eric Brown, Jeff Moehlis, Philip Holmes, and Jonathan D. Cohen. 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. https://doi.org/10.1037/0033-295X.113.4.700.

Drugowitsch, Jan, Rubén Moreno-Bote, Anne K. Churchland, Michael N. Shadlen, and Alexandre Pouget. 2012. "The Cost of Accumulating Evidence in Perceptual Decision Making." *The Journal of Neuroscience* 32 (11):3612–28. https://doi.org/10.1523/JNEUROSCI.4010-11.2012.

Gold, Joshua I., and Michael N. Shadlen. 2007. "The Neural Basis of Decision Making." *Annual Review of Neuroscience* 30 (1):535–74. https://doi.org/10.1146/annurev.neuro.29.051605.113038.

Hawkins, Guy E., Birte U. Forstmann, Eric-Jan Wagenmakers, Roger Ratcliff, and Scott D. Brown. 2015. "Revisiting the Evidence for Collapsing Boundaries and Urgency Signals in Perceptual Decision-Making." *The Journal of Neuroscience* 35 (6):2476–84. https://doi.org/10.1523/JNEUROSCI.2410-14.2015.

Kang, Yul H. R., Frederike H. Petzschner, Daniel M. Wolpert, and Michael N. Shadlen. 2017. "Piercing of Consciousness as a Threshold-Crossing Operation." *Current Biology* 27 (15):2285–2295.e6. https://doi.org/10.1016/j.cub.2017.06.047.

Palmer, John, Alexander C. Huk, and Michael N. Shadlen. 2005. "The Effect of Stimulus Strength on the Speed and Accuracy of a Perceptual Decision." *Journal of Vision* 5 (5). https://doi.org/10.1167/5.5.1.

Shadlen, Michael N., T.D. Hanks, A.K. Churchland, R. Kiani, and T. Yang. 2006. “The Speed and Accuracy of a Simple Perceptual Decision: A Mathematical Primer.” In *Bayesian Brain: Probabilistic Approaches to Neural Coding*, 209–237.

Shadlen, Michael N., and William T. Newsome. 2001. “Neural Basis of a Perceptual Decision in the Parietal Cortex (Area LIP) of the Rhesus Monkey.” *J Neurophysiol* 86 (4):1916–36.

Smith, Philip L. 2000. “Stochastic Dynamic Models of Response Time and Accuracy: A Foundational Primer.” *Journal of Mathematical Psychology* 44 (3):408–63. https://doi.org/10.1006/jmps.1999.1260.

Wald, A. 1973. *Sequential Analysis*. Dover Books on Advanced Mathematics. Dover Publications.

Wald, A., and J. Wolfowitz. 1948. “Optimum Character of the Sequential Probability Ratio Test.” *The Annals of Mathematical Statistics* 19 (3):326–39.