

Psychological Review

Optimal Decision Making in Neural Inhibition Models

Don van Ravenzwaaij, Han L. J. van der Maas, and Eric-Jan Wagenmakers

Online First Publication, November 21, 2011. doi: 10.1037/a0026275

CITATION

van Ravenzwaaij, D., van der Maas, H. L. J., & Wagenmakers, E.-J. (2011, November 21). Optimal Decision Making in Neural Inhibition Models. *Psychological Review*. Advance online publication. doi: 10.1037/a0026275

Optimal Decision Making in Neural Inhibition Models

Don van Ravenzwaaij, Han L. J. van der Maas, and Eric-Jan Wagenmakers
University of Amsterdam

In their influential *Psychological Review* article, Bogacz, Brown, Moehlis, Holmes, and Cohen (2006) discussed optimal decision making as accomplished by the drift diffusion model (DDM). The authors showed that neural inhibition models, such as the leaky competing accumulator model (LCA) and the feedforward inhibition model (FFI), can mimic the DDM and accomplish optimal decision making. Here we show that these conclusions depend on how the models handle negative activation values and (for the LCA) across-trial variability in response conservativeness. Negative neural activations are undesirable for both neurophysiological and mathematical reasons. However, when negative activations are truncated to 0, the equivalence to the DDM is lost. Simulations show that this concern has practical ramifications: The DDM generally outperforms truncated versions of the LCA and the FFI, and the parameter estimates from the neural models can no longer be mapped onto those of the DDM in a simple fashion. We show that for both models, truncation may be avoided by assuming a baseline activity for each accumulator. This solution allows the LCA to approximate the DDM and the FFI to be identical to the DDM.

Keywords: drift diffusion model, leaky competing accumulator model, feedforward inhibition model, phase planes, model equivalence

Supplemental materials: <http://dx.doi.org/10.1037/a0026275.supp>

From the end of the 19th century until the present day, experimental psychology has been dominated by the choice response time, or choice RT, paradigm. In this paradigm, researchers seek to draw substantive conclusions about psychological processes from the measurement of response speed and accuracy in relatively simple tasks. Traditionally, performance in these tasks is summarized by two straightforward measures: mean RT and proportion correct responses. While these measures are intuitively appealing, they ignore important findings such as the shape of the RT distribution and the speed–accuracy tradeoff. Also, measures such as mean RT and proportion correct do not allow a direct estimation of the latent psychological processes that drive performance. In an attempt to examine the underlying psychological processes of performance, a number of choice RT models have been developed (Luce, 1986; Ratcliff, 1978; Townsend & Ashby, 1983).

One of the most popular choice RT models is the drift diffusion model (DDM; Ratcliff, 1978; Ratcliff & McKoon, 2008). In the DDM, a decision process with two response alternatives is conceptualized as the gradual accumulation of noisy evidence. The DDM has been successfully applied to a wide range of experimental paradigms and subject populations. For example, Ratcliff,

Thapar, Gomez, and McKoon (2004) applied the DDM to RT data from a recognition memory task in order to assess the cognitive components affected by aging. In line with earlier research, Ratcliff, Thapar, et al. found that older participants responded more slowly than college-age participants. A DDM analysis revealed that this difference in performance was due to two factors: Older participants were relatively cautious, and they were slower in the peripheral, nondecision component of processing (e.g., speed of perceptual and motor processes). Importantly, the DDM analysis showed that older participants processed information at the same rate as did the college-age participants. This example shows how the DDM allows researchers to draw detailed conclusions about cognitive processes, something that is impossible from observed data alone.

The most parsimonious version of the DDM without across-trial variability in its parameters is a continuous version of the sequential probability ratio test (e.g., Bogacz, Brown, Moehlis, Holmes, & Cohen, 2006; Laming, 1968; Wald & Wolfowitz, 1948). This means that the DDM makes optimal decisions in the sense that for a given percentage correct, the model will have the lowest associated mean RT. Similarly, for a given mean RT, the model will produce the highest associated percentage correct.¹ Thus, not only has the DDM been shown to be useful as a tool to measure latent cognitive processes, the existence of the model can also be motivated from the perspective of optimal information processing.

The DDM is one of the most successful models in mathematical psychology and has been applied to data many times (see the next section for a summary). However, Usher and McClelland (2001) argued that a biologically plausible model should feature leakage

Don van Ravenzwaaij, Han L. J. van der Maas, and Eric-Jan Wagenmakers, Department of Psychology, University of Amsterdam, Amsterdam, the Netherlands.

This research was supported by a Vidi grant from the Dutch Organization for Scientific Research (NWO).

Correspondence concerning this article should be addressed to Don van Ravenzwaaij, University of Amsterdam, Department of Psychology, Roetersstraat 15, 1018 WB Amsterdam, the Netherlands. E-mail: d.vanravenzwaaij@uva.nl

¹ This definition is only one of many possible definitions of optimality, a point to which we return in the Conclusion.

of activation (which the DDM does not have; see also Wang, 2002); Usher and McClelland also noted that the DDM does not easily extend to decision problems with more than two alternatives. In a biologically plausible model, decisions with more than two alternatives are accommodated by means of feedforward, pooled, or lateral inhibition (e.g., Ditterich, 2010).

The search for a more biologically plausible model for speeded decision making has resulted in the development of several neural inhibition models. While neural inhibition models are similar to the DDM in some respects (e.g., gradual integration of noisy evidence to a threshold), they differ from the DDM in others (e.g., the presence of leakage and inhibition). Two prominent examples of neural inhibition models are the leaky competing accumulator model (LCA; Usher & McClelland, 2001) and the feedforward inhibition model (FFI; Ditterich, 2010; Shadlen & Newsome, 2001); for models inspired by LCA or FFI, see for instance Brown and Heathcote (2008), Ratcliff and Starns (2009), and Purcell et al. (2010).

An important question is whether the neural inhibition models are able, just as the DDM is, to make decisions in an optimal way. If so, the neural inhibition models would have two desirable properties: biological plausibility and optimality. This issue was addressed with mathematical rigor by Bogacz et al. (2006), who identified conditions under which the DDM and a number of neural inhibition models are equivalent. The goal of this article is to demonstrate when the conclusions from Bogacz et al. hold and when they do not. Specifically, we show that model equivalence is compromised by truncating negative neural activations to zero and by across-trial variability in the response conservativeness. We show in this article that the first of these complications, the necessity of truncation, may be resolved by assuming a baseline activity for neurons.

The outline of this article is as follows. In the first part we introduce the DDM and the LCA. We then introduce phase planes, a graphical method to represent data from two accumulators in a two-dimensional plane. Bogacz et al. (2006) used phase planes to demonstrate conditions of model equivalence. The next two sections discuss two challenges that have to be overcome before there can be model equivalence between the DDM and the LCA. The first challenge concerns the potential for negative activations in the LCA. The second challenge is that the LCA (but not the DDM) allows response conservativeness to vary over trials. We then present two different sets of simulations that show how both the LCA with truncation and the LCA without truncation perform differently from the DDM in practice.

In the second part of this article, we introduce the FFI and the model equivalence solution between the DDM and the FFI as presented by Bogacz et al. (2006). This solution is also challenged by the existence of negative neural activations. We then present simulations that show that the truncated FFI performs differently from the DDM, whereas the nontruncated FFI performs identically to the DDM.

In the third part of this article, we offer a potential solution to the problem of truncation for both the LCA and the FFI: baseline accumulator activity. We discuss the ramifications and plausibility of the assumption of neural baseline firing rates. In addition, we discuss optimality and the many ways in which this may be defined.

The DDM

In the DDM (Ratcliff, 1978; Ratcliff & Rouder, 2000; van Ravenzwaaij & Oberauer, 2009; Wagenmakers, 2009), a decision process with two response alternatives is conceptualized as the accumulation of noisy evidence over time. Evidence is represented by a single accumulator, so that evidence in favor of one alternative is evidence against the other alternative. A response is initiated when the accumulated evidence reaches one of two predefined thresholds. For instance, in a *lexical decision task*, participants have to decide whether a letter string is an English word, such as TANGO, or a nonword, such as TANAG (Figure 1).

The model assumes that the decision process starts at the starting point z , from which point evidence is accumulated with a signal-to-noise ratio that is governed by drift rate v and Wiener noise. Without trial-to-trial variability in drift rate, the change in evidence x is described by the following stochastic differential equation:

$$dx(t) = v \cdot dt + s \cdot dW(t), \quad (1)$$

where W represents the Wiener noise process (i.e., idealized Brownian motion). Parameter s represents the standard deviation of $dW(t)$. Values of v near zero produce long RTs and high error rates.

Evidence accumulation stops and a decision is initiated once the accumulator hits one of two response boundaries. The difference between these boundaries, boundary separation a , determines the speed-accuracy tradeoff; lowering a leads to faster RTs at the cost of a higher error rate. The starting point, z , is usually fixed at $a/2$, which indicates no a priori bias in the decision process. Together, these parameters generate a distribution of decision times (DTs). The observed RT, however, also consists of stimulus-nonspecific components such as response preparation and motor execution, which together make up nondecision time T_{er} . The model assumes that T_{er} simply shifts the distribution of DT, such that $RT = DT + T_{er}$ (Luce, 1986).

Thus, the four key parameters of the DDM are (a) speed of information processing, quantified by drift rate v ; (b) response caution, quantified by boundary separation a ; (c) a priori bias, quantified by starting point z ; and (d) nondecision time, quantified by T_{er} . In addition to these four parameters, the full DDM also includes parameters that specify across-trial variability in drift rate, starting point, and nondecision time (Ratcliff & Tuerlinckx, 2002), but our analytical work involves the simpler four-parameter version.

Both the optimal and the full DDMs have been successfully applied to a number of experimental paradigms, including perceptual discrimination, letter identification, lexical decision, recognition memory, signal detection, preference, and decision making. Applications for the optimal DDM include the following: Kamienkowski, Pashler, Dehaene, and Sigman (2011); Schmiedek, Lövdén, and Lindenberger (2009); Schmiedek, Oberauer, Wilhelm, Süß, and Wittmann (2007); and van Ravenzwaaij, Dutilh, and Wagenmakers (2011). Applications for the full DDM include the following: Busemeyer and Townsend (1993); Cavanagh et al. (2011); Klauer, Voss, Schmitz, and Teige-Mocigemba (2007); Philiastides, Ratcliff, and Sajda (2006); Ratcliff (1978); Ratcliff, Gomez, and McKoon (2004); Ratcliff, Hasegawa, Hasegawa, Smith, and Segraves (2007); Ratcliff, Thapar, and McKoon




Figure 1. The drift diffusion model and its key parameters, illustrated for a lexical decision task. Evidence accumulation begins at starting point z , proceeds over time guided by drift rate v , but subject to random noise, and stops when either the upper or the lower boundary is reached. Boundary separation a quantifies response caution. The predicted response time equals the accumulation time plus the time required for nondecision processes T_{er} (i.e., stimulus encoding and response execution).

(2010); Ratcliff and van Dongen (2009); Roe, Busemeyer, and Townsend (2001); van Ravenzwaaij, van der Maas, and Wagenmakers (2011); Wagenmakers, Ratcliff, Gomez, and McKoon (2008); and White, Ratcliff, Vasey, and McKoon (2010).

The LCA

The LCA (Usher & McClelland, 2001) assumes that each choice alternative is represented by a single accumulator that collects noisy evidence over time. A decision is made as soon as one of the accumulators reaches a boundary Z (Figure 2).




Figure 2. Evidence accumulation in the leaky competing accumulator model, illustrated for a lexical decision task. Evidence accumulation for each response alternative begins at decision time $DT = 0$, proceeds over time guided by drift rates I_1 and I_2 , but subject to random noise, and stops when the response threshold Z is reached. Both accumulators are subject to decay k and mutual inhibition w .

In contrast to the DDM, both of the accumulators lose a constant proportion of their current activation in each unit of time, so that their growth is negatively accelerated. The proportional leakage is a free parameter k . The accumulators for different alternatives also inhibit each other, and the strength of inhibition is a free parameter w .² Thus, for a choice between two alternatives, the accumulators' change in activation over time (cf. Bogacz et al., 2006, Equation 14) is given by

$$dx_1(t) = [-k \cdot x_1(t) - w \cdot x_2(t) + I_1]dt + s \cdot dW_1(t),$$

$$dx_2(t) = [-k \cdot x_2(t) - w \cdot x_1(t) + I_2]dt + s \cdot dW_2(t), \quad (2)$$

where dx_1 and dx_2 are the change in activation per unit time dt for the two accumulators. Parameters I_1 and I_2 are the drift rates of both accumulators. Quantities $s \cdot dW_1$ and $s \cdot dW_2$ are white noise added at each time step with $M = 0$ and variance = $s^2 \cdot dt$, analogous to the noise process of the DDM. Activation starts at zero (i.e., $x_1(0) = x_2(0) = 0$) and in the original model, negative values of x_1 and x_2 are truncated to zero (Usher & McClelland, 2001, p. 557).

More recently, some versions of the LCA have been fit to empirical data. For instance, an LCA without inhibition has been applied to data from a brightness discrimination task (Ratcliff et al., 2007); an LCA without leakage has been applied to data from a stop-signal task (Boucher, Palmeri, Logan, & Schall, 2007); and the original LCA has been applied to preferential choice tasks (Usher & McClelland, 2004).

² Usher and McClelland (2001) called this parameter β , but we follow the terminology in Bogacz et al. (2006).

LCA Phase Plane Analysis

In an effort to combine the optimality from the DDM with the biological plausibility of the LCA, Bogacz et al. (2006) identified conditions under which the LCA reduces to the DDM. To bring about this reduction, the authors made use of *phase planes* (see Figure 3).

The phase plane in Figure 3 represents the data from Figure 2. In a phase plane, x_2 is plotted against x_1 . The time dimension is now represented by the line, so that three dimensions are now depicted in a two-dimensional space. With the phase plane in place,

x_1 and x_2 can now be transformed to x_1^* and x_2^* by means of the Pythagorean theorem:

$$\begin{aligned} x_1^* &= \frac{x_1 - x_2}{\sqrt{2}}, \\ x_2^* &= \frac{x_1 + x_2}{\sqrt{2}}. \end{aligned} \quad (3)$$

In this reparameterization, x_2^* is proportional to the sum of the accumulated evidence of both accumulators, whereas x_1^* is proportional to the difference. The stochastic differential equations of x_1^* and x_2^* are given by

$$\begin{aligned} dx_1^*(t) &= \left[(w - k) \cdot x_1^*(t) + \frac{I_1 - I_2}{\sqrt{2}} \right] dt + s \cdot dW'_1(t), \\ dx_2^*(t) &= \left[(-k - w) \cdot x_2^*(t) + \frac{I_1 + I_2}{\sqrt{2}} \right] dt + s \cdot dW'_2(t), \end{aligned} \quad (4)$$

where $s \cdot dW'_1$ and $s \cdot dW'_2$ are white noise added at each time step with $M = 0$ and variance = $s^2 \cdot dt$ (cf. Bogacz et al., 2006, Equations 19, 20, 21, 22).

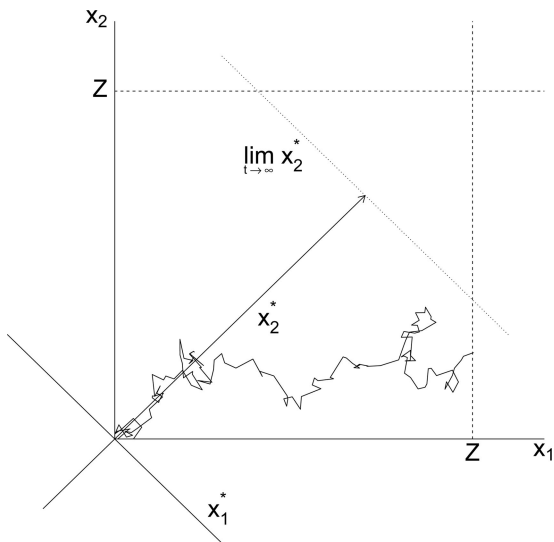


Figure 3. Phase plane of the leaky competing accumulator model from Figure 2; accumulators x_1 and x_2 are transformed to x_1^* and x_2^* .

Bogacz et al. (2006) demonstrated that the LCA process approximates the DDM process if x_2^* , which is proportional to the sum of the accumulated evidence, rapidly approaches its asymptote, $\lim_{t \rightarrow \infty} x_2^*$ (cf. Figure 3, the dotted line), while x_1^* , which is proportional to the difference in accumulated evidence, remains approximately zero. Once x_2^* is reached, x_1^* moves away from zero, ultimately yielding a response when response threshold Z is reached. This way, the choice between the two accumulators is determined mainly by movement along the x_1^* -axis, that is, through considering only the difference in activation between the two accumulators. Bogacz et al. showed this dynamic to hold when $k = w \gg 0$ and input values for both accumulators are similar. Thus, under these conditions the LCA mimics the DDM and behaves optimally.

If the DDM and the LCA are indeed equivalent, we would have the best of two worlds: a formal model for decision making that is biologically plausible and yields optimal decisions. However, examination of the model equations proposed by Bogacz et al. (2006) and the phase plane dynamics suggest two discrepancies. The first discrepancy is that the version of the LCA used in the equations by Bogacz et al. allows the accumulators to have negative activation values. Negative accumulator activation is undesirable, because it implies negative neural firing rates. However, when truncating negative accumulator activation to zero, the derivations presented by Bogacz et al. no longer hold.

The second discrepancy is that the model equivalence solution of Bogacz et al. (2006) leads to between-trial variability in the boundary separation parameter of the DDM. Although such variability may be expected in real data, it is not present in the four-parameter DDM, or in the full DDM, compromising a strict equality between the LCA and the DDM. In the next two sections, we deal with each of these discrepancies in turn.

LCA Discrepancy 1: Negative Activations

The derivation of model equivalence presented in the previous section ignored the fact that the activation of accumulators potentially has to be truncated at zero. Put differently, equivalence was demonstrated for LCA versions in which the losing accumulator is either allowed to cross the zero boundary (i.e., have a negative activation) or never hits the zero boundary. To illustrate the problem, consider a phase plane with a truncated accumulator, shown in Figure 4.

In the figure it can be seen that negative activations are prevented by means of reflecting boundaries along the x_1 and x_2 axes (e.g., Zhang & Bogacz, 2010). Negative activations are unwanted because they are biologically implausible (see also Bogacz, Usher, Zhang, & McClelland, 2007; Usher & McClelland, 2004); activations are meant to represent a neuron's firing rate, and this cannot be negative.

The interpretational problems of negative activation were originally avoided by truncating activations at zero as illustrated in Figure 4 (Grossberg, 1988; Usher & McClelland, 2001, 2004). When truncating, however, the accumulation process is no longer described by the dynamics of Equation 4 as soon as one of the accumulators hits the lower boundary. This compromises the model equivalence results from Bogacz et al. (2006).

Moreover, there is a qualitative difference between nontruncated and truncated accumulation paths for the LCA, as shown in Figure




Figure 4. Phase plane of the leaky competing accumulator model with reflecting lower boundaries (bold lines).

5. In the left panel, accumulators are not truncated at zero. The winning accumulator will grow unboundedly, because the inhibition of the losing accumulator turns into excitation (i.e., negative inhibition) when its activation is below zero. Therefore, the increasing loss of activation from the winning accumulator as a result of leakage is counteracted by an increasing gain as a result of negative inhibition from the losing accumulator. The DDM behaves similarly, as it allows an unbounded growth of evidence.

In the right panel, a reflecting boundary is imposed so that accumulators are truncated at zero. When the losing accumulator is truncated at zero, the winning accumulator will move stochastically toward an attractor, because at a certain level of activation the loss of activation that results from leakage cancels the gain of activation from the input.

Assuming that the system is deterministic (i.e., $s = 0$ in Equation 4), the attractor for the winning accumulator, $\lim_{t \rightarrow \infty} x_1(t)$, can be derived as follows:

$$\begin{aligned}
 t \gg 1 &\Rightarrow x_2(t) = 0 \\
 \Rightarrow \frac{dx_1}{dt} &= -k \cdot x_1(t) - w \cdot x_2(t) + I_1 = -k \cdot x_1(t) + I_1 \\
 \Rightarrow x_1(t) &= c \cdot e^{-kt} + \frac{I_1}{k} \\
 \Rightarrow \lim_{t \rightarrow \infty} x_1(t) &= \frac{I_1}{k}. \tag{5}
 \end{aligned}$$

The dynamics are a little more complicated for $s \neq 0$, but for typical values of s , x_1 continues to evolve close to the attractor line. The first line of Equation 5 shows that for very large t the losing accumulator approaches zero. The remaining lines of the equation show that when the losing accumulator is close to zero, the asymptote of the winning accumulator is given by I_1/k . If the LCA response threshold Z is sufficiently higher than the value of $\lim_{t \rightarrow \infty} x_1$, the asymptote of the winning accumulator lies below the

response threshold. It may then take a very long time before the stochastic process causes the accumulator to hit the response threshold. In terms of the phase planes, if Z is larger than $\lim_{t \rightarrow \infty} x_1$, the gravitation point of the accumulators lies somewhere to the left and below the decision thresholds Z , thereby preventing the execution of a response (cf. Bogacz et al., 2006, Figure 6a). The dynamics are that of a system in which there is only leakage and no inhibition, instead of a system in which the inhibition equals the leakage.

In sum, if the losing accumulator hits the response threshold, truncation is necessary to prevent negative activations from occurring (see the left panel of Figure 6). This compromises the proposed model equivalence between the DDM and the LCA. We now introduce the second discrepancy of model equivalence between the LCA and the DDM.

LCA Discrepancy 2: Changing Boundaries

The second discrepancy with the model equivalence between the LCA and the DDM as proposed by Bogacz et al. (2006) is that an LCA response threshold that is fixed over trials may correspond to across-trial variability in DDM boundary separation. Consider the phase plane in Figure 7, in which the dark accumulation path is the same as the one in Figure 3. The DDM boundary separation a_1 is determined by the location on the LCA threshold Z where the accumulator hits, as shown in Figure 7. On another trial, the evidence accumulation may correspond to the lighter accumulation path. This means that for this trial, the DDM boundary separation a_2 is different from a_1 .

Across-trial variability in boundary separation is a source of discrepancy for two reasons. The first reason is that optimal behavior for the DDM has been demonstrated only for situations in which boundary separation is fixed over trials. Unpublished simulations confirmed that allowing for across-trial variability in boundary separation increases error rate while leaving mean RT

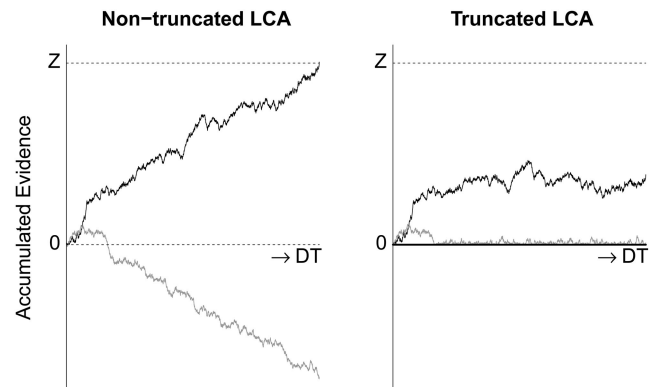


Figure 5. The effect of truncation on the dynamics of the leaky competing accumulator model (LCA). The left panel shows accumulation paths that are not truncated at zero. This allows the winning accumulator to grow unboundedly, as the increasing leakage is compensated by the increasing excitation (i.e., negative inhibition) of the losing accumulator. The right panel shows the same accumulators, now truncated at zero. The winning accumulator grows toward an asymptote and may take a very long time to reach the response threshold Z . DT = decision time.




Figure 6. The two discrepancies of model equivalence between the leaky competing accumulator model and the drift diffusion model. If the losing accumulator hits the zero boundary before the winning accumulator hits the response threshold, the discrepancy of truncation occurs (left panel). If the losing accumulator never hits the zero boundary, the discrepancy of across-trial variability in boundary separation occurs (right panel). Perfect correspondence between both models happens only when the losing accumulator hits the zero boundary just as the winning accumulator hits the response threshold (middle panel). See text for details. DT = decision time.

largely intact. In other words, the DDM with across-trial variability in boundary separation does not perform optimally.

The second reason is that a fixed across-trial LCA response threshold corresponds to a variable DDM boundary separation. There is no theoretical reason why such a correspondence should exist. Thus, transforming the LCA to the DDM leads to across-trial variability in DDM boundary separation and as such compromises a strict equality between the LCA and the DDM.

The across-trial variability in DDM boundary separation can be visualized as the difference in accumulated evidence between the winning and the losing accumulators of the LCA over repeated trials. Across-trial variability in DDM boundary separation is avoided if the losing accumulator has the same amount of activation on each trial at the time of a decision. In theory, optimal model correspondence may be accomplished only if the difference between the losing accumulator and the winning accumulator is the

same on each trial at the moment when the winning accumulator hits the response threshold on every trial (see the middle panel of Figure 6). This situation is unlikely to occur, as both the winning and the losing accumulators are noisy.

In sum, the model equivalence between the LCA and the DDM as proposed by Bogacz et al. (2006) leads to a DDM boundary separation that changes over trials (see the right panel of Figure 6). The variable boundary separation and the issue of truncation constitute two theoretical discrepancies between the DDM and the LCA. However, it is quite possible that both the DDM and the LCA show indistinguishable performance in practice. In the next two sections, we investigate how both models fare when pitted against each other.

LCA Simulation 1: Performance

In order to examine whether the LCA matches the optimal performance of the DDM, we

1. analytically calculated mean decision times for sets of DDM parameters that correspond to three different mean percentages correct;
2. used these DDM parameters to analytically calculate the corresponding LCA parameters for a range of values of decay and inhibition; and
3. generated LCA data based on the calculated LCA parameters and then compared the mean decision time and percentage correct for both models.

The next three paragraphs discuss each of these steps in more detail.

For the first step of this set of simulations, we calculated boundary separation values that correspond to percentages correct of 80%, 90%, and 95% based on drift rate $\nu = 0.2$.³ Practice has shown this to be a plausible value (e.g., Matzke & Wagenmakers, 2009). For this and following simulations, we set DDM noise parameter $s = 0.1$. Boundary separations were calculated using the relation $a = s^2 \times \ln(P_c / (1 - P_c)) / \nu$ (cf. Wagenmakers, van der Maas, & Grasman, 2007, Equation 5). The mean decision time was

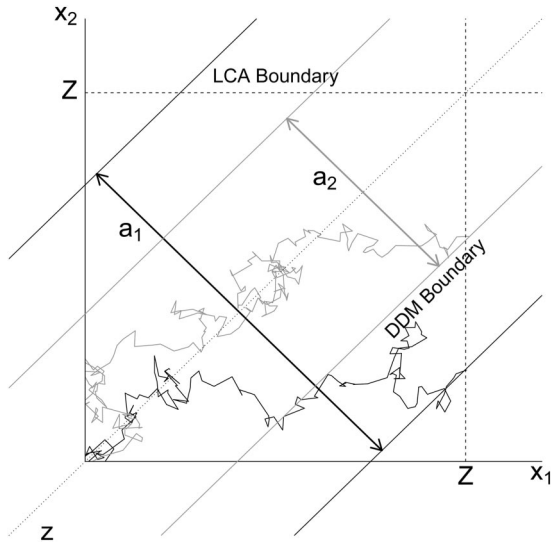


Figure 7. Phase plane of the leaky competing accumulator model (LCA), with the drift diffusion model (DDM) boundary separation a overlaid. The dark line is identical to the one displayed in Figure 3 and has DDM boundary separation a_1 . The light line represents an accumulation path for a different trial and has DDM boundary separation a_2 .

³ Simulation results for $\nu = 0.3$ are shown in the section Higher Input LCA Simulation 1 in the supplemental material.

then calculated using the relation $MDT = a/(2v) \times (1 - e^{-va/s^2}) / (1 + e^{-va/s^2})$ (cf. Wagenmakers et al., 2007, Equation 9).

For the second step of this set of simulations, we used the relation $I_1 = v \times \sqrt{2} + I_2$ (cf. Bogacz et al., 2006, Equation 24) to calculate the corresponding value for the LCA inputs (we set $I_2 = 1$). We then applied $Z = [a/2 + (I_1 + I_2) / (\sqrt{2} \times (k + w))] / \sqrt{2}$ (cf. Bogacz et al., 2006, Equation 26) to calculate corresponding values for the LCA threshold for a range of values of decay parameter k and inhibition parameter w : $k = w = \{1, 2, \dots, 30\}$.

For the third step of this set of simulations, we generated 20,000 RT trials from the LCA with LCA parameters that correspond to the DDM parameters. Finally, we compared the mean decision time and the percentage correct of the DDM to the LCA. Figure 8 shows the simulation results, with the left, middle, and right panels displaying results for DDM percentages correct of 80%, 90%, and 95%, respectively.

The top panels show that for all three levels of DDM accuracy, the truncated LCA consistently has a higher mean decision time than the DDM. The truncated LCA seems to have an optimum that depends on the values of $k = w$. For lower values of $k = w$, the discrepancy between the truncated LCA and the DDM increases, supporting the observation by Bogacz et al. (2006) that $k = w$ needs to be much larger than zero for model equivalence to hold. For higher values of $k = w$, the discrepancy between the models also increases. This occurs because higher decay and inhibition lead to more truncation as the losing accumulator is pushed down more strongly. For higher levels of accuracy, truncation starts to dominate the process for relatively low values of $k = w$, because the response threshold is set higher (compare the right panels to the left panels). For the range of $k = w$ where the truncated LCA

approximates the optimal mean decision time, the corresponding accuracy (shown in the bottom panels) is consistently lower for all three levels of DDM accuracy.

The nontruncated LCA performs somewhat better than the truncated LCA. All panels show that for lower values of $k = w$, performance of the nontruncated LCA approximates the performance of the truncated LCA. This is expected, as no truncation takes place for low values of $k = w$. The difference between the truncated and the nontruncated LCAs emerges for values of $k = w$ that are higher than the values that corresponded to an optimal mean decision time for the truncated LCA. Where performance of the truncated LCA clearly deteriorates for higher values of $k = w$, performance of the nontruncated LCA appears to be largely unaffected. In general, the DDM is now slightly slower than the LCA, but more accurate. The lower percentage correct of the nontruncated LCA is most likely caused by the changing boundary separations (cf. the section LCA Discrepancy 2: Changing Boundaries), as a variable boundary separation leads to more errors. We ran additional simulations in which we simulated LCA and DDM data for the same accuracy level, eliminating the speed-accuracy tradeoff present in Figure 8 (see the section Matched Accuracy LCA Simulation 1 in the supplemental material). These simulations confirmed that the DDM outperforms the nontruncated LCA.

In sum, for a relatively large range of parameter values, the DDM outperforms the truncated LCA. This underscores the fact that the DDM and the truncated LCA are not identical and, as a consequence, that the truncated LCA does not make optimal decisions. The advantage of the DDM over the nontruncated LCA is smaller than the advantage of the DDM over the truncated LCA, but a discrepancy remains.

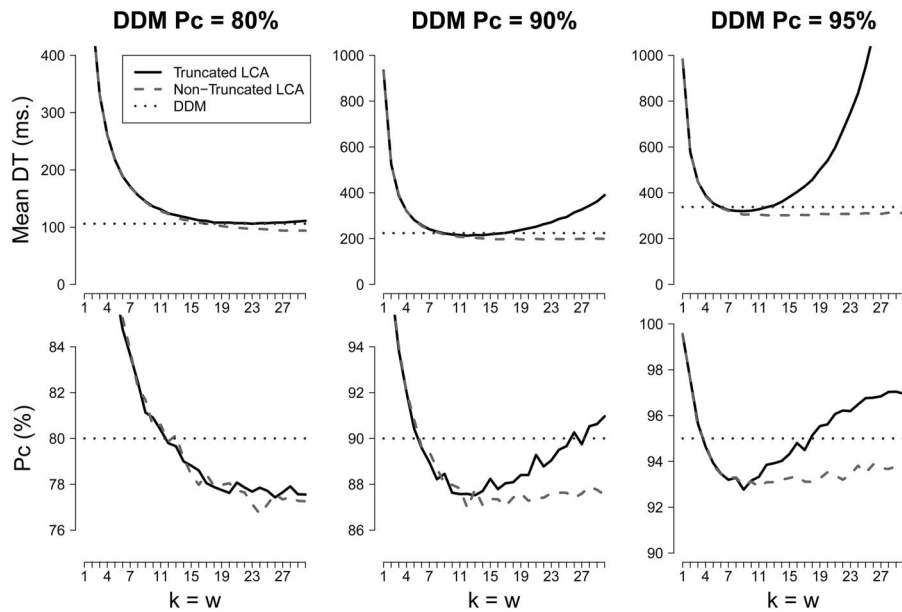


Figure 8. Simulation results for drift diffusion model (DDM) drift rate $v = 0.2$. Top panels: mean decision time (DT). Bottom panels: percentage correct (Pc). DDM boundary separations were set to match mean percentages correct of 80% (left panels), 90% (middle panels), and 95% (right panels). LCA = leaky competing accumulator model.

LCA Simulation 2: Parameters

In the previous section we established that for a relatively large parameter range the LCA performs worse than the DDM. In this section we examine whether the two models draw the same conclusions about latent psychological processes. We do so by fitting the DDM to data simulated with the LCA for input values $I_1 = 1.28$ and $I_2 = 1$, two relatively high values of $k = w$ (i.e., 15 and 30), a nondecision time of 400 ms, and mean percentages correct of 80%, 90%, and 95%. If the models are equivalent, the DDM estimates should be very close to the corresponding LCA parameters that generated the data. We used $v = (I_1 - I_2)/\sqrt{2}$ and $a = 2\sqrt{2}Z - (\sqrt{2}(I_1 + I_2))/(k + w)$ (cf. Bogacz et al., 2006, Equations 24 and 26) to calculate the DDM equivalents for LCA drift rate and LCA boundary separation parameters.

To assess systematic biases in the parameter estimates, we generated 20,000 RTs for both the truncated and the nontruncated LCAs. To quantify the small remaining uncertainty in the parameter estimates, we used a bootstrap procedure and repeatedly drew a random sample of 20,000 RTs with replacement from the full data set. To each subset, we applied the simple EZ algorithm (Wagenmakers et al., 2007) to calculate the DDM drift rate v , DDM boundary separation a , and DDM nondecision time T_{er} .⁴ We repeated this procedure 1,000 times for both the truncated and the nontruncated LCA data. On the basis of the previous section, we expect the DDM estimates and the calculated LCA equivalents to be more similar for the nontruncated LCA than for the truncated LCA.

First, we present simulation results for a mean percentage correct of 80%. Results for $k = w = 15$ are presented in the first column of Figure 9. The figure shows that for both the truncated and the nontruncated LCAs, there is a slight overestimation of drift rate v and nondecision time T_{er} . This result mirrors our earlier finding that for a mean percentage correct of 80% and $k = w = 15$, the truncated and the nontruncated LCAs perform very similarly. It also confirms that both models do not correspond to the DDM.

Results for $k = w = 30$ are presented in the second column of Figure 9. For the truncated LCA, the increase in k and w leads to a negative bias for drift rate v and a positive bias for boundary separation a . The model does this to compensate for the fact that the LCA produces a higher RT than the DDM for the same accuracy level. The nontruncated LCA corresponds more closely to the DDM.

Second, we present simulation results for a mean percentage correct of 90%. Results for $k = w = 15$ are presented in the third column of Figure 9. Results are very similar to those found for the mean percentage correct of 80% presented in the first column, with slight biases in drift rate v and nondecision time T_{er} .

Results for $k = w = 30$ are shown in the fourth column of Figure 9. Compared to the mean percentage correct of 80% presented in the second column, the biases for drift rate v and boundary separation a reappear more strongly. In addition, there is now a small negative bias for nondecision time T_{er} . These findings reflect the fact that the model discrepancy between the LCA and the DDM increases for this accuracy level.

Third, we present simulation results for a mean percentage correct of 95%. Results for $k = w = 15$ are presented in the fifth column of Figure 9. For lower mean percentages correct, serious biases are present only for $k = w = 30$. However, for an accuracy

level of 95%, a negative bias for drift rate v and a positive bias for boundary separation a already appear for $k = w = 15$. Once again, this bias is there to compensate for the fact that the LCA produces a higher RT than the DDM for the same accuracy level.

Results for $k = w = 30$ are shown in the sixth column of Figure 9. Compared to the mean percentage correct of 90% presented in the fourth column, the biases for drift rate v , boundary separation a , and nondecision time T_{er} have increased substantially. Again, these findings reflect the fact that the model discrepancy between the LCA and the DDM increases for this accuracy level.

In sum, nontruncated LCA parameters correspond reasonably well to DDM parameters for different levels of accuracy and different values of decay k and inhibition w . For data generated with the truncated LCA, DDM parameters do not match the LCA parameters used to generate the data for all parameter settings. The simulations presented above show that the nonequivalence is not only a theoretical concern but has practical ramifications; when $k = w \gg 0$, parameters change substantially when truncation is imposed. Moreover, this problem is compounded for higher levels of accuracy.

Interim Conclusion

Is the LCA with $k = w \gg 0$ equivalent to the DDM? The results presented here suggest that for a large range of parameters, both the truncated and the nontruncated LCAs are not equivalent to the DDM. The lack of model equivalence is caused by the fact that the LCA needs to be truncated at zero and by the fact that the derivations for model equivalence imply across-trial variability in boundary separation for the DDM. We now turn to a different neural inhibition model: the feedforward inhibition model (FFI).

The Feedforward Inhibition Model

The FFI (Shadlen & Newsome, 2001) is a multiple accumulator model just as the LCA is (Figure 10, right panel). However, the FFI differs from the LCA in two principal ways. First, the FFI has no leakage, so each accumulator is a perfect evidence integrator. Second, inhibition between accumulators is not affected by their current activation (i.e., their output) but by their drift rates (i.e., their input). Taking once again the two-alternative version as an example, the differential equations for the FFI are given by

$$\begin{aligned} dx_1(t) &= I_1 \cdot dt + s \cdot dW_1(t) - u[I_2 \cdot dt + s \cdot dW_2(t)], \\ dx_2(t) &= I_2 \cdot dt + s \cdot dW_2(t) - u[I_1 \cdot dt + s \cdot dW_1(t)]. \end{aligned} \quad (6)$$

As in the LCA, dx_1 and dx_2 refer to the changes per unit time dt in both accumulators. The drift rate for accumulator 1 is I_1 , whereas the drift rate of the competing accumulator is I_2 . Quantities $s \cdot dW_1$ and $s \cdot dW_2$ are white noise added at each time step with $M = 0$ and variance $= s^2 \cdot dt$, analogous to the DDM and LCA. For the same reason of biological plausibility that applies to the LCA, we

⁴ Results in which the full DDM was fit to the data are shown in section Full DDM Estimates LCA Simulation 2 in the supplemental material.




Figure 9. Drift diffusion model (DDM) estimates for leaky competing accumulator model (LCA) data. Dots represent the mean of the 1,000 bootstrap parameter estimates, with boxes containing 50% and whiskers extending to 90% of these estimates. The left two columns represent data for a mean percentage correct (Pc) of 80%, the middle two columns represent data for a mean percentage correct of 90%, and the right two columns represent data for a mean percentage correct of 95%. The three panel rows represent estimates for DDM parameters v , a , and T_{er} , respectively. For each panel, the left boxes represent the truncated LCA (Trunc.) and the right boxes represent the nontruncated LCA (N-trunc.).

assume that negative values of x_1 and x_2 have to be truncated to zero (e.g., Ditterich, 2006, p. 998). The starting point of the accumulators is the same as for the LCA: $x_1(0) = x_2(0) = 0$, and evidence accumulation stops once an accumulator reaches a deci-

sion threshold Z . The changes in activation for both accumulators, dx_1 and dx_2 , are assumed to be perfectly anticorrelated. Parameter u is a constant that represents the weight of the feedforward inhibition.

The FFI has also been applied to empirical data (e.g., Ditterich, 2006, 2010). For instance, Ditterich (2006) applied the FFI to both behavioral and neurophysiological data from a visual decision-making task.

Using phase planes and a reparameterization equivalent to the one in Equation 4, Bogacz et al. (2006) derived expressions for x_2^* (which is proportional to the sum of the accumulated evidence of both accumulators) and x_1^* (which is proportional to the difference in the accumulated evidence of both accumulators). The stochastic differential equations of x_1^* and x_2^* are given by

$$dx_1^*(t) = (1 + u) \left[\frac{I_1 - I_2}{\sqrt{2}} dt + s \cdot dW_1'(t) \right],$$

$$dx_2^*(t) = (1 - u) \left[\frac{I_1 - I_2}{\sqrt{2}} dt + s \cdot dW_2'(t) \right], \quad (7)$$

where $s \cdot dW_1'$ and $s \cdot dW_2'$ are white noise added at each time step with $M = 0$ and variance = $s^2 \cdot dt$.

Bogacz et al. (2006) reasoned that if $u = 1$, the sum of the accumulated evidence, x_2^* , remains zero. Both accumulators would then be perfectly anticorrelated. Quantity x_1^* and, by extension, the

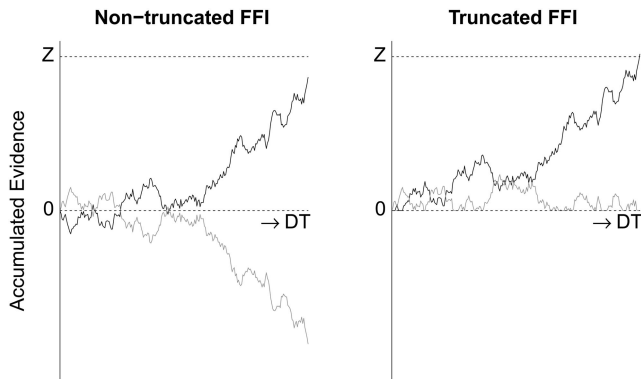


Figure 10. The effect of truncation on the dynamics of the feedforward inhibition model (FFI). The left panel shows accumulation paths that are not truncated at zero. This leads to perfectly anticorrelated accumulation paths. The right panel shows the same accumulators, now truncated at zero. The winning accumulator does not need to overcome any potential negative activation, leading to responses that are faster, but more error prone, than for the nontruncated FFI. DT = decision time.

difference in accumulated evidence, determines the decision process, and it can be seen from the top part of Equation 7 that for $u = 1$ this corresponds to a diffusion drift of $v = \sqrt{2} \cdot (I_1 - I_2)$.

If the DDM and the FFI are equivalent, we would have a model for decision making that is biologically plausible and optimal. However, the FFI used in the equations by Bogacz et al. (2006) also suffers from the discrepancy that the accumulators are allowed to have negative activation values.

FFI Discrepancy: Negative Activations

The discrepancy with the model equivalence between the DDM and the FFI as proposed by Bogacz et al. (2006) is that, just like the LCA, the FFI requires that negative accumulator activation values are truncated to zero for biological plausibility. Consider the accumulators in Figure 10, in which the left panel shows two accumulation paths without truncation and the right panel shows the same accumulation paths with truncation. The left panel shows that without truncation, both accumulators are perfectly anticorrelated, consistent with the dynamics of the DDM. However, negative activations are undesirable, because they imply a negative neural firing rate. To prevent negative activations, the FFI is truncated at zero. The right panel shows the same accumulators with a reflecting boundary at zero.

There is a discrepancy between the truncated and the nontruncated FFIs. Namely, truncation accelerates the decision process if the accumulator that eventually wins has been truncated at zero along the way. An example of this is shown in Figure 10. The winning accumulator in the left panel (corresponding to nontruncated activation) has not yet reached the decision threshold Z , whereas the same accumulator in the right panel (corresponding to truncated activation) has. The reason for this is that in the left panel, the winning accumulator has to overcome its initial negative course, leading to a slower RT. In the right panel, however, negative activation does not occur. In a sense this means that the accumulator has “forgotten” its initial negative evidence.

In addition, truncation also affects the proportion of correct responses. For instance, it is conceivable that the lighter accumulator in the left panel of Figure 10 makes a comeback and eventually wins the race, even though it has already lost in the right panel with truncated activation. Such comebacks are more frequent in the FFI than in the LCA; in the FFI, accumulators inhibit each other only through their inputs, and hence there is no increase in inhibitory pressure on the losing accumulator over time. Thus, the truncated and nontruncated FFIs do not necessarily produce the same response. Only the left panel with nontruncated activation is equivalent to the DDM, whereas the right panel with truncated activation is more biologically plausible.

In sum, the truncated FFI permits decisions that are faster and have a different outcome than the nontruncated FFI. Whereas the truncated FFI is biologically plausible, only the nontruncated FFI is mathematically equivalent to the DDM and as such behaves as an optimal decision maker. The problem for the truncated FFI appears to be less serious than for the truncated LCA, as the FFI accumulators can still reach arbitrary heights, whereas the LCA accumulators approach asymptotic values.

We have shown the truncation problem for model equivalence between the DDM and the FFI in theory. However, it is quite possible that both the FFI and the DDM show indistinguishable performance

in practice. In the next two sections, we investigate how both models fare when pitted against each other.

FFI Simulation 1: Performance

In order to examine whether the FFI matches the optimal performance of the DDM, we

1. analytically calculated mean decision times for sets of DDM parameters that correspond to three different mean percentages correct;
2. used these DDM parameters to analytically calculate the corresponding FFI parameters; and
3. generated FFI data based on the calculated FFI parameters, and then compared the mean decision time and percentage correct for both models.

The next three paragraphs discuss each of these steps in more detail.

For the first step of this set of simulations, we calculated boundary separation values that correspond to percentages correct of 80%, 90%, and 95% based on drift rate $v = \{0.15, 0.20, \dots, 0.40\}$. Practice has shown this to be a plausible range of values (e.g., Matzke & Wagenmakers, 2009). Boundary separations were calculated as before (cf. Wagenmakers et al., 2007, Equations 5 and 9).

For the second step of this set of simulations, we used the relation $I_1 = v/\sqrt{2} + I_2$ (cf. Bogacz et al., 2006, Equation 32, first part) to calculate the corresponding values for the FFI inputs (we set $I_2 = 1$). We then applied $Z = a/(2\sqrt{2})$ (cf. Bogacz et al., 2006, Equation 32, third part) to calculate corresponding values for the FFI thresholds. Finally, we applied $s_{\text{FFI}} = s_{\text{DDM}}/2$ (cf. Bogacz et al., 2006, Equation 32, second part) to calculate a corresponding value for the FFI noise parameter.

For the third step of this set of simulations, we generated 20,000 RT trials from the FFI with FFI parameters that correspond to the DDM parameters. Finally, we compared the mean decision time and the percentage correct of the DDM to the FFI. Figure 11 shows the simulation results, with the left, middle, and right panels displaying results for percentages correct of 80%, 90%, and 95%, respectively.

Figure 11 shows that the DDM is indeed slightly slower than the truncated FFI, as anticipated in the section FFI Discrepancy: Negative Activations. The results also confirm that the slightly lower decision time of the truncated FFI comes at the expense of an increase in error rate of between 5% and 10% for each level of DDM accuracy. We ran additional simulations in which we simulated FFI and DDM data for the same accuracy level, eliminating the speed–accuracy tradeoff present in Figure 11 (see the section Matched Accuracy FFI Simulation 1 in the supplemental material). These additional simulations confirmed that for the same percentage correct, the mean decision time of the DDM is faster than the mean decision time of the FFI, with an advantage of 20 to 100 ms (or 16% to 26%) depending on the value of the input parameter.

Performance of the nontruncated FFI cannot be distinguished from that of the DDM. This is not surprising; without the necessity for truncation, the FFI and the DDM are mathematically identical.

In sum, we see that compared to the DDM, the truncated FFI sacrifices a lot of accuracy for a little speed. Performance of the nontruncated FFI and the DDM is equivalent.

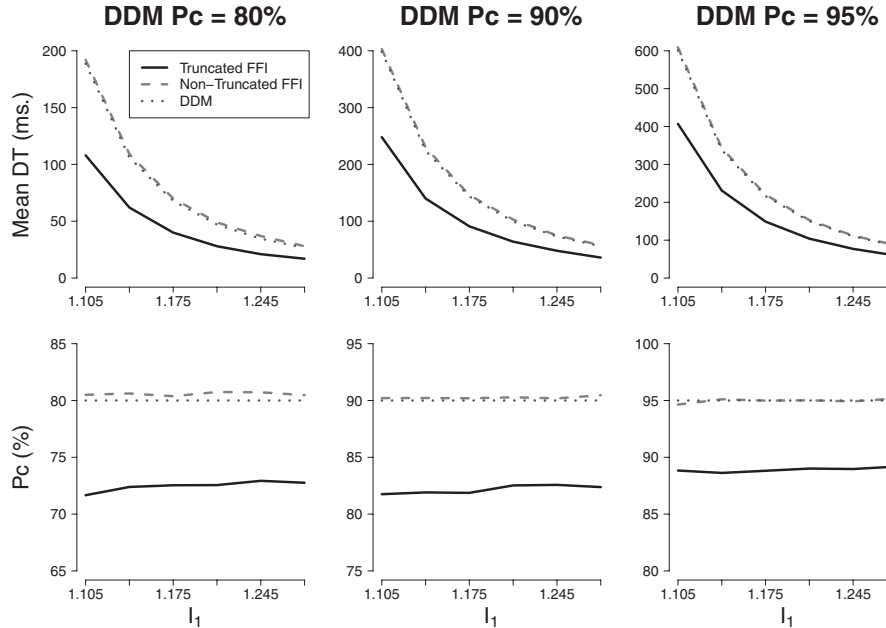


Figure 11. Simulation results for drift diffusion model (DDM) drift rates $v = \{0.15, 0.20, \dots, 0.40\}$. Top panels: mean decision time (DT). Bottom panels: percentage correct (Pc). DDM boundary separations were set to match mean percentages correct of 80% (left panels), 90% (middle panels), and 95% (right panels). FFI = feedforward inhibition model.

FFI Simulation 2: Parameters

In the previous section we established that the truncated FFI performs worse than the DDM, whereas the nontruncated FFI performs identically to the DDM. In this section we examine whether the two models draw the same conclusions about latent psychological processes as the DDM. We do so by fitting the DDM to data simulated with the FFI for input values $I_1 = 1.28$ and $I_2 = 1$. If the models are equivalent, the DDM estimates should be very close to the corresponding FFI parameters that generated the data. We used $v = \sqrt{2} \cdot (I_1 - I_2)$, $a = 2/\sqrt{2} \cdot Z$, and $s_{\text{DDM}} = 2 \cdot s_{\text{FFI}}$ (cf. Bogacz et al., 2006, Equation 32) to calculate the DDM equivalents for FFI drift rate, FFI boundary separation, and FFI noise parameters, respectively.

To assess systematic biases in the parameter estimates, we generated 20,000 RTs for both the truncated and the nontruncated FFIs. To quantify the small remaining uncertainty in the parameter estimates, we used a bootstrap procedure and repeatedly drew a random sample of 20,000 RTs with replacement from the full data set. To each subset, we applied the simple EZ algorithm (Wagenmakers et al., 2007) to calculate the DDM drift rate v , DDM boundary separation a , and DDM nondecision time T_{er} .⁵ We repeated this procedure 1,000 times for both the truncated and the nontruncated FFI data. On the basis of the previous section, we expect the DDM estimates and the calculated FFI equivalents to be more similar for the nontruncated FFI than for the truncated FFI.

The left, middle, and right columns of Figure 12 display results for mean percentages correct of 80%, 90%, and 95%, respectively. The right box in each of the nine panels shows that parameter correspondence between the nontruncated FFI and the DDM is near perfect for all three levels of accuracy. This result was expected, since the models are mathematically equivalent.

The left boxes in each of the three columns show that for the truncated FFI, drift rate v is systematically overestimated, boundary separation a is underestimated, and nondecision time T_{er} is overestimated. An overestimation of nondecision time T_{er} causes higher RTs, an effect that is partially counteracted by a higher drift rate v and a lower boundary separation a . Parameter recovery for the full DDM with across-trial variability parameters may be found in the section Full DDM Estimates FFI Simulation 2 in the supplemental material.

In sum, only for data generated with the nontruncated FFI can the corresponding DDM parameters be recovered with reasonable accuracy.

A Possible Solution to the Problem of Truncation: Baseline Activation

We have shown that model equivalence between the DDM and two neural inhibition models—the LCA and the FFI—are complicated by the potential for negative activation values. Negative activation values may be resolved by truncating them to zero, but this jeopardizes the proposed model equivalence by Bogacz et al. (2006).

There is a way around the necessity of truncation: having accumulators start at a *baseline activation level*. Baseline activity is implemented relatively easily in the FFI. Instead of having each accumulator start at zero and collecting evidence toward response threshold Z , accumulators now start at the old value of response

⁵ Results in which the full DDM was fit to the data are shown in the section Full DDM Estimates FFI Simulation 2 in the supplemental material.




Figure 12. Drift diffusion model (DDM) estimates for feedforward inhibition model (FFI) data. Dots represent the mean of the 1,000 bootstrap parameter estimates, with boxes containing 50% and whiskers extending to 90% of these estimates. The left, middle, and right columns represent data for mean percentages (Pc) correct of 80%, 90%, and 95%, respectively. The three panel rows represent estimates for DDM parameters v , a , and T_{er} , respectively. For each panel, the left boxes represent the truncated FFI (Trunc.) and the right boxes represent the nontruncated FFI (N-trunc.).

threshold Z and collect evidence toward twice the value of Z . In other words, the accumulators have a baseline activation level that allows them to start halfway toward the response threshold. Because the threshold is twice as high, but the distance between the starting point and the threshold is still intact, the winning accumulator still has to collect the same amount of evidence. The distance between the starting point and the zero boundary is now identical to the distance between the starting point and the response threshold, which guarantees that no truncation is ever needed. Such a model is mathematically equivalent to the DDM.

The idea of a nonzero starting point for the FFI is not new. For instance, Ditterich (2010) chose a starting point of $Z/3$, a choice informed by recordings of neural activity in the lateral intraparietal area.

The implementation of baseline activation levels for the LCA involves a little more work than it does for the FFI. For the LCA, simply increasing the starting point and the response threshold, while leaving the other parameters intact, will have an effect only on the starting trajectories of the accumulators. Over time, the

combined force of the decay k and initial mutual inhibition w will lead both accumulators to revolve to the same trajectory they would have followed had they started out at the zero boundary; consequently, truncation may still take place.⁶ In order to prevent this from happening, it is necessary to augment the inputs of each accumulator with a *baseline input*, or I_B (e.g., Bogacz et al., 2007; Bussey, Townsend, Diederich, & Barkan, 2005). Simulations by Bogacz et al. (2007) showed that the addition of high baseline inputs may lead to performance for the truncated LCA that is very similar to that of the nontruncated LCA (Bogacz et al., 2007, Figures 5a and 5b).

Specifically, prior to presentation of the stimulus, both accumulators have the same baseline input I_B , causing their activation values to move to their natural equilibrium of $I_B/(w + k)$ (see Marshall, Bogacz, & Gilchrist, 2011). With a starting point of $I_B/(w + k)$ and a response threshold of $Z + I_B/(w + k)$ (so that the distance between the starting point and the response threshold is still Z), the increased force of the decay and inhibition is exactly canceled out by the added baseline input. LCA accumulation paths with a baseline input I_B and a starting point of $I_B/(w + k)$ are identical to nontruncated LCA accumulation paths without baseline input and a starting point of zero, provided that the losing accumulator never hits the zero boundary. The only difference between the baseline input LCA and the nontruncated LCA is the absolute amount of accumulated evidence. Since the response is determined by the amount of accumulated evidence relative to the starting point, this difference has no effect on model performance.

Some neurophysiological data support the assumption that neurons responsible for the stochastic accumulation of evidence have a baseline firing rate that is at least half of the threshold firing rate. For instance, work by Roitman and Shadlen (2002, Figure 7), Huk and Shadlen (2005, Figure 5), and Churchland, Kiani, and Shadlen (2008, Figures 4a and 5a) indicated that the baseline firing rate of neurons in the lateral intraparietal area is approximately half of the maximum firing rate. However, there is also evidence that suggests that neural baseline firing rates in the frontal eye field are approximately 20% of the maximum firing rate (e.g., Bruce & Goldberg, 1985; Hanes, Patterson, & Schall, 1998; Hanes & Schall, 1996; Pouget et al., 2011; Schall, 1991). Also, research on neural spike rates in the superior colliculus suggests that the baseline firing rates in this brain area are roughly between 10% and 30% of the maximum firing rate (e.g., Basso & Wurtz, 1998; Mays & Sparks, 1980; Munoz & Wurtz, 1995; Ratcliff, Chierian, & Segraves, 2003). Finally, baseline firing rates in the medial superior temporal area seem to be much lower than half of the maximum firing rate (e.g., Duffy & Wurtz, 1991).⁷

In conclusion, some neurophysiological evidence supports the notion that neural baseline firing rates are at least half the amount of neural threshold firing rates, although there is also much evidence to the contrary. At least from a modeler's perspective, both the LCA and the FFI can be implemented in such a way that truncation is not necessary without quantitatively altering their performance.

⁶ For a simulation that demonstrates this effect, see the section LCA Truncation With Starting Points Above Zero in the supplemental material.

⁷ We thank two anonymous reviewers for alerting us to some of the studies mentioned in this paragraph.

Conclusion

We have highlighted a tension between two desiderata for neural network models: optimality and biological plausibility. We examined an attempt by Bogacz et al. (2006) to resolve this tension between the DDM and two neural network models: the LCA and the FFI. We have shown that for both the LCA and the FFI there are challenges that have to be addressed before each of these models can be reduced to the DDM.

For the LCA, our simulations showed that the DDM outperforms the LCA with truncation at zero. With the assumption of baseline activity of neurons that is at least half the amount of the threshold firing rate of neurons, the necessity of truncation disappears for the LCA. However, even the nontruncated LCA cannot be completely reduced to the DDM by means of the equations presented by Bogacz et al. (2006). Specifically, these derivations allow for DDM boundary separations that vary over trials, a condition that the DDM does not allow. Our simulations showed that the DDM achieves higher accuracy than the nontruncated LCA.

Although we showed performance deficits for the truncated LCA relative to both the nontruncated LCA and the DDM, both Bogacz et al. (2007, Figures 5a and 5b) and Tsetsos, Usher, and McClelland (2011) have demonstrated that for choice tasks with more than two alternatives, the truncated LCA can outperform the nontruncated LCA. Tsetsos et al. argued that in order to account for data from multialternative experiments, truncation is in fact required. It is important to remember that, in contrast to the LCA, the DDM is meant for choice tasks with only two response alternatives. For tasks with more alternatives, truncated LCA may lead to performance that is superior to that of nontruncated LCA.

For the FFI, our simulations showed that the truncated model is slightly faster than the DDM, but this increase in speed comes at the expense of a substantially higher error rate. The necessity to truncate at the zero boundary for the FFI can be overcome by assuming baseline activity of at least half the threshold firing rate. The resulting nontruncated version of the FFI is identical to the DDM.

In the introduction, we highlighted one of the main selling points of the DDM: It acts as an optimal decision maker. The DDM is optimal in the sense that it achieves the lowest possible mean RT for any given accuracy level. However, many other definitions of optimality are possible. In the context of human categorization, Anderson (1991) discussed three crucial questions that relate to the definition of optimality:

1. What is the decision maker trying to optimize?
2. What is the structure of the environment (e.g., to what extent and in what way is optimal behavior learned by the decision maker)?
3. What are the costs associated with achieving optimal behavior?

For a participant in a two-choice RT experiment, the first question may pertain to time spent on the task (e.g., Hawkins, Brown, Steyvers, & Wagenmakers, 2011), monetary reward (e.g., Simen, Cohen, & Holmes, 2006), or minimizing the number of avoidable errors (e.g., Starns & Ratcliff, 2010). The second question deals with the transparency of the task. For instance, when a decision maker wants to maximize reward rate, the speed–accuracy tradeoff needs to be adjusted depending on the difficulty

of the stimuli, the nature and timing of feedback messages, and so forth. The change in a decision maker's belief about the environment can be modeled with Bayesian updating functions (see Berger, 1985, for more on Bayesian decision theory). For our RT example, the third question may imply that pursuing an optimal speed–accuracy trade means forfeiting spending an hour outside in the sun (something that could have been accomplished by answering very quickly at the cost of some errors). Shenoy and Yu (2011) used Bayesian utility functions to define optimal decision-making behavior in the stop-signal task. The extensive literature on reward rate maximization and utility functions attests to the fact that optimality is a construct that depends on context and goals that are set by the decision maker.

The simple DDM ensures the lowest possible mean RT for any given accuracy level. However, in practice an extended version of the DDM is often used that also includes parameters that specify across-trial variability in drift rate, starting point, and nondesired time (Ratcliff & Tuerlinckx, 2002). The full DDM is capable of dealing with a number of empirical phenomena the simple DDM cannot handle (i.e., fast or slow errors; see Ratcliff & Rouder, 1998). However, the full DDM is not identical to the simple DDM and therefore does not generate optimal decision-making behavior.

For both the LCA and the FFI, truncation may be avoided by assuming a baseline firing rate of neurons that is at least half the amount of their threshold firing rates. With these assumptions, performance of the neural inhibition models approximates that of the optimal DDM (for the FFI, performance is even identical to that of the DDM). However, the neurophysiological evidence suggests that in many brain areas, neural baseline firing rates are relatively low. Thus, the quest for a neurologically based, multi-alternative, optimally behaving decision model continues. With new and exciting biological models based on Bayesian principles (Ma, Beck, Latham, & Pouget, 2006; Ma, Beck, & Pouget, 2008; Soltani & Wang, 2010) and studies that examine decision models that behave optimally for multiple alternatives (Bogacz & Gurney, 2007; McMillen & Holmes, 2006), there is hope for a decision model that has the best of all worlds.

References

- Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological Review*, *98*, 409–429. doi:10.1037/0033-295X.98.3.409
- Basso, M. A., & Wurtz, R. H. (1998). Modulation of neuronal activity in superior colliculus by changes in target probability. *Journal of Neuroscience*, *18*, 7519–7534.
- Berger, J. O. (1985). *Statistical decision theory and Bayesian analysis*. New York, NY: Springer.
- Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (2006). The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review*, *113*, 700–765. doi:10.1037/0033-295X.113.4.700
- Bogacz, R., & Gurney, K. (2007). The basal ganglia and cortex implement optimal decision making between alternative actions. *Neural Computation*, *19*, 442–477. doi:10.1162/neco.2007.19.2.442
- Bogacz, R., Usher, M., Zhang, J., & McClelland, J. L. (2007). Extending a biologically inspired model of choice: Multi-alternatives, nonlinearity and value-based multidimensional choice. *Philosophical Transactions of the Royal Society, Series B: Biological Sciences*, *362*, 1655–1670. doi: 10.1098/rstb.2007.2059
- Boucher, L., Palmeri, T. J., Logan, G. D., & Schall, J. D. (2007). Inhibitory control in mind and brain: An interactive race model of countermanding

- saccades. *Psychological Review*, *114*, 376–397. doi:10.1037/0033-295X.114.2.376
- Brown, S. D., & Heathcote, A. (2008). The simplest complete model of choice reaction time: Linear ballistic accumulation. *Cognitive Psychology*, *57*, 153–178. doi:10.1016/j.cogpsych.2007.12.002
- Bruce, C. J., & Goldberg, M. E. (1985). Primate frontal eye fields: I. Single neurons discharging before saccades. *Journal of Neurophysiology*, *53*, 603–635.
- Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: A dynamic-cognitive approach to decision making. *Psychological Review*, *100*, 432–459. doi:10.1037/0033-295X.100.3.432
- Busemeyer, J. R., Townsend, J. T., Diederich, A., & Barkan, R. (2005). Contrast effects or loss aversion? Comment on Usher and McClelland (2004). *Psychological Review*, *112*, 253–255. doi:10.1037/0033-295X.112.1.253
- Cavanagh, J. F., Wiecki, T. V., Cohen, M. X., Figueroa, C. M., Samanta, J., Sherman, S. J., & Frank, M. J. (2011). Subthalamic nucleus stimulation reverses mediofrontal influence over decision threshold. *Nature Neuroscience*. Advance online publication. doi:10.1038/nn.2925
- Churchland, A. K., Kiani, R., & Shadlen, M. N. (2008). Decision-making with multiple alternatives. *Nature Neuroscience*, *11*, 693–702. doi:10.1038/nn.2123
- Ditterich, J. (2006). Stochastic models of decisions about motion direction: Behavior and physiology. *Neural Networks*, *19*, 981–1012. doi:10.1016/j.neunet.2006.05.042
- Ditterich, J. (2010). A comparison between mechanisms of multi-alternative perceptual decision making: Ability to explain human behavior, predictions for neurophysiology, and relationship with decision theory. *Frontiers in Decision Neuroscience*, *4*, 184. doi:10.3389/fnins.2010.00184
- Duffy, C. J., & Wurtz, R. H. (1991). Sensitivity of MST neurons to optic flow stimuli: I. A continuum of response selectivity to large-field stimuli. *Journal of Neurophysiology*, *65*, 1329–1345.
- Grossberg, S. (Ed.). (1988). *Neural networks and natural intelligence*. Cambridge, MA: MIT Press.
- Hanes, D. P., Patterson, W. F., & Schall, J. D. (1998). Role of frontal eye fields in countermanding saccades: Visual, movement, and fixation activity. *Journal of Neurophysiology*, *79*, 817–834.
- Hanes, D. P., & Schall, J. D. (1996, October 18). Neural control of voluntary movement initiation. *Science*, *274*, 427–430. doi:10.1126/science.274.5286.427
- Hawkins, G., Brown, S. D., Steyvers, M., & Wagenmakers, E.-J. (2011). *An optimal adjustment procedure to minimize experiment time in decisions with multiple alternatives*. Manuscript submitted for publication.
- Huk, A. C., & Shadlen, M. N. (2005). Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. *Journal of Neuroscience*, *25*, 10420–10436. doi:10.1523/JNEUROSCI.4684-04.2005
- Kamienkowski, J. E., Pashler, H., Dehaene, S., & Sigman, M. (2011). Effects of practice on task architecture: Combined evidence from interference experiments and random-walk models of decision making. *Cognition*, *119*, 81–95. doi:10.1016/j.cognition.2010.12.010
- Klauer, K. C., Voss, A., Schmitz, F., & Teige-Mocigemba, S. (2007). Process components of the Implicit Association Test: A diffusion-model analysis. *Journal of Personality and Social Psychology*, *93*, 353–368. doi:10.1037/0022-3514.93.3.353
- Laming, D. R. J. (1968). *Information theory of choice-reaction times*. London, England: Academic Press.
- Luce, R. D. (1986). *Response times*. New York, NY: Oxford University Press.
- Ma, W. J., Beck, J. M., Latham, P. E., & Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nature Neuroscience*, *9*, 1432–1438. doi:10.1038/nn1790
- Ma, W. J., Beck, J. M., & Pouget, A. (2008). Spiking networks for Bayesian inference and choice. *Current Opinion in Neurobiology*, *18*, 217–222. doi:10.1016/j.conb.2008.07.004
- Marshall, J. A. R., Bogacz, R., & Gilchrist, I. D. (2011). *Consistent implementation of decision in the brain*. Manuscript submitted for publication.
- Matzke, D., & Wagenmakers, E.-J. (2009). Psychological interpretation of ex-Gaussian and shifted Wald parameters: A diffusion model analysis. *Psychonomic Bulletin & Review*, *16*, 798–817. doi:10.3758/PBR.16.5.798
- Mays, L. E., & Sparks, D. L. (1980). Dissociation of visual and saccade-related responses in superior colliculus neurons. *Journal of Neurophysiology*, *43*, 207–232.
- McMillen, T., & Holmes, P. (2006). The dynamics of choice among multiple alternatives. *Journal of Mathematical Psychology*, *50*, 30–57. doi:10.1016/j.jmp.2005.10.003
- Munoz, D. P., & Wurtz, R. H. (1995). Saccade-related activity in monkey superior colliculus: I. Characteristics of burst and buildup cells. *Journal of Neurophysiology*, *73*, 2313–2333.
- Philastides, M. G., Ratcliff, R., & Sajda, P. (2006). Neural representation of task difficulty and decision-making during perceptual categorization: A timing diagram. *Journal of Neuroscience*, *26*, 8965–8975. doi:10.1523/JNEUROSCI.1655-06.2006
- Pouget, P., Logan, G. D., Palmeri, T. J., Boucher, L., Par, M., & Schall, J. D. (2011). Neural basis of adaptive response time adjustment during saccade countermanding. *Journal of Neuroscience*, *31*, 12604–12612. doi:10.1523/JNEUROSCI.1868-11.2011
- Purcell, B. A., Heitz, R. P., Cohen, J. Y., Schall, J. D., Logan, G. D., & Palmeri, T. J. (2010). Neurally constrained modeling of perceptual decision making. *Psychological Review*, *117*, 1113–1143. doi:10.1037/a0020311
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*, 59–108. doi:10.1037/0033-295X.85.2.59
- Ratcliff, R., Cheria, A., & Segraves, M. (2003). A comparison of macaque behavior and superior colliculus neuronal activity to predictions from models of simple two-choice decisions. *Journal of Neurophysiology*, *90*, 1392–1407. doi:10.1152/jn.01049.2002
- Ratcliff, R., Gomez, P., & McKoon, G. (2004). Diffusion model account of lexical decision. *Psychological Review*, *111*, 159–182. doi:10.1037/0033-295X.111.1.159
- Ratcliff, R., Hasegawa, Y. T., Hasegawa, Y. P., Smith, P. L., & Segraves, M. A. (2007). Dual diffusion model for single-cell recording data from the superior colliculus in a brightness-discrimination task. *Journal of Neurophysiology*, *97*, 1756–1774. doi:10.1152/jn.00393.2006
- Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation*, *20*, 873–922. doi:10.1162/neco.2008.12-06-420
- Ratcliff, R., & Rouder, J. N. (1998). Modeling response times for two-choice decisions. *Psychological Science*, *9*, 347–356. doi:10.1111/1467-9280.00067
- Ratcliff, R., & Rouder, J. N. (2000). A diffusion model account of masking in two-choice letter identification. *Journal of Experimental Psychology: Human Perception and Performance*, *26*, 127–140. doi:10.1037/0096-1523.26.1.127
- Ratcliff, R., & Starns, J. J. (2009). Modeling confidence and response time in recognition memory. *Psychological Review*, *116*, 59–83. doi:10.1037/a0014086
- Ratcliff, R., Thapar, A., Gomez, P., & McKoon, G. (2004). A diffusion model analysis of the effects of aging in the lexical-decision task. *Psychology and Aging*, *19*, 278–289. doi:10.1037/0882-7974.19.2.278
- Ratcliff, R., Thapar, A., & McKoon, G. (2010). Individual differences, aging, and IQ in two-choice tasks. *Cognitive Psychology*, *60*, 127–157. doi:10.1016/j.cogpsych.2009.09.001
- Ratcliff, R., & Tuerlinckx, F. (2002). Estimating parameters of the diffusion model: Approaches to dealing with contaminant reaction times and

- parameter variability. *Psychonomic Bulletin & Review*, 9, 438–481. doi:10.3758/BF03196302
- Ratcliff, R., & van Dongen, H. P. A. (2009). Sleep deprivation affects multiple distinct cognitive processes. *Psychonomic Bulletin & Review*, 16, 742–751. doi:10.3758/PBR.16.4.742
- Roe, R. M., Busemeyer, J. R., & Townsend, J. T. (2001). Multi-alternative decision field theory: A dynamic artificial neural network model of decision-making. *Psychological Review*, 108, 370–392. doi:10.1037/0033-295X.108.2.370
- Roitman, J. D., & Shadlen, M. N. (2002). Responses of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *Journal of Neuroscience*, 22, 9475–9489.
- Schall, J. D. (1991). Neuronal activity related to visually guided saccades in the frontal eye fields of rhesus monkeys: Comparison with supplementary eye fields. *Journal of Neurophysiology*, 66, 559–579.
- Schmiedek, F., Lövdén, M., & Lindenberger, U. (2009). On the relation of mean reaction time and intraindividual reaction time variability. *Psychology and Aging*, 24, 841–857. doi:10.1037/a0017799
- Schmiedek, F., Oberauer, K., Wilhelm, O., Süß, H.-M., & Wittmann, W. W. (2007). Individual differences in components of reaction time distributions and their relations to working memory and intelligence. *Journal of Experimental Psychology: General*, 136, 414–429. doi:10.1037/0096-3445.136.3.414
- Shadlen, M. N., & Newsome, W. T. (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *Journal of Neurophysiology*, 86, 1916–1936.
- Shenoy, P., & Yu, A. J. (2011). Rational decision-making in inhibitory control. *Frontiers in Human Neuroscience*, 5, 48. doi:10.3389/fnhum.2011.00048
- Simen, P., Cohen, J. D., & Holmes, P. (2006). Rapid decision threshold modulation by reward rate in a neural network. *Neural Networks*, 19, 1013–1026. doi:10.1016/j.neunet.2006.05.038
- Soltani, A., & Wang, X.-J. (2010). Synaptic computation underlying probabilistic inference. *Nature Neuroscience*, 13, 112–119. doi:10.1038/nn.2450
- Starns, J. J., & Ratcliff, R. (2010). The effects of aging on the speed–accuracy compromise: Boundary optimality in the diffusion model. *Psychology and Aging*, 25, 377–390. doi:10.1037/a0018022
- Townsend, J. T., & Ashby, F. G. (1983). *Stochastic modeling of elementary psychological processes*. London, England: Cambridge University Press.
- Tsetsos, T., Usher, M., & McClelland, J. L. (2011). Testing multi-alternative decision models with non-stationary evidence. *Frontiers in Decision Neuroscience*, 5, 63. doi:10.3389/fnins.2011.00063
- Usher, M., & McClelland, J. L. (2001). On the time course of perceptual choice: The leaky competing accumulator model. *Psychological Review*, 108, 550–592. doi:10.1037/0033-295X.108.3.550
- Usher, M., & McClelland, J. L. (2004). Loss aversion and inhibition in dynamical models of multialternative choice. *Psychological Review*, 111, 757–769. doi:10.1037/0033-295X.111.3.757
- van Ravenzwaaij, D., Dutilh, G., & Wagenmakers, E.-J. (2011). A diffusion model decomposition of the effects of alcohol on perceptual decision making. *Psychopharmacology*. Advance online publication. doi:10.1007/s00213-011-2435-9
- van Ravenzwaaij, D., & Oberauer, K. (2009). How to use the diffusion model: Parameter recovery of three methods: EZ, fast-dm, and DMAT. *Journal of Mathematical Psychology*, 53, 463–473. doi:10.1016/j.jmp.2009.09.004
- van Ravenzwaaij, D., van der Maas, H. L. J., & Wagenmakers, E.-J. (2011). Does the Name–Race Implicit Association Test measure racial prejudice? *Experimental Psychology*, 58, 271–277. doi:10.1027/1618-3169/a000093
- Wagenmakers, E.-J. (2009). Methodological and empirical developments for the Ratcliff diffusion model of response times and accuracy. *European Journal of Cognitive Psychology*, 21, 641–671. doi:10.1080/09541440802205067
- Wagenmakers, E.-J., Ratcliff, R., Gomez, P., & McKoon, G. (2008). A diffusion model account of criterion shifts in the lexical decision task. *Journal of Memory and Language*, 58, 140–159. doi:10.1016/j.jml.2007.04.006
- Wagenmakers, E.-J., van der Maas, H. J. L., & Grasman, R. P. P. P. (2007). An EZ-diffusion model for response time and accuracy. *Psychonomic Bulletin & Review*, 14, 3–22. doi:10.3758/BF03194023
- Wald, A., & Wolfowitz, J. (1948). Optimal character of the sequential probability ratio test. *Annals of Mathematical Statistics*, 19, 326–339. doi:10.1214/aoms/1177730197
- Wang, X.-J. (2002). Probabilistic decision making by slow reverberation in cortical circuits. *Neuron*, 36, 955–968. doi:10.1016/S0896-6273(02)01092-9
- White, C. N., Ratcliff, R., Vasey, M. W., & McKoon, G. (2010). Using diffusion models to understand clinical disorders. *Journal of Mathematical Psychology*, 54, 39–52. doi:10.1016/j.jmp.2010.01.004
- Zhang, J., & Bogacz, R. (2010). Bounded Ornstein–Uhlenbeck models for two-choice time controlled tasks. *Journal of Mathematical Psychology*, 54, 322–333. doi:10.1016/j.jmp.2010.03.001

Received April 18, 2011

Revision received October 12, 2011

Accepted October 12, 2011 ■