

Accumulating Advantages: A New Conceptualization of Rapid Multiple Choice

Don van Ravenzwaaij^{1,2}, Scott D. Brown¹, A. A. J. Marley^{3,4}, and Andrew Heathcote⁵

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

Independent racing evidence-accumulator models have proven fruitful in advancing understanding of rapid decisions, mainly in the case of binary choice, where they can be relatively easily estimated and are known to account for a range of benchmark phenomena. Typically, such models assume a one-to-one mapping between accumulators and responses. We explore an alternative independent-race framework where more than one accumulator can be associated with each response, and where a response is triggered when a sufficient number of accumulators associated with that response reach their thresholds. Each accumulator is primarily driven by the difference in evidence supporting one vs. another response (i.e., that response’s “advantage”), with secondary inputs corresponding to the total evidence for both responses and a constant term. We use Brown and Heathcote’s (2008) LBA to instantiate the framework in a mathematically tractable measurement model (i.e., a model whose parameters can be successfully recovered from data). We show this “Advantage LBA” model provides a detailed quantitative account of a variety of benchmark binary and multiple choice phenomena that traditional independent accumulator models struggle with; in binary choice the effects of additive versus multiplicative changes to input values, and in multiple choice the effects of manipulations of the strength of lure (i.e., non-target) stimuli and Hick’s Law. We conclude that the Advantage LBA provides a tractable new avenue for understanding the dynamics of decisions among multiple choices.

Keywords: Evidence accumulation models, RT tasks, Hick’s Law, lateral inhibition, max-next.

In everyday life, we are constantly confronted with tasks that require choosing one among many options. These decisions often become more difficult as the number of alternatives increase, leading to slowed response time (RT) and decreases in choice accuracy. It is attractive to model the dynamics of such multiple-choice decisions with racing evidence-accumulation processes as such models can be applied to choosing among any number of options by simply allocating one accumulator to each option. These models assume that once the relevant information is perceptually encoded and/or extracted from memory, each accumulator accrues evidence favoring its option. The first accumulator to satisfy a stopping rule (e.g. a threshold on its evidence total) leads to the response with which it is associated. Notable recent examples include the leaky competing accumulator (LCA; Usher & McClelland, 2001), the “max-next” (Brown, Steyvers, & Wagenmakers, 2009; McClelland, Usher, & Tsetsos, 2011; McMillen & Holmes, 2006), the ballistic accumulator (Brown & Heathcote, 2005), and the linear ballistic accumulator (LBA; Brown & Heathcote, 2008). The LBA model differs from the others in that the “stopping rule” which determines when a response is chosen depends only on whether accumulated evidence has exceeded a threshold, and in that accumulation is independent. Throughout this manuscript, we use independence to refer to the relationship among accumulators during accumulation (see Discussion for discussion of this and alternative definitions). These assumptions make it functionally and computationally simple, mathematically tractable, and easily extended to more complex decision paradigms (e.g., Eidels, Donkin, Brown, & Heathcote, 2010; Holmes, Trueblood, & Heathcote, 2016; Trueblood, Brown, & Heathcote, 2014).

A typical multiple-choice experiment either: 1) presents one of N possible stimuli on each trial, with each stimulus associated with a single correct response (e.g., Lacouture & Marley, 1995; Leite & Ratcliff, 2010; Pachella & Fisher, 1972), or 2) simultaneously presents N stimuli on each trial, again with each stimulus associated with a single correct response (e.g., Brown et al., 2009; Dassonville, Lewis, Foster, & Ashe, 1999; Kveraga, Boucher, & Hughes, 2002; K.-M. Lee, Keller, & Heinen, 2005; ten Hoopen, Akerboom, & Raaymakers, 1982; Vickrey & Neuringer, 2000). A well-known problem in the application of independent racing accumulator models to both multiple-choice paradigms is that the conventional one-to-one mapping between stimuli and accumulators leads to *faster* decisions with more accumulators (i.e., “statistical facilitation”, Raab, 1962), whereas in practice decisions slow down. One proposed solution is to relax the assumption that accumulation is independent, as occurs in the LCA via lateral inhibitory interactions. Another is a stopping rule that depends on the moment-to-moment evidence totals in more than one accumulator, as occurs in the max-next model through requiring a minimum difference between the largest and second largest evidence totals to initiate a response. A third solution involves adjusting response thresholds for increasing number of choices to counteract the increase in response time. Here we explore an alternative framework, which applies to both types of multiple-choice paradigm, and which maintains independence in accumulation, but relaxes the assumption that each response is represented by only one accumulator. The stopping rule for the framework we propose can depend on more than one accumulator, but only through threshold-crossing events and not through the evidence levels in each accumulator. This makes the framework mathematically tractable.

In our proposed framework the rate of evidence accumulation for each unit is primarily based on relative rather than absolute inputs (see Marley, 1991 and Tversky & Simonson,

1993, for relative evidence models of choice probabilities, and Usher & McClelland, 2004 and Trueblood et al., 2014, for relative advantage models of both RT and choice). Specifically, we propose that alternatives are evaluated in pairs, so, when there are more than two alternatives, more than one accumulator is associated with each response. The input for each accumulator is a weighed sum of: (1) the difference or “advantage” in evidence for the alternative associated with the accumulator over the other alternative, (2) the total evidence for both alternatives, and (3) a bias term (see Blavatsky, 2012, for a related formulation). Because our fits to data show the first term has the dominant effect we describe this as an *advantage input* scheme. We explore the mathematically tractable situation where these pairwise comparisons run independently and in parallel. When this parallel independent race model is instantiated using linear ballistic accumulators, as we do here, we call the resulting model the “Advantage LBA” (ALBA).

One new contribution of our modelling framework is the idea that a response option may be associated with more than one accumulator. This occurs when there are more than two response options, whereas when there are only two options a one-to-one mapping applies. Thus, we are able to release independent accumulator models of multiple-alternative choice from the traditional one-to-one mapping between accumulators and responses while remaining consistent with traditional approaches to binary choice. We first show that our advantage-input scheme enables good fits of the ALBA to data from a two-alternative forced-choice paradigm that have been problematic for independent racing accumulator models but consistent with dependent accumulation as instantiated in the LCA model (Teodorescu, Moran, & Usher, 2016, Experiment 1). We then extend the two-alternative ALBA to choices among more than two response alternatives, and demonstrate that it provides good fits to data from both types of multiple-choice paradigm that are problematic for existing independent-race models. The mathematical properties of the accumulators and input scheme in the multiple-alternative ALBA are identical to those in the two-alternative ALBA, but an extension to the idea of a stopping rule is required to account for the association of each response to more than one accumulator.

A second new contribution of our work is an exploration of stopping rules. In the main body of the paper we focus on a “Win-All” stopping rule, with details of alternative stopping rules reported in supplementary materials. We report fits of the Win-All ALBA to a task requiring choice among four simultaneously presented alternatives (Teodorescu & Usher, 2013, Experiment 1a) in which effects of the relative strengths of non-target (lure) response options were best fit by the max-next model, and which were taken to be incompatible with independent accumulation. We show that the Win-All ALBA, whose stopping rule is conceptually related to the max-next stopping rule, provides an accurate and detailed account of this data. We then extend the Win-All ALBA to address a data set that exemplifies a long-standing benchmark phenomenon for multiple-choice paradigms when assigning a single stimulus into one of many classes, Hick’s Law (van Maanen et al., 2012). Hick’s Law states that the mean RT and the logarithm of the number of choice alternatives are linearly related (Hick, 1952; Hyman, 1953). We demonstrate that the Win-All ALBA naturally provides an account of Hick’s Law.

In both of the applications of the ALBA to multiple-choice data, we show that all the parameters of the Win-All ALBA are identifiable by performing parameter-recovery simulations. These successful recoveries underline a significant improvement in the utility of

our approach for behavioral applications compared to non-independent race models. Non-independent models tend to be mathematically intractable, and so it is difficult to compute a key quantity required to fit them to data, their likelihood functions. Miletic, Turner, Forstmann, and Van Maanen (2017) explored a computationally intensive simulation-based method to obtain the LCA’s likelihood, but found that it was “extremely difficult to faithfully recover the parameters of the LCA model ” (p.25.). When parameter recovery is not possible it is difficult to interpret estimated parameter values as they may not be psychologically meaningful. Note that we are not implying that the parameters themselves are meaningless, only their estimates. Further, even if this is the case it does not mean that such models are of no use. Psychological questions can still be addressed through model selection techniques, as was shown with reference to the LCA by Evans, Holmes, and Trueblood (in press). Parameter recovery may also be possible for restricted versions of the LCA (see Miletic et al., 2017, for further discussion). For the Win-All ALBA, in contrast, we can safely interpret parameter estimates, and so we present and discuss them in each application, particularly highlighting the consistency of estimated weights for the sums and difference components of the advantage coding scheme that hold over the variety of paradigms we examine. In the next section we begin by defining this coding scheme for the simple binary-choice case.

Advantage-input coding for binary choice

The standard LBA model for binary choice (Brown & Heathcote, 2008) has two accumulators, each of which starts from an independently sampled and uniformly distributed point between 0 and $A_i > 0$, $i = 1, 2$, after which evidence is accumulated linearly for each response option if the stimulus input remains fixed throughout the trial. Each evidence accumulator has a drift rate d_i , and for each trial each drift rate is independently drawn from a normal distribution truncated at zero (Heathcote & Love, 2012), with means v_i , and standard deviations s_i .¹ Thresholds $b_i > A_i$ determine a speed–accuracy trade off; smaller values lead to faster decisions at the cost of a higher error rate. Sometimes the thresholds and/or maximum starting points are assumed to be the same for both accumulators, in which case the subscript can be dropped. Usually, rather than directly estimating the threshold the distance from the maximum starting point (A) to the response threshold (b), $B = b - A$, is estimated. This makes it easy to fulfill the assumption that an accumulator cannot start above its threshold (i.e, $b > A$) by enforcing $B > 0$. Manipulations affecting the a-priori plausibility of responses (say, a cue that predicts the correct response 80% of the time, see Teodorescu & Usher, 2013) can be expected to elevate the mean starting point of the compatible stimulus and/or depress the mean starting point of the incompatible stimulus. This is equivalent to an equal but opposite effect on the threshold in terms of response time and probability (Heathcote, Holloway, & Sauer, in press).

Together, the accumulator (A and B) and input (v and s) parameters define a distribution of decision times, DT . Response times also include the time taken for processes such as stimulus encoding and response production, which together make up the non-decision

¹The original 2008 model assumed an unbounded normal distribution. Other drift rate distributions also yield tractable models (see e.g. Terry et al., 2015), but most recent applications of the LBA assume a normal distribution truncated at zero.

time (Luce, 1986). We assume non-decision time is a constant, $t_0 \geq 0$, that shifts the distribution of DT such that $RT = DT + t_0$.

For binary choice based on perceptual properties, stimulus i has a physical value O_i , $i = 1, 2$, and these determine the drift rates for evidence accumulation. For example in Experiment 1 of Teodorescu et al. (2016), which we analyze with the ALBA, the luminance (in lumens) of the visual stimuli are linear with respect to a measure (vis., MATLAB RGB values) for which 0 (resp., 1) represents the minimum (resp., maximum) screen luminance. We assume that those objective values, in the interval $(0, 1)$, are logarithmically transformed to subjective brightness values, $S_i = \log(O_i)$ (Fechner, Boring, Howes, & Adler, 1966). The advantage-input rate for each accumulator is then an additive combination of the difference between the subjective brightness values, $S_1 - S_2$ (resp., $S_2 - S_1$), with weight w_D , and their sum, $S_1 + S_2$, with weight w_S , plus a bias parameter, $v_0 > 0$; see Equations 1 and 2 below.

To clearly differentiate this type of input scheme from that used in past applications of the standard LBA (where objective values and/or their mapping to subjective values were often not known and so rates were freely estimated) we denote the mean rate for the accumulator associated with the advantage of stimulus 1 over 2 (and hence also associated with a response favoring stimulus 1) as v_{1-2} , and similarly v_{2-1} for the other accumulator.

$$v_{1-2} = v_0 + w_D(S_1 - S_2) + w_S(S_1 + S_2) \quad (1)$$

$$v_{2-1} = v_0 + w_D(S_2 - S_1) + w_S(S_1 + S_2) \quad (2)$$

The bias parameter, v_0 , can take on values that ensure that each accumulator has a non-negative drift rate and hence eventually reaches its threshold, which in turn ensures that a response is made in finite time (for a similar mechanism see e.g., Bogacz, Usher, Zhang, & McClelland, 2007; Busemeyer, Townsend, Diederich, & Barkan, 2005; van Ravenzwaaij, van der Maas, & Wagenmakers, 2012). Different schemes for ensuring that some or all drift rates are positive are also possible, such as by taking the ratio rather than difference of positive subjective brightness values (e.g., Hawkins et al., 2014). These possibilities may have practical and conceptual advantages, but we leave their investigation to future work. The “difference weight”, w_D , is constrained to be non-negative and therefore the drift rate v_{1-2} (resp., v_{2-1}) increases (resp., decreases) as the brightness difference $S_1 - S_2$ increases. We constrain the “sum weight” w_S to non-negative values and therefore the drift rate increases with the overall magnitude of the pair.

We describe this as an “advantage” input coding scheme as typically $w_D \gg w_S$, and so the difference term dominates in determining the drift rate. A large difference effect makes sense as it means the rates favor the correct response. However, a non-zero sum term is also necessary in order to account for effects of the absolute strength of the stimuli. In the framing given by Teodorescu et al. (2016), whose work inspired this formulation and whose data we fit in the next section, these rates are partially absolute but mostly relative.

Each of v_0 , w_S , and w_D , are estimated from the data, and so the units used to measure the stimuli do not matter up to a linear transformation - i.e., the stimulus measures are interval scales. We assume a common variance, s , for the drift rate distribution of all advantage accumulators within a condition, and we assume that the inputs to the accumulators are uncorrelated. An illustration of the two-alternative ALBA for a brightness




Figure 1. The ALBA and its parameters for a two alternative brightness identification task. Evidence accumulation begins at a start point drawn randomly from a uniform distribution on the interval $[0, A]$. Evidence accumulation is governed by drift rates d_{1-2} and d_{2-1} , drawn across trials from a normal distribution with means v_{1-2} and v_{2-1} and standard deviation s , truncated to positive values. A response is given as soon as one accumulator reaches the threshold $b = A + B$. Observed RT is an additive combination of the time during which evidence is accumulated and non-decision time t_0 .

identification task with two response options is given in Figure 1. In the next section, we test this model by fitting data that test the relative influences of the sum and difference components of the inputs.

Absolute versus Relative Input

Teodorescu et al.’s (2016) Experiment 1 compared two-alternative forced choice of the brightest stimulus in a baseline condition with luminance values of $\{.4 \text{ vs. } .3\}$, against performance in an “additive boost” condition, in which luminance values were elevated through the addition of 0.2 to $\{.6 \text{ vs. } .5\}$, and a “multiplicative boost” condition, in which they were elevated through multiplication by 1.5 to $\{.6 \text{ vs. } .45\}$. The two boosts were chosen such that the correct stimuli have identical objective values (.6). As a result, the additive and multiplicative conditions differ only in the luminance values of their incorrect stimuli. Although the task required a judgment about relative brightness, the authors found that both accuracy and response time were also sensitive to the absolute values of luminance relative to the baseline condition, both when the absolute value of the difference in luminance between stimuli was the same as in the baseline condition (i.e., in the additive condition) and when the ratio of luminance was the same (i.e., in the multiplicative condition; see Teodorescu et al., 2016, Figure 1c, bottom panel, see also Figure 2 of this paper). The authors attributed this pattern either to non-independent accumulation of absolute values,

Table 1

Median parameter values, with a 95% credible intervals for two-alternative ALBA model fit to Teodorescu et al. (2016) Experiment 1. Rows correspond to participants (Pp), except the bottom row, which is the average of the corresponding values above).

Pp	A	B_R	t_0	v_0	s	w_D	w_S
1	1.38 (0.78, 2.12)	0.87 (0.80, 0.94)	0.22 (0.16, 0.27)	2.73 (2.01, 3.73)	1.39 (0.99, 1.92)	5.43 (3.52, 8.19)	0.76 (0.45, 1.25)
2	4.81 (3.40, 6.97)	1.18 (1.04, 1.37)	0.42 (0.39, 0.45)	6.56 (4.82, 9.10)	2.28 (1.65, 3.26)	11.14 (7.56, 17.18)	1.56 (0.97, 2.47)
3	1.27 (0.76, 1.86)	0.97 (0.91, 1.02)	0.17 (0.12, 0.22)	2.96 (2.24, 3.82)	1.11 (0.84, 1.46)	2.92 (2.03, 4.07)	0.33 (0.16, 0.57)
4	4.20 (3.07, 5.64)	0.90 (0.79, 1.03)	0.30 (0.28, 0.32)	7.72 (5.83, 10.21)	3.75 (2.87, 4.9)	14.31 (10.17, 20.08)	1.31 (0.60, 2.24)
5	1.87 (1.43, 2.65)	1.28 (1.19, 1.41)	0.12 (0.10, 0.17)	3.15 (2.63, 4.15)	1.26 (1.08, 1.68)	3.16 (2.52, 4.58)	0.61 (0.41, 0.95)
6	1.63 (1.19, 2.22)	1.07 (1.01, 1.14)	0.21 (0.13, 0.27)	2.05 (1.65, 2.58)	0.66 (0.51, 0.86)	2.03 (1.48, 2.84)	0.17 (0.07, 0.30)
7	2.62 (2.15, 3.16)	0.93 (0.85, 1.02)	0.10 (0.10, 0.12)	2.11 (1.81, 2.44)	0.77 (0.70, 0.85)	2.52 (2.19, 2.93)	0.35 (0.22, 0.50)
Mean	2.54 (1.82, 3.52)	1.03 (0.94, 1.13)	0.22 (0.18, 0.26)	3.90 (3.00, 5.15)	1.60 (1.23, 2.13)	5.93 (4.21, 8.55)	0.73 (0.41, 1.18)

due to lateral inhibition as in the LCA, or independent accumulation of differences with activation dependent processing noise. Here we show that the latter mechanism can be replaced in an independent accumulation model by allowing the sum of the subjective brightness values over stimuli to have a small effect on drift rates.

In Teodorescu et al.’s (2016) experiment, each participant performed 1,200 trials, 400 in each condition, with the conditions randomly intermixed. As described in the previous section, we assumed subjective brightness to be the logarithm of the luminance values and these subjective brightness values were entered into Equations 1 and 2 to calculate drift rates. Note that the logarithmic transformation means the baseline and multiplicative conditions have equal subjective differences, which are larger than the subjective difference for the additive condition, whereas the subjective sum increases from baseline to multiplicative to additive conditions. There are no parameters which are free to vary between conditions in the ALBA model for these data. Instead, the sum and difference values entirely account for condition effects, with the same seven estimated parameters applied to the objective brightness inputs from each condition: baseline drift rate (v_0), sum (w_S) and difference (w_D) weights, non-decision time (t_0), rate variability (s), start-point variability (A) and the right-response accumulator threshold (B_R). The left-response accumulator threshold was fixed at $B_L = 1$ to make the model identifiable (Donkin, Brown, & Heathcote, 2009) and different thresholds for each accumulator allowed for response bias.

Details of the estimation methods are given in subsection “B1. Estimation Details: Absolute versus Relative Input” in Appendix B. Table 1 reports posterior median parameter estimates. For all participants the difference component of the rates had a much higher weight than the sum component, on average by approximately an order of magnitude, but the sum component was non-negligible. This resulted in mean drift rates for the target advantage accumulator of 4.06, 4.65 and 4.1 for baseline, multiplicative and additive conditions, respectively, and 0.65, 1.24 and 1.94, respectively, for the lure advantage accumulator. The small sum component does not change the equal target-lure differences in subjective brightness for baseline and multiplicative conditions (both 3.41), with a much smaller difference in the additive condition (2.16) reflecting the smaller difference in subjective brightness. However, the sum component is sufficient to account for the small absolute effects in the data.

Figure 2 shows the model fits the data well, not only in terms of accuracy and average RT but also RT distribution. The ALBA parameter estimates are consistent with

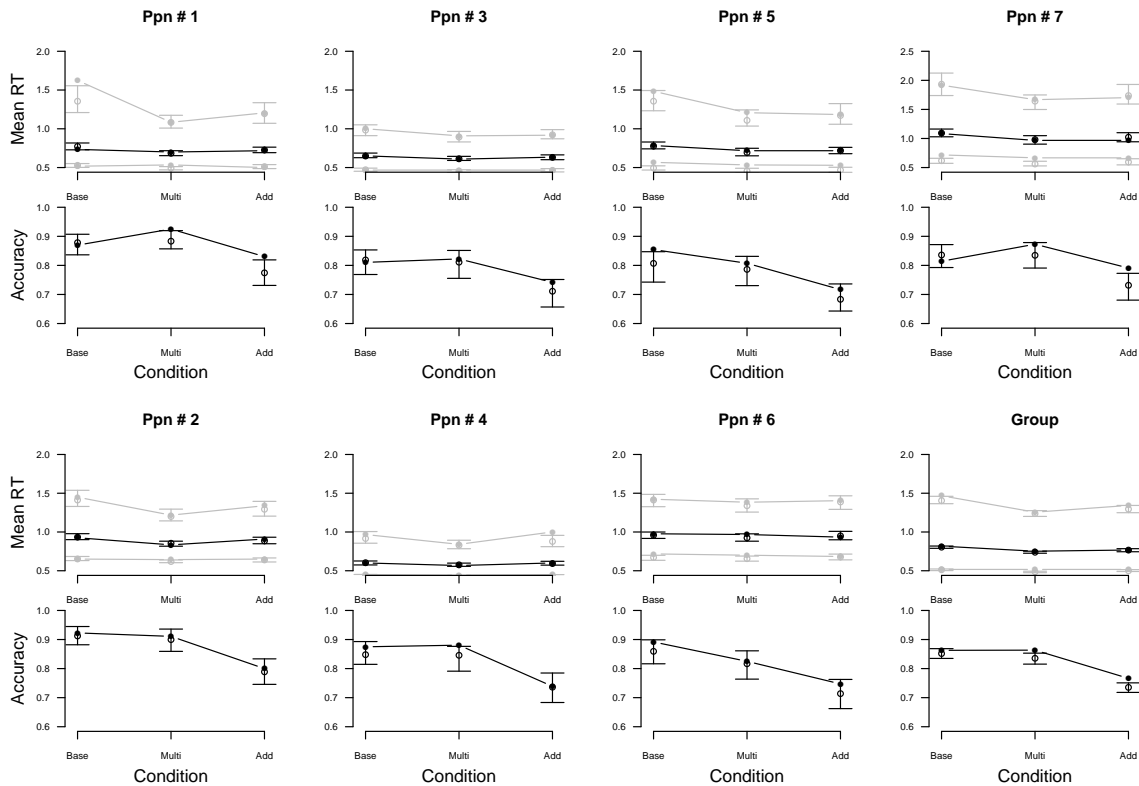


Figure 2. Posterior predictive data for fits to the Experiment 1 data of Teodorescu et al. (2016). RTs for the .5 (black), .1, and .9 (gray) deciles calculated for the baseline (Base), multiplicative (Mult), and additive (Add) conditions, and the proportion of correct responses for the respective conditions, both at the individual level (left three columns and top of right column) and for aggregate data (bottom right column). For all panels, error bars represent posterior predictive data simulated from model fits (the bar extends to the middle 95% of generated summary statistics, with the dot in the middle indicating the median) and lines represent data.

Teodorescu et al.’s (2016) conclusion that accumulation is partially absolute (the sum component of the ALBA) and partially relative (the difference component of the ALBA). Our model fit is at least as good as their fit with the LCA. In the next section, we show how ALBA can be generalized to multiple alternatives.

The multiple-alternative ALBA

The multiple-alternative ALBA maintains the same underlying type of accumulation as the two-alternative ALBA, but decisions are made when each of a pre-specified *set* of accumulators has crossed its threshold, as opposed to a single accumulator crossing a threshold (for a similar approach see Eidels et al., 2010). The combination stopping rules may be thought of as being realized by counters, with one counter for each possible response, although other conceptualizations are also possible (e.g., logic gates). Counts are incre-

mented 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 stimuli is the brightest: 1, 2, 3, or 4. For this decision, a standard accumulator model, such as the LBA, would assume a one-to-one mapping between accumulators and choices. This leads to four accumulators, which we denote 1, 2, 3, and 4 with corresponding drift rates $d(1)$, $d(2)$, $d(3)$, and $d(4)$. For the same decision, the ALBA has a total of 12 advantage accumulators, each taking as input a difference between the evidence values for an ordered pair of stimuli. We denote these accumulators: 1-2, 2-1, 1-3, 3-1, 1-4, 4-1, 2-3, 3-2, 2-4, 4-2, 3-4, and 4-3. In general for n responses there are $n(n-1)/2$ comparisons that can be made and hence $n(n-1)$ accumulators, half for comparisons in one direction and half for comparisons in the other direction (e.g., 1-2 and also 2-1).

Even though the ALBA model has more accumulators than the standard LBA model, all of the ALBA drift rates are produced from stimulus inputs via the same set of base parameters as in the two-choice example above. To illustrate, consider a trial on which stimulus 1 is brightest and the other stimuli, all less bright than stimulus 1, are equally bright to one another. In the traditional LBA, this stimulus set provides a strong “matching” subjective input value S_M to accumulator 1 and a smaller “mismatching” subjective input value of S_m to all other accumulators. In the corresponding case for the ALBA, each “matching” advantage accumulator (i.e., 1-2, 1-3, and 1-4, where the matching term is first) would have an advantage drift rate value of $v_0 + w_D(S_M - S_m) + w_S(S_M + S_m)$; each “mismatching” advantage accumulator (i.e., 2-1, 3-1, and 4-1, where the matching term is second) would have an advantage drift rate value of $v_0 + w_D(S_m - S_M) + w_S(S_M + S_m)$; and each of the remaining six “unrelated” accumulators (i.e., 2-3, 2-4, 3-2, 3-4, 4-2, and 4-3, where the matching term does not appear) would have an advantage drift rate value of $v_0 + w_S(S_m + S_M)$, as the difference term is zero. These values serve as the mean drift rates for their respective advantage accumulators.

Unless stated otherwise, we assume that the standard deviation for the drift rate distribution of all advantage accumulators is the same. We also assume that the inputs to all accumulators are uncorrelated. These assumptions correspond to the case where, on each trial for each accumulator, an independent random sample drawn from the same distribution is added to the mean drift rate of the accumulator.²

In summary, the mean drift rates for all advantage accumulators are determined by only three free parameters, the baseline rate, v_0 , and the sum, w_S , and difference, w_D , weights. 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 multiple-choice drift rates are constructed from differences). For our applications here, other standard LBA parameters A , B , and t_0 are assumed to be identical across ad-

²If the drift rate standard deviation was in part due to variability in each input, and that variability could differ between inputs, then only equality of drift rate standard deviation between advantage accumulators with the same inputs (e.g., 2-1 and 1-2) follows. That is, correlations would arise among accumulators that share inputs, which would make the model less mathematically tractable. However, systematic differences in rate variability across accumulators that are not a function of inputs do not affect tractability, and were implemented in some of the model fits of the Hick’s Law data set below.

vantage accumulators and are free parameters to be estimated from the data. However, situations likely exist where these restrictions must be relaxed. For instance, to accommodate response bias, different values of B could be allowed for the different sets of accumulators associated with each response. A lower value of B would make it quicker for accumulators in the set to finish, and hence bias responding towards the associated response.

With the details of the advantage accumulators established, the last thing is to determine a stopping rule: which (set of) accumulator(s) needs to finish before a response is initiated? Here, we focus on one stopping rule, which we call Win-All, that is conceptually closest to a max-next model. We investigated two other stopping rules, Lose-All, and Lose-One, both of which are discussed in Appendix A. Note that for the two-alternative case, all these stopping rules collapse to the same end-result, as there are only two advantage accumulators.

Win-All

The Win-All rule assumes that a response is made as soon as each of the accumulators associated with one of the response options has reached its threshold. For example, a Win-All rule will choose option 1 from $\{1,2,3\}$ if and only if:

1. Accumulators 1-2 and 1-3 have reached their thresholds, and:
2. At least one of the accumulators in each of the sets $\{2-1, 2-3\}$ and $\{3-1, 3-2\}$ has *not* reached its threshold.

Put simply, response option 1 is chosen if it is the first option to have beaten every other response option. This rule could be instantiated by linking each response with a counter having two inputs (e.g., from 1-2 and 1-3 for a 1 response) and requiring two counts to trigger its response. An illustration for a brightness identification task with three response options is given in Figure 3.

With the Win-All rule, it is mathematically possible for accumulator termination (i.e., threshold-crossing) sequences to occur which give rise to responses in a way that appears counter intuitive. For example, the termination sequence 2-1, 3-1, 1-2, and then 1-3 would result in choosing option 1, as it is the first option to have beaten all of its competitors. This may appear counter intuitive, because option 1 has also been beaten by each of its competitors. With reasonable parameter settings such sequences are exceedingly unlikely, because they would require opposite pairs to reach threshold close together in a sequence, which 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 with choice 1 at time t is:

$$p_1(t) = \sum_{I \neq 1} \left[PDF_{1-I}(t) \times \prod_{J \neq [1,I]} CDF_{1-J}(t) \right] \times \prod_{I \neq 1} \left[1 - \prod_{K \neq I} CDF_{I-K}(t) \right] \quad (3)$$

I is an option in the set $\{2, 3\}$, J is an option in the same set that is not I , and K is an option in the set $\{1, 2, 3\}$ that is not I . The cumulative distribution functions

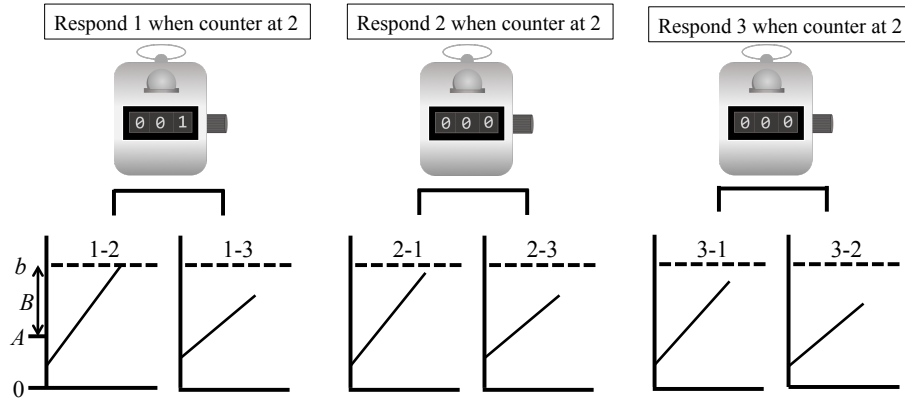


Figure 3. The Win-All version of ALBA for a three alternative task. Only the first counter to reach a count of 2 triggers a response.

(CDFs) and probability density functions (PDFs) are those of the standard LBA model (see Terry et al., 2015, also Appendix A). The derivation for Equation 3 may also be found in subsection “A1. Win-All Derivation” of Appendix A.

In the max-next model a decision is made as soon as the difference between the most active and the next most active accumulator exceeds a given threshold. The Win-All model is similar in that a response is made once the winning accumulator has beaten all of its competitors – that is, all relevant accumulators corresponding to pairwise comparisons have exceeded a given threshold. With this rule, the last advantage accumulator to cross its threshold will – on average – represent a contrast between the winner and the next best response option. The Win All ALBA and max-next models are also similar in terms of computational complexity, as for the latter model a full evaluation of the stopping rule must be made at each moment during accumulation. One possible serial algorithm for the max-next stopping rule involves first identifying the accumulator with the highest evidence total, then the one with the second highest, then comparing the difference to a threshold. A possible parallel algorithm could involve evaluating the same set of advantages (in this context differences in momentary evidence totals) as in the ALBA, with a response initiated when an accumulator has both the maximum advantage (and hence must have the maximum evidence total) and a minimum advantage greater than a threshold amount.

The max-next model does not have an easily computed likelihood, so requires the same simulation methods as the LCA to be fit to data in an optimal way, but its computational complexity, like that of the LCA (whose number of lateral inhibitory connections increase

with the square of n), makes that practically difficult as the number of options increases. In contrast, ALBA does have an easily computed likelihood, which makes it straightforward to fit data from choices among many options, as illustrated below in an application requiring choice among up to 9 options. Before reporting that application, we report fits of data from forced choices among four simultaneously presented options where the max-next and LCA models were preferred over independent accumulation (Teodorescu & Usher, 2013).

Strong versus weak lures

In Experiment 1A reported by Teodorescu and Usher (2013) participants made a forced choice about which of four patches was brightest. The key comparison was between trials that had one relatively attractive incorrect answer and two very unattractive incorrect answers (from here on, a “difficult trial”) and trials with a set of three relatively unattractive incorrect answers (from here on, an “easy trial”). Teodorescu and Usher theorized that, due to the comparatively elevated input of the attractive incorrect answer in the difficult trial, an independent race model will always predict a speed-up for correctly answered difficult trials compared to easy trials, due to statistical facilitation. In contrast, they found correct responses on difficult trials were actually slower than on easy trials.

Eight participants performed between 1,000 and 1,200 trials. Half of these trials constituted the easy condition with luminance values of $\{.4, .2, .2, .2\}$, respectively, for the target and three lures. The other half of the trials constituted the difficult condition with brightness values of $\{.4, .3, .15, .15\}$ (see Figure 4 for example stimuli, adapted from Teodorescu & Usher, 2013, Figure 4). Trials from the easy and difficult condition were randomly mixed within each block. In each condition, the sum of the brightness values is the same, so that normalizing these values by dividing them by the sum preserves the ratios between values, a feature which was used to rule out independent race models with sum-normalized feed-forward input competition. The ALBA is another kind of input-competition model, but with a different architecture and stopping rule.

As described in section “Rethinking stimulus inputs: the two-alternative ALBA”, we assume luminance values are log-transformed to obtain subjective brightness values. Advantage accumulators for each pair are dictated by Equations 1 and 2. Unfortunately, due to a computer error, the data for this experiment only recorded whether the response was correct or incorrect (Teodorescu, personal communication). As a result, in the case of an incorrect response it is unknown which of the incorrect options was chosen. To respect this, we aggregated the model’s log-likelihoods for all three error response options in our fits to the data.

We constrained parameters A , B , t_0 , v_0 , w_S , w_D , and s to be identical across the two conditions. We fixed the value of $s = 1$ and estimated the remaining 6 parameters.³ Details of the estimation methods are given in subsection “B2. Estimation Details: Strong versus Weak Distractors” in Appendix B. We confirmed the model was identifiable with a parameter-recovery study (for details see subsection “C1. Parameter Recovery Strong versus Weak Distractors” in Appendix C).

³We also fit a model that relaxed the assumption of equal rate variability for the easy and difficult condition, estimating it for one and fixing it to $s = 1$ for the other. Model fit did not qualitatively improve (see subsection “D1. Additional Fits Strong versus Weak Distractors” in Appendix D), so we report the more parsimonious model here.




Figure 4. Example stimuli from the easy condition (left) and the difficult condition (right). In the actual task, the numbers were not presented.

Parameter estimates for the Win-All ALBA fit can be found in Table 2. As with our fits to Teodorescu et al. (2016)’s binary choice data, the difference component of the rates had a higher weight than the sum component for all participants, again on average by approximately an order of magnitude. Taking the first participant as a representative example, the mean drift rates in the easy condition that follow from the median parameter estimates in the table are 3.1 for the target accumulator and -1.6 for the lure accumulators. Mean drift rate estimates for the difficult condition involving easy and hard lures bracket these values: for the target relative to the hard and easy lures 1.8 and 4.0, respectively, and for the hard and easy lures -0.1 and -2.6 respectively.

Posterior predictive data from the Win-All ALBA are compared against the observed data in Figure 5, showing that the model fits relatively well. The number of free parameters required to obtain this fit (six) is no larger than the number of free parameters (six or seven) in the several models that were fit by Teodorescu and Usher (2013). The model misfits accuracy for the difficult condition for some participants, although the aggregate posterior predictive data (right-most column) captures the data at least as well as the best (max-next) model reported by Teodorescu and Usher (2013).

To confirm the shortcomings of a conventional race model with one-to-one accumulator-to-response mapping, we also fit a regular LBA to the data. This LBA had one correct drift rate and three error drift rates for each of the two conditions, along with the standard A , B , and t_0 parameters. This parametrization makes the model very flexible, and includes any potential input competition or feedforward inhibition model as a special case. Despite this flexibility, it did not fit as well as the Win-All ALBA; in order to capture the pattern in RTs between the easy and difficult condition it somewhat over-predicts error rates in both conditions (see subsection “D1. Additional Fits Strong versus Weak Distractors” in Appendix D). We performed model selection using the Deviance Informa-

Table 2

Posterior median parameter values, with a 95% credible intervals for the Win-All ALBA model of Teodorescu and Usher (2013) Experiment 1A data. Rows correspond to participants (P_p), except the top row, which contains parameters of the group-level distributions (*Hyper*).

P_p	B	A	t_0	v_0	w_S	w_D
Hyper	0.18 (0.01, 0.53)	1.03 (0.21, 1.79)	0.51 (0.11, 0.64)	1.26 (0.43, 1.71)	0.17 (0.01, 0.90)	1.61 (0.45, 3.55)
1	0.05 (0.00, 0.19)	0.87 (0.74, 1.02)	0.66 (0.62, 0.69)	1.33 (0.80, 2.00)	0.23 (0.02, 0.59)	3.38 (2.86, 3.80)
2	0.09 (0.01, 0.32)	1.11 (0.89, 1.35)	0.68 (0.61, 0.72)	1.02 (0.32, 1.70)	0.31 (0.03, 0.78)	3.91 (3.25, 4.62)
3	0.01 (0.00, 2.13)	0.58 (0.15, 0.70)	0.64 (0.27, 0.65)	1.41 (0.91, 10.25)	0.17 (0.00, 3.26)	3.13 (1.32, 3.45)
4	0.11 (0.00, 0.37)	4.07 (3.39, 4.81)	0.48 (0.37, 0.56)	1.28 (0.86, 1.75)	0.09 (0.00, 0.34)	3.69 (3.30, 4.04)
5	0.39 (0.07, 0.75)	1.19 (1.00, 1.38)	0.58 (0.50, 0.65)	1.40 (0.91, 3.42)	0.24 (0.02, 1.49)	4.16 (3.64, 5.48)
6	0.55 (0.23, 1.21)	1.84 (1.52, 2.22)	0.63 (0.50, 0.70)	1.73 (1.24, 2.32)	0.05 (0.00, 0.26)	3.58 (3.10, 4.14)
7	0.35 (0.11, 0.93)	2.34 (1.98, 2.74)	0.60 (0.47, 0.67)	1.57 (1.08, 13.33)	0.21 (0.03, 4.43)	3.95 (1.77, 4.54)
8	0.00 (0.00, 0.07)	0.91 (0.78, 1.07)	0.64 (0.60, 0.66)	0.96 (0.12, 11.39)	0.47 (0.10, 4.15)	4.05 (2.02, 4.41)

tion Criterion (DIC; Spiegelhalter, Best, Carlin, & van der Linde, 2002), a measure that balances goodness-of-fit against model complexity. A smaller DIC for the ALBA (2854) than the LBA (3089) suggests it is the superior model. Aside from a better fit, the ALBA model is more parsimonious with six free parameters compared to eleven free parameters for the LBA. These results suggest that it is one-to-one assumption of the traditional LBA, rather than the way in which stimuli map to drift rates, that is problematic.

In this section, we have demonstrated that the ALBA model can account for the strong versus weak lure data. This result suggests instead of independence, it is the assumption of a one-to-one mapping of accumulators to responses and the associated response rule that is problematic for the class of input-competition models. Next, we turn to another challenging empirical pattern for a multiple-alternative accumulator model: Hick’s Law.

Hick’s Law

Hick’s Law is a long-standing benchmark result for multiple-alternative decisions (Hick, 1952; Teichner & Krebs, 1974). It states that the mean RT and the logarithm of the number of choice alternatives are approximately linearly related. A well-known problem with independent race models with a one-to-one accumulator to response mapping is that they produce the opposite trend to Hick’s Law, *faster* decisions with more accumulators, because of statistical facilitation (Raab, 1962). Usher, Olami, and McClelland (2002) note that competitive accumulation (i.e., lateral inhibition among accumulators) can produce increasing RT with the number of options (see also Usher & McClelland, 2001), but at least in the LCA they found this was not sufficient to quantitatively account for Hick’s Law. They then showed that both in the LCA and an independent racing accumulator model, Hick’s Law can be accommodated if evidence thresholds are increased with set size in order to compensate for a decrease in accuracy that otherwise occurs as the number of choices increases.

Like the LCA, the Win-All ALBA naturally predicts longer response times as number of options (n) increases. This is because at least $n - 1$ accumulators need to reach threshold before a decision can be triggered. Effectively this means decision time increases as the maximum of a set of random variables (the times for accumulators to each threshold), where the size of that set increases in proportion to n . Simulations with a range of different

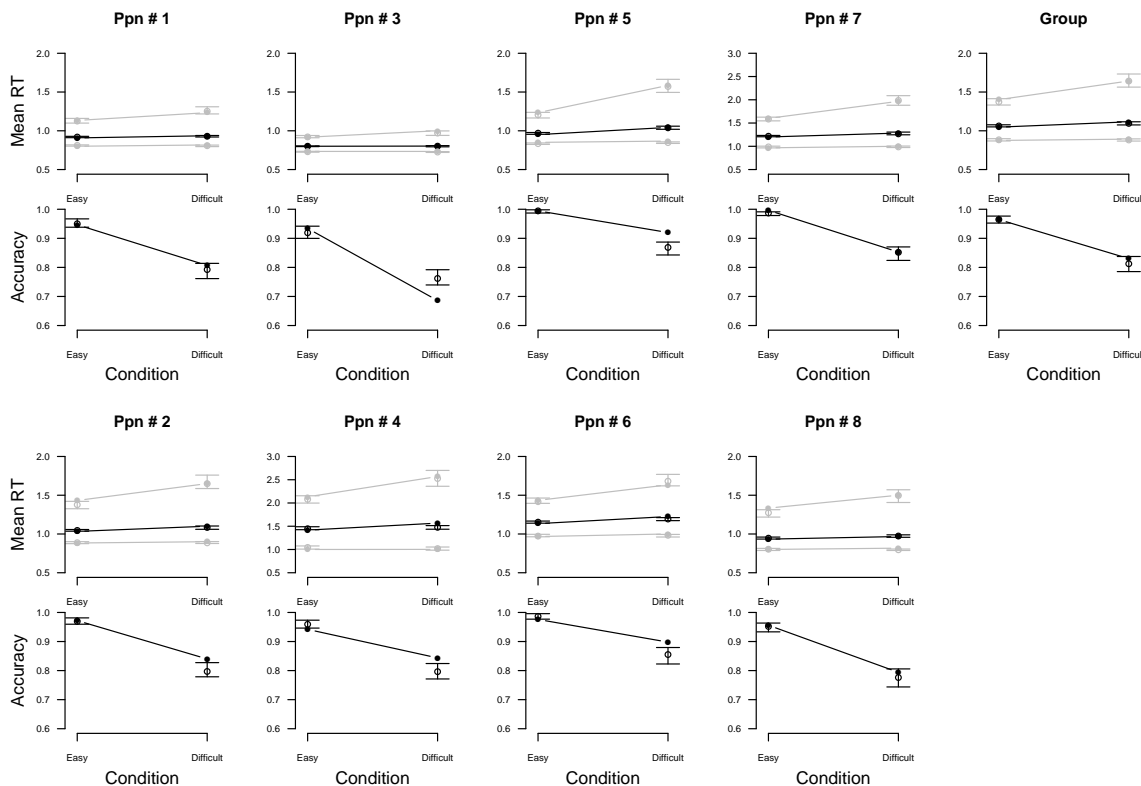


Figure 5. Posterior predictive data for fits to the Experiment 1A data of Teodorescu and Usher (2013). RTs for the .5 (black), .1, and .9 (gray) deciles calculated for the easy (top-left) and difficult (top-right) condition, and the proportion of correct responses for the easy (bottom-left) and difficult (bottom-right) condition, both at the individual level (left four columns) and for aggregate data (right column). For all panels, error bars represent posterior predictive data simulated from model fits (the bar extends to the middle 95% of generated summary statistics, with the dot in the middle indicating the median) and lines represent data. See text for details.

random variables indicate that this increase is approximately linear in the logarithm of n . However, the question remains whether the ALBA can quantitatively account for the fine details of RT and accuracy changes as a function of the number of response options due to this feature of its architecture alone, or whether evidence thresholds or other parameters also need to change with set size.

We took advantage of the tractability of the ALBA to directly fit 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, 2008a; Hawkins, Brown, Steyvers, & Wagenmakers, 2012a, 2012b).

van Maanen et al. (2012) had participants view displays consisting mostly of ran-

domly moving dots with a subset that move coherently in one direction (Britten, Shadlen, Newsome, & Movshon, 1992). Each trial involved either 3, 5, 7 or 9 directions, with the corresponding number of responses. There were eight blocks of trials, and within each block, trials were pseudo-randomized, such that no more than two consecutive trials had the same number of response options. In the “clustered” condition, which we address here, the angular spacing between adjacent stimulus directions was the same for all set-sizes, and hence the range of the stimulus directions increased with set-size, in an attempt to equate perceptual discriminability across set-sizes. All four conditions were administered within all five subjects, and there were 144 trials per condition.

Assume there are n stimuli, and therefore n responses matched to stimuli in a 1-to-1 fashion. Let stimulus k , $k \in \{1, \dots, n\}$, have subjective value s_k . In the experiment we consider, we assume the stimuli are subjectively equally spaced; that is, there is a subjective stimulus value s such $s_i - s_j = (i - j)s$ for all $i, j \in \{1, \dots, n\}$. We assume that the subject has a (referent) memory of the subjective value of each stimulus that is presented in the current task. Let $S_{j|i}$ denote the ‘strength’ of response j when stimulus i is presented. Then we assume $S_{j|i}$ has the form:⁴

$$S_{j|i} = \left(\frac{1}{1 + \frac{|s_j - s_i|}{s}} \right)^\alpha = \left(\frac{1}{1 + |j - i|} \right)^\alpha.$$

with a constant $\alpha > 0$. To provide some intuition about this function, consider the condition with five choice options, and a trial in which stimulus 2 is presented. For $\alpha = 1$, this leads to the set of input values $\{0.5, 1, 0.5, 0.33, 0.25\}$, reflecting the fact that nearby options are more plausible than options further removed. For $\alpha = \infty$, this leads to the set of input values $\{0, 1, 0, 0, 0\}$, reflecting no difference in the input values for competitors (i.e., no effect of proximity). Calculation of drift rates for each advantage accumulator followed a slightly modified version of Equations 1 and 2 to account for the fact that inputs depend on the angular distance from the correct response:

$$v_{j-k|i} = v_0 + w_D(S_{j|i} - S_{k|i}) + w_S(S_{j|i} + S_{k|i}) \quad (4)$$

$$v_{k-j|i} = v_0 + w_D(S_{k|i} - S_{j|i}) + w_S(S_{j|i} + S_{k|i}) \quad (5)$$

Details of model fitting can be found in the subsection “B3. Estimation Details: Hick’s Law” in Appendix B. We confirmed the model was identifiable with a parameter-recovery study (for details see subsection “C2. Parameter Recovery Hick’s Law” in Appendix C). In order to see if the Win-All ALBA naturally produces Hick’s Law we fit a model that constrained all parameters to be equal across set-size conditions (i.e., B , A , t_0 , v_0 , w_S , w_D , and α), with a fixed value of $s = 1$ (ALBA-1). Estimated parameters for the resulting model are given in Table 3. The pattern of weight parameters follows that found in earlier fits with the difference weight more than an order of magnitude greater than the sum weight. Although estimates of α are relatively small, mean rates change monotonically with the

⁴The presented form is for stimuli that are subjectively equally spaced and, as we see later, does not fit certain data for stimuli at, or near, the ends of the range of presented stimuli well. A complete theory, building on the current assumptions, might include a *rehearsal component*, similar to that in SAMBA (Brown et al., 2008a).

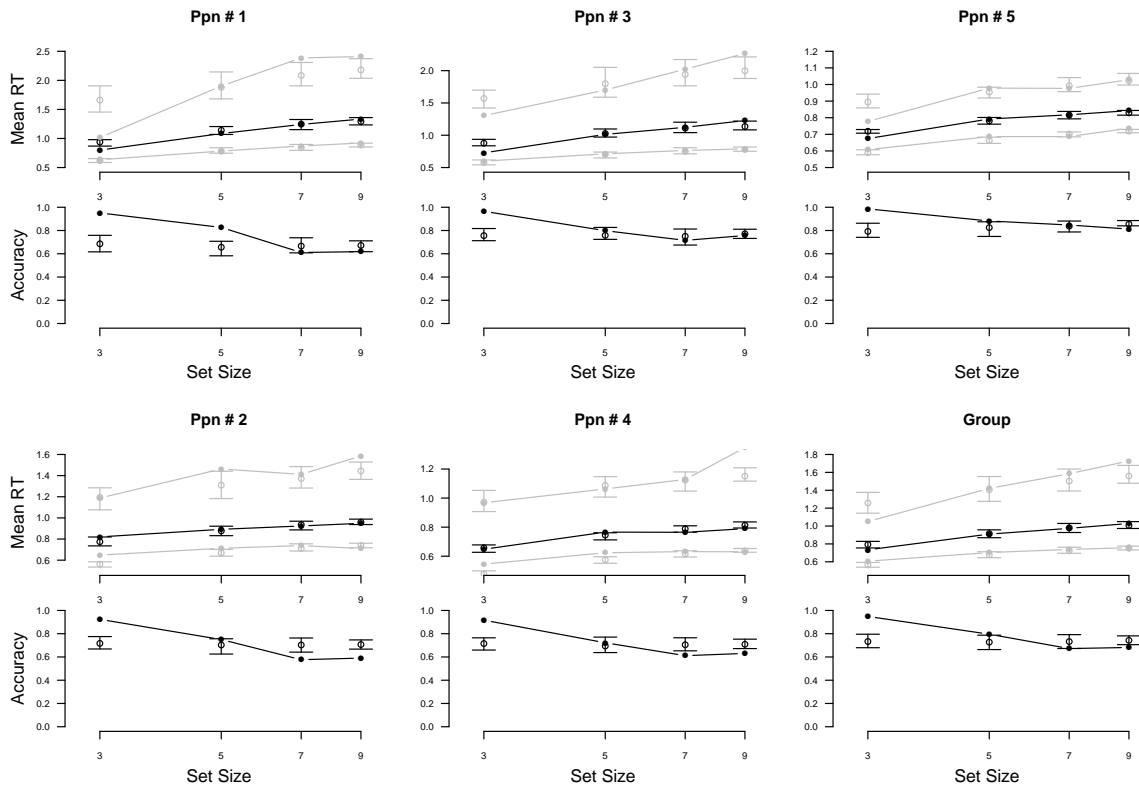


Figure 6. Posterior predictive data for the ALBA-1 fit to the van Maanen data. RTs for the .5 (black), .1, and .9 (gray) deciles (top) and the proportion of correct responses (bottom) as a function of set-size (N) on a logarithmic scale. Posterior predictives are presented at the individual level and for aggregate data (bottom-right panel). For all panels, box-and-whiskers represent posterior predictive data (the box contains 95% of the simulated data, with a bar across the middle indicating the median, and whiskers extend to the data extremes) and lines represent data. See text for details.

distance between inputs. For example, based on the median posterior parameter estimates for the first participant, mean rates for the accumulator associated with the advantage of the correct choice over options 1, 2, and 3 spaces removed are 1.1, 1.7, and 2.1, respectively. Similarly, for the advantage accumulator associated with choice options 1, 2, and 3 spaces removed from the correct choice mean rates were -0.9, -1.5, and -1.9 respectively. Estimates of A were quite large, indicating strong effects of factors like response biases due to carry-over effects from previous responses (Heathcote, Suraev, Curley, Gong, & Love, 2015), In comparison B estimates were small, although they were in most cases clearly greater than zero, indicating that participants exercised a small degree of response caution.

As shown in Figure 6, the model fit the median RT data well, consistent with the ALBA architecture accommodating the logarithmically increasing effect of set-size. It also fit effects on fast RTs, but did not fit the increase in error rates with set size and RTs in the slow tail of the distributions for higher set sizes. Given the misfit we explored models

Table 3

Estimated parameters of the ALBA-1 model for the van Maanen data set. Displayed are the median parameter values, with a 95% credible interval of the posterior presented in parentheses. Rows correspond to participants (Pp), except the top row, which contains parameters of the group-level distributions (Hyper).

Pp	B	A	t_0	v_0	w_S	w_D	α
Hyper	0.11 (0.01, 0.32)	1.18 (0.94, 1.45)	0.34 (0.19, 0.41)	0.20 (0.01, 0.80)	0.34 (0.04, 0.73)	10.00 (4.58, 15.17)	0.16 (0.07, 0.48)
1	0.11 (0.01, 0.32)	1.22 (1.02, 1.51)	0.33 (0.25, 0.39)	0.13 (0.01, 0.43)	0.13 (0.01, 0.34)	10.21 (3.95, 19.08)	0.14 (0.08, 0.40)
2	0.05 (0.00, 0.17)	1.14 (0.94, 1.34)	0.37 (0.32, 0.42)	0.32 (0.02, 0.86)	0.27 (0.02, 0.51)	16.56 (10.64, 26.53)	0.09 (0.05, 0.14)
3	0.14 (0.03, 0.38)	1.25 (1.03, 1.56)	0.30 (0.20, 0.36)	0.12 (0.01, 0.42)	0.11 (0.00, 0.36)	12.17 (4.03, 21.38)	0.15 (0.09, 0.44)
4	0.05 (0.00, 0.17)	1.03 (0.82, 1.22)	0.33 (0.27, 0.37)	0.28 (0.02, 0.99)	0.42 (0.06, 0.65)	14.88 (9.89, 23.64)	0.09 (0.06, 0.14)
5	0.19 (0.06, 0.40)	1.19 (1.02, 1.44)	0.43 (0.37, 0.46)	0.20 (0.01, 1.38)	1.12 (0.27, 1.42)	4.49 (3.02, 15.19)	0.38 (0.12, 0.58)

that allowed selected parameters to change with set size n . These analyses demonstrate how the ALBA's easily computed likelihood makes it practical to fit and evaluate a range of alternative model parameterizations. DIC values for all models can be found in Table 4. The table also reports the two components of DIC, one of which quantifies the model misfit, and the other that determines the penalty for model complexity.

Following Usher et al. (2002), we first examined a model that allowed thresholds to vary across set-size conditions. Varying threshold with set size could occur because set-size was manipulated between blocks of trials so participants could implement a trade off between speed and accuracy (model ALBA-4B). Although there were small improvements in both DIC and the account of accuracy effects, there was still clear misfit (see Appendix D, Figure D3).

Table 4

DIC summed over participants for ALBA fits to the van Maanen et al. (2012) data. Parameter(s) varying with set-size, ALBA-1: None, ALBA-4B: B, ALBA-4A: A, ALBA-4BA: both B and A, ALBA- β : s, ALBA- β 4B: both s and B. # Pars = number of free parameters per participant for all four conditions; Misfit = -2 times the likelihood of the median parameter estimate; Complexity = -4 times the median likelihood of the overall model + 4 times the likelihood of the median parameter estimate; DIC = Misfit + Complexity

Model	# Pars	Misfit	Complexity	DIC
ALBA-1	7	6439	36	6475
ALBA-4B	10	6399	-15	6384
ALBA-4A	10	6303	56	6359
ALBA-4BA	13	6020	66	6085
ALBA- β	8	6045	-8	6037
ALBA- β 4B	11	5925	17	5943
LBA	17	33213	71	33355

As parameters determining the level of trial-to-trial variability (i.e., start-point noise, A , and rate variability, s) affect error rates, it seems likely that this aspect of the misfit might be addressed by allowing one or more of these parameters to change with set size. We first considered changes in A . Start-point noise is usually attributed to factors like response

Table 5

Estimated parameters of the ALBA- β model for the van Maanen data set. Displayed are the median parameter values, with a 95% credible interval of the posterior presented in parentheses. Rows correspond to participants (Pp), except the top row, which contains parameters of the group-level distributions (Hyper).

Pp	B	A	t_0	v_0	w_S	w_D	α	β
Hyper	0.01 (0.00, 0.05)	1.76 (0.84, 2.48)	0.41 (0.28, 0.46)	1.21 (0.35, 1.80)	0.20 (0.01, 0.61)	9.08 (4.67, 13.37)	0.32 (0.10, 1.00)	0.07 (0.01, 0.13)
1	0.01 (0.00, 0.04)	2.61 (2.05, 3.31)	0.42 (0.37, 0.46)	1.87 (1.18, 2.48)	0.08 (0.00, 0.31)	3.30 (2.79, 5.73)	1.32 (0.41, 1.70)	0.13 (0.07, 0.17)
2	0.01 (0.00, 0.05)	1.48 (1.23, 1.80)	0.42 (0.39, 0.44)	1.19 (0.43, 1.58)	0.12 (0.01, 0.51)	14.49 (7.44, 25.39)	0.13 (0.07, 0.26)	0.05 (0.04, 0.06)
3	0.01 (0.00, 0.07)	1.79 (1.46, 2.21)	0.42 (0.38, 0.45)	0.92 (0.17, 1.40)	0.15 (0.01, 0.56)	11.60 (4.16, 20.97)	0.21 (0.11, 0.68)	0.10 (0.07, 0.12)
4	0.01 (0.00, 0.03)	1.21 (1.02, 1.45)	0.39 (0.37, 0.41)	1.28 (0.50, 1.73)	0.16 (0.01, 0.57)	13.30 (8.01, 24.51)	0.15 (0.08, 0.26)	0.06 (0.05, 0.08)
5	0.01 (0.00, 0.10)	1.80 (1.43, 2.28)	0.49 (0.46, 0.51)	1.78 (0.45, 3.03)	0.57 (0.03, 1.38)	10.80 (6.12, 20.18)	0.30 (0.16, 0.56)	0.03 (0.02, 0.04)

biases due to carry-over effects from previous responses (Heathcote et al., 2015) and so could plausibly vary with the number of responses. We allowed a different value of A for every set-size (ALBA-4A), with all other parameters constrained to be equal across set-size conditions. However, although DIC and the fit were again slightly improved, substantial misfit was still evident (see Appendix D, Figure D4).

We next considered rate variability (s), and inspired by the work of Ratcliff, Voskuilen, and Teodorescu (2018), we fit a model (ALBA- β) that assumed it increased linearly with the set size and in proportion to the mean rate. This was achieved by estimating one additional free (slope) parameter, β , where $s_n = 1 + \beta \times (n - 3) \times v$. This equation fixes $s = 1$ for the smallest set-size ($N=3$), which makes the model identifiable. We bounded the value of s_n below by 0.01 to enforce the necessary non-negativity of a standard deviation. Note that a more complex model with a different value of s estimated for each set size did not fit much better. Again all other parameters were constrained to be equal across set-size conditions. Despite requiring the estimation of only one extra parameter, there was a very substantial reduction in misfit and improvement in DIC. As shown in Figure 7, this model produced a good fit to almost all aspects of the data, including the decrease in accuracy with increasing set-size, with only accuracy for set-size 3 being underestimated.

Estimated parameters for the ALBA- β model are given in Table 5. Estimates of β were positive for all participants, which forced drift rate variability (s) to increase with mean drift rate, although the increase was modest. Overall, mean rates were more extreme than those for the baseline (ALBA-1) model. For example, based on the median parameter estimates for the first participant, mean rates were 4.8, 6.5, and 7.8 for the correct choice over options 1, 2, and 3 spaces removed, respectively, and -1.2, -3.0, and -4.3 for choice options 1, 2, and 3 spaces removed relative to the correct choice. This occurred because s estimates were greater than $s = 1$ for larger set sizes, which leads to more errors. For set-size 5, for example, s values associated with the correct choice over options 1, 2, and 3 spaces removed were 2.3, 2.7, and 3.0, respectively, although this was partially compensated for by decreased variability for choice options 1, 2, and 3 spaces removed from the correct choice, with values of 0.7, 0.2, and 0.01, respectively. The other parameter values shared with the ALBA-1 model were similar, except that α was larger, producing a shallower decrease in rates with distance from the stimulus direction, and B was close to zero, indicating that participants exercised minimal response caution.

Finally, we examined two models that allowed threshold (B) to vary with set size in addition to a between-trial variability parameter. For the case where A also varies with

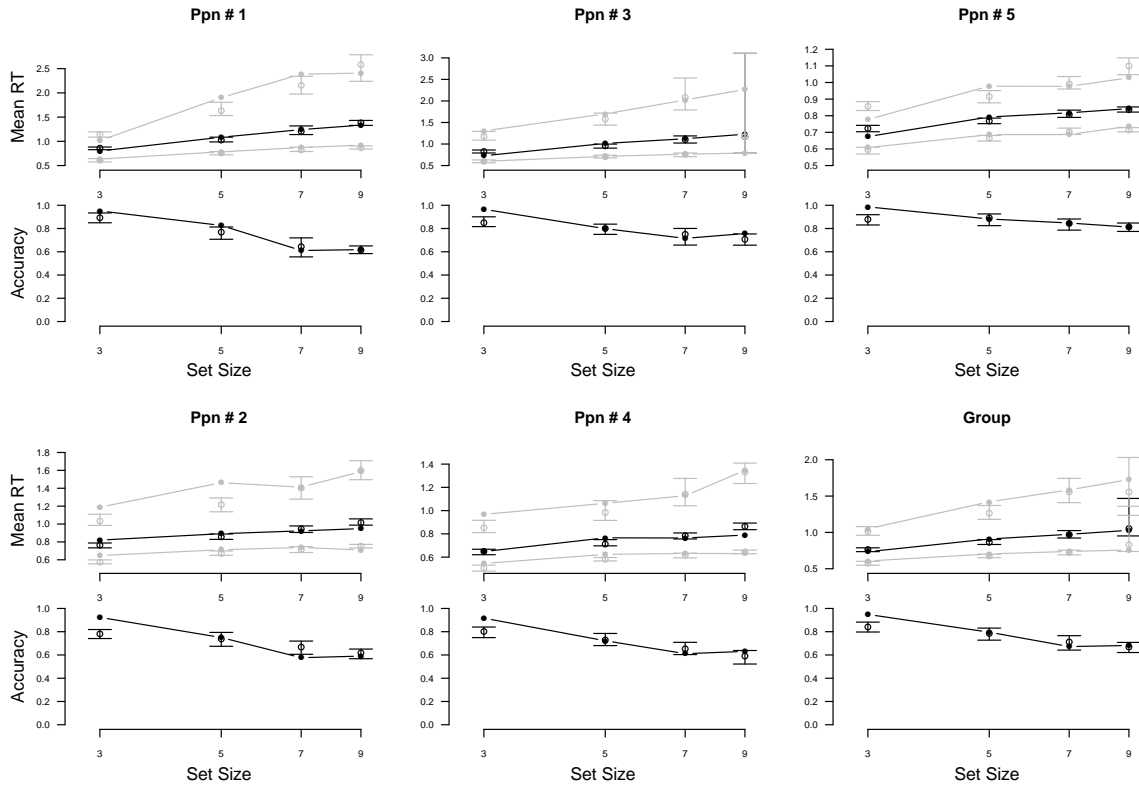


Figure 7. Posterior predictive data for the ALBA- β fit to the van Maanen data. RTs for the .5 (black), .1, and .9 (gray) deciles (top) and the proportion of correct responses (bottom) as a function of set-size (N) on a logarithmic scale. Posterior predictives are presented at the individual level and for aggregate data (bottom-right panel). For all panels, box-and-whiskers represent posterior predictive data (the box contains 95% of the simulated data, with a bar across the middle indicating the median, and whiskers extend to the data extremes) and lines represent data. See text for details.

set size (model ALBA-4BA) there was a very large improvement in DIC, although this was still not sufficient to be selected over the much simpler ALBA- β model. The ALBA-4BA model also under-predicts accuracy for the smallest set size and over-predicts accuracy for the two largest set-sizes (Appendix D, Figure D5).

The case where B and s vary with set size (model ALBA- β 4B) produced the lowest DIC of any model in Table 4 but the improvement compared to the ALBA- β model was modest. Figure 8 shows that accuracy for set-size 3 is now captured slightly better than the ALBA- β model, but is still somewhat under-predicted. Estimated parameters for the ALBA- β 4B model are given in Table 6. Most parameters shared with the ALBA- β follow a similar pattern. The B parameters for the ALBA- β 4B generally decrease as set size increases, starting at values similar to the ALBA-1 model for smaller set sizes with the values for $n = 9$ being similar to the single estimate for the ALBA- β at close to zero, indicating a very low level of response caution.

Finally, we also fit a standard LBA model, in which we let A , B , v_c (corresponding to mean drift rate matching the correct direction), and v_e (corresponding to mean drift rate not matching the correct direction) all vary freely with set size, but constrained t_0 to be equal across set size. Despite its complexity, this model, with 17 free parameters, failed to fit the data satisfactorily, because it overestimated the increase in error rate for increasing set sizes (for details, see Appendix D). Note that the fits presented in van Maanen et al. (2012) were based on even larger numbers of free parameters.

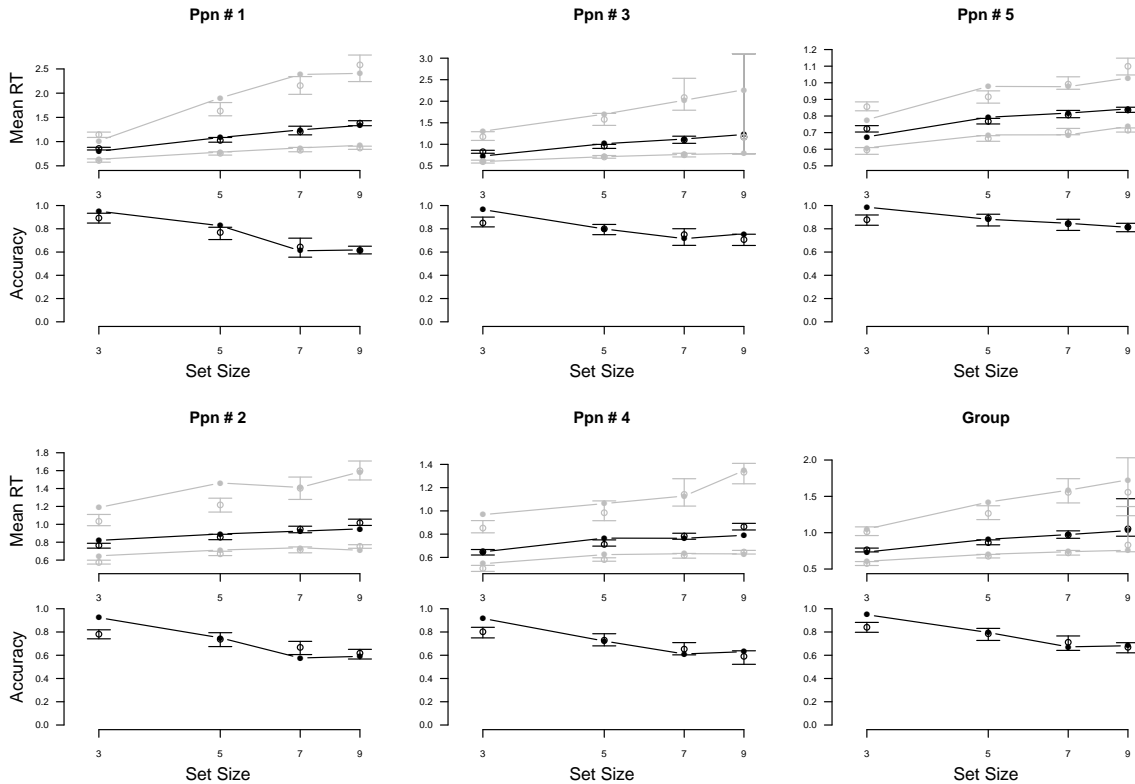


Figure 8. Posterior predictive data for the ALBA- β 4B fit to the van Maanen data. RTs for the .5 (black), .1, and .9 (gray) deciles (top) and the proportion of correct responses (bottom) as a function of set-size (N) on a logarithmic scale. Posterior predictives are presented at the individual level and for aggregate data (bottom-right panel). For all panels, box-and-whiskers represent posterior predictive data (the box contains 95% of the simulated data, with a bar across the middle indicating the median, and whiskers extend to the data extremes) and lines represent data. See text for details.

In summary, these analyses clearly show that the Win-All ALBA naturally predicts Hicks Law in terms of the central tendency of RT, and is able to capture most fine-grained effects of set size not only on the full distribution of RT, but also on accuracy, at least when some of its parameters are allowed to change with set size in reasonable ways. In the data set examined here (van Maanen et al., 2012) there was strong support for a parsimonious account in terms of a linear effect of set size on a proportional relationship between the

Table 6

Estimated parameters of the ALBA- β_4B model for the van Maanen data set. Displayed are the median parameter values, with a 95% credible interval of the posterior presented in parentheses. Rows correspond to participants (Pp), except the top row, which contains parameters of the group-level distributions (Hyper).

Pp	B_3	B_5	B_7	B_9	A	t_0
Hyper	0.12 (0.01, 0.35)	0.16 (0.03, 0.31)	0.05 (0.00, 0.14)	0.01 (0.00, 0.07)	1.81 (1.20, 2.41)	0.39 (0.22, 0.45)
1	0.06 (0.00, 0.19)	0.21 (0.12, 0.34)	0.07 (0.01, 0.17)	0.01 (0.00, 0.03)	2.40 (1.80, 3.09)	0.41 (0.35, 0.45)
2	0.30 (0.21, 0.41)	0.19 (0.12, 0.28)	0.06 (0.01, 0.12)	0.01 (0.00, 0.03)	1.58 (1.30, 1.90)	0.39 (0.35, 0.42)
3	0.09 (0.01, 0.18)	0.11 (0.02, 0.18)	0.03 (0.00, 0.09)	0.02 (0.00, 0.11)	1.83 (1.52, 2.22)	0.39 (0.36, 0.43)
4	0.18 (0.10, 0.28)	0.13 (0.06, 0.21)	0.04 (0.00, 0.10)	0.01 (0.00, 0.04)	1.43 (1.19, 1.77)	0.35 (0.32, 0.37)
5	0.03 (0.00, 0.11)	0.19 (0.11, 0.31)	0.06 (0.01, 0.17)	0.02 (0.00, 0.12)	1.81 (1.49, 2.34)	0.48 (0.45, 0.50)
Pp	v_0	w_S	w_D	α	β	
Hyper	1.10 (0.31, 1.73)	0.35 (0.03, 0.77)	8.19 (4.57, 11.88)	0.37 (0.14, 1)	0.07 (0.01, 0.12)	
1	1.73 (1.16, 2.28)	0.07 (0.00, 0.34)	3.33 (2.69, 6.20)	1.19 (0.37, 1.61)	0.12 (0.07, 0.15)	
2	1.10 (0.18, 1.73)	0.31 (0.02, 0.82)	11.27 (7.06, 17.63)	0.17 (0.11, 0.30)	0.06 (0.05, 0.07)	
3	0.91 (0.13, 1.43)	0.21 (0.01, 0.64)	9.63 (5.17, 18.86)	0.26 (0.12, 0.55)	0.10 (0.08, 0.12)	
4	1.22 (0.23, 1.95)	0.39 (0.03, 0.97)	12.10 (7.42, 20.98)	0.17 (0.10, 0.30)	0.06 (0.05, 0.08)	
5	1.30 (0.23, 2.93)	1.06 (0.19, 1.79)	8.19 (5.08, 13.01)	0.41 (0.25, 0.65)	0.03 (0.03, 0.04)	

mean and standard deviation of variability in rates, and some evidence for a decrease in response caution as set size increased. Whether such effects apply to other instances of Hick’s Law remains to be seen.

It is possible that the remaining misfit, under-prediction of accuracy for the smallest set size of 3, may not be due to the Win-All ALBA itself but instead because of our specification of the way mean rates change as a function of distance from the correct response. Although the function we specified is flexible, it does not take account of “edge effects” – improved discriminability for stimuli at the extremes of the stimulus set – which are known to be prevalent in absolute identification tasks, such as the present one, that require classification of stimuli along a single dimension (Brown, Marley, Donkin, & Heathcote, 2008b). For $n = 3$ the majority of the stimuli are at the edges, whereas this proportion drops off rapidly as n increases, consistent with a pronounced under-prediction of accuracy for $n = 3$.

Discussion

We have proposed a theory of multiple choice decisions in terms of *advantages*, directed pairwise comparisons among the subjective values of response alternatives. We instantiated this theory through a linear scheme for mapping subjective values to the inputs for linear evidence accumulation processes that race independently to determine a choice. Together these assumptions are required for the validity of the simple race equation that we use to instantiate the theory in the “Advantage LBA” (ALBA) model, making it sufficiently mathematically tractable to support an easily computed likelihood. We exploited this likelihood to explore the ability of the model to provide comprehensive fits to both choice probabilities and the full distribution of RT. We addressed tasks requiring either identification or forced choice among sets of responses ranging in size from two to nine, with a focus on phenomena that have been claimed to rule out independent race models (Teodorescu & Usher, 2013). Contrary to these claims, the ALBA provided a good account of these data in a parsimonious and parametrically plausible and coherent manner.

We first focused on a task requiring two-alternative forced choices based on brightness (Teodorescu et al., 2016, Experiment 1). We exploited the known luminance values and research supporting a logarithmic mapping to subjective brightness values (Fechner et al., 1966) to test a linear mapping to the rate of evidence accumulation in terms of three estimated parameters, an intercept and weights on the sum of and difference between the subjective brightness values for the two options being compared by each advantage accumulator. We described this as an “advantage-input” coding scheme because it is the difference component that determines whether responses are accurate. Consistent with this nomenclature, the difference weight was estimated as an order of magnitude greater than the sum weight. This finding was replicated in our two subsequent applications of the ALBA, for choices among more than two brightness values and movement directions in forced choice and identification respectively, bolstering the plausibility of the advantage-input coding scheme. In all cases the sum weight, although smaller, was non-negligible, consistent with Teodorescu et al.’s conclusion, that forced choice has both absolute and relative components.

Although we focused on cases in which objective stimulus values are known and a mapping assumed that produces corresponding subjective values for each stimulus, the advantage-input coding scheme also enables subjective values to be directly estimated, at least when there are sufficiently many stimuli. In binary choice, for example, with only two stimuli the corresponding two subjective values cannot be identified because a total of 5 parameters must be estimated (i.e., the two subjective values and three advantage-input parameters) in order to specify four rates (i.e., inputs for each of the two accumulator for each stimulus). However, with three stimuli identification is possible because the required number of 6 estimated parameters is commensurate with the six required rates. As the number of stimuli (S) increases estimates become increasingly constrained as only $S + 3$ parameters are required to calculate $2 \times S$ rates. Thus, our approach provides a method to estimate a scaling of subjective values for a set of three or more stimuli based on binary responses that, for the first time to our knowledge, takes account of RT as well as choices. This approach can be applied when objective values are unknown (e.g., for items in a recognition memory experiment) and also when they are known to infer an unknown mapping to subjective values.

The same logic applies to estimating subjective values from choices among more than two options. This offers potential efficiencies above standard methods of obtaining a scaling based on testing all possible binary comparisons among a set of stimuli, as all such binary comparisons are assumed to occur as part of the ALBA architecture. Clearly further work is needed to determine the best designs to realize this potential. Our applications here focused on cases like brightness judgments where a unidimensional scaling of subjective values is likely to apply. However, our approach could be applied more broadly to cases where multidimensional scalings might be required, such as in multiattribute choice. In this case subjective coordinates would be estimated and differences taken according to an assumed distance metric (e.g., Euclidean or City Block), and goodness of fit used to adjudicate among potential choices of dimensionality and metric (see e.g., M. D. Lee, 2001).

In more complex situations that violate simple scalability, such as the multiattribute choice context effects studied by Trueblood et al. (2014), a potential approach is an architecture in which there is a separate ALBA for each attribute, so each attribute is treated like a separate stimulus, and an appropriate stopping rule applied to combine the outputs of each

ALBA. In Appendix E we investigated one such Win-All type stopping rule, choosing the option whose entire set of advantages (for both attributes) finishes first. We obtained the attraction effect (Huber, Payne, & Puto, 1982), where an indifferent forced binary choice is tipped towards one option in a trinary choice by adding a stimulus that is equal on one attribute and slightly inferior on the other to the now favored stimulus. We also obtained the compromise effect (Simonson, 1989), where the added stimulus is clearly dominated on one attribute and dominates on the other attribute relative to the new favored stimulus, but less so than the now disfavored stimulus. Qualitative and quantitative details on why ALBA is able to produce these effects are provided in Appendix E.

We offer these findings as a demonstration that the ALBA is not necessarily incompatible with, and offers a potential alternative approach to, phenomena that violate simple scalability. We are certainly not claiming that this version of the Win-All ALBA provides the same sort of comprehensive account of context effects on RT and choice as the currently leading models, the MLBA (Trueblood et al., 2014), LCA (Tsetsos, Usher, & Chater, 2010) and MDFT (Roe, Busemeyer, & Townsend, 2001). However, it might offer an attractive alternative route to pursue in modeling these effects as it has recently been shown that not only the LCA, but also MDFT and the MLBA are not measurement models (i.e., a model whose parameters can be successfully recovered from data, Evans et al., in press). Of course, it remains to be shown if the multiattribute model we have proposed, with one ALBA per attribute, is a measurement model.

Our second application of the ALBA focused on a four-alternative forced-choice task, again requiring selection of the brightest stimulus (Teodorescu & Usher, 2013), so we used the same method of determining rates from objective stimulus values. With $n = 4$ possible responses there are $n(n - 1) = 12$ advantage accumulators, with each response being associated with a set of $n - 1$ advantages. We showed that a Win-All decision rule was able to accommodate the “near-competitor” effect, whereby decisions are slower and less accurate when a lure stimulus is close in value to the correct stimulus. This occurs under the Win-All rule because decision time is determined by the *slowest* member of the winner’s set of $n - 1$ advantages, which for an accurate response will typically correspond to the contrast between the correct stimulus (with the maximum stimulus value) and the near competitor (with the next-to-maximum value). As the later contrast has a slow rate (i.e., the average advantage is small) the near-competitor effect arises, and as we noted this also makes the Win-All ALBA similar to the max-next model.

Our final application was to an identification task that varied the number of potential responses between $n = 3$ and $n = 9$ (van Maanen et al., 2012), with a focus on Hick’s Law, a linear increase in mean RT with the logarithm of n . Hick’s Law is problematic for independent race models with a one-to-one mapping between stimuli and accumulators, because decision time corresponds to the *minimum* of the n accumulator completion times, which, all other things being equal, decreases with n . For the Win-All ALBA the same logic about a minimum time applies, but to the counters that require all of their $n - 1$ accumulators to complete before they complete. Hence, all other things are not equal, as the completion time of a counter depends on the *maximum* completion time over its $n - 1$ accumulators, which increases with n . We showed that Hick’s Law naturally emerges from this “minimum-of-maxima” setup.

The Win-All stopping rule is a key component in successfully extending the ALBA to

tasks with more than two choices. However, it is only one of a variety of potential stopping rules. In the appendix, we detail two alternative stopping rules. The Lose-All rule assumes that the decision maker responds as soon as all but one of the response options have been beaten by every other contrasting alternative. The Lose-One rule 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. We focused on the Win-All rule because of its conceptual similarity to max-next models, because of the relatively transparent way in which it explains the near-competitor effects and Hick’s Law, and because it provided the best fit to the data we examined. However, we do not believe it would be prudent to conclude that the later finding will always hold, that the same stopping rule necessarily applies in all situations, or indeed that the three we considered are the only possibilities. That said, at present we recommend the Win-All rule as the default choice for applications of the ALBA.

Our further analysis of van Maanen et al.’s data (2012) showcased the power afforded by the ALBA’s easily computed likelihood in terms of our ability to relatively easily fit and evaluate a range of alternative model parameterizations despite the computational complexity associated with $n = 9$ choices, and hence $n(n - 1) = 72$ accumulators. We were able to explore six models with Bayesian methods that enabled us to thoroughly evaluate and compare them. This also demonstrated the application of a flexible functional method for determining subjective values from objective stimulus values. We did not intend either the proposed mapping or the model exploration to be definitive, but rather as illustrative of the potential for future applications of the ALBA. Even so, we were able to show that the ALBA was able to fit fine-grained detail in a complex data set.

Throughout this manuscript we have made a number of simplifying assumptions in the interests of parsimony or tractability. However, it is important to acknowledge that in some of these cases there is no in principle objection to relaxing these assumptions. First, we generally attempted to minimize differences in trial-to-trial variability. We always assumed the start-point range (A) to be fixed between accumulators within trials. Rate variability (s) was fixed between accumulators for most applications, but was allowed to change linearly with the mean rate (v) and set-size to account for the reduction in accuracy with increasing set-size for the Hick’s law data set. It is possible that future applications will require further relaxation of assumptions about trial-to-trial variability.

Throughout this manuscript we also assumed independence in two senses. The first is in the relationship between accumulators *during* accumulation (i.e., that the value in one accumulator does not affect the value in another). The second sense refers to the values of starting points and drift rates on a given trial being sampled independently over accumulators. However, we note that there are a number of ways to define independence (Teodorescu & Usher, 2013), and that we did not assume independence in the relationship of the inputs that are fed into different accumulators, which is the type of independence Teodorescu & Usher, 2013 focused on. It is possible that future work aiming to maintain mathematical simplicity may also relax the second sense in which we used independence, as the resulting race models remain tractable in some cases (see Heathcote & Love, 2012). This may be desirable as plausible mechanisms exist that could result in positively or negatively correlated thresholds and/or rates over accumulators. For example positive correlations could result from trial-to-trial fluctuations in attention that have a common effect on the rates of both accumulators, or fluctuations in response caution that have a common effect

on the thresholds of both accumulators.

The ALBA shares with the LCA and max-next models a high degree of computational complexity that either scales with the square of the number of possible responses or involves some serial components. In the LCA this increase in computational complexity occurs because, although there are only n accumulators, $n(n - 1)$ lateral inhibitory connections are required. For the max-next model this increase in computational complexity occurs either because several serial operations are required at each moment during accumulation or because a parallel version requires the same order of advantage comparisons as the ALBA. A common element here is the requirement to base decisions on some sort of pairwise comparison among potential responses, which naturally leads to a polynomial increase in computational cost with n . Although ways to avoid this have been proposed, such as normalizing inputs, it has been argued that this approach is inadequate (Teodorescu & Usher, 2013). A reviewer pointed out that in the LCA the number of inhibitory connections, and hence computational complexity, could be reduced to a linear function of n by outputs from each accumulator projecting to a single unit that then sends back the same inhibitory value to each accumulator. This possibility was explicitly explored by Wang (2002, see also; Bogacz, Brown, Moehlis, Holmes, & Cohen, 2006, Fig. 3). Further, even in the standard LCA in which all response units mutually inhibit, typically all connections are assumed to have equal strength, and so only a single free parameter is added. These considerations show that computational simplifications are possible for the LCA, whereas analogous simplifications to the ALBA's architecture are much more difficult to envisage. The ALBA expresses this cost in terms of a complex architecture, but we would argue that it is plausible given the brain's massively parallel architecture.

In conclusion, this paper presents a new framework for modeling multialternative speeded choice data. The framework is based on racing accumulators corresponding to binary advantages of choice options. It can be used to instantiate tractable independent accumulator model with an explicit likelihood function that supports comprehensive and efficient fitting to data. Further, the model can account for a number of benchmark data sets in perceptual decision making in term of psychologically interpretable parameters. On a broader scale, this framework provides a general way of dealing with key phenomena for multiple choice, such as response competition, the effect of number of choice options, and simultaneous absolute and relative effects among choice options, that are potentially important beyond perceptual decision making. When combined with the fact that it supports a measurement model, which in turn allows it to address cases where objective input values, or their mapping to subjective values, must be inferred, we believe this framework could be applied more widely than perceptual choice, potentially providing detailed quantitative characterizations performance in areas ranging from memory and psycho-linguistics to judgment and decision making and thinking and reasoning.

Acknowledgements

This research was supported by Australian Research Council grants to DvR, SDB, and AH (DE140101181, FT120100244, DP12102907, DP160101891 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 Adjunct Research Professor (part-time) at the Institute for Choice, University of South Australia Business School.

Appendix A
ALBA Stopping Rules

All stopping rules for the ALBA require access to an analytical expression for the probability density function and cumulative density function of a single accumulator. 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:

$$CDF(t) = 1 + \left(\frac{tZ(t) - b}{A} \right) G\left(\frac{b}{t}\right) + \left(\frac{b - A - tZ(t)}{A} \right) G\left(\frac{b - A}{t}\right) \quad (6)$$

with

$$Z(t) = \frac{1}{G\left(\frac{b}{t}\right) - G\left(\frac{b-A}{t}\right)} \int_{\frac{b-A}{t}}^{\frac{b}{t}} ug(u) du \quad (7)$$

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 (6) with respect to t .⁵ Assuming $Z(t)$ is differentiable for all $t > 0$, and denoting its derivative by $Z'(t)$, we get

$$PDF(t) = \left(\frac{Z(t) + tZ'(t)}{A} \right) G\left(\frac{b}{t}\right) + \left(\frac{tZ(t) - b}{A} \right) \left(-\frac{b}{t^2} \right) g\left(\frac{b}{t}\right) - \left(\frac{Z(t) + tZ'(t)}{A} \right) G\left(\frac{b - A}{t}\right) + \left(\frac{b - A - tZ(t)}{A} \right) \left(-\frac{b - A}{t^2} \right) g\left(\frac{b - A}{t}\right) \quad (8)$$

Note that Equations 6 and 8 only involve the expressions for the PDF and CDF of the drift rate distribution (g and G respectively). This results in expressions for the PDF and CDF for a single accumulator that are analogous to those presented as Equations (1) and (2) of Brown and Heathcote (2008).

A1. Win-All Derivation

In this section, we unpack the equations for the Win-All version of the ALBA model, (3) for a decision trial with three response options (i.e., 1, 2, and 3). It can be written out as:

$$p_1(t) = PDF_{1-2}(t) \times CDF_{1-3}(t) \times \prod_{I \neq A} \left[1 - \prod_{K \neq I} CDF_{I-K}(t) \right] + PDF_{1-3}(t) \times CDF_{1-2}(t) \times \prod_{I \neq A} \left[1 - \prod_{K \neq I} CDF_{I-K}(t) \right] \quad (9)$$

where I is an option in the set $\{2, 3\}$ and K is an option in the set $\{1, 2, 3\}$ that is not I . The first line represents the scenario where accumulator 1 – 2 is the terminating accumulator that prompts the response ($PDF_{1-2}(t)$), accumulator 1 – 3 had finished before

⁵Note that there is a typo in Equation 3 of Terry et al. (2015), the form presented here is the corrected version.

($CDF_{1-3}(t)$), and at least one accumulator out of each of the sets $\{2-1, 2-3\}$, and $\{3-1, 3-2\}$ had not yet finished ($\prod_{I \neq 1} [1 - \prod_{K \neq I} CDF_{I-K}(t)]$).

Similarly, the second line represents the scenario where accumulator $1-3$ is the terminating accumulator that prompts the response ($PDF_{1-3}(t)$), accumulator had finished before ($CDF_{1-2}(t)$), and at least one accumulator out of each of the sets $\{2-1, 2-3\}$, and $\{3-1, 3-2\}$ had not yet finished ($\prod_{I \neq 1} [1 - \prod_{K \neq I} CDF_{I-K}(t)]$).

The PDF for response 1 is completed by summing the expressions on both of these lines.

A2. Lose-All Stopping Rule

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; thus, the Lose-All model is a “last man standing” algorithm, and the conceptual inverse of the Win-All model. For example, the decision maker chooses 1 from three response options (i.e., 1, 2, and 3) if and only if

1. All of the accumulators in each of the sets $\{1-2, 3-2\}$, and $\{1-3, 2-3\}$ have reached their threshold, and:
2. At least one of the accumulators in the set $\{2-1, 3-1\}$ has *not* reached its threshold.

Specifically, response option 1 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 six inputs (e.g., all but $1-2$, and $1-3$ for a 1 response) and requiring six 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 $\{1-2, 1-3\}$, followed by accumulators $\{2-1, 3-1\}$, and $\{2-3\}$. Then response 2 is made, despite the fact that 1 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 1 at time t is given by

$$p_1(t) = \sum_{I \neq 1} \left[\sum_{J \neq I} \left(PDF_{J-I}(t) \times \prod_{K \neq I, J} CDF_{K-I}(t) \times \prod_{L \neq I, 1} \prod_{M \neq L} CDF_{M-L}(t) \right) \right] \times \left(1 - \prod_{I \neq 1} CDF_{I-1}(t) \right) \quad (10)$$

I is an option in the set $\{2, 3\}$, J is an option in the set $\{1, 2, 3\}$, K is an option in the set $\{1, 2, 3\}$ that is neither I nor J , L is an option in the set $\{2, 3\}$ that is not I , and M is an option in the set $\{1, 2, 3\}$ that is not L . Each CDF is obtained by applying (6) to the respective advantage accumulator, and each PDF is obtained by applying (8) to the respective advantage accumulator.

For a decision trial with three response options (i.e., 1, 2, and 3), (10) can be expanded as:

$$\begin{aligned}
p_1(t) = & \sum_{J \neq 2} \left(PDF_{J-2}(t) \times \prod_{K \neq 2, J} CDF_{K-2}(t) \times \prod_{L \neq 2, 1} \prod_{M \neq L} CDF_{M-L}(t) \right) \times \left(1 - \prod_{I \neq 1} CDF_{I-1}(t) \right) + \\
& \sum_{J \neq 3} \left(PDF_{J-3}(t) \times \prod_{K \neq 3, J} CDF_{K-3}(t) \times \prod_{L \neq 3, 1} \prod_{M \neq L} CDF_{M-L}(t) \right) \times \left(1 - \prod_{I \neq 1} CDF_{I-1}(t) \right) \quad (11)
\end{aligned}$$

where I and L are options in the set $\{2, 3\}$, J is an option in the set $\{1, 2, 3\}$, K is an option in the set $\{1, 2, 3\}$ that is not J , and M is an option in the set $\{1, 2, 3\}$ that is not L . The first line represents the sum of all scenarios of J where accumulator $J-2$, where J is not 2, is the terminating accumulator that prompts the response ($PDF_{J-2}(t)$), all accumulators $K-2$, where K is not 2 or J , had finished before ($\prod_{K \neq 2, J} CDF_{K-2}(t)$), all accumulators out of the set $\{1-3, 2-3\}$ had finished before ($\prod_{L \neq 2, 1} \prod_{M \neq L} CDF_{M-L}(t)$), and at least one accumulator out of the set $\{2-1, 3-1\}$ had not yet finished ($1 - \prod_{I \neq 1} CDF_{I-1}(t)$).

Similarly, the second line represents the sum of all scenarios of J where accumulator $J-3$, where J is not 3, is the terminating accumulator that prompts the response ($PDF_{J-3}(t)$), all accumulators $K-3$, where K is not 3 or J , had finished before ($\prod_{K \neq 3, J} CDF_{K-3}(t)$), all accumulators out of the set $\{1-2, 3-2\}$ had finished before ($\prod_{L \neq 3, 1} \prod_{M \neq L} CDF_{M-L}(t)$), and at least one accumulator out of the set $\{2-1, 3-1\}$ had not yet finished ($1 - \prod_{I \neq 1} CDF_{I-1}(t)$).

The PDF for response 1 is completed by summing the expressions on both of these lines.

A3. Lose-One Stopping Rule

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. That is, the decision maker chooses 1 from three response options (i.e., 1, 2, and 3) if and only if

1. At least one of the accumulators in each of the sets $\{1-2, 3-2\}$, and $\{1-3, 2-3\}$ have reached their threshold, and:
2. None of the accumulators in the set $\{2-1, 3-1\}$ has reached their threshold.

Specifically, response option 1 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. Three counters in the first layer take input from the sets of two just described, each requiring only one count to be triggered. Three counters in the second layer each correspond to a response. They take inputs from two counters in the previous layer (e.g. the counter corresponding to the 1 response takes inputs from the two sets in (1) above), and require two 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 1 without any accumulator that favors 1 having reached threshold (e.g., the sequence 2-3, 3-2 will trigger a response 1). Again, with sensible drift rates, this set of events is exceedingly unlikely.

The probability density function for the distribution of responses 1 at time t is given by

$$p_1(t) = \sum_{I \neq 1} \left[\sum_{J \neq I} \left(PDF_{J-I}(t) \times \prod_{K \neq I, J} [1 - CDF_{K-I}(t)] \times \prod_{L \neq I, 1} \left[1 - \prod_{M \neq L} [1 - CDF_{M-L}(t)] \right] \right) \right] \times \prod_{I \neq 1} [1 - CDF_{I-1}(t)] \quad (12)$$

I is an option in the set $\{2, 3\}$, J is an option in the set $\{1, 2, 3\}$, K is an option in the set $\{1, 2, 3\}$ that is neither I nor J , L is an option in the set $\{2, 3\}$ that is not I , and M is an option in the set $\{1, 2, 3\}$ that is not L . Each CDF is obtained by applying (6) to the respective advantage accumulator, and each PDF is obtained by applying (8) to the respective advantage accumulator.

For a decision trial with three response options (i.e., 1, 2, and 3), (12) can be expanded as:

$$p_1(t) = \sum_{J \neq 2} \left(PDF_{J-2}(t) \times \prod_{K \neq 2, J} [1 - CDF_{K-2}(t)] \times \prod_{L \neq 2, 1} \left[1 - \prod_{M \neq L} [1 - CDF_{M-L}(t)] \right] \right) \times \prod_{I \neq 1} [1 - CDF_{I-1}(t)] + \sum_{J \neq 3} \left(PDF_{J-3}(t) \times \prod_{K \neq 3, J} [1 - CDF_{K-3}(t)] \times \prod_{L \neq 3, 1} \left[1 - \prod_{M \neq L} [1 - CDF_{M-L}(t)] \right] \right) \times \prod_{I \neq 1} [1 - CDF_{I-1}(t)] \quad (13)$$

where I and L are options in the set $\{2, 3\}$, J is an option in the set $\{1, 2, 3\}$, K is an option in the set $\{1, 2, 3\}$ that is not J , and M is an option in the set $\{1, 2, 3\}$ that is not L . The first line represents the sum of all scenarios of J where accumulator $J - 2$, where J is not 2, is the terminating accumulator that prompts the response ($PDF_{J-2}(t)$), all accumulators $K - 2$, where K is not 2 or J , had not yet finished ($\prod_{K \neq 2, J} [1 - CDF_{K-2}(t)]$), at least one accumulator out of the set $\{1 - 3, 2 - 3\}$ had finished before ($\prod_{L \neq 2, 1} [1 - \prod_{M \neq L} [1 - CDF_{M-L}(t)]]$), and no accumulator out of the set $\{2 - 1, 3 - 1\}$ had finished yet ($\prod_{I \neq 1} [1 - CDF_{I-1}(t)]$).

Similarly, the second line represents the sum of all scenarios of J where accumulator $J - 3$, where J is not 3, is the terminating accumulator that prompts the response ($PDF_{J-3}(t)$), all accumulators $K - 3$, where K is not 3 or J , had not yet finished ($\prod_{K \neq 3, J} [1 - CDF_{K-3}(t)]$), at least one accumulator out of the set $\{1 - 2, 3 - 2\}$ had finished before ($\prod_{L \neq 3, 1} [1 - \prod_{M \neq L} [1 - CDF_{M-L}(t)]]$), and no accumulator out of the set $\{2 - 1, 3 - 1\}$ had finished yet ($\prod_{I \neq 1} [1 - CDF_{I-1}(t)]$).

The PDF for response 1 is completed by summing the expressions on both of these lines.

Appendix B

Estimation Details

B1. Estimation Details: Absolute versus Relative Input

The model was fit to each participant’s data separately using Bayesian Markov-chain Monte Carlo methods in R with the DMC software (Heathcote, Lin, et al., in press).⁶ All scripts, RData files, and plotting code of these and subsequent fits are available on <https://osf.io/2s6ax/>.

Vague normal priors were used, truncated at zero for all parameters except t_0 , which was bounded between 0.1s and 1s, with the following means: $A = 25$, $B_R = 1$, $v_0 = 5$, $w_S = 5$, $w_D = 100$, $s = 5$, and $t_0 = 0.3$. Prior standard deviations had the same values, except for t_0 , where it was 0.2. After burn in, 21 chains of 500 samples, thinned to retain every 10th sample, were used for analysis, with convergence supported by multivariate scale-reduction factors (Brooks & Gelman, 1998) of less than 1.01 in all cases, and confirmed visually, as was dominance of the posterior by the prior.

B2. Estimation Details: Strong versus Weak Distractors

We estimated parameters for a hierarchical version of the Win-All model using a differential evolution Markov-chain Monte Carlo procedure (ter Braak, 2006; Turner, Sederberg, Brown, & Steyvers, 2013).

Starting values for the MCMC chains for individual parameters were drawn from the following distributions: $B \sim N(0.5, 0.05)|(0,)$, $A \sim N(1, 0.1)|(0,)$, $t_0 \sim N(0.2, 0.02)|(0,)$, $v_0 \sim N(1, 0.1)|(0,)$, $w_S \sim N(1, 0.1)|(0,)$, and $w_D \sim N(1, 0.1)|(0,)$.

Priors for all group level mean parameters were normal distributions, with $B_\mu \sim N(0.5, 0.2)|(0,)$, $A_\mu \sim N(1, 0.5)|(0,)$, $t_{0\mu} \sim N(0.2, 0.1)|(0,)$, $v_{0\mu} \sim N(1, 0.5)|(0,)$, $w_{S\mu} \sim N(1, 0.5)|(0,)$, and $w_{D\mu} \sim N(1, 0.5)|(0,)$. Priors for all group level standard deviation parameters were 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.

Following burn-in, sampling chains that were at least 3 across-chain standard deviations removed from the mean were reset to the mean in an iterative procedure for each estimated parameter (2.6% of all chains were reset this way). Parameter con-

⁶We provide an RStudio project containing the data, DMC functions, and scripts to fit and check the model as a special case of a more general modeling framework that allows a power transformation of objective to subjective values and for rate variability to increase with the mean rate.

vergence was assessed visually and considered satisfactory (trace plots are available on <https://osf.io/2s6ax/>).

B3. Estimation Details: Hick’s Law

We used the same Bayesian hierarchical estimation methods as in the previous section. Again, following burn-in, each parameter chain that was at least 3 across-chain standard deviations removed from the mean was reset to the mean in an iterative procedure (2.8% of all chains were reset this way). Parameter convergence was assessed visually and considered satisfactory (trace plots are available on <https://osf.io/2s6ax/>).

Starting values for the MCMC chains for individual parameters were drawn from the following distributions: $B \sim N(0.5, 0.05)|(0,)$, $A \sim N(2, 0.2)|(0,)$, $t_0 \sim N(0.25, 0.025)|(0,)$, $v_0 \sim N(1, 0.1)|(0,)$, $w_S \sim N(0.5, 0.05)|(0,)$, $w_D \sim N(6, 0.6)|(0,)$, $\beta \sim N(0.1, 0.01)|(0,)$, and $\log(\alpha) \sim N(0, 0.3)|(-3,)$.

Priors for all group level mean parameters were normal distributions, with $B_\mu \sim N(0.5, 0.25)|(0,)$, $A_\mu \sim N(2, 1)|(0,)$, $t_{0\mu} \sim N(0.25, 0.1)|(0,)$, $v_{0\mu} \sim N(1, 0.5)|(0,)$, $w_{S\mu} \sim N(0.5, 0.25)|(0,)$, $w_{D\mu} \sim N(6, 3)|(0,)$, $\beta_\mu \sim N(0.1, 0.05)|(0,)$, and $\log(\alpha_\mu) \sim N(3, 1)|(-3,)$. Priors for all group level standard deviation parameters were 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 2,000 burn-in iterations followed by 2,000 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.

Appendix C

Parameter Recovery

C1. Parameter Recovery Strong versus Weak Distractors

Parameter Recovery was performed by generating data from the median parameter estimates of a Win-All fit to the Teodorescu et al. (2016) data set. The Win-All model with the same parameter constraints that were used on the empirical data set was then fit to this generated data set. The resulting parameter estimates (95% credible interval in parentheses) were then compared to the true parameters. Parameter recovery was excellent, details are shown in Table C1.

C2. Parameter Recovery Hick’s Law

Parameter Recovery was performed by generating data from the median parameter estimates of the ALBA-1 fit to the van Maanen data set. The resulting parameter esti-

Table C1

Estimated parameters of the ALBA model for the generated Teodorescu data set. Displayed are the true parameter values, with a 95% credible interval of the posterior for the recovered parameters presented in parentheses. Columns represent parameters and rows represent different participants (hyper = parameters of the group-level distributions).

Pp	B	A	t_0	v_0	w_S	w_D
Hyper	0.18 (0.03, 0.44)	1.03 (0.23, 1.83)	0.51 (0.16, 0.62)	1.26 (0.61, 2.14)	0.17 (0.06, 0.79)	1.61 (0.36, 2.97)
1	0.07 (0.06, 0.35)	0.88 (0.72, 0.95)	0.65 (0.58, 0.66)	1.41 (1.43, 2.62)	0.21 (0.09, 0.62)	3.28 (2.48, 3.20)
2	0.11 (0.07, 0.40)	1.12 (0.97, 1.33)	0.67 (0.58, 0.68)	1.08 (1.04, 12.11)	0.29 (0.18, 4.21)	3.84 (2.30, 4.48)
3	0.01 (0.00, 0.08)	0.59 (0.54, 0.70)	0.64 (0.61, 0.64)	1.58 (1.45, 8.97)	0.20 (0.04, 2.75)	3.07 (1.48, 3.17)
4	0.12 (0.01, 0.42)	4.08 (3.39, 4.32)	0.47 (0.37, 0.52)	1.31 (1.36, 2.70)	0.10 (0.21, 0.84)	3.65 (3.50, 4.24)
5	0.43 (0.09, 0.43)	1.19 (1.02, 1.35)	0.57 (0.57, 0.66)	1.50 (0.83, 2.40)	0.25 (0.16, 1.13)	4.09 (4.31, 5.78)
6	0.63 (0.24, 0.94)	1.86 (1.77, 2.30)	0.61 (0.53, 0.68)	1.86 (1.86, 3.31)	0.06 (0.05, 0.74)	3.51 (3.22, 4.01)
7	0.39 (0.17, 0.77)	2.34 (2.27, 2.92)	0.59 (0.49, 0.63)	1.67 (1.34, 2.39)	0.23 (0.01, 0.48)	3.88 (3.41, 4.25)
8	0.01 (0.03, 0.28)	0.91 (0.81, 1.05)	0.64 (0.55, 0.63)	1.05 (0.61, 1.95)	0.46 (0.14, 0.75)	4.01 (3.12, 4.15)

Table C2

Estimated parameters of the ALBA- β model for the generated van Maanen data set. Displayed are the true parameter values, with a 95% credible interval of the posterior for the recovered parameters presented in parentheses. Columns represent parameters and rows represent different participants (hyper = parameters of the group-level distributions).

Pp	B	A	t_0	v_0	w_S	w_D	α
Hyper	0.11 (0.01, 0.34)	1.16 (1.02, 1.27)	0.35 (0.21, 0.39)	0.21 (0.03, 0.90)	0.41 (0.03, 0.68)	11.66 (4.81, 12.78)	0.15 (0.07, 0.50)
1	0.11 (0.11, 0.34)	1.22 (1.00, 1.21)	0.33 (0.26, 0.36)	0.13 (0.01, 0.47)	0.13 (0.01, 0.27)	10.21 (6.46, 15.28)	0.14 (0.09, 0.22)
2	0.05 (0.02, 0.14)	1.14 (1.03, 1.19)	0.37 (0.33, 0.39)	0.32 (0.04, 0.85)	0.27 (0.02, 0.45)	16.56 (9.86, 25.10)	0.09 (0.06, 0.15)
3	0.14 (0.07, 0.25)	1.25 (1.12, 1.33)	0.30 (0.26, 0.34)	0.12 (0.01, 0.39)	0.11 (0.00, 0.20)	12.17 (6.91, 14.06)	0.15 (0.13, 0.29)
4	0.05 (0.01, 0.12)	1.03 (1.01, 1.17)	0.33 (0.29, 0.34)	0.28 (0.06, 1.10)	0.42 (0.02, 0.59)	14.88 (10.28, 22.82)	0.09 (0.06, 0.14)
5	0.19 (0.16, 0.37)	1.19 (1.11, 1.27)	0.43 (0.38, 0.44)	0.20 (0.08, 1.17)	1.12 (0.64, 1.25)	4.49 (3.17, 5.06)	0.38 (0.32, 0.55)

mates (95% credible interval in parentheses) were then compared to the true parameters. Parameter recovery was excellent, details are shown in Table C2.

Appendix D Additional Fits

D1. Additional Fits Strong versus Weak Distractors

The posterior predictive data for the WA fit with rate variability s free to vary are shown in Figure D1. Figure D1 shows that this model fits the data well, but offers no qualitative improvement over the model with rate variability s constrained between conditions.

The posterior predictive data for the LBA fit can be found in Figure D2. Figure D2 shows that this model fits the RT data well, but overestimates error rates in both conditions.

D2. Additional Fits Hick's Law

The posterior predictive data for the Win-All version of the ALBA model with parameter B free to vary with set-size conditions can be found in Figure D3. Figure D3 shows that relaxing B to vary across set-sizes does not yield a noticeable improvement over the more constrained model, it cannot pick up the increasing error rates for higher set-sizes.

The posterior predictive data for the Win-All version of the ALBA model with parameter A free to vary with set-size conditions can be found in Figure D4. Figure D4 shows

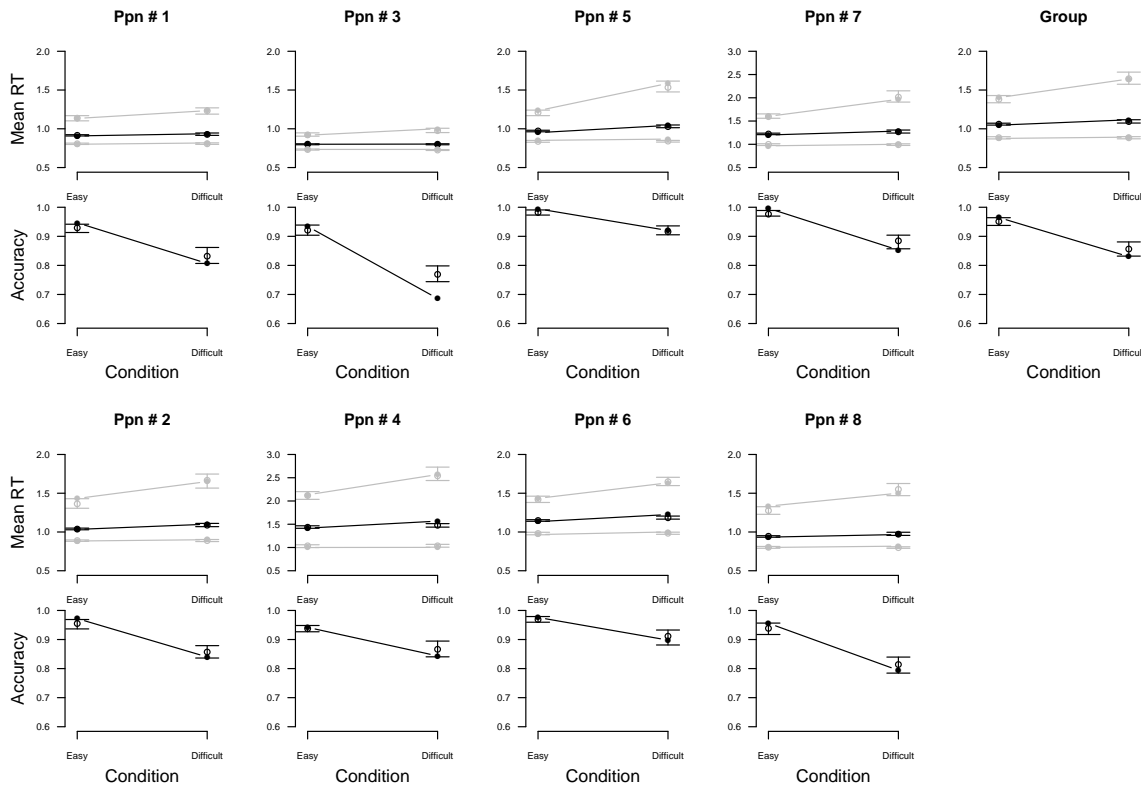


Figure D1. Posterior predictive data for fits to the Experiment 1A data of Teodorescu and Usher (2013) with rate variability s free to vary between the two conditions. RTs for the .5 (black), .1, and .9 (gray) deciles calculated for the easy (top-left) and difficult (top-right) condition, and the proportion of correct responses for the easy (bottom-left) and difficult (bottom-right) condition, both at the individual level (left four columns) and for aggregate data (right column). For all panels, error bars represent posterior predictive data simulated from model fits (the bar extends to the middle 95% of generated summary statistics, with the dot in the middle indicating the median) and lines represent data. See text for details.

that relaxing A to vary across set-sizes does not yield a noticeable improvement over the more constrained model, it cannot pick up the increasing error rates for higher set-sizes.

The posterior predictive data for the Win-All version of the ALBA model with parameters B and A free to vary with set-size conditions can be found in Figure D5. Figure D5 shows that relaxing both B and A allows the model to pick up both the RT and proportion correct data fairly well. Compared to the model that varies rate variability S presented in the main text, the model with B and A free to vary struggles to pick up proportion correct data for set-size 3 and struggles to capture some of the slower RT quantiles.

The posterior predictive data for the LBA fit can be found in Figure D6. Figure D6 shows that this model fits the RT data well, but overestimates the increase in error rates for higher set-sizes.

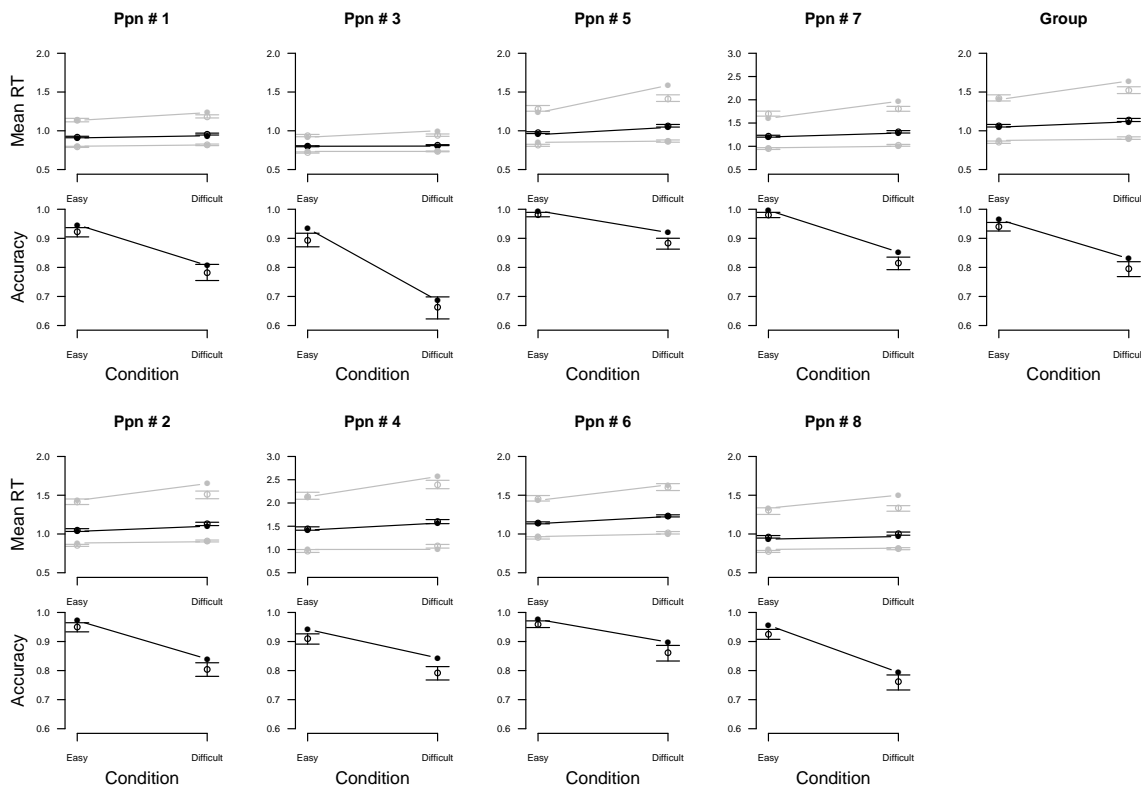


Figure D2. LBA posterior predictive data for the Teodorescu data. RTs for the .1, .5, and .9 deciles (top) and the proportion of correct responses (bottom) as a function of set-size (N) on a logarithmic scale. Posterior predictives are presented at the individual level and for aggregate data (bottom-right panel). For all panels, box-and-whiskers represent posterior predictive data (the box contains 95% of the simulated data, with a bar across the middle indicating the median, and whiskers extend to the data extremes) and lines represent data. See text for details.

Appendix E Context Effects

Here, we show that the ALBA can produce two multiattribute context effects: the attraction effect (Huber et al., 1982) (also called the asymmetric dominance effect) and the compromise effect (Simonson, 1989). There is a third well-known multiattribute context effect, the similarity effect (Tversky, 1972), which the ALBA does not produce. Each of these effects are about two stimuli, S1 and S2, that differ on two attributes, A1 and A2. S1 is preferable on A1, but S2 is preferable on A2, such that in a binary choice S1 and S2 are indifferent. The attraction effect occurs when a third stimulus, S3, is introduced that is slightly inferior to S1 on both A1 and A2, resulting in a preference for S1 over S2. The compromise effect occurs when a third stimulus, S4, is introduced that is even more preferable on A1 and even less preferable on A2 than S1, resulting in a preference for S1, the intermediate option. The similarity effect occurs when a third stimulus, S5, is introduced that is very similar to S2

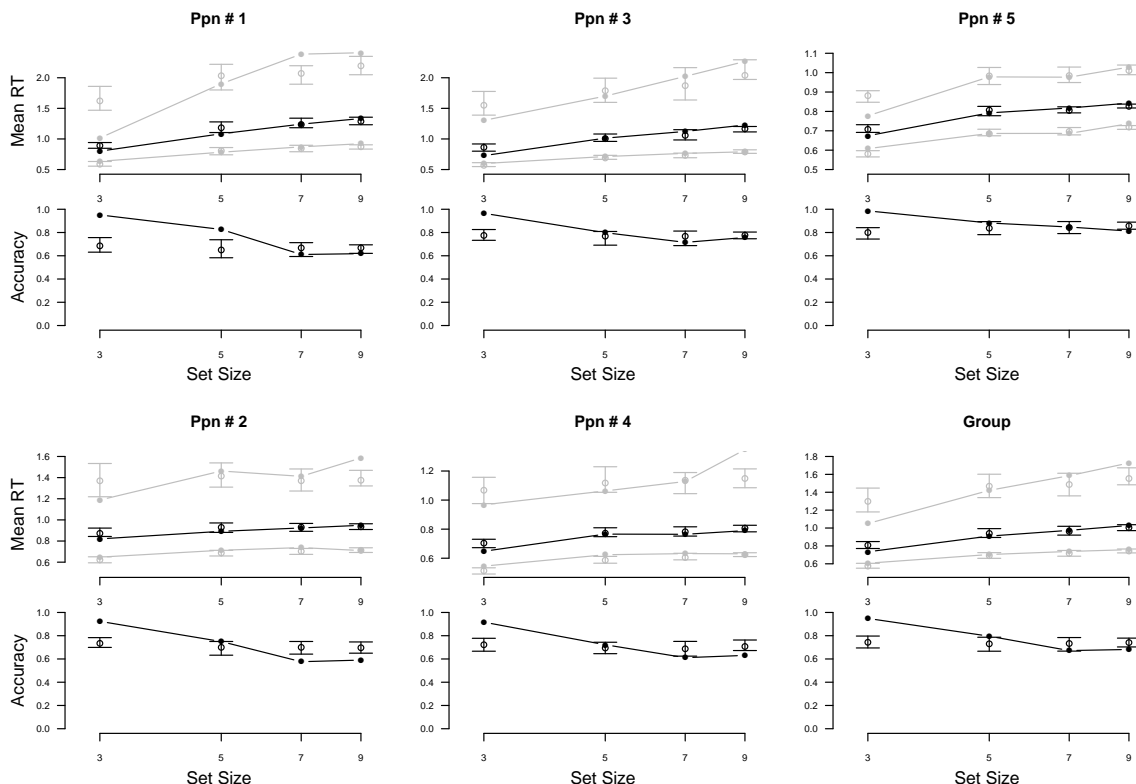


Figure D3. Posterior predictive data for the ALBA-4B fit to the van Maanen data. RTs for the .5 (black), .1, and .9 (gray) deciles (top) and the proportion of correct responses (bottom) as a function of set-size (N) on a logarithmic scale. Posterior predictives are presented at the individual level and for aggregate data (bottom-right panel). For all panels, box-and-whiskers represent posterior predictive data (the box contains 95% of the simulated data, with a bar across the middle indicating the median, and whiskers extend to the data extremes) and lines represent data. See text for details.

on both attributes, resulting in a preference for S1, the dissimilar option.

The purpose of this section is not to provide a detailed ALBA of context effects. Researchers interested in quantitatively capturing these phenomena should use a specialized model like the one developed by Trueblood et al. (2014). Rather, this section intends to explain in a qualitative way how the architecture of ALBA naturally produces some context effects.

For this simulation, context effects were modeled as follows. All three effects were examined using two basic stimuli that varied on two attributes: S1 had subjective input values {4, 6} for A1 and A2 respectively, S2 had subjective input values {6, 4} for A1 and A2 respectively. We investigated the consequences of adding an extra stimulus to this pair. For the attraction effect, the extra stimulus (“S3”) had subjective input values {3, 6} for A1 and A2 respectively. The attraction effect posits that the presence of S3 should lead decision makers to choose S1, because S1 “dominates” S3. For the compromise effect, the

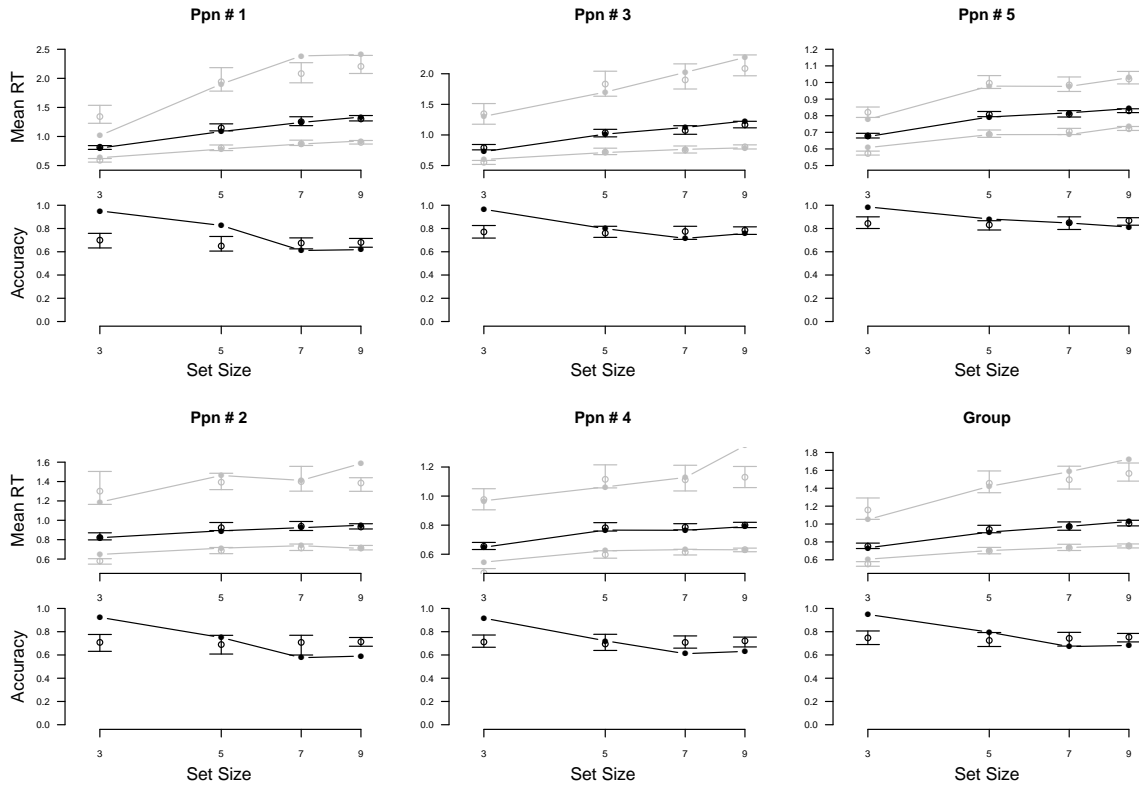


Figure D4. Posterior predictive data for the ALBA-4A fit to the van Maanen data. RTs for the .5 (black), .1, and .9 (gray) deciles (top) and the proportion of correct responses (bottom) as a function of set-size (N) on a logarithmic scale. Posterior predictives are presented at the individual level and for aggregate data (bottom-right panel). For all panels, box-and-whiskers represent posterior predictive data (the box contains 95% of the simulated data, with a bar across the middle indicating the median, and whiskers extend to the data extremes) and lines represent data. See text for details.

extra stimulus (“S4”) had subjective input values {2, 8} for A1 and A2 respectively. The compromise effect posits that the presence of S4 should lead decision makers to choose S1, because it is a compromise between S2 and S4. For the similarity effect, the extra stimulus (“S5”) had subjective input values {5.5, 4.5} for A1 and A2 respectively. The similarity effect posits that the presence of S5 should lead decision makers to S1, because it is different from the highly similar S2 and S5.

All effects were simulated in the ALBA using Equations 1 and 2 by modeling two separate ALBA processes, one for A1 and one for A2. A decision was made once the stopping rule was satisfied for both processes. For the Win-All stopping rule, this meant S1 needed to beat the other stimuli on both A1 and A2 before a response in favor of S1 was executed.

For every effect, we ran 10,000 individual simulations. Individual parameters were $B = 0.2$, $A = 1$, $t_0 = 0.5$, $v_0 = 1.3$, $w_S = 0.2$, and $w_D = 3.5$. Results of the simulations are

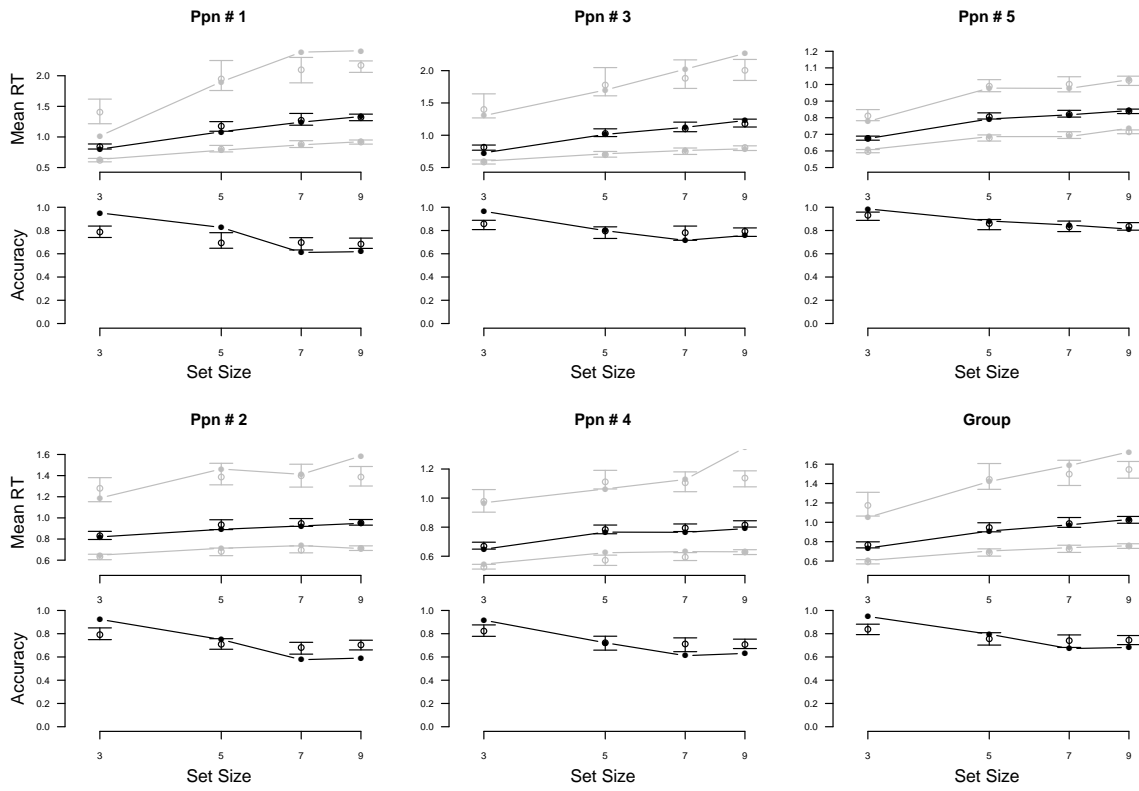


Figure D5. Posterior predictive data for the ALBA-4BA fit to the van Maanen data. RTs for the .5 (black), .1, and .9 (gray) deciles (top) and the proportion of correct responses (bottom) as a function of set-size (N) on a logarithmic scale. Posterior predictives are presented at the individual level and for aggregate data (bottom-right panel). For all panels, box-and-whiskers represent posterior predictive data (the box contains 95% of the simulated data, with a bar across the middle indicating the median, and whiskers extend to the data extremes) and lines represent data. See text for details.

presented in Table E1. The table presents the proportion of times each of the three stimuli was chosen. The probability of choosing S1 (resp., S2) from the set $\{1, 2\}$ is $1/2$ (resp., $1/2$).

The attraction effect is shown by the fact that S1 is chosen from the set $\{1, 2, 3\}$ more often than either S2 or S3. The intuition for this result is as follows: the Win-All stopping rule should produce equal probabilities of choosing S1 and S2 without the presence of S3. Adding S3 adds 1 – 3 advantage accumulators for A1 and A2 that need to finish for S1 to be chosen, and 2 – 3 advantage accumulators for A1 and A2 that need to finish for S2 to be chosen. The 1 – 3 advantage accumulators have inputs 4-3 and 6-6 for A1 and A2 respectively; of these two the slowest one will usually be 6-6. The 2 – 3 advantage accumulators have inputs 6-3 and 4-6 for A1 and A2 respectively; of these two the slowest one will usually be 4-6. As the Win-All stopping rule hinges on the last advantage accumulator finishing, S2 (with the inclusion of 4-6) will become less popular compared to S1 (with the

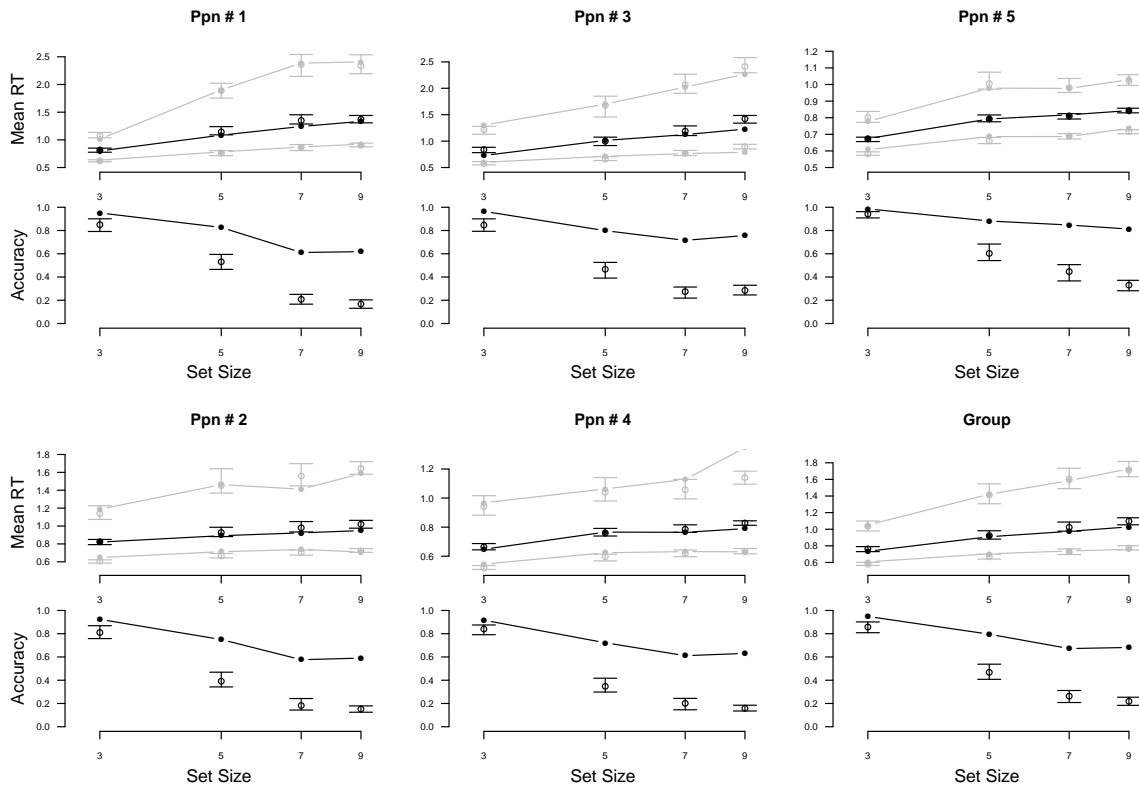


Figure D6. LBA posterior predictive data for the van Maanen data. RTs for the .1, .5, and .9 deciles (top) and the proportion of correct responses (bottom) as a function of set-size (N) on a logarithmic scale. Posterior predictives are presented at the individual level and for aggregate data (bottom-right panel). For all panels, box-and-whiskers represent posterior predictive data (the box contains 95% of the simulated data, with a bar across the middle indicating the median, and whiskers extend to the data extremes) and lines represent data. See text for details.

inclusion of 6-6).

The compromise effect is shown by the fact that S1 is chosen from the set $\{1, 2, 4\}$ more often than either S2 or S4. The intuition for this result is as follows: the Win-All stopping rule should produce equal probabilities of choosing S1 and S2 without the presence of S4. Adding S4 adds 1 – 4 advantage accumulators for A1 and A2 that need to finish for S1 to be chosen, and 2 – 4 advantage accumulators for A1 and A2 that need to finish for S2 to be chosen. The 1 – 4 advantage accumulators have inputs 4-2 and 6-8 for A1 and A2 respectively; of these two the slowest one will usually be 6-8. The 2 – 4 advantage accumulators have inputs 6-2 and 4-8 for A1 and A2 respectively; of these two the slowest one will usually be 4-8. As the Win-All stopping rule hinges on the last advantage accumulator finishing, S2 (with the inclusion of 4-8) will become less popular compared to S1 (with the inclusion of 6-8).

The similarity effect would be shown by the fact that S1 is chosen from the set

Table E1

Attraction and compromise effects, as indicated by the proportion of times stimulus 1 was chosen. Attribute values for stimulus 3 (attraction effect) and stimulus 4 (compromise effect) are displayed in the rows below attribute values for stimulus 1 and 2. The most common choice is printed bold.

Stimulus	Attribute 1	Attribute 2	Attraction	Compromise	Similarity
1	4	6	0.51	0.57	0.26
2	6	4	0.30	0.32	0.32
3	3	6	0.19	–	–
4	2	8	–	0.11	–
5	5.5	4.5	–	–	0.43

{1, 2, 5} more often than either S2 or S5, but the ALBA does not produce the similarity effect. The intuition for this result is as follows: the Win-All stopping rule should produce equal probabilities of choosing S1 and S2 without the presence of S5. Adding S5 adds 1 – 5 advantage accumulators for A1 and A2 that need to finish for S1 to be chosen, and 2 – 5 advantage accumulators for A1 and A2 that need to finish for S2 to be chosen. The 1 – 5 advantage accumulators have inputs 4-5.5 and 6-4.5 for A1 and A2 respectively; of these two the slowest one will usually be 4-5.5. The 2 – 5 advantage accumulators have inputs 6-5.5 and 4-4.5 for A1 and A2 respectively; the slowest one will usually be 4-4.5. As the Win-All stopping rule hinges on the last advantage accumulator finishing, S1 (with the inclusion of 4-5.5) will become less popular compared to S2 (with the inclusion of 4-4.5). Most popular is S5, similar to the result for the compromise effect, as it presents the intermediate option for both attributes.

The fact that the ALBA produces the attraction and compromise effects, but not similarity effect is consistent with recent work that shows that people who exhibit the attraction and compromise effects often do not exhibit the similarity effect (Berkowitsch, Scheibehenne, & Rieskamp, 2014).

References

- Berkowitsch, N. A. J., Scheibehenne, B., & Rieskamp, J. (2014). Rigorously testing multialternative decision field theory against random utility models. *Journal of Experimental Psychology: General*, *143*, 1331–1348.
- Blavatsky, P. (2012). Probabilistic choice and stochastic dominance. *Economic Theory*, *50*, 59 – 83.
- 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.
- 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.
- Brooks, S. P., & Gelman, A. (1998). General methods for monitoring convergence of iterative simulations. *Journal of computational and graphical statistics*, *7*(4), 434–455.
- 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. (2008a). An integrated model of choices and response times in absolute identification. *Psychological Review*, *115*, 396–425.
- Brown, S. D., Marley, A. A. J., Donkin, C., & Heathcote, A. (2008b). 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*, 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.
- Dassonville, P., Lewis, S. M., Foster, H. E., & Ashe, J. (1999). Choice and stimulus–response compatibility affect duration of response selection. *Cognitive Brain Research*, *7*, 235–240.
- Donkin, C., Brown, S. D., & Heathcote, A. (2009). The overconstraint of response time models: Rethinking the scaling problem. *Psychonomic Bulletin & Review*, *16*(6), 1129–1135.
- Eidels, A., Donkin, C., Brown, S. D., & Heathcote, A. (2010). Converging measures of workload capacity. *Psychonomic bulletin & review*, *17*(6), 763–771.
- Evans, N. J., Holmes, W. R., & Trueblood, J. S. (in press). Response-time data provide critical constraints on dynamic models of multi-alternative, multi-attribute choice. *Psychonomic Bulletin & Review*.

- Fechner, G. T., Boring, E. G., Howes, D. H., & Adler, H. E. (1966). *Elements of psychophysics*. Holt, Rinehart and Winston.
- 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). An optimal adjustment procedure to minimize experiment time in decisions with multiple alternatives. *Psychonomic Bulletin & Review*, *19*, 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., Holloway, E., & Sauer, J. (in press). Confidence and varieties of bias. *Journal of Mathematical Psychology*.
- Heathcote, A., Lin, Y.-S., Reynolds, A., Strickland, L., Gretton, M., & Matzke, D. (in press). Dynamic models of choice. *Behavior Research Methods*.
- Heathcote, A., & Love, J. (2012). Linear deterministic accumulator models of simple choice. *Frontiers in Psychology*, *3*.
- Heathcote, A., Suraev, A., Curley, S., Gong, Q., & Love, J. (2015). Decision processes and the slowing of simple choices in schizophrenia. *Journal of Abnormal Psychology*, *124*, 961–974.
- Hick, W. E. (1952). On the rate of gain of information. *Quarterly Journal of Experimental Psychology*, *4*, 11–26.
- Holmes, W. R., Trueblood, J. S., & Heathcote, A. (2016). A new framework for modeling decisions about changing information: The piecewise linear ballistic accumulator model. *Cognitive Psychology*, *85*, 1–29.
- Huber, J., Payne, J. W., & Puto, C. (1982). Adding asymmetrically dominated alternatives: Violations of regularity and the similarity hypothesis. *Journal of Consumer Research*, *9*, 90–98.
- Hyman, R. (1953). Stimulus information as a determinant of reaction time. *Journal of Experimental Psychology*, *45*, 188–196.
- Kveraga, K., Boucher, L., & Hughes, H. C. (2002). Saccades operate in violation of Hick's law. *Experimental Brain Research*, *146*, 307–314.
- Lacouture, Y., & Marley, A. A. J. (1995). A mapping model of bow effects in absolute identification. *Journal of Mathematical Psychology*, *39*, 383–395.
- Lee, K.-M., Keller, E. L., & Heinen, S. J. (2005). Properties of saccades generated as a choice response. *Experimental Brain Research*, *162*, 278–286.
- Lee, M. D. (2001). Determining the dimensionality of multidimensional scaling representations for cognitive modeling. *Journal of Mathematical Psychology*, *45*, 149–166.
- Leite, F. P., & Ratcliff, R. (2010). Modeling reaction time and accuracy of multiple-alternative decisions. *Attention, Perception, & Psychophysics*, *72*, 246–273.
- 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. (2017). Parameter Recovery for the Leaky Competing Accumulator Model. *Journal of Mathematical Psychology, 76*, 25–50.
- Pachella, R. G., & Fisher, D. (1972). Hick's law and the speed-accuracy trade-off in absolute judgment. *Journal of Experimental Psychology, 92*, 378.
- Raab, D. H. (1962). Statistical facilitation of simple reaction times. *Transactions of the New York Academy of Sciences, 24*, 574–590.
- Ratcliff, R., Voskuilen, C., & Teodorescu, A. (2018). Modeling 2-alternative forced-choice tasks: Accounting for both magnitude and difference effects. *Cognitive Psychology, 103*, 1–22.
- 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*(2), 370–392.
- Simonson, I. (1989). Choice based on reasons: The case of attraction and compromise effects. *Journal of Consumer Research, 16*, 158–174.
- 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.
- ten Hoopen, G., Akerboom, S., & Raaymakers, E. (1982). Vibrotactile choice reaction time, tactile receptor systems and ideomotor compatibility. *Acta Psychologica, 50*, 143–157.
- Teodorescu, A. R., Moran, R., & Usher, M. (2016). Absolutely relative or relatively absolute: Violations of value invariance in human decision making. *Psychonomic Bulletin & Review, 23*, 22–38.
- Teodorescu, A. R., & Usher, M. (2013). Disentangling decision models: From independence to competition. *Psychological Review, 120*, 1–38.
- ter Braak, C. J. F. (2006). A Markov chain Monte Carlo version of the genetic algorithm Differential Evolution: easy Bayesian computing for real parameter spaces. *Statistics and Computing, 16*, 239–249.
- 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. *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.
- Tsetsos, K., Usher, M., & Chater, N. (2010). Preference reversal in multiattribute choice. *Psychological Review, 117*(4), 1275–1291.
- 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.

- Tversky, A. (1972). Choice by elimination. *Journal of Mathematical Psychology*, *9*, 341–367.
- Tversky, A., & Simonson, I. (1993). Context-dependent preferences. *Management Science*, *39*, 1179–1189.
- 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., & McClelland, J. L. (2004). Loss aversion and inhibition in dynamical models of multialternative choice. *Psychological Review*, *111*, 757–769.
- Usher, M., Olami, Z., & McClelland, J. L. (2002). Hick’s Law in a Stochastic Race Model with Speed–Accuracy Tradeoff. *Journal of Mathematical Psychology*, *46*, 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., van der Maas, H. L. J., & Wagenmakers, E.-J. (2012). Optimal decision making in neural inhibition models. *Psychological Review*, *119*, 201–215.
- Vickrey, C., & Neuringer, A. (2000). Pigeon reaction time, Hick’s law, and intelligence. *Psychonomic Bulletin & Review*, *7*, 284–291.
- Wang, X.-J. (2002). Probabilistic decision making by slow reverberation in cortical circuits. *Neuron*, *36*, 955–968.