Accumulating Advantages: A New Approach to Multialternative Forced Choice Tasks

Don van Ravenzwaaij[1,2], Scott D. Brown[1], Anthony A. J. Marley[3,4], and Andrew Heathcote[1,5]

[1]School of Psychology, University of Newcastle [2]Faculty of Behavioral and Social Sciences, University of Groningen [3]Department of Psychology, University of Victoria [4]Institute for Choice, University of South Australia [5]School of Psychology, University of Tasmania

Correspondence concerning this article should be addressed to:
Don van Ravenzwaaij
University of Groningen, Department of Psychology
Grote Kruisstraat 2/1, Heymans Building, room 169
9712 TS Groningen, The Netherlands
Ph: (+31) 50 363 7021
E–mail should be sent to d.van.ravenzwaaij@rug.nl.

Abstract

Accumulator models have a relatively long history of application to choice and response time, especially for binary choice. When the number of response options is greater than two, these models usually posit one accumulator per response option. Here, we propose a theoretical framework in which there is one accumulator for every ordered pairwise difference between choices, each quantifying the evidence for one response over another ("advantages"). We instantiate and test this framework in the computationally tractable *advantage linear ballistic accumulator model* (ALBA), for which we derive explicit likelihoods that support efficient Bayesian estimation. In the first part of the paper, we present three model architectures that differ in terms of stopping rules: conditions on accumulator completions that have to be met for a response to be chosen. In the second part of the paper we present fits of each model architecture to a multiple-choice data set, and select the model architecture that provides the best account of Hick's Law. In the third part of the paper, we discuss a recent claim by Teodorescu and Usher (2013) that accumulator models require response competition in order to account for the effect of strong distractors in their multialternative forced choice data. Counter to that assertion, we show that this effect naturally emerges from the independent accumulation of advantages, and that the data of Teodorescu and Usher (2013) are at least as well fit by the ALBA as by any of the response-competition models they considered.

Keywords: Evidence accumulation models, RT tasks, Hick's Law, mutual inhibition.

In everyday life, we are constantly confronted with decision problems that require a choice of one option from many. An example is choosing a pasta sauce in the supermarket out of a dozen potential options. A decision like this is complex in nature, because you are typically weighing your options on a number of dimensions: price, flavor, perhaps number of calories, and so on. Magnitude estimation is a simpler, but related, kind of decision problem. Say, you are in your local bakery and you have just paid for a croissant. The server behind the counter lets you choose your own croissant, but a cursory glance shows you that the five croissants that you can choose from are not all equal in size. Whether it is because you are particularly hungry or want your money's worth, you are bent on selecting the largest of the available options; however, you also want to make your decision in a timely manner. Such speed–accuracy trade-offs have been studied extensively in psychology (e.g., Wickelgren, 1977) and are often modeled by racing evidence accumulators.

Accumulator models assume that once the relevant stimulus information is perceptually encoded and/or extracted from memory, evidence for each response option is accumulated towards a response criterion. The first accumulator to reach its criterion leads to a response in favor of the option the accumulator represents. Notable examples include the diffusion decision model (DDM; Ratcliff, 1978), the leaky competing accumulator model (LCA; Usher & McClelland, 2001), the Ising Decision Maker (IDM; Verdonck & Tuerlinckx, 2014), the ballistic accumulator model (BA; Brown & Heathcote, 2005), and the linear ballistic accumulator model (LBA; Brown & Heathcote, 2008). For decision problems with more than two response options, these models assume that evidence accumulation for response alternatives is accrued in an absolute manner.[1] Evidence accumulators may or may not influence one another (the LCA, for instance, includes a mutual inhibition mechanism), but the evidence accrued by an accumulator typically counts as absolute evidence for just one response option.

However, returning to the croissants example, it could be argued that an approach based on relative judgments is more realistic than one based on absolute judgments. For example, one might select two of the croissants at random, decide which has the advantage (i.e., which is larger), and repeat the process with the other three croissants with the winner at each stage being carried forward to the next comparison[2]. This example suggests a serial process, but another option, which we explore here, is that the comparisons run in parallel. Whether the process is considered to be serial or parallel, the general framework is that evidence for each option is always *relative* to the evidence for each of the other available options (see e.g., Marley, 1991, and Trueblood, Brown, & Heathcote, 2014, for relative evidence models of choice probabilities, with the Trueblood et al., 2014 results derived from a model of choice and response time.

We propose and test a specific model within this parallel relative judgment framework: the Advantage LBA (ALBA). The mathematical properties of the accumulators in the ALBA are identical to those in the LBA, but the decision criterion or *stopping rule* is very different. As a result – and in contrast to response-competition models like the LCA, IDM, and BA – the ALBA retains the tractability of the LBA and has a closed-form likelihood expression. A new contribution of this framework is the idea that different response options may be represented by *more than one accumulator* – there may be accumulators correspond-

[1]Note that the diffusion model is typically applied to two response options, though it can be reconceptualized and extended to more than two response options (e.g., van Ravenzwaaij, van der Maas, & Wagenmakers, 2012).

[2]Blavatskyy (2012) develops a variant of this process, which can lead to various context effects

ing to evidence in favor of an option and others corresponding to evidence against that same option, for example. This approach releases accumulator models of multialternative choice from the traditional one-to-one constraint between accumulators and responses.

The next section summarizes the properties of the LBA, and the following one introduces the core framework of the ALBA, as well as three specific versions of the stopping rule: Win–All, Lose–All, and Lose–One. We show that all of the parameters of the ALBA are identifiable with every stopping rule by performing a parameter-recovery study with realistic sample sizes. This underlines a significant advantage of our approach compared to the LCA, where recovery of the full set of model parameters is not possible in practice (Miletic, Turner, Forstmann, & Van Maanen, 2016).

We then investigate to what extent each model can account for data that are consistent with Hick's Law (Hick, 1952; Hyman, 1953). Hick's Law is a benchmark result for multialternative response time, or RT, paradigms; it states that the mean RT and the logarithm of the number of choice alternatives are linearly related. We fit all three model architectures to an archival Hick's Law data set (van Maanen et al., 2012), and compare the results via formal model comparison against full RT distributions. This allows us to identify the stopping rule that is able to most parsimoniously and accurately account for the data.

In the final part of the paper we examine an important claim by Teodorescu and Usher (2013) that accumulator models require response competition in order to account for the effect of a strong distractor (i.e., a stimulus that closely matches the stimulus corresponding to the correct choice) in multialternative forced choice. We demonstrate by example that this claim fails when one considers the accumulation of advantages. In this framework, their finding of slowing in correct responses caused by a strong distractor emerges naturally even though each accumulator races independently. We also show that the ALBA model, in particular, can provide an accurate and detailed account of Teodorescu and Usher's data.

The LBA

The LBA decision model (Brown & Heathcote, 2008) conceptualizes choice processing as the accumulation of information over time, with a response initiated when the accumulated evidence reaches a predefined threshold. An illustration for a brightness identification task with two response options is given in Figure 1.

The LBA assumes that each evidence accumulator starts from an independently sampled and uniformly distributed point between 0 and a fixed $A > 0$, after which evidence is accumulated linearly for each response option. Each of the N evidence accumulators has a drift rate d_i, $i = 1, \ldots, N$. For each trial, each drift rate is independently drawn from a Normal distribution – which we assume here is truncated at zero – with means ν_i, $i = 1, \ldots, N$, and a common standard deviation s, which we fix to 1 for all further applications of the model in this paper.[3] A common threshold b determines a speed–accuracy tradeoff; smaller b leads to faster decisions at the cost of a higher error rate.

Together, these parameters generate a distribution of decision times, DT. The observed RT, however, also includes components such as stimulus encoding and response

[3] Other drift rate distributions also yield tractable models (see e.g. Terry et al., 2015), but the most commonly used version of the LBA assumes a truncated normal distribution.

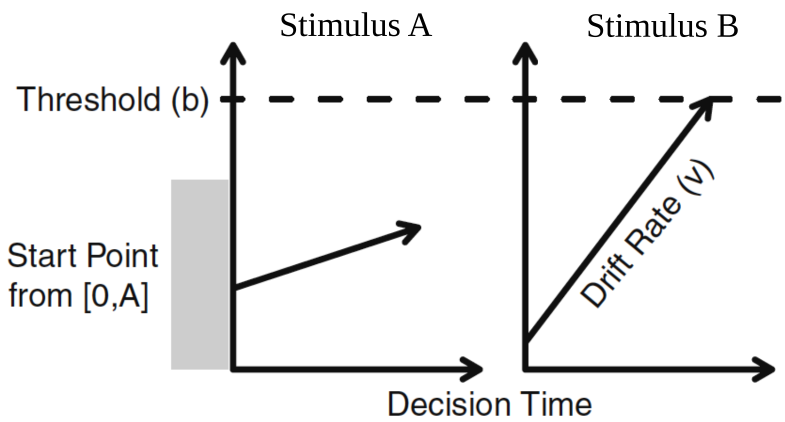

Figure 1. The LBA and its parameters for a brightness identification task with two response options ("stimulus A" and "stimulus B"). Evidence accumulation begins at a start point drawn randomly from a uniform distribution on the interval $[0, A]$. Evidence accumulation is governed by drift rate d, drawn across trials from a normal distribution with mean ν and standard deviation s, truncated to positive values. A response is given as soon as one accumulator reaches threshold b. Observed RT is an additive combination of the time during which evidence is accumulated and non–decision time t_0.

production, which together make up the non-decision time (Luce, 1986). We assume non-decision time is a constant, t_0, that shifts the distribution of DT such that $RT = DT + t_0$. Hence, the three key components of the LBA are (1) the rate of information processing, quantified by mean drift rate ν; (2) response caution, quantified by distance from start-point to threshold, with mean $b - A/2$; and (3) non–decision time t_0.

As shown in Terry et al. (2015), the cumulative distribution function (CDF) for the random variable associated with the decision time ($t \geq 0$) of a single accumulator is given by:

$$CDF(t) = 1 + \left(\frac{tZ(t) - b}{A}\right) G\left(\frac{b}{t}\right) + \left(\frac{b - A - tZ(t)}{A}\right) G\left(\frac{b - A}{t}\right) \quad (1)$$

with

$$Z(t) = \frac{1}{G(\frac{b}{t}) - G(\frac{b-A}{t})} \int_{\frac{b-A}{t}}^{\frac{b}{t}} ug(u)du \quad (2)$$

Here, G and g represent the cumulative distribution function and probability density function for the distribution of drift rates, respectively. The PDF for finishing times of a single accumulator is obtained by differentiation of (1) with respect to t. Assuming Z(t) is differentiable for all $t > 0$, and denoting its derivative by Z'(t), we get

$$PDF(t) = \left(\frac{Z(t) + tZ'(t)}{A}\right)\left[G\left(\frac{b}{t}\right) - G\left(\frac{b-A}{t}\right)\right] + \left(\frac{tZ(t) - b}{A}\right) g\left(\frac{b}{t}\right)\left(\frac{b - A - tZ(t)}{A}\right) g\left(\frac{b-A}{t}\right) \quad (3)$$

Note that 1 and 3 only involve the expressions for the PDF and CDF of the drift rate distribution (g and G respectively). For the remainder of this paper, we will choose the truncated normal distribution for drift rates. This results in expressions for the PDF and CDF for a single accumulator that are analogous to those presented as Equations (1) and (2) by Brown and Heathcote (2008). For the LBA, one is typically not interested in the termination time of a single accumulator. However, equations 1 and 3 are the building blocks for the ALBA, to which we now turn.

The Advantage LBA (ALBA)

The ALBA assumes the same underlying type of processing as the LBA: a set of accumulators each moving towards a threshold. The crucial differences between the ALBA and the LBA are that (1) the evidence to be accumulated comes from *differences* between evidence for ordered pairs of stimuli (i.e, from the *advantage* for one stimulus over another), and that (2) decisions are made when a combination of accumulators have all crossed threshold, as opposed to a single accumulator crossing a threshold (see also Eidels, Donkin, Brown, & Heathcote, 2010). These combination stopping rules can be thought of as being realized by counters, with one counter for each possible response, although there may be other interpretations (e.g., logic gates). Counts are incremented by threshold-crossing events in a set of accumulators connected to the counter. The response associated with the counter is initiated as soon as a criterion number of counts is achieved.

As an example, consider a task in which a participant has to decide which of four presented stimuli is the brightest: A, B, C, or D. For this decision, a standard LBA with one-to-one mapping between accumulators and choices has four accumulators, which we denote A, B, C, and D with corresponding drift rates d(A), d(B), d(C), and d(D). The ALBA for the same decision has a total of 12 advantage accumulators, each taking as input a difference between the evidence values for a different ordered pair of stimuli. We denote these accumulators: $A - B$, $A - C$, $A - D$, $B - A$, $B - C$, $B - D$, $C - A$, $C - B$, $C - D$, $D - A$, $D - B$, and $D - C$.

Consider a stimulus that in the traditional LBA provides a "matching" input value M to its corresponding accumulator A and a smaller "mismatching" input value of m to all other accumulators. In the corresponding case for the ALBA, each "matching" advantage accumulator (e.g., $A - B$, $A - C$, and $A - D$) would have an advantage input value of $M - m = I$; each "mismatching" advantage accumulator (e.g., $B - A$, $C - A$, and $D - A$) would have an advantage input value of $m - M = -I$; and each of the remaining six "unrelated" accumulators would have an advantage input value of $m - m = 0$. We describe these differences as stimulus contrasts.

The next step of the ALBA is to multiply all stimulus contrasts by a discrimination parameter ν. The discrimination parameter has a function that is very similar to mean drift rate ν in the LBA, ans so we retain ν as the parameter symbol. The resulting products, which we call discrimination contrasts, index a subject's ability to discriminate a stimulus from its competitors. For example, consider a scenario where all mismatching input values are m, $M - m = 1.5$, and $\nu = 2$.[4] In this scenario, a subject's matching discrimination contrast is 3, their mismatching discrimination contrast is -3, and their unrelated accumulators' discrimination contrast is 0.

Our implementation of the ALBA adds two values to all discrimination contrasts. Firstly, the maximum discrimination contrast across all advantage accumulators is added to the discrimination contrasts of each advantage accumulator. In the example above, the set of discrimination contrasts {3, 0, -3} becomes {6, 3, 0}. Secondly, the same positive baseline input value, ι, is added to the discrimination contrasts for all advantage accumulators. The value of ι is a free parameter. For instance, for $\iota = 2$, the final set of stimulus contrasts becomes {8, 5, 2}. These values serve as the mean drift rates for their respective advantage accumulators. For simplicity, we assume a common variance, $s = 1$, for the drift rate distribution of all advantage accumulators, and we assume that the inputs to all accumulators are uncorrelated. That is, we assume that on each trial an independent random sample is added to the mean drift rate of the accumulator. This set–up ensures that most accumulators have a positive non–zero drift rate and hence eventually reach their threshold, so that a response is guaranteed (for a similar mechanism see e.g., van Ravenzwaaij et al., 2012; Bogacz, Usher, Zhang, & McClelland, 2007; Busemeyer, Townsend, Diederich, & Barkan, 2005).

Note that this whole process requires only two free parameters: discrimination parameter ν and baseline input ι. Each of the advantage accumulators has an input, and hence mean drift rate, determined by the dimensions of the stimuli (see Trueblood et al., 2014, for another approach where drift rates are constructed from differences). Other standard LBA parameters A, b, and t_0 are assumed to be identical across advantage accumulators and are free parameters to be estimated from the data.

Different schemes for ensuring that some or all drift rates are positive are possible (e.g., exponentiation of the stimulus contrasts). Here we calculate stimulus contrasts by taking input differences, but one could also quantify them as the ratio of all positive values (e.g., Hawkins et al., 2014). These schemes may have practical and conceptual advantages, but we leave investigation of such possibilities to future work. Here, instead, we focus on different stopping rules that could be used with an advantage-accumulation architecture. In particular, in the following sections, we define and investigate three possible stopping rules, which we call Win-All, Lose-All, and Lose-One.

Win-All (WA)

The Win-All (WA) rule assumes that a decision is made as soon as one of the response options has beaten each of the other alternatives in a pairwise comparison. For example, the decision maker chooses option A from A,B,C, D if and only if:

1. Accumulators $A - B$, $A - C$, and $A - D$ have reached their thresholds, and:

[4] In general, in both the LBA and ALBA, mismatching input values can differ among different accumulators. For example, suppose the stimuli corresponding to accumulators B and C are similar and dissimilar, respectively, to the stimulus corresponding to accumulator A. When the stimulus corresponding to accumulator A is presented, there will be a lesser mismatch, and hence a larger input, to accumulator B than to accumulator C.

2. At least one of the accumulators in each of the sets $\{B-A, B-C, B-D\}$, $\{C-A, C-B, C-D\}$, and $\{D-A, D-B, D-C\}$ has *not* reached its threshold.

In words, response option is chosen A if is the first option to have beaten every other response option. This rule could be instantiated by linking each response with a counter having three inputs (e.g., from $A-B$, $A-C$, and $A-D$ for an A response) and requiring three counts to trigger its response.

For this stopping rule, accumulator termination (i.e., threshold-crossing) sequences can arise that appear counter-intuitive. For example, the termination sequence $B-A$, $C-A$, $D-A$, $A-B$, $A-C$, and then $A-D$ would result in choosing option A, as it is the first option to have beaten all its competitors, even though option A has also been beaten by each of its competitors. However, with reasonable parameter settings, such sequences are exceedingly unlikely. The example just given is unlikely because it requires opposite pairs to complete close together in a sequence, but this will only happen if they have similar inputs. However, in this case the difference between inputs will be small, and so they are unlikely to complete early in the sequence.

Under the Win-All rule, probability of responding A at time t is:

$$p_A(t) = \sum_{I \neq A} \left[PDF_{A-I}(t) \times \prod_{J \neq [A,I]} CDF_{A-J}(t) \right] \times \prod_{I \neq A} \left[1 - \prod_{K \neq I} CDF_{I-K}(t) \right] \tag{4}$$

I is an option in the set $\{B, C, D\}$, J is an option in the set $\{B, C, D\}$ that is not I, and K is an option in the set $\{A, B, C, D\}$ that is not I. Each CDF is obtained by applying (1) to the respective advantage accumulator, and each PDF is obtained by applying (3) to the respective advantage accumulator. The derivation for this Equation may be found in subsection "WA Derivation" of Appendix A.

The WA model is conceptually similar to a max-minus-next accumulator model (see e.g., McMillen & Holmes, 2006; McClelland, Usher, & Tsetsos, 2011; Brown, Steyvers, & Wagenmakers, 2009). In the max-next diffusion model, a decision is made as soon as the difference in activation between the most active accumulator and the *next* most active accumulator crosses a given threshold. The WA model is similar in that a response is made once the winning accumulator has beaten all of its competitors – that is, all the relevant accumulators have exceeded a given threshold. With this rule, the last advantage accumulator to cross its threshold will on average contrast the winner and the next best response option. Although the decision criterion is similar, the process by which the winning and next best accumulator accrue evidence is quite different between the two models. On top of that, the max-minus-next diffusion model fails to specify an algorithm for computing the decision rule. Exactly how is the model to "know" which is the second-best accumulator, and thus which activation to compare with the winning accumulator and monitor for threshold difference? The WA version of the ALBA specifies some of this missing detail, at least to the level of explanation which assumes that events can be triggered by threshold crossing, because the WA model specifies an algorithm for selecting which accumulators must be monitored for threshold crossing.

Lose-All (LA)

The Lose-All (LA) model assumes that the decision maker responds as soon as all but one of the response options have been beaten by every other contrasting alternative; thus, the LA model is a "last man standing" algorithm, and the conceptual inverse of the WA model. For example, the decision maker chooses A from A,B,C,D if and only if

1. All of the accumulators in each of the sets $\{A-B, C-B, D-B\}$, $\{A-C, B-C, D-C\}$, and $\{A-D, B-D, C-D\}$ have reached their threshold, and:
2. At least one of the accumulators in the set $\{B-A, C-A, D-A\}$ has *not* reached its threshold.

In words, response option A is the sole remaining option that has not yet been beaten by *every* competitor. This rule could be instantiated by linking each response with a counter having nine inputs (e.g., all but $A-B$, $A-C$, and $A-D$ for an A response) and requiring nine counts to trigger its response.

Again, accumulator termination sequences can arise that look somewhat contradictory. For example, consider the following sequence of accumulators reaching threshold: all of the accumulators in $\{A-B, A-C, A-D\}$, followed by all accumulators in $\{B-A, C-A, D-A\}$, $\{B-C, D-C\}$, and $\{B-D, C-D\}$. Then response B is made, despite the fact that A started out beating every competitor. Again, with sensible drift rates, this set of events is exceedingly unlikely.

The probability density function for the distribution of responses A at time t is given by

$$p_A(t) = \sum_{I \neq A} \left[\sum_{J \neq I} \left(PDF_{J-I}(t) \times \prod_{K \neq I,J} CDF_{K-I}(t) \times \prod_{L \neq I,A} \prod_{M \neq L} CDF_{M-L}(t) \right) \right] \times \left(1 - \prod_{I \neq A} CDF_{I-A}(t) \right) \tag{5}$$

I is an option in the set $\{B, C, D\}$, J is an option in the set $\{A, B, C, D\}$, K is an option in the set $\{A, B, C, D\}$ that is neither I nor J, L is an option in the set $\{B, C, D\}$ that is not I, and M is an option in the set $\{A, B, C, D\}$ that is not L. Each CDF is obtained by applying (1) to the respective advantage accumulator, and each PDF is obtained by applying (3) to the respective advantage accumulator. The derivation for this Equation may be found in subsection "LA Derivation" in Appendix A.

Lose-One (LO)

The Lose-One (LO) model assumes that the decision maker responds as soon as all but one of the response options have been beaten by at least one contrasting alternative. That is, the decision maker chooses A from A,B,C,D if and only if

1. At least one of the accumulators in each of the sets $\{A-B, C-B, D-B\}$, $\{A-C, B-C, D-C\}$, and $\{A-D, B-D, C-D\}$ has reached threshold, and:

2. None of the accumulators in the set $\{B-A, C-A, D-A\}$ has reached threshold.

In words, response option A is the last remaining option which has not been beaten by any competitor (another version of "last man standing"). Two layers of counters are required to instantiate this stopping rule. Four counters in the first layer take input from the sets of three just described, each requiring only one count to be triggered. Four counters in the second layer each correspond to a response. They take inputs from three counters in the previous layer (e.g. the counter corresponding to the A response takes inputs from the three sets in (1) above), and require three counts to trigger their response.

An advantage of this model is that it cannot produce sequences in which the winning response has ever lost in a direct comparison. However, it is possible to respond A without any accumulator that favors A having reached threshold (e.g., the sequence $D-B$, $B-C$, $C-D$ will trigger a response A). Again, with sensible drift rates, this set of events is exceedingly unlikely.

The probability density function for the distribution of responses A at time t is given by

$$p_A(t) = \sum_{I\neq A}\left[\sum_{J\neq I}\left(PDF_{J-I}(t)\times\prod_{K\neq I,J}[1-CDF_{K-I}(t)]\times\prod_{L\neq I,A}\left[1-\prod_{M\neq L}[1-CDF_{M-L}(t)]\right]\right)\right] \times\prod_{I\neq A}[1-CDF_{I-A}(t)] \tag{6}$$

I is an option in the set $\{B, C, D\}$, J is an option in the set $\{A, B, C, D\}$, K is an option in the set $\{A, B, C, D\}$ that is neither I nor J, L is an option in the set $\{B, C, D\}$ that is not I, and M is an option in the set $\{A, B, C, D\}$ that is not L. Each CDF is obtained by applying (1) to the respective advantage accumulator, and each PDF is obtained by applying (3) to the respective advantage accumulator. The derivation for this Equation may be found in subsection "LO Derivation" in Appendix A.

Detailed Model Specification and Parameter Recovery

An important element for the success of any scientific model is practical usability. The most widely cited, widely adopted, and perhaps influential decision-making models are those that can be used by any interested scientist, and can be compared against data using modern statistical techniques. For example, the diffusion model (Ratcliff, 1978), the EZ diffusion (Wagenmakers, van der Maas, & Grasman, 2007), and the LBA (Brown & Heathcote, 2008) all owe much of their success to the existence of simple mathematical expressions and widely-available software solutions. Without these elements, it can be difficult to be confident about the meaning of parameters estimated by the models.

Hence, our first step in assessing the ALBA models is to examine parameter recovery (Heathcote, Brown, & Wagenmakers, 2014). This requires a demonstration that the three proposed ALBA models, WA, LA, and LO, can recover their own parameters in reasonably sized data samples. This demonstration is made by simulating synthetic data from the models, with known parameters, and then estimating model parameters from those data,

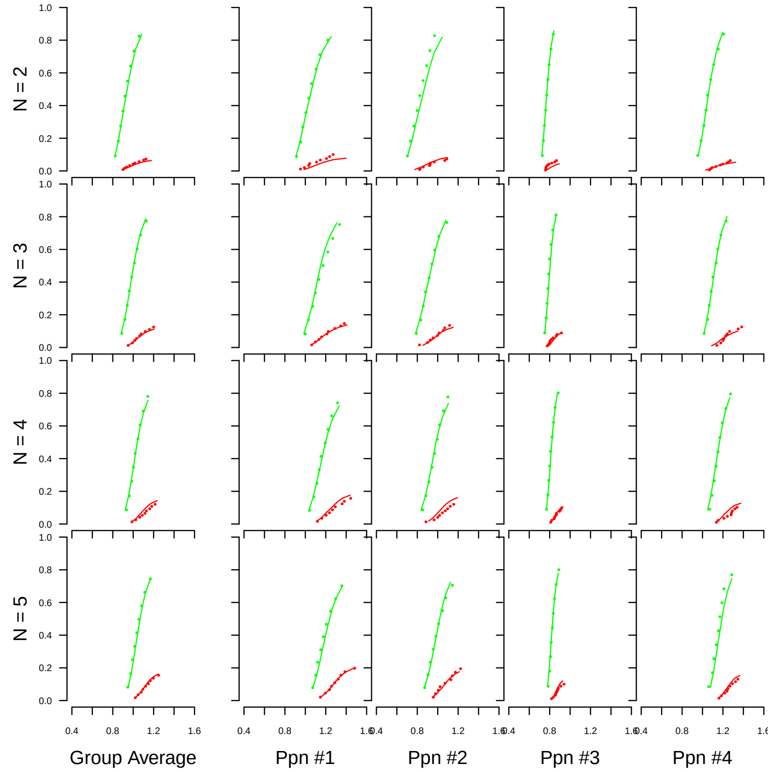

Figure 2. Posterior predictive data from the WA model. Left column: group data. Right four columns: individual data. Proportions (y–axis) plotted against RT percentiles (x–axis) for the .1, .2, . . . , .9 percentiles calculated from correct RTs (green) and error RTs (red). For all panels, dots represent simulated data to be fit, and lines represent posterior predictive data from the fitted model.

with standard techniques. The estimated parameters are then compared with the known, data-generating parameters.

For each of the three models, we simulated four participants that performed 200 trials each with response set size two, three, four, and five, for a total of 800 trials. All individual parameter values were generated from the following group level distributions: $b \sim N(1, 0.5)|(0, \infty)$, $A \sim N(1, 0.5)|(0, \infty)$, $\nu \sim N(1.5, 0.5)|(0, \infty)$, $t_0 \sim N(0.5, 0.1)|(0, \infty)$, and $\iota \sim N(2, 0.5)|(0, \infty)$. The notation $\sim N(,)$ indicates that values were drawn from a normal distribution with mean and standard deviation parameters given by the first and second number between parentheses, respectively. The notation $|(,)$ indicates that the values sampled from the normal distribution were truncated between the first and second numbers in parentheses. Except for drift rates, all parameters were identical across accumulators.

The sampled discrimination parameter ν determines the relative processing speeds of different advantages for each participant, and thus their accuracy. A high value of ν leads to the correct (set of) accumulator(s) finishing relatively fast compared to the competing accumulators, and so the response being more often correct. The baseline input ι can be

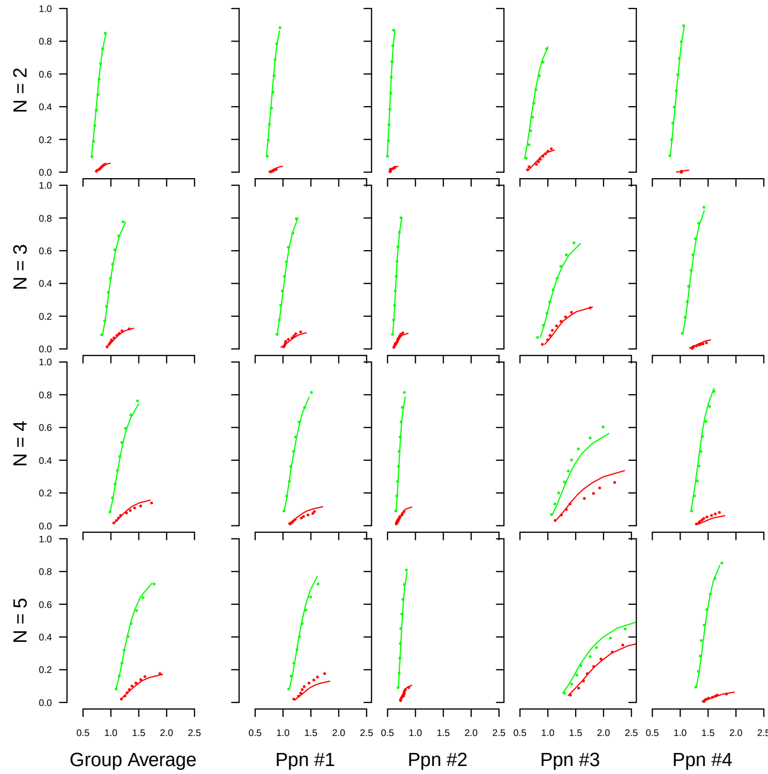

Figure 3. Posterior predictive data from the LA model. Left column: group data. Right four columns: individual data. Proportion correct (y–axis) plotted against RT percentiles (x–axis) for the .1, .2, ...9 percentiles calculated from correct RTs (green) and error RTs (red). For all panels, dots represent simulated data to be fit, and lines represent posterior predictive data from the fitted model.

thought of as controlling the absolute speed of information processing; a high value of ι leads to fast average terminations for all accumulators.

We fit each of the three models to their respective simulated data and compared the resulting parameter estimates with the original values that were used to generate the data. We also generated posterior predictive data, using the posterior distributions over parameter values, and compared them to the original simulation data.

All ALBA model fits presented in this paper were realized in a hierarchical Bayesian fashion (see for details Turner, Sederberg, Brown, & Steyvers, 2013). The Bayesian set-up allows for using MCMC sampling, which is an efficient way of finding the optimal set of parameters (Gamerman & Lopes, 2006; Gilks, Richardson, & Spiegelhalter, 1996; van Ravenzwaaij, Cassey, & Brown, in press). Details on starting values, prior distributions, and number of iterations may be found in subsection "Recovery Distributions" in Appendix B.

Parameter convergence was satisfactory in all cases. Figures 2, 3 and 4 show that all models fit their data well. The simulated data are shown as a dot for each decile, with the posterior predictive data shown by lines, and the two correspond closely.

Results of the parameter recovery are displayed in Table 1. The table provides the

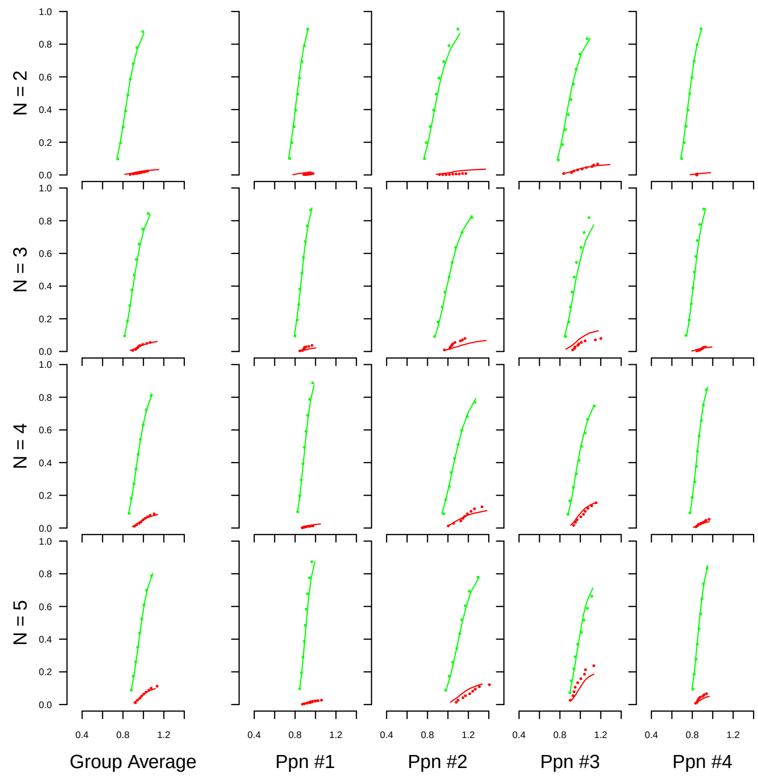

Figure 4. Posterior predictives of the LO model. Left column: group data. Right four columns: individual data. Proportion correct (y–axis) plotted against RT percentiles (x–axis) for the .1, .2, ...9 percentiles calculated from correct RTs (green) and error RTs (red). For all panels, dots represent simulated data to be fit, and lines represent posterior predictive data from the fitted model.

population parameter values as well as the parameter values sampled from the population distribution for each participant. For each true value the 50% credible interval from the fitted models is presented in parentheses. The top five rows display results for the WA model, the middle five rows the results for the LA model, and the bottom five rows the results for the LO model.

Table 1 shows that all models generally recover the individual participant parameters well with the simulated sample size, 800 observations per participant, with no evidence of any systematic bias. Also, in all cases the 50% credible intervals contained true values at the nominal rate or a little better (overall 14/20, 11/20, and 13/20 intervals cover the true value for WA, LA and LO respectively, and in all but one case two or more out of four for individual parameters). The same is true for the population parameters; despite the small number of participants at least four out of five intervals cover the true value for each model.

Table 1
Parameter recovery results. True parameter values, with a 50% credible interval of the posterior presented in parentheses. Columns represent parameters and rows represent different participants ("hyper" indicates parameters of the group-level distributions). The top four rows display results for the WA model, the middle three rows display results for the LA model, and the bottom three rows display results for the LO model.

Model	Pp	ν	b	A	t_0	ι
WA	Hyper	1.50 (1.18, 1.31)	1.00 (1.01, 2.48)	1.00 (0.71, 1.39)	0.50 (0.42, 0.64)	2.00 (1.88, 3.06)
	1	1.14 (1.10, 1.20)	1.02 (0.77, 1.24)	1.30 (1.11, 1.29)	0.64 (0.62, 0.70)	1.74 (1.47, 1.96)
	2	1.21 (1.21, 1.32)	1.29 (1.07, 1.67)	1.44 (1.45, 1.68)	0.44 (0.37, 0.46)	2.31 (2.07, 2.67)
	3	1.39 (1.22, 1.39)	0.31 (0.36, 2.35)	0.44 (0.32, 0.44)	0.65 (0.43, 0.64)	1.30 (1.40, 4.61)
	4	1.24 (1.16, 1.28)	1.50 (1.53, 3.09)	0.94 (0.50, 0.98)	0.68 (0.50, 0.68)	2.58 (2.46, 3.68)
LA	Hyper	1.50 (1.48, 2.06)	1.00 (0.89, 1.61)	1.00 (1.16, 1.61)	0.50 (0.34, 0.53)	2.00 (2.08, 2.79)
	1	1.57 (1.63, 1.75)	1.16 (0.78, 1.07)	1.13 (1.17, 1.34)	0.49 (0.50, 0.55)	2.17 (1.87, 2.13)
	2	1.96 (1.96, 2.15)	0.73 (0.60, 0.82)	0.98 (0.95, 1.08)	0.39 (0.38, 0.40)	3.33 (3.03, 3.46)
	3	0.94 (0.84, 0.91)	0.85 (0.98, 1.26)	1.35 (1.02, 1.28)	0.31 (0.26, 0.31)	1.56 (1.66, 1.80)
	4	2.25 (2.13, 2.31)	1.80 (1.33, 1.83)	1.88 (1.74, 1.98)	0.54 (0.54, 0.60)	2.61 (2.29, 2.66)
LO	Hyper	1.50 (1.55, 2.14)	1.00 (0.72, 1.53)	1.00 (0.85, 1.09)	0.50 (0.48, 0.58)	2.00 (1.85, 2.57)
	1	2.00 (2.09, 2.29)	0.73 (0.67, 0.77)	0.76 (0.87, 1.05)	0.60 (0.59, 0.61)	1.74 (1.86, 2.46)
	2	1.35 (1.31, 1.38)	1.73 (1.51, 1.70)	0.53 (0.59, 0.91)	0.40 (0.40, 0.43)	2.62 (2.64, 3.02)
	3	1.32 (1.28, 1.37)	0.53 (0.53, 0.61)	0.88 (0.78, 0.93)	0.61 (0.59, 0.61)	1.20 (1.13, 1.58)
	4	1.95 (1.86, 1.99)	0.73 (0.69, 0.76)	0.93 (0.88, 1.03)	0.53 (0.52, 0.54)	1.92 (2.09, 2.58)

Hick's Law

Hick's Law, a benchmark result for RT models of multialternative decisions, states that the mean RT and the logarithm of the number of choice alternatives are approximately linearly related, when accuracy is high. Hick's Law has been encountered in empirical data sets on numerous occasions (Hick, 1952; Teichner & Krebs, 1974). Thus, a crucial test for any new multialternative RT model is that it should be able to produce Hick's Law. A well-known problem with independent racing accumulator models is that they produce almost the opposite of Hick's Law. The conventional one-to-one mapping between choices and accumulators leads to *faster* decisions with more accumulators, because the winning accumulator is likely to be faster when a larger set is presented (i.e., "statistical facilitation", Raab, 1962).

Usher, Olami, and McClelland (2002) discuss how Hick's Law can be accommodated by an independent race model if participants adjust their evidence thresholds in order to compensate for the fact that the probability of an error also increases with the number of choices. They showed that maintaining constant accuracy in this way approximates Hick's Law. Usher and McClelland (2001) showed that this can also be the case in the LCA model, where an increased number of choices can be associated with slowing, even when the threshold is held constant, due to competition (i.e., lateral inhibition) among accumulators. The ALBA displays similar slowing because larger choice sets require larger numbers of accumulators to reach threshold before a decision is triggered.

We took advantage of the tractability of the ALBA to directly fit it to an archival

Hick's Law data set (van Maanen et al., 2012). This approach allows us to go beyond the conventional formulation of Hick's Law in terms of mean RT, expanding our test of the ALBA to its ability to account for the effects of choice-set-size simultaneously on both accuracy and the full distribution of RT (see also Brown, Marley, Donkin, & Heathcote, 2008; Hawkins, Brown, Steyvers, & Wagenmakers, 2012a, 2012b). In these fits we used model selection to investigate the effects of adjustments in parameters associated with each accumulator, such as the response threshold, as a function of choice set size. Discrimination rates ν were obtained in the same manner as in the simulation study, with the same parameter values applying for all set sizes.

van Maanen et al. (2012) had participants judge the direction of a random dot kinematogram (Britten, Shadlen, Newsome, & Movshon, 1992), which is a pattern of dots on a computer screen, most of which move in random directions. A small proportion of the dots move coherently in one direction, and it is this direction that the participant must identify. van Maanen et al. had participants choose between either 3, 5, 7, or 9 alternative directions. In their "clustered" condition, which we address here, the angular spacing between adjacent alternatives remained constant across the different set sizes. All four conditions were administered within all five subjects, and there were 144 trials per condition.

Table 2
DIC values summed over participants for all nine models fit to van Maanen's Hick's law data set. The best model for these data is the LA model with parameters b and A free to vary between conditions. Type = the specific parameterization of the model (see text for details); P = number of free parameters per participant for all four conditions; Deviance = -2 times the likelihood of the mean parameter estimate; pD = -2 times the mean likelihood of the overall model + 2 times the likelihood of the mean parameter estimate; DIC = Deviance+2pD.

Model	Type	P	Deviance	pD	DIC
WA	1b1A	5	8024	17	8058
	4b1A	8	7731	31	7794
	4b4A	11	7482	38	7558
LA	1b1A	5	9366	-73	9220
	4b1A	8	7000	117	7233
	4b4A	11	6854	35	6924
LO	1b1A	5	8872	21	8914
	4b1A	8	8820	35	8891
	4b4A	11	7680	43	7766

For each stopping rule, we fit a version that restricted the b (threshold) and the A (start-point noise) parameter to be fixed between set-size conditions (we call this "1b1A"), a version that restricted the A parameter to be fixed between conditions but allowed the b parameter to vary freely between conditions (4b1A), and a version that allowed both the b and the A parameters to vary freely between conditions (4b4A). Threshold and start point noise adjustments are reasonable prima facie, because the experiment included a cue indicating the number of alternatives 500 msec prior to each stimulus. For all nine resulting

Table 3
Estimated parameters of the LA 4b4A model of the van Maanen data set. Displayed are the median parameter values, with a 50% credible interval of the posterior presented in parentheses. Columns represent parameters and rows represent different participants (hyper = parameters of the group level distributions).

Pp	v	b.3	b.5	b.7	b.9	A.3
Hyper	1.49 (1.39, 1.58)	1.49 (1.36, 1.57)	0.2 (0.12, 0.33)	0.06 (0.03, 0.1)	0.08 (0.04, 0.14)	0.29 (0.16, 0.47)
1	1.4 (1.36, 1.43)	1.47 (1.36, 1.56)	0.27 (0.16, 0.39)	0.07 (0.04, 0.12)	0.08 (0.04, 0.14)	0.23 (0.12, 0.35)
2	1.25 (1.22, 1.28)	1.53 (1.38, 1.64)	0.23 (0.14, 0.34)	0.06 (0.03, 0.11)	0.08 (0.04, 0.13)	0.55 (0.35, 0.8)
3	1.67 (1.63, 1.71)	1.51 (1.4, 1.61)	0.15 (0.08, 0.23)	0.05 (0.02, 0.1)	0.12 (0.06, 0.19)	0.31 (0.19, 0.46)
4	1.3 (1.26, 1.33)	1.52 (1.37, 1.6)	0.18 (0.11, 0.29)	0.06 (0.03, 0.11)	0.1 (0.05, 0.15)	0.34 (0.2, 0.51)
5	1.85 (1.81, 1.89)	1.44 (1.31, 1.55)	0.27 (0.16, 0.38)	0.09 (0.05, 0.15)	0.09 (0.05, 0.14)	0.25 (0.12, 0.41)
Pp	A.5	A.7	A.9	t_0	ι	
Hyper	1.76 (1.66, 1.89)	1.57 (1.5, 1.66)	1.34 (1.24, 1.42)	0.14 (0.1, 0.18)	1.82 (1.64, 2)	
1	1.73 (1.58, 1.87)	1.58 (1.51, 1.67)	1.32 (1.21, 1.4)	0.12 (0.1, 0.14)	1.43 (1.38, 1.49)	
2	1.73 (1.62, 1.87)	1.55 (1.46, 1.63)	1.26 (1.16, 1.35)	0.13 (0.11, 0.15)	1.93 (1.87, 1.98)	
3	1.84 (1.72, 1.96)	1.55 (1.47, 1.64)	1.25 (1.14, 1.35)	0.08 (0.06, 0.1)	1.26 (1.19, 1.32)	
4	1.79 (1.68, 1.92)	1.56 (1.48, 1.65)	1.32 (1.24, 1.41)	0.09 (0.07, 0.1)	2.16 (2.12, 2.22)	
5	1.78 (1.66, 1.9)	1.64 (1.55, 1.74)	1.52 (1.42, 1.64)	0.26 (0.24, 0.27)	2.4 (2.31, 2.47)	

models, the discrimination rate (ν), non-decision time (t_0), and baseline (ι) parameters were fixed over conditions.

Details on starting values, prior distributions, and number of iterations may be found in subsection "van Maanen Distributions" in Appendix B. In all cases, parameter convergence was satisfactory. We performed model selection using the Deviance Information Criterion (DIC; Spiegelhalter, Best, Carlin, & van der Linde, 2002), a measure that balances goodness of fit against model complexity. In this sense, DIC is similar to the well-known BIC and AIC measures, which use the number of free parameters to quantify complexity. In hierarchical models, however, the effective number of free parameters does not equal the nominal value due to constraints imposed by the population model. DIC addresses this issue by providing an estimate of the effective number of parameters, and hence model complexity, that takes into account the hierarchical constraint. Lower values of DIC indicate better support for a model. DIC values for all nine versions of the model can be found in Table 2. These values suggest that the best model for these data is the LA 4b4A model with parameters b and A free to vary between conditions. Further, within each model architecture (WA, LA, LO) the best model was the 4b4A variant of that architecture. Estimated parameters can be found in Table 3.

The posterior predictive data for the LA 4b4A model is compared with the aggregated data the left panel of Figure 5. For reference, posterior predictive data for the LA 1b1A model are compared against the same aggregated data in the right panel. The original data are shown by points joined by lines, and distributions of posterior predictive data are shown by box–and–whiskers. Every observation contained in each box–and–whiskers is an average value across participants based on data generated from a sample from the joint posterior. Boxes contain 50% of the observations, and tails extend to 100%. The top row shows correct RTs, the middle row shows error RTs, and the bottom row shows proportion correct. Deciles of .1, .5, and .9 are displayed.

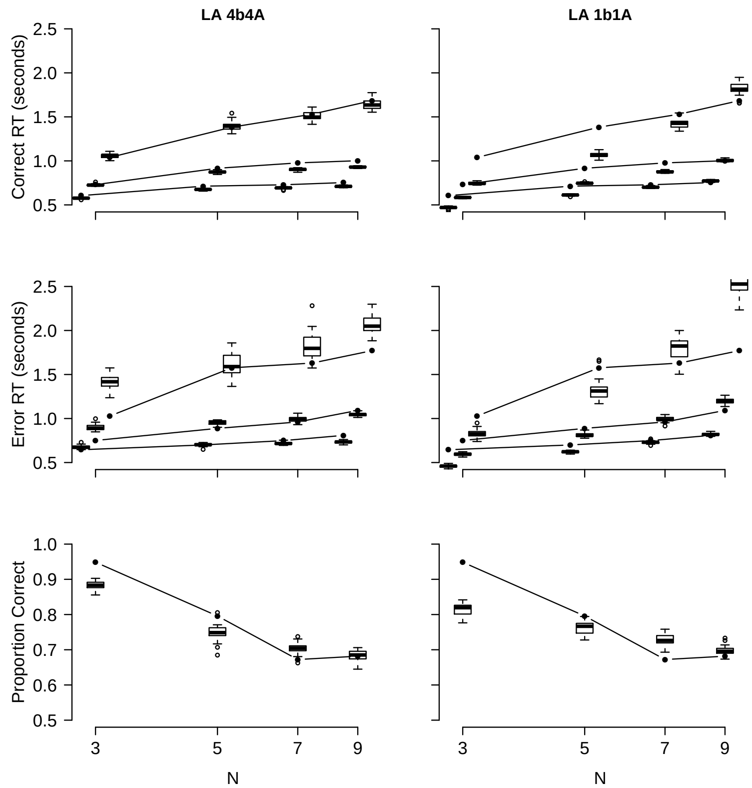

Figure 5. Posterior predictive data for the van Maanen data. RT .1, .5, and .9 deciles (y–axis) calculated for correct RTs (top row), error RTs (middle row), and the proportion of correct responses (bottom row), all as functions of set size (N) on a logarithmic scale. For all panels, box-and-whiskers represent posterior predictive data (the box contains 50% of the simulated data, with a bar across the middle indicating the median, and whiskers extend to the data extremes) and lines represent data. Left panels: LA 4b4A; right panels: LA 1b1A. See text for details.

The plot of the 1b1A model shows that the combination of the ALBA architecture and LA stopping rule naturally produce Hick's Law: an approximately linear increase in mean RT with the logarithm of set size, as well as as an increase in variability with set size. The same is also true for the WA stopping rule. However, the 1b1A model overestimates the difference between smaller set sizes, particularly between 3 and 5 alternatives, in this data set. Table 3 shows that to compensate the threshold and start-point noise are markedly higher for set size three in model 4b4A. This result indicates participants engaged in a speed-accuracy tradeoff, displaying less response caution for the smallest set size relative to larger set sizes. Hawkins et al. (2012b) came to a remarkably similar conclusion, using an entirely different decision paradigm, and modeling approach. Figure 5 shows that, overall, the LA 4b4A model fits the data well, except for the error responses of the smallest set size, where it under-estimates the slower error responses and overestimates accuracy.

Mutual Inhibition

As part of a larger investigation, Teodorescu and Usher (2013) examined the ability of different kinds of sequential accumulator models to fit a variety of data from multi-alternative forced choice tasks. They divided models into three sub-types based on how accumulators and their inputs interact: independent models, input competition models, and response competition models. Independent models assume no direct competition between the accumulators. All accumulators race independently towards their threshold, after which a response is executed. The LBA (Brown & Heathcote, 2008) is an example of an independent model. Input competition models incorporate competition at the input level. Specifically, if a stimulus provides some input in favor of response A, it will also provide some input against response B. Note that such competition is not between the accumulators themselves, but between what is fed into the accumulators (e.g. drift rates). The diffusion model (Ratcliff, 1978; van Ravenzwaaij et al., 2012), as well as the ALBA, are examples of input-competition models. Response competition models incorporate competition at the output level. Competition is implemented as a function of the activation levels of each accumulator. The leaky, competing accumulator model (Usher & McClelland, 2001) is an example of a response competition model.

Teodorescu and Usher (2013) report data from several experiments indicating that response competition models are capable of capturing an empirical phenomenon that is not adequately captured by either independent models or input competition models. Specifically, Teodorescu and Usher claim that independent models predict a decrease in response time in conditions where the correct response has a near competitor, whereas response competition models predict an increase in response time. They also present empirical results showing that decision makers have longer response times in such conditions, and through model fitting rule out input-competition, concluding that some form of response competition is necessary to capture their results. In particular, they examined two types of input-competition, "normalized" models, where the input to each accumulator is divided by the sum of inputs to all accumulators, and "feedforward" inhibition, where evidence supporting each accumulator can also subtract from the evidence supporting other accumulators.

In this section, we re-evaluate the conclusions of Teodorescu and Usher (2013) from the point of view of the ALBA model. In each trial of Teodorescu and Usher's Experiment 1A participants were presented with four patches of different brightness, and had to decide which patch was the brightest. In the easy condition, one patch was much brighter than all the others, whereas in the difficult condition, one of the other three patches was close in brightness to the brightest patch (see Teodorescu & Usher, 2013, Figure 4). The authors fit different models to the data and conclude: "This study demonstrated that both the independent and normalized race as well as the FFI [feedforward inhibition] model speed up in the difficult condition for all parameter values, while response competition models are able to account well for the observed slowdown in the data." (pp. 14).

Here, we present fits of the three ALBA models to the data from Experiment 1A reported by Teodorescu and Usher (2013). We demonstrate that the LA model fits the data well and overturn the claim that response competition is a necessary requirement to fit multialternative response data. This result suggests instead that it was the assumption of a one-to-one mapping of accumulators to responses that was problematic for the class of input-competition models.

Brightness Data

Eight participants provided data for the brightness discrimination experiment reported by Teodorescu and Usher (2013). Each participant performed between 1,000 and 1,200 trials. Half of these trials constituted the easy condition with brightness input values of {8, 5, 3, 3}, the other half of the trials constituted the difficult condition with brightness input values of {8, 7, 2, 2}. Note that in each condition the sum of the brightness values if the same, so that normalizing these values by dividing them by the sum preserves the ratios between values.

Following Teodorescu and Usher (2013), we assume a simple identity psychophysical function, directly using nominal stimulus brightness values as inputs to the ALBA. More psychologically plausible psychophysical functions, such as sigmoidal functions, are discussed later. Stimulus contrasts were calculated as all the pairwise differences between input values (e.g., for the difficult condition: 8-7, 8-2, 8-2, 7-8, 7-2, 7-2, 2-8, 2-7, 2-2, 2-8, 2-7, 2-2). Next, the maximum stimulus contrast was added to each of these stimulus contrasts. The resulting values were multiplied by discrimination parameter ν, after which baseline parameter ι was added. Unfortunately, due to a programming error, the recorded data only contain a log of whether the response was correct or incorrect (Teodorescu, personal communication). Which of the incorrect options was chosen is unknown. As such, we were forced to average the model's log-likelihoods for all three error response options in our fits to the data.

Table 4
DIC values for the three models fit to the Teodorescu data set. The best model for these data is the LA model. Deviance = -2 times the likelihood of the mean parameter estimate; pD = -2 times the mean likelihood of the overall model + 2 times the likelihood of the mean parameter estimate; DIC = BestFit+2pD.

Model	Deviance	pD	DIC
WA	3024	-56	2911
LA	2304	48	2399
LO	2841	34	2909

All participants' data were fit with the WA, the LA, and the LO models. For each model, we constrained parameters b, A, and t_0 to be identical across conditions. The restriction on b and A is different from our fits of van Maanen et al.'s (2012) Hick's law data set. The reason lies in the nature of the manipulation: different numbers of alternatives were cued before stimulus presentation in van Maanen et al.'s experiment, allowing the potential for changes in the amount of evidence required to make a decision. In Teodorescu and Usher's experiment, differences in stimulus difficulty only become apparent during the decision–making process. As such, differences in thresholds or starting points are not reasonable. We estimated ν and ι parameters for each of the two conditions, as these parameters reflect the stimulus differences that distinguish these conditions.

Table 5
Estimated parameters of the LA model for the Teodorescu data set. Displayed are the median parameter values, with a 50% credible interval of the posterior presented in parentheses. Columns represent parameters and rows represent different participants (hyper = parameters of the group-level distributions).

Pp	v(Easy)	v(Hard)	b	A	t0	iota(Easy)	iota(Hard)
Hyper	0.91 (0.87, 0.95)	1.09 (0.85, 1.24)	0.25 (0.18, 0.33)	1.04 (0.73, 1.33)	0.55 (0.52, 0.57)	0.3 (0.17, 0.42)	0.35 (0.25, 0.46)
1	0.77 (0.73, 0.81)	1.17 (1.11, 1.24)	0.26 (0.2, 0.33)	0.84 (0.79, 0.9)	0.6 (0.58, 0.61)	0.71 (0.53, 0.9)	1.31 (1.19, 1.43)
2	0.89 (0.85, 0.92)	1.37 (1.28, 1.45)	0.26 (0.22, 0.31)	1.16 (1.08, 1.24)	0.63 (0.61, 0.64)	0.16 (0.07, 0.28)	0.87 (0.76, 0.98)
3	0.87 (0.83, 0.9)	0.75 (0.71, 0.79)	0.01 (0, 0.02)	0.72 (0.69, 0.76)	0.62 (0.62, 0.62)	0.55 (0.43, 0.67)	1.67 (1.58, 1.76)
4	1.06 (1.02, 1.1)	1.29 (1.23, 1.35)	0.05 (0.02, 0.09)	4.39 (4.15, 4.65)	0.47 (0.44, 0.49)	0.31 (0.18, 0.44)	1.31 (1.21, 1.41)
5	1 (0.97, 1.04)	1.94 (1.85, 2.04)	0.37 (0.31, 0.43)	1.18 (1.11, 1.24)	0.58 (0.56, 0.59)	0.22 (0.11, 0.34)	0.66 (0.53, 0.79)
6	1.02 (0.97, 1.07)	1.77 (1.68, 1.86)	0.51 (0.42, 0.61)	1.85 (1.74, 1.96)	0.63 (0.61, 0.65)	0.71 (0.52, 0.91)	1.3 (1.17, 1.43)
7	1.11 (1.07, 1.15)	1.43 (1.36, 1.5)	0.56 (0.49, 0.65)	2.32 (2.19, 2.46)	0.55 (0.53, 0.57)	0.16 (0.07, 0.28)	1.44 (1.34, 1.53)
8	0.8 (0.77, 0.83)	1.23 (1.13, 1.32)	0.14 (0.11, 0.18)	0.97 (0.91, 1.03)	0.58 (0.57, 0.59)	0.1 (0.04, 0.19)	0.78 (0.67, 0.9)

Details on starting values, prior distributions, and number of iterations may be found in subsection "Teodorescu Distributions" in Appendix B. Parameter convergence was satisfactory. Mean DIC values for all three versions of the model can be found in Table 4. The Table shows that that DIC selects the LA model, consistent with our finding for van Maanen et al.'s experiment. Parameter estimates for the LA fit of the Teodorescu and Usher data set can be found in Table 5.

The aggregate posterior predictive data for the LA model with ν and ι free between conditions, labeled $2\nu2\iota$ are compared against the data in the left panels of Figure 6. Figure 6 shows that this model fits the data well. This good fit is achieved with seven estimated parameters, which is comparable to the seven or eight free parameters in the models that were fit by Teodorescu and Usher (2013). The model has some difficulty with the tail of the distribution of the errors, in particular for the easy condition. This is common with RT model fits and is the result of very sparse data (i.e., few errors).

Posterior predictive data for the LA model with discrimination rate ν fixed between conditions, labeled $1\nu1\iota$, are compared against data in the right panels of Figure 6. In this model differences between the easy and hard conditions are purely a function of the model architecture. Although restricting parameters ν and ι to be equal between both conditions results in poorer model fit (DIC = 3607), it still captures the qualitative pattern of results without the need for response competition. In particular, RTs in the hard condition are slower and accuracy in the hard condition is lower. All of this is done with only five parameters free to vary.

Table 5 shows that the extra freedom in the $2\nu2\iota$ model was used to increase the discrimination parameter and the baseline parameter in the hard condition relative to the easy condition. Fit was improved roughly equally by allowing two ι estimates but only one ν estimate (DIC = 2522) and by allowing two ν estimates but only one ι estimate (DIC = 2534). We speculate that these model variants will only be able to be separated with certainty if greater detail, and constraint, is built into the psychophysical front-end of the model. Without such detail, it is difficult to know whether one model or another is being unduly handicapped by the implausibly strict assumption that internal representations of brightness are linearly related to physical brightness.

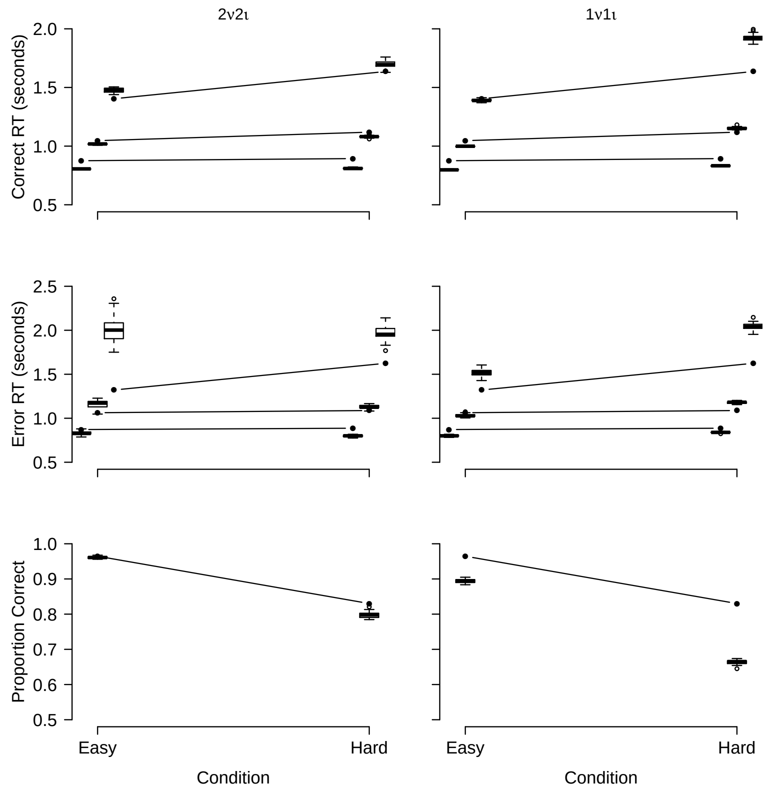

Figure 6. Posterior predictive data for fits to the Teodorescu data. RT .1, .5, and .9 deciles (y–axis) calculated for correct RTs (top row), error RTs (middle row), and the proportion of correct responses (bottom row), all as functions of difficulty condition. For all panels, box-and-whiskers represent posterior predictive data (the box contains 50% of the simulated data, with a bar across the middle indicating the median, and whiskers extend to the data extremes) and lines represent data. Left panel: LA $2\nu2\iota$ model; right panel: LA $1\nu1\iota$ model. See text for details.

Conclusion

Accumulator models of choice response time include versions that differ depending on the level of competition between separate accumulators. When these models are applied to decisions with more than two response options, almost all have one thing in common: accumulation is conceived in terms of absolute evidence for each response option (but see Marley, 1991, and Trueblood et al., 2014, for relative evidence models of choice probabilities, with Trueblood et al.'s results derived from a model of choice and response times). All of the existing accumulator models for multialternative choice make the basic assumption that there is one and only one accumulator representing each choice option.

We have developed and tested the advantage linear ballistic accumulator (ALBA) model, in which evidence for each option is relative to every other option, and in which many accumulators contribute to decisions about each response option. We presented the core architecture of this new model along with three different stopping rules[5], each of which combines the information quantified by the accumulators. The Win-All model assumes that the decision maker responds as soon as one of the response options has beaten every other contrasting alternative. The Lose-All model assumes that the decision maker responds as soon as all but one of the response options have been beaten by every other contrasting alternative. Finally, the Lose-One model assumes that the decision maker responds as soon as all but one of the response options have been beaten by at least one contrasting alternative.

Parameter recovery simulations demonstrated that, with all three stopping rules, the ALBA can satisfactorily recover its parameters, a feat that has proved difficult for other models of the type of data modeled here (Miletic et al., 2016). Then, we showed that each of these model architectures is capable of reproducing a ubiquitous phenomenon in multialternative decision making situations: Hick's Law. According to Hick's Law, the mean RT of correct responses and the logarithm of the number of choice alternatives are linearly related. The Win–All, the Lose–All, and the Lose–One models all qualitatively produce Hick's Law even when auxiliary parameters, such as the level of evidence at the start of a trial and response thresholds, are fixed across set-size conditions. We then subjected the three different stopping rules to a quantitative test, by fitting three different parameterizations of each to a data set which exhibits Hick's Law (van Maanen et al., 2012). The ALBA variant that was selected as providing the best tradeoff between simplicity and goodness-of-fit used the Lose-All stopping rule and allowed start points and thresholds to vary with set size. It also provided a good fit in an absolute sense, not only in terms of the central tendency of RT for correct responses, but also in terms of accuracy and the full distribution of RT, with the exception of error related data at the smallest set size where errors data were sparse. Our ability to provide this fine-grained test using Bayesian hierarchical methods, and associated model-selection techniques, was facilitated by the analytic tractability of the ALBA.

In the last part of the paper, we challenged the claim by Teodorescu and Usher (2013) that only models with response competition are capable of accurately accounting for the effect of a near competitor (e.g., an incorrect alternative that is closely matched to the correct alternative in terms of brightness) in multialternative forced-choice data. Relative to the easier displays where the nearest competitor is less similar to the target, the authors state that independent race models predict a speedup in the harder condition where there is a near competitor, whereas response competition models predict a slowdown. A slowdown was observed by Teodorescu and Usher (2013), and they found the best fit for their data was provided by the max-minus-next diffusion model, which is closely related to an ALBA model with a Win-All stopping rule. A crucial difference between the two models concerns where competition occurs. The max-minus-next diffusion model has competition at the response level, whereas the Win-All model, in common with all ALBA models, has competition at the input level. Despite this, all of the ALBA models we considered are able to qualitatively produce both the increased error rate and decreased speed associated with near competitors, even when easy and hard conditions had the same parameters.[6] Hierarchical Bayesian fitting

[5]Note that, under an appropriate parameterization, when there are only two choices the ALBA is equivalent to the standard LBA, and in this case all stopping rules are also equivalent.

[6]See supplementary materials, available at `http://www.donvanravenzwaaij.com/Papers`, for demonstrations of this fact for the Win-All and Lose-One stopping rules.

revealed that the Lose-All stopping rule provided the best trade off between simplicity and goodness-of-fit among the ALBA models we considered. Once again the fit to both accuracy and RT distributions was good in an absolute sense, with the exception of over predicting slower error responses, again in cells where data were sparse (rare errors for easy decisions). This close fit required slightly larger input parameters for the hard than easy condition, potentially because inputs are based on a transformation of objective brightness values.

The relative evidence framework we have presented in this paper incorporates two features that are non-standard. The first feature is the conceptualization of relative evidence accumulation. This allows the ALBA to fit Hick's Law data without the need for additional parameters. It also creates mutual inhibition in a structural way, and so does not need an explicit mutual inhibition parameter, which can be hard to estimate (Miletic et al., 2016). The second feature is the idea of using input values to determine the drift rates for accumulators. The idea of some front-end process that determines drift rate is not new; for example, a different but related approach has been taken by Trueblood et al. (2014). Where their approach was tailored specifically to multiattribute context effects, our approach is applicable to multialternative decision data more generally. The result of such integrative approaches are customizable sequential accumulator models that can account for complex phenomena such as the attraction effect (Trueblood et al., 2014) or Hick's Law and the near competitor effect, without sacrificing model tractability.

We considered three variants of the ALBA model, defined by three different stopping rules. Of the three, the Lose-All model provided superior fits to both data sets that we examined. At this stage, we are cautious to conclude from this that the Lose-All variant is the best of the three; more data sets need to be examined to confidently support that conclusion. However, the Lose-All architecture shows a promise and is likely a reasonable starting point for researchers interested in applying the ALBA to their data.

Some of the choices we made when fitting the ALBA are admittedly arbitrary. For instance, for the Hick's Law data set, we chose to give the correct response an input value of 1 and the incorrect responses all an input value of 0. A different, more gradual, distribution of input values based on the distance from the correct response is conceivable as well. Error responses directly adjacent to the correct response would then have higher input values than more distal error responses. Our main motivation in this paper was to highlight the capacity of the ALBA framework to account for a range of empirical data patterns without the need for ad hoc assumptions. As such, our specification was mainly guided by simplicity, and should not be taken as a necessary feature of future applications.

This paper presents a framework for conceiving of multialternative decision making as a process of relative, rather than absolute, evidence accumulation. This framework leads to specific model instantiations which are tractable and come with tractable, closed-form expressions for the associated likelihood functions. These expressions are extremely important for practical model usage, including via maximum-likelihood and Bayesian approaches. We showed that the ALBA model can fit empirical Hick's Law data, and can fit data that were previously claimed to be hard to account for without a specific mutual inhibition parameter. We believe that recent developments that attempt to separate complex decisions into a back-end and a front-end part are exciting, because they provide explicit and parsimonious accounts of the components that lead to complex decisions. We hope that the advantage linear ballistic accumulator model provides one more building block towards a general account of such processes.

Acknowledgements

This research was supported by Australian Research Council grants to DvR, SDB, and AH (DE140101181, FT120100244, DP12102907 and DP110100234) and a Natural Science and Engineering Research Council Discovery Grant (8124-98) to AAJM. The work was carried out, in part, while Marley was a Distinguished Professor (part-time) at the Institute for Choice, University of South Australia Business School.

Appendix A
Derivations

WA Derivation

In this section, we unpack the equations for the WA version of the ALBA model. (4) for a decision trial with four response options (i.e., A, B, C, and D) can be written out as:

$$\begin{aligned} p_A(t) =& PDF_{A-B}(t) \times CDF_{A-C}(t) \times CDF_{A-D}(t) \times \prod_{I \neq A} \left[1 - \prod_{K \neq I} CDF_{I-K}(t)\right] + \\ & PDF_{A-C}(t) \times CDF_{A-B}(t) \times CDF_{A-D}(t) \times \prod_{I \neq A} \left[1 - \prod_{K \neq I} CDF_{I-K}(t)\right] + \\ & PDF_{A-D}(t) \times CDF_{A-B}(t) \times CDF_{A-C}(t) \times \prod_{I \neq A} \left[1 - \prod_{K \neq I} CDF_{I-K}(t)\right] \end{aligned} \tag{7}$$

where I is an option in the set $\{B, C, D\}$ and K is an option in the set $\{A, B, C, D\}$ that is not I. The top line represents the scenario where accumulator $A-B$ is the terminating accumulator that prompts the response ($PDF_{A-B}(t)$), accumulators $A-C$ and $A-D$ had finished before ($CDF_{A-C}(t) \times CDF_{A-D}(t)$), and at least one accumulator out of each of the sets $\{B-A, B-C, B-D\}$, $\{C-A, C-B, C-D\}$, and $\{D-A, D-B, D-C\}$ had not yet finished ($\prod_{I \neq A} \left[1 - \prod_{K \neq I} CDF_{I-K}(t)\right]$).

Similarly, the middle line represents the scenario where accumulator $A-C$ is the terminating accumulator that prompts the response ($PDF_{A-C}(t)$), accumulators $A-B$ and $A-D$ had finished before ($CDF_{A-B}(t) \times CDF_{A-D}(t)$), and at least one accumulator out of each of the sets $\{B-A, B-C, B-D\}$, $\{C-A, C-B, C-D\}$, and $\{D-A, D-B, D-C\}$ had not yet finished ($\prod_{I \neq A} \left[1 - \prod_{K \neq I} CDF_{I-K}(t)\right]$).

Finally, the bottom line represents the scenario where accumulator $A-D$ is the terminating accumulator that prompts the response ($PDF_{A-D}(t)$), accumulators $A-B$ and $A-C$ had finished before ($CDF_{A-B}(t) \times CDF_{A-C}(t)$), and at least one accumulator out of each of the sets $\{B-A, B-C, B-D\}$, $\{C-A, C-B, C-D\}$, and $\{D-A, D-B, D-C\}$ had not yet finished ($\prod_{I \neq A} \left[1 - \prod_{K \neq I} CDF_{I-K}(t)\right]$).

The PDF for response A is completed by summing the expressions on all three of these lines.

LA Derivation

In this section, we unpack the equations for the LA version of the ALBA model. (5) for a decision trial with four response options (i.e., A, B, C, and D) can be written out as:

$$\begin{aligned}p_A(t) = &\sum_{J\neq B}\left(PDF_{J-B}(t)\times\prod_{K\neq B,J}CDF_{K-B}(t)\times\prod_{L\neq B,A}\prod_{M\neq L}CDF_{M-L}(t)\right)\times\left(1-\prod_{I\neq A}CDF_{I-A}(t)\right)+\\&\sum_{J\neq C}\left(PDF_{J-C}(t)\times\prod_{K\neq C,J}CDF_{K-C}(t)\times\prod_{L\neq C,A}\prod_{M\neq L}CDF_{M-L}(t)\right)\times\left(1-\prod_{I\neq A}CDF_{I-A}(t)\right)+\\&\sum_{J\neq D}\left(PDF_{J-D}(t)\times\prod_{K\neq D,J}CDF_{K-D}(t)\times\prod_{L\neq D,A}\prod_{M\neq L}CDF_{M-L}(t)\right)\times\left(1-\prod_{I\neq A}CDF_{I-A}(t)\right)\end{aligned} \tag{8}$$

where I and L are options in the set $\{B, C, D\}$, J is an option in the set $\{A, B, C, D\}$, K is an option in the set $\{A, B, C, D\}$ that is not J, and M is an option in the set $\{A, B, C, D\}$ that is not L. The top line represents the sum of all scenarios of J where accumulator $J - B$, where J is not B, is the terminating accumulator that prompts the response ($PDF_{J-B}(t)$), all accumulators $K - B$, where K is not B or J, had finished before ($\prod_{K\neq B,J} CDF_{K-B}(t)$), all accumulators out of the sets $\{A - C, B - C, D - C\}$ and $\{A - D, B - D, C - D\}$ had finished before ($\prod_{L\neq B,A}\prod_{M\neq L} CDF_{M-L}(t)$), and at least one accumulator out of the set $\{B-A, C-A, D-A\}$ had not yet finished ($1-\prod_{I\neq A} CDF_{I-A}(t)$).

Similarly, the middle line represents the sum of all scenarios of J where accumulator $J - C$, where J is not C, is the terminating accumulator that prompts the response ($PDF_{J-C}(t)$), all accumulators $K - C$, where K is not C or J, had finished before ($\prod_{K\neq C,J} CDF_{K-C}(t)$), all accumulators out of the sets $\{A - B, C - B, D - B\}$ and $\{A - D, B - D, C - D\}$ had finished before ($\prod_{L\neq C,A}\prod_{M\neq L} CDF_{M-L}(t)$), and at least one accumulator out of the set $\{B-A, C-A, D-A\}$ had not yet finished ($1-\prod_{I\neq A} CDF_{I-A}(t)$).

Finally, the bottom line represents the sum of all scenarios of J where accumulator $J - D$, where J is not D, is the terminating accumulator that prompts the response ($PDF_{J-D}(t)$), all accumulators $K - D$, where K is not D or J, had finished before ($\prod_{K\neq D,J} CDF_{K-D}(t)$), all accumulators out of the sets $\{A - B, C - B, D - B\}$ and $\{A - C, B - C, D - C\}$ had finished before ($\prod_{L\neq D,A}\prod_{M\neq L} CDF_{M-L}(t)$), and at least one accumulator out of the set $\{B-A, C-A, D-A\}$ had not yet finished ($1-\prod_{I\neq A} CDF_{I-A}(t)$).

The PDF for response A is completed by summing the expressions on all three of these lines.

LO Derivation

In this section, we unpack the equations for the LO version of the ALBA model. (6) for a decision trial with four response options (i.e., A, B, C, and D) can be written out as:

$$\begin{aligned} p_A(t) = &\sum_{J \neq B} \left(PDF_{J-B}(t) \times \prod_{K \neq B,J} [1 - CDF_{K-B}(t)] \times \prod_{L \neq B,A} \left[1 - \prod_{M \neq L} [1 - CDF_{M-L}(t)] \right] \right) \times \prod_{I \neq A} [1 - CDF_{I-A}(t)] + \\ &\sum_{J \neq C} \left(PDF_{J-C}(t) \times \prod_{K \neq C,J} [1 - CDF_{K-C}(t)] \times \prod_{L \neq C,A} \left[1 - \prod_{M \neq L} [1 - CDF_{M-L}(t)] \right] \right) \times \prod_{I \neq A} [1 - CDF_{I-A}(t)] + \\ &\sum_{J \neq D} \left(PDF_{J-D}(t) \times \prod_{K \neq D,J} [1 - CDF_{K-D}(t)] \times \prod_{L \neq D,A} \left[1 - \prod_{M \neq L} [1 - CDF_{M-L}(t)] \right] \right) \times \prod_{I \neq A} [1 - CDF_{I-A}(t)] \end{aligned} \tag{9}$$

where I and L are options in the set $\{B, C, D\}$, J is an option in the set $\{A, B, C, D\}$, K is an option in the set $\{A, B, C, D\}$ that is not J, and M is an option in the set $\{A, B, C, D\}$ that is not L. The top line represents the sum of all scenarios of J where accumulator $J - B$, where J is not B, is the terminating accumulator that prompts the response $(PDF_{J-B}(t))$, all accumulators $K - B$, where K is not B or J, had not yet finished $(\prod_{K \neq B,J} [1 - CDF_{K-B}(t)])$, at least one accumulator out of the sets $\{A - C, B - C, D - C\}$ and $\{A - D, B - D, C - D\}$ had finished before $(\prod_{L \neq B,A} \left[1 - \prod_{M \neq L} [1 - CDF_{M-L}(t)]\right])$, and no accumulator out of the set $\{B - A, C - A, D - A\}$ had finished yet $(\prod_{I \neq A} [1 - CDF_{I-A}(t)])$.

Similarly, the middle line represents the sum of all scenarios of J where accumulator $J - C$, where J is not C, is the terminating accumulator that prompts the response $(PDF_{J-C}(t))$, all accumulators $K - C$, where K is not C or J, had not yet finished $(\prod_{K \neq C,J} [1 - CDF_{K-C}(t)])$, at least one accumulator out of the sets $\{A - B, C - B, D - B\}$ and $\{A - D, B - D, C - D\}$ had finished before $(\prod_{L \neq C,A} \left[1 - \prod_{M \neq L} [1 - CDF_{M-L}(t)]\right])$, and no accumulator out of the set $\{B - A, C - A, D - A\}$ had finished yet $(\prod_{I \neq A} [1 - CDF_{I-A}(t)])$.

Finally, the bottom line represents the sum of all scenarios of J where accumulator $J - D$, where J is not D, is the terminating accumulator that prompts the response $(PDF_{J-D}(t))$, all accumulators $K - D$, where K is not D or J, had not yet finished $(\prod_{K \neq D,J} [1 - CDF_{K-D}(t)])$, at least one accumulator out of the sets $\{A - B, C - B, D - B\}$ and $\{A - C, B - C, D - C\}$ had finished before $(\prod_{L \neq D,A} \left[1 - \prod_{M \neq L} [1 - CDF_{M-L}(t)]\right])$, and no accumulator out of the set $\{B - A, C - A, D - A\}$ had finished yet $(\prod_{I \neq A} [1 - CDF_{I-A}(t)])$.

The PDF for response A is completed by summing the expressions on all three of these lines.

Appendix B
Distributional Choices

Recovery Distributions

Starting values for the MCMC chains for individual parameters were drawn from the following distributions: $b \sim N(3, 0.3)|(0, \infty)$, $A \sim N(3, 0.3)|(0, \infty)$, all νs $\sim N(3, 0.3)|(0, \infty)$, $t_0 \sim N(1, 0.1)|(0, \infty)$, and $\iota \sim N(3, 0.3)|(0, \infty)$.

The hierarchical set–up prescribes that all individual parameters come from a truncated Gaussian group–level distribution (truncated to positive values only). Thus, for each parameter to be estimated, we are estimating a group level mean parameter and a group level standard deviation parameter. All group level mean parameters are normally distributed, with $b_\mu \sim N(3, 1.5)|(0,)$, $A_\mu \sim N(3, 1.5)|(0, \infty)$, all v_μs $\sim N(3, 1.5)|(0, \infty)$, $t0_\mu \sim N(1, 0.5)|(0, \infty)$, and $\iota_\mu \sim N(3, 1.5)|(0, \infty)$. All group level standard deviation parameters are exponentially distributed with mean values of 1. Starting values for the MCMC chains for group level μ parameters were drawn from the same distributions as those for the individual parameters, and starting values for group level σ parameters were derived from starting value distributions for the individual parameters by dividing the mean by 10 and the standard deviation by 2. These prior settings are quite uninformative, and are based on previous experience with parameter estimation for the LBA model. As a result, the specific settings will not have a large influence on the shape of the posterior. For more details on distributional choices for the priors, we refer the reader to Turner et al. (2013).

For sampling, we used 32 interacting Markov chains for every run, and drew 1,000 burn-in iterations of each chain followed by another 1,000 iterations after convergence. The two tuning parameters of the differential evolution proposal algorithm were set to standard values used in previous work: random perturbations drawn uniformly from the interval [-.001,.001] were added to all proposals; and the scale of the difference added for proposal generation was set to $\gamma = 2.38 \times (2K)^{-0.5}$, where K is the number of parameters per participant. The MCMC chains blocked proposals separately for each participant's parameters, and also blocked the group–level parameters in $\{\mu, \sigma\}$ pairs.

van Maanen Distributions

For all models, starting values for the MCMC chains for individual parameters were drawn from the following distributions: all bs $\sim N(1, 0.1)|(0,)$, all As $\sim N(2, 0.2)|(0,)$, all νs $\sim N(1.5, 0.15)|(0,)$, all t_0s $\sim N(0.25, 0.025)|(0,)$, and all ιs $\sim N(2, 0.2)|(0,)$.

For each parameter to be estimated, we estimated a group level mean parameter and a group level standard deviation parameter. Priors for all group level mean parameters were normal distributions, with all b_μs $\sim N(1, 0.5)|(0,)$, all A_μs $\sim N(2, 1)|(0,)$, all v_μs $\sim N(1.5, 0.75)|(0,)$, all $t0_\mu$s $\sim N(0.25, 0.1)|(0,)$, and all ι_μs $\sim N(2, 1)|(0,)$. Priors for all group level standard deviation parameters were exponential distributions with mean values of 1. Starting values for the MCMC chains for group level μ parameters were drawn from the same distributions as those for the individual parameters, and starting values for group level σ parameters were derived from starting value distributions for the individual parameters by dividing the mean by 10 and the standard deviation by 2.

For sampling, we used 32 interacting Markov chains for all runs, and ran each for 1,000 burn-in iterations followed by 1,000 iterations after convergence. The two tuning parameters of the differential evolution proposal algorithm were set to standard values used in previous work: random permutations drawn uniformly from the interval [-.001, .001] were added to all proposals; and the scale of the difference added for proposal generation was set to $\gamma = 2.38 \times (2K)^{-0.5}$, where K is the number of parameters per participant. The MCMC chains blocked proposals separately for each participant's parameters, and also blocked the group–level parameters in $\{\mu, \sigma\}$ pairs.

Teodorescu Distributions

Starting values for the MCMC chains for individual parameters were drawn from the following distributions: $b \sim N(0.5, 0.05)|(0,)$, $A \sim N(1, 0.1)|(0,)$, $v_{easy} \sim N(1.4, 0.14)|(0,)$, $v_{difficult} \sim N(0.6, 0.06)|(0,)$, $t_0 \sim N(0.2, 0.02)|(0,)$, $\iota_{easy} \sim N(1, 0.1)|(0,)$, and $\iota_{difficult} \sim N(0.3, 0.03)|(0,)$.

Priors for all group level mean parameters were normal distributions, with $b_\mu \sim N(0.5, 0.2)|(0,)$, $A_\mu \sim N(1, 0.5)|(0,)$, $v_{\mu-easy} \sim N(1.4, 0.7)|(0,)$, $v_{\mu-difficult} \sim N(0.6, 0.3)|(0,)$, $t0_\mu \sim N(0.2, 0.1)|(0,)$, $\iota_{\mu-easy} \sim N(1, 0.5)|(0,)$, and $\iota_{\mu-difficult} \sim N(0.3, 0.15)|(0,)$. Priors for all group level standard deviation parameters exponential distributions with a mean of 1. Starting values for the MCMC chains for group level μ parameters were drawn from the same distributions as those for the individual parameters, and starting values for group level σ parameters were derived from starting value distributions for the individual parameters by dividing the mean by 10 and the standard deviation by 2.

For sampling, we used 32 interacting Markov chains for all runs, and ran each for 1,000 burn-in iterations followed by 1,000 iterations after convergence. The two tuning parameters of the differential evolution proposal algorithm were set to standard values used in previous work: random perturbations were added to all proposals drawn uniformly from the interval [-.001, .001]; and the scale of the difference added for proposal generation was set to $\gamma = 2.38 \times (2K)^{-0.5}$, where K is the number of parameters per participant. The MCMC chains blocked proposals separately for each participant's parameters, and also blocked the group–level parameters in $\{\mu, \sigma\}$ pairs.

References

Blavatskyy, P. (2012). Probabilistic choice and stochastic dominance. *Economic Theory*, *50*, 59 – 83.

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*, *362*, 1655–1670.

Britten, K. H., Shadlen, M. N., Newsome, W. T., & Movshon, J. A. (1992). The analysis of visual motion: A comparison of neuronal and psychophysical performance. *Journal of Neuroscience*, *12*, 4745–4765.

Brown, S. D., & Heathcote, A. (2005). A ballistic model of choice response time. *Psychological Review*, *112*, 117–128.

Brown, S. D., & Heathcote, A. (2008). The simplest complete model of choice reaction time: Linear ballistic accumulation. *Cognitive Psychology*, *57*, 153–178.

Brown, S. D., Marley, A. A. J., Donkin, C., & Heathcote, A. (2008). An integrated model of choices and response times in absolute identification. *Psychological Review*, *115*(2), 396–425.

Brown, S. D., Steyvers, M., & Wagenmakers, E.-J. (2009). Observing evidence accumulation during multi-alternative decisions. *Journal of Mathematical Psychology*, *53*(6), 453–462.

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.

Eidels, A., Donkin, C., Brown, S. D., & Heathcote, A. (2010, December). Converging measures of workload capacity. *Psychonomic bulletin & review*, *17*(6), 763–771.

Gamerman, D., & Lopes, H. F. (2006). *Markov chain Monte Carlo: Stochastic simulation for Bayesian inference.* Boca Raton, FL: Chapman & Hall/CRC.

Gilks, W. R., Richardson, S., & Spiegelhalter, D. J. (Eds.). (1996). *Markov chain Monte Carlo in practice.* Boca Raton (FL): Chapman & Hall/CRC.

Hawkins, G. E., Brown, S. D., Steyvers, M., & Wagenmakers, E.-J. (2012a). Context effects in multi–alternative decision making: Empirical data and a Bayesian model. *Cognitive Science*, *36*, 498–516.

Hawkins, G. E., Brown, S. D., Steyvers, M., & Wagenmakers, E.-J. (2012b, February). An optimal adjustment procedure to minimize experiment time in decisions with multiple alternatives. *Psychonomic Bulletin & Review*, *19*(2), 339–348.

Hawkins, G. E., Marley, A. A. J., Heathcote, A., Flynn, T. N., Louviere, J. J., & Brown, S. D. (2014). Integrating cognitive process and descriptive models of attitudes and preferences. *Cognitive Science*, *38*, 701–735.

Heathcote, A., Brown, S. D., & Wagenmakers, E.-J. (2014). An introduction to good practices in cognitive modeling. In B. Forstmann & E.-J. Wagenmakers (Eds.), *An introduction to model-based cognitive neuroscience* (pp. 25–48). New York: Springer.

Hick, W. E. (1952). On the rate of gain of information. *Quarterly Journal of Experimental Psychology*, *4*, 11–26.

Hyman, R. (1953). Stimulus information as a determinant of reaction time. *Journal of Experimental Psychology*, *45*, 188–196.

Luce, R. D. (1986). *Response times.* New York: Oxford University Press.

Marley, A. A. J. (1991). Context dependent probabilistic choice models based on measures of binary advantage. *Mathematical Social Sciences*, *21*, 201–231.

McClelland, J. L., Usher, M., & Tsetsos, K. (2011). Testing multi–alternative decision models with non–stationary evidence. *Frontiers in Neuroscience*, *5*, 1662–4548.

McMillen, T., & Holmes, P. (2006). The dynamics of choice among multiple alternatives. *Journal of Mathematical Psychology*, *50*, 30–57.

Miletic, S., Turner, B. M., Forstmann, B., & Van Maanen, L. (2016). Parameter Recovery for the Leaky Competing Accumulator Model. Manuscript submitted for publication.

Raab, D. H. (1962). Statistical facilitation of simple reaction times. *Transactions of the New York Academy of Sciences*, *24*, 574–590.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*, 59–108.

Spiegelhalter, D. J., Best, N. G., Carlin, B. P., & van der Linde, A. (2002). Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society B*, *64*, 583–639.

Teichner, W. H., & Krebs, M. J. (1974). Laws of visual choice reaction time. *Psychological Review*, *81*, 75–98.

Teodorescu, A. R., & Usher, M. (2013). Disentangling decision models: From independence to competition. *Psychological Review*, *120*, 1–38.

Terry, A., Marley, A. A. J., Barnwal, A., Wagenmakers, E.-J., Heathcote, A., & Brown, S. D. (2015). Generalising the drift rate distributions for Linear Ballistic Accumulators. Manuscript submitted for publication. *Journal of Mathematical Psychology*, *68*, 49–58.

Trueblood, J. S., Brown, S. D., & Heathcote, A. (2014). The multiattribute linear ballistic accumulator model of context effects in multialternative choice. *Psychological Review*, *121*, 179–205.

Turner, B. M., Sederberg, P. B., Brown, S. D., & Steyvers, M. (2013). A method for efficiently sampling from distributions with correlated dimensions. *Psychological Methods*, *18*, 368–384.

Usher, M., & McClelland, J. L. (2001). On the time course of perceptual choice: The leaky competing accumulator model. *Psychological Review*, *108*, 550–592.

Usher, M., Olami, Z., & McClelland, J. L. (2002, December). Hick's Law in a Stochastic Race Model with Speed–Accuracy Tradeoff. *Journal of Mathematical Psychology*, *46*(6), 704–715.

van Maanen, L., Grasman, R. P. P. P., Forstmann, B. U., Keuken, M. C., Brown, S. D., & Wagenmakers, E.-J. (2012). Similarity and number of alternatives in the random–dot motion paradigm. *Attention, Perception, & Psychophysics*, *74*, 739–753.

van Ravenzwaaij, D., Cassey, P., & Brown, S. D. (in press). A simple introduction to Markov Chain Monte–Carlo. *Psychonomic Bulletin & Review*.

van Ravenzwaaij, D., van der Maas, H. L. J., & Wagenmakers, E.-J. (2012). Optimal decision making in neural inhibition models. *Psychological Review*, *119*, 201–215.

Verdonck, S., & Tuerlinckx, F. (2014). The Ising Decision Maker: A binary stochastic network for choice response time. *Psychological Review*, *121*, 422–462.

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.

Wickelgren, W. A. (1977). Speed–accuracy tradeoff and information processing dynamics. *Acta Psychologica*, *41*, 67–85.