American Psychological Association

Decision

2325-9965/19/$12.00

2019, Vol. 6, No. 3, 201–222
http://dx.doi.org/10.1037/dec0000096

A Quantitative Test of Computational Models of Multialternative Context Effects

Jared M. Hotaling
University of New South Wales

Jörg Rieskamp
University of Basel

When people choose between multiple alternatives it has been observed that the probability of choosing one option over another often depends on the additional options available. For instance, one's preference for a high-end notebook computer over a cheaper midrange model can reverse when a third, low price, but bulkier, computer is also considered. Importantly, such *context effects* violate fundamental choice axioms of most standard economic choice theories. To account for these effects various computational models have been proposed. The present work examines two prominent models that have been proposed to simultaneously account for the *similarity*, *compromise*, and *attraction* effects: multialternative decision field theory (MDFT; Roe, Busemeyer, & Townsend, 2001) and the multiattribute linear ballistic accumulator (MLBA; Trueblood, Brown, & Heathcote, 2014). Because the models imply quite different psychological processes, we focus our comparison on the various mechanisms the models use to produce each effect. We test the models using data from a large within-subjects consumer choice experiment in which all three effects occurred. Overall, we find that both MLBA and MDFT have valuable mechanisms for understanding multialternative context effects. When considering each context effect in isolation MDFT gave the better account of the similarity and attraction effects, while MLBA better described the compromise effect. We conclude that three mechanisms—*attention switching* and *distance-dependent inhibition* from MDFT and *extremeness aversion* from MLBA—appear crucial. These results provide insights into the psychological mechanisms underlying human decision making and should be considered in further development of decision theories.

Keywords: context effects, computational models, consumer choice, MDFT, MLBA

Supplemental materials: http://dx.doi.org/10.1037/dec0000096.supp

People often make choices among multiple alternatives. For example, when buying a new laptop computer a decision maker is faced with many options, varying along attributes such as *price*, *portability*, and *computation power*. How people make such multialternative choices is a topic of interest to researchers across many disciplines, from psychology to economics to marketing. In particular, much work has focused on how decision makers perceive the relative strengths and weaknesses of each alternative, and how these context-dependent differences between alternatives influence choices. A key motivation for this research is to contrast naturalistic choice behavior with that prescribed by normative standards. In economics, the *independence of irrelevant alternatives* (IIA) principle states that if option *A* is preferred to option *B* when these are the only available options, A should also be preferred to *B* when a new alter-

This article was published Online First July 12, 2018.
Jared M. Hotaling, School of Psychology, University of New South Wales; Jörg Rieskamp, Department of Psychology, University of Basel.
We thank Jennifer Trueblood, Jerome Busemeyer, Nicolas Berkowitsch, and Benjamin Scheibehenne for helpful comments to previous versions of our manuscript.
Correspondence concerning this article should be addressed to Jared M. Hotaling, School of Psychology, University of New South Wales, Sydney 2052, Australia. E-mail: jaredhotaling@gmail.com

native, C, is added to the choice set (cf., Rieskamp, Busemeyer, & Mellers, 2006). The related *regularity* principle states that the probability of choosing an option cannot increase as result of adding new alternatives to the choice set. In contradiction to these axioms, it has long been observed that people's preferences are context dependent and change systematically as a function of the available choice alternatives (e.g., Huber, Payne, & Puto, 1982; Simonson, 1989; Tversky, 1972). These context effects have led many to theorize that individuals do not evaluate alternatives independently, but instead compare them with each other; and these comparisons can result in violations of IIA or regularity.

At present, only a few theories are capable, in contrast to standard economic choice theories, of simultaneously explaining the various context effects. One of the first attempts to explain all of these context effects simultaneously was by Roe et al. (2001). Subsequently, alternative models were presented by Usher and McClelland (2004); Bhatia (2013), and most recently by Trueblood et al. (2014). While Usher and McClelland (2004) and Bhatia (2013) made qualitative arguments in favor of their models, Trueblood et al. (2014) went further with arguments that their newly proposed model actually fits context effects data better than the model by Roe et al. (2001). In the present article, we examine this assertion. Importantly, we do so using data from a preferential consumer choice study, in contrast to Trueblood et al. (2014), who reported results from inference and perceptual decision-making experiments. The goal of the present work is to rigorously compare two theories against each other. We follow Trueblood et al. (2014) and consider multialternative decision field theory (MDFT; Roe et al., 2001) and the multiattribute linear ballistic accumulator (MLBA; Trueblood et al., 2014). Rather than attempt a comparison of all possible models, we opt to focus on these two prominent models. This allows us to examine each model in greater detail, with a special emphasis on the psychological mechanisms that each posits. Understanding how these mechanisms act (and interact) to produce each context effect is a central aim of this article. The work we present is novel because it involves fitting the models to actual response data from a consumer choice experiment, at the individual level. Thus, we can test how well each model is able to account for the behavior of each individual, rather than the group.

Multialternative Context Effects

People's evaluation of choice alternatives can be substantially influenced by the choice context. Among different context effects three effects have received substantial attention in the literature. To understand each, first imagine a scenario with two options, A and B, that are chosen with equal probability. Figure 1 graphically represents these options in a two-dimensional attribute space. The *similarity effect* refers to a situation where the introduction of an additional option, S, which is more similar to one existing option, A, increases the likelihood of choosing the dissimilar option, B, relative to A (Tversky, 1972).

In contrast, the *attraction effect* refers to a situation where the added alternative, D, is similar but dominated by one of the existing options (see Figure 1). This increases the absolute probability of choosing the dominating option, A, representing a violation of regularity (Huber et al., 1982). Finally, the *compromise effect* denotes a situation where a third alternative, C, is added to the choice set A and B and has attribute values such that B appears as a compromise between the two other more extreme options A

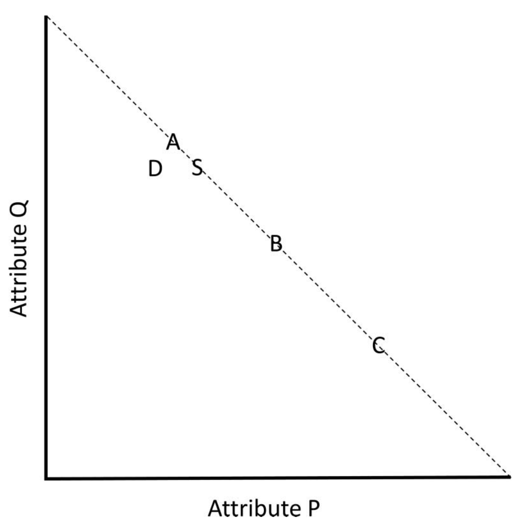

Figure 1. Graphical representation of choice alternatives in two-dimensional attribute space.

and *C*. This makes *B* more attractive, so that the choice proportion of *B* relative to *A* increases, representing another violation of the independence principle (Simonson, 1989).

Context effects have been demonstrated in various domains (see Heath & Chatterjee, 1995 for a review). Recently, some studies have demonstrated that all three effects can also occur within a single person (Berkowitsch, Scheibehenne, & Rieskamp, 2014; Noguchi & Stewart, 2014; Trueblood et al., 2014), while others have challenged this assertion (Liew, Howe, & Little, 2016; Trueblood, Brown, & Heathcote, 2015). Together these results call for a theory to explain why and when a decision maker might exhibit the context effects.

Standard economic choice theories, such as the probit or logit model cannot explain these effects, as they assume that the subjective value of an option is solely a function of the option's attribute values, ignoring other options in the choice set. Thus, new theories were needed. For instance, to account for the similarity effect, Tversky (1972) proposed the theory of *elimination by aspects*, in which the decision maker is assumed to use a sequential process of elimination to gradually reduce the number choice options to one. The model posits that people's *attention switches* probabilistically between the different attributes of the alternatives. Once an attribute receives the attention of the decision maker, each alternative's attribute value is compared with a predetermined aspiration level. Alternatives that fail to meet the aspiration level are eliminated, and attention switches to another attribute to eliminate further alternatives until only one alternative remains. To account for the compromise effect, Tversky and Simonson (1993) proposed the *componential context theory*, which posits that alternatives are valued according to the combination of background context—represented as the weighted sum of attribute values—and local context—represented with a binary comparison process that sums the attribute-wise advantages and disadvantages of each alternative (see Simonson, 1989 for an alternative account). Both theories are limited in their ability to explain all three context effects, with elimination by aspects unable to explain the compromise and attraction effects, while the componential context theory cannot explain the similarity effect. Despite these limitations, the success of both theories in explaining violations of IIA shows the value of incorporating comparison processes into theories of decision making.

Computational models aim to explain violations of the IIA or regularity principles by positing psychological mechanisms for how people process information (Bhatia, 2013; Usher & McClelland, 2001). Despite the growing number of models designed to capture context effects, relatively little work has been done to compare the models. In the present article we present a rigorous comparison of two prominent computational models: MDFT and the MLBA. Below, we test the models against each other using data from a consumer choice task. Our approach is twofold. First, we compare the models rigorously using quantitative model comparison criteria. Second, we illustrate the qualitative predictions of the models to gain a better understanding of when and why the models make different predictions, and how these differences affect their accounts of context effects. Since we focus on the underlying psychological mechanisms of choice behavior, we use our findings to infer which mechanisms are most useful in explaining context effects.

MLBA

Trueblood et al. (2014) introduce a model comprised of two largely independent components. The "front-end" component specifies how the different options are compared against each other, and produces a subjective value for each option. These subjective values serve as input to the "back-end" component. The back-end component is defined by the linear ballistic accumulator (LBA) model (Brown & Heathcote, 2008) that determines the model's choice probabilities and response times (RTs). What follows is a brief summary of MLBA, for more details please see Trueblood et al. (2014). We use variables *i*, *j*, and *k* to represent three choice alternatives, and *P* and *Q* to denote the two attribute dimensions, defined on the same scale, along which the alternatives vary.

According to the front-end process of MLBA the attribute values (Q_i,P_i) of option *i* are transformed into subjective values (u_{Pi},u_{Qi}; see Appendix C in Trueblood et al., 2014 for a full description of MLBA's subjective value function). On the basis of Chernev's (2004) finding of extremeness aversion, MLBA's value function can be tuned to favor options with moderate attribute values. The function has one free parameter, m, that determines the curvature of the attribute space according to

$$\left(\frac{x}{a}\right)^m + \left(\frac{y}{b}\right)^m = 1, \quad (1)$$

where a and b are the x- and y-intercepts, respectively, of the line connecting two objectively indifferent options in the attribute space. Importantly, when $m > 1$, if choice options i and j are matched such that $Q_i + P_i = Q_j + P_j$, the alternative with lower variance across attributes will have the higher total subjective value. Next, the model evaluates alternatives relative to each other in a pairwise fashion using the weighted difference equation,

$$V_{ij} = w_{Pij}(u_{Pi} - u_{Pj}) + w_{Qij}(u_{Qi} - u_{Qj}). \quad (2)$$

Attention weights, w, are an exponentially increasing function of similarity,

$$\begin{aligned} w_{Pij} &= \exp(-\lambda|u_{Pi} - u_{Pj}|) \\ w_{Qij} &= \exp(-\lambda|u_{Qi} - u_{Qj}|), \end{aligned} \quad (3)$$

such that comparisons between similar attributes receive greater attention. In situations where options are described by two attributes, the model allows that one attribute receives greater attention, in which case the formula for w_{Qij} becomes

$$w_{Qij} = \exp(-\lambda\beta|u_{Qi} - u_{Qj}|) \quad (4)$$

where $\beta > 0$. When $\beta > 1$ attribute P is favored, whereas with $\beta < 1$ attribute Q is favored. MLBA also assumes that similarity may be asymmetric, in that $w_{Pij} \neq w_{Pji}$ and $w_{Qij} \neq w_{Qji}$. This is achieved by using λ_+ when the value difference is positive and λ_- when the value difference is negative. Finally, pairwise comparisons are combined to produce a drift rate for each option according to

$$\begin{aligned} d_i &= V_{ij} + V_{ik} + I_0 \\ d_j &= V_{ji} + V_{jk} + I_0 \\ d_k &= V_{ki} + V_{kj} + I_0 \end{aligned} \quad (5)$$

where $I_0 > 0$ is a free parameter assuring that at least one drift rate is positive. Thus, according to MLBA, all competition between alternatives occurs in the front-end process and defines the input of the accumulation process. The end result of the front-end is a set of context-dependent drift rates.

These drift rates are then passed to a back-end process comprised of the LBA model. LBA uses linear, deterministic accumulators to represent the accrual of information in favor of each option. Accumulators independently race toward a decision threshold, with the winner determining the chosen option and its response time. The model was originally proposed as a simplified approximation to choice models that posit noisy accumulation of information within trial, such as decision field theory (DFT; Busemeyer & Townsend, 1993) and the leaky competing accumulator model (Usher & McClelland, 2001). Its linear determinism lends it greater mathematical tractability compared with stochastic accumulator models, and uniquely allows for simple analytic solutions for the distribution of RTs in situations with three or more choice options. Interestingly, despites its simplifying assumptions the LBA model nevertheless captures most established RT phenomena and can potentially produce any RT distribution (Jones & Dzhafarov, 2014).

In the present work we use the internally controlled stopping rule of MLBA because it most closely models the choice scenarios we consider, and because it is the version of the model to which Trueblood et al. (2014) devote the most attention. This stopping rule specifies that a decision is made as soon as the value of one accumulator reaches the threshold. Also, as with all previous research into multialternative context effects, we focus our analyses on choices. Trueblood et al. (2014, p. 187) state that the model's decision threshold and noise parameters should be fixed when fitting choices alone. Therefore, the version of MLBA we use has a total of five free parameters: β, m, λ_+, λ_-, and I_0, which represent attribute weight, curvature of the subjective value space, the rate parameters in the similarity function for positive and negative differences, and a constant used to assure that at least one drift rate is positive, respectively (Table 1).

MLBA's similarity effect. In the case of the similarity effect, MLBA uses an *asymmetric*

Table 1
Parameters of MLBA and MDFT

Model	Parameter	Description
MLBA	β	Attribute weight
	m	Exponent transforming objective to subjective values
	λ_+	Rate parameter for attention weights with positive differences
	λ_-	Rate parameter for attention weights with negative differences
	I_0	Baseline input
MDFT	w	Attention weight
	b	Dominance dimension weight
	φ_1	Sensitivity relating distance to similarity
	φ_2	Memory decay
	σ	Input noise

Note. MLBA = multiattribute linear ballistic accumulator; MDFT = multialternative decision field theory.

attention weighting mechanism to explain the effect. It posits that decision makers place greater weight on positive evidence than negative evidence ($\lambda_+ > \lambda_-$) when judging one option relative to another. This asymmetry grows as the options become more dissimilar, as implemented in the exponential attention weighting function. Consequently, the dissimilar option, B, receives very high valuation because it is furthest away from it competitors and benefits most from this overweighting of positive comparisons. Thus, according to MLBA, the similarity effect occurs because individuals overweight positive information relative to negative information and this bias increases with the distance between alternatives.

MLBA's attraction effect. To explain the attraction effect, MLBA posits a *distance-dependent comparison* mechanism in which comparisons between option A and option D receive greater attention because they are similar and therefore difficult to discriminate. Since A dominates D, comparisons generally favor A, and because the options are similar, these are weighted heavily. B is also compared favorably to D, but because it is so easy to discriminate from D, these comparisons receive little weight and A becomes the favored alternative. Thus, according to MLBA, decision makers attend to comparisons between similar options more because they are difficult to discriminate. Ironically, people are thought to overcompensate for this difficulty to the point where they overweight these hard-to-judge comparisons.

MLBA's compromise effect. MLBA accounts for the compromise effect through two mechanisms. The first is extremeness aversion and is represented using a subjective value function. Although MLBA's m parameter can be any positive value, $m > 1$ is typically used to increase the subjective value of options near the middle of the scale. For example, all of the estimated m values reported by Trueblood et al. (2014) are greater than 1. Our analyses yield similar results, with 81% of best fitting m parameters being greater than 1.[1] Thus, although the model does not strictly assume extremeness aversion, in practice this mechanism appears crucial for producing context effects. In this respect, by employing a value function that encompasses preference for intermediate options the compromise effect can be partially "built in" to the model for robust predictions. Second, distance-dependent comparisons help to produce the compromise effect because the compromise option, B, benefits from being closer to its competitors. Since it has a higher value (and because positives are always weighted more than negatives), comparisons with A and C are generally favorable, and because the distances involved in these comparisons are relatively small, these comparisons receive considerable attention. In contrast, when the two extreme options are compared with each other their dissimilarity results in relatively little attention. As Tsetsos, Chater, and Usher (2015) point out, MLBA requires a "fine balance" to simultane-

[1] This included fits to the full dataset and to the compromise-specific subset. This does not include best fitting parameters for the other effect-specific subsets.

ously predict both similarity and compromise effects because in the former case extreme options benefit from their extremeness—due to the asymmetric importance of positive and negative evidence—and in the latter case they are hurt by their extremeness—due to the greater attention paid to similar comparisons. Increasing the curvature of the value function to favor the compromise, that is, increasing extremeness aversion, helps alleviate this fine balance, but moves the model away from a psychological process account of the compromise effect.

In sum, for MLBA to explain the similarity effect, the asymmetric attention weighting mechanism is essential. For the attraction effect distance-dependent comparisons are necessary. For the compromise effect extremeness aversion and distance-dependent comparisons are crucial.

MDFT

DFT (Busemeyer & Townsend, 1993) is a dynamical model—built on cognitive principles—describing a noisy, attention-driven process of deliberation. It has been applied to a wide range of decision-making phenomena including preference reversals under time pressure (Diederich, 2003), multiattribute choice (Diederich, 1997), pricing (Johnson & Busemeyer, 2005), and planning in multistage choice (Hotaling, 2018; Hotaling & Busemeyer, 2012). MDFT is a generalization of DFT designed to account for multialternative context effects. Like MLBA, MDFT assumes that individuals accumulate information in favor of each choice alternative until a decision threshold is reached. However, MDFT describes a temporally extended sampling process that drives information accumulation. At each moment the evidence in favor of each alternative is represented in a preference state vector, $\mathbf{P}_t$, containing the preference values for all options. During deliberation, attribute information is sampled to compute a vector of momentary valences representing the advantages or disadvantages of each alternative over the others, according to

$$V_t = \mathbf{CMW}_t + \varepsilon. \tag{6}$$

Here $\mathbf{M}$ is a matrix containing the attribute values for each option, $\mathbf{C}$ is a contrast matrix for computing advantages and disadvantages across alternatives, and $\mathbf{W}_t$ is a vector containing the attention weights associated with each attribute (see Busemeyer & Diederich, 2002, for a more complete specification). These valences accumulate over time, such that

$$\mathbf{P}_t = \mathbf{SP}_{t-1} + V_t. \tag{7}$$

$\mathbf{S}$ is a feedback matrix with diagonal elements representing the decay of previous preference values over time, and off-diagonal element describing the interactions between alternatives. $\mathbf{S}$ is defined by the equation

$$\mathbf{S} = \delta - \varphi_2 \times \exp(-\varphi_1 \times D^2) \tag{8}$$

where δ is an identity matrix, φ_2 is a decay parameter, and φ_1 is a sensitivity parameter used to calculate the similarity between alternatives as a function of their distance, D, in the attribute space. We use the generalized distance function development by Berkowitsch, Scheibehenne, Rieskamp, and Matthäus (2015). This is based on the idea that psychological distance changes less as an alternative moves along the dimension of preferential indifference (the dotted line in Figure 1) compared with when an alternative moves along the dimension of preferential dominance (a line with a slope of 1 in Figure 1; Hotaling, Busemeyer, & Li, 2010). The parameter $b > 1$ controls the relative weight placed on the dominance versus indifference dimensions.

On an intuitive level, these equations describe a process involving variable input wherein alternatives compete by suppressing activation for one another via distance-dependent lateral inhibition. Two simple assumptions govern the competition between alternatives: (a) more similar options exert a stronger negative effect on each other, and (b) similarity between alternatives decreases faster along the dominance dimension than the indifferences dimension (see Hotaling et al., 2010).

We follow Berkowitsch et al. (2014) in using the internally controlled stopping version of the model where a choice is made only after preferences stabilize. This allows for analytic solutions for choice probabilities, but sacrifices the model's ability to make response time predictions. Although we fit MDFT to some of the same data as Berkowitsch et al. (2014) we have redone the model fitting in an effort to equate the fitting procedures for MLBA and MDFT. We use a version of MDFT with five free parameters. The attention weight, w, represents the relative importance of each attribute to the decision maker. The dominance dimension weight, b, is used to represent the greater sensitivity to differences along the indifference dimension when calculating the distance between alternatives. Sensitivity parameter, φ_1, governs the relationship between distance and similarity. Memory decay parameter, φ_2, controls the influence of previous valences on the preference state. Finally, the noise parameter, σ, represents random input to the system caused by error. See Table 1 for a summary of MDFT's parameters.

MDFT's similarity effect. MDFT's explanation of the similarity effect is quite different from MLBA's. Because *A* and *S* are similar, their activation is correlated. When attention focuses on attribute *Q*, both options will have high valences and *B* will have a low valence. When attention focuses on attribute *P*, both *A* and *S* will have low valences and *B* will have a high valence. As a result *A* and *S* compete with each other, suppress each other's activation, and are chosen less often than the dissimilar option, *B*. Thus, according to MDFT, attention switching and distance-dependent inhibition between alternatives produces the similarity effect. Similar items moderate preference for one another through correlated activation and strong lateral inhibition. Whereas MLBA posits that the comparison between similar options *A* and *S* is difficult and has little impact on deliberation, MDFT proposes that these options compete fiercely and diminish each other over time to produce the effect.

MDFT's attraction effect. In the case of the attraction effect, both models tell a similar story, with MDFT using its distance-dependent inhibition mechanism to explain the effect. Here, *D* will tend to have negative valences due to its relatively low attribute values. This negative activation boosts the activation of the competing alternatives just as positive activation would have inhibited it. However, competition is strongest among similar options, so *A* benefits more than *B*, and is the favored option.

It is worth noting that the models make very different claims about the roles of attention and discriminability in producing the attraction effect. According to MLBA, decision makers attend to comparisons between similar options more because they are difficult to discriminate. Ironically, people are thought to overcompensate for this difficulty to the point where they overweight these hard-to-judge comparisons. MDFT suggests quite the opposite, and posits that decision makers place greater importance on competition between similar options because their similarity makes differences obvious.

MDFT's compromise effect. According to MDFT, the compromise effect also arises due to distance-dependent inhibition because option *B* has an inhibitory effect on both *A* and *C*. Since all three alternatives lie on the equal preference contour (and are chosen equally often in binary choice) the mean valences are equal. However, due to random noise caused by attention to irrelevant information each option will occasionally receive higher activation than the others. For option *B*, the effects of momentary fluctuations in *A* and *C* balances out over time. On the contrary, since *B* is relatively close to both extreme options, it suppresses activation in both whenever its momentary valence is high. Over time this process favors the compromise and causes it to be chosen more often.

In sum, for MDFT to explain the similarity effect the attention switching and distance-dependent inhibition mechanisms are essential. To explain the attraction and compromise effect the distance-dependent inhibition mechanism is crucial.

Quantitative Model Comparison

In principle, MLBA and MDFT are able to predict the three context effects as described above. However, it is an open question whether they describe empirical choice data equally well. To compare the models on the basis of empirical data we rely on the consumer choice experiment by Berkowitsch et al. (2014, Study 2). The experiment involved repeated choices from a set of three consumer products. The sets were selected from six different product domains (e.g., bicycles, vacuum cleaners, cameras, etc.). For each decision, participants saw three products and chose the one they liked best. The set of three options were calibrated to each individual's subjective preferences and attribute importances ahead of time using a matching task where participants created indifference pairs. This allowed the construction of choice triplets corresponding to the classic similarity, attraction, and compromise effects, in terms of their positioning in participants' subjective attribute spaces. Additional details regarding the Berkowitsch et al. (2014) dataset can be found in the Appendix A and at osf.io/26u7z.

Our model comparison has two components. First, we applied each model to the entire dataset to test its ability to simultaneously account for behavior across all three effects. Second, to better examine how and whether the models can explain the specific context effects, we also estimated the models' parameters separately for each effect. A third component—a cross-validation analysis to assess each model's ability to generalize and its susceptibility to overfitting—was unsuccessful for reasons explained below. These analyses provide valuable insights into the abilities and limitations of each model, as well as the tradeoff and interactions between the models' context effect mechanisms. To preview our conclusions, we find that: (a) both MLBA and MDFT have useful mechanisms for producing and explaining multialternative context effects, (b) MDFT's mechanisms allow it greater flexibility to accommodate different patterns of behavior, and (c) MLBA's ability to generalize is significantly limited by negative drift rate errors stemming from its I_0 parameter.

Simultaneously Accounting for Similarity, Attraction, and Compromise Effects

The data from Berkowitsch et al. (2014) are presented in Figure 2. These data are challenging to model for several reasons. First, they involve attraction, similarity, and compromise effects, therefore necessitating that the models' context effect mechanisms work together harmoniously to simultaneously produce each effect. Second, these data aggregate choices across trials from several product conditions. There is good reason to believe that some model parameters would change across product domain because people may distribute their attention or process information differently depending on the domain. For example, MDFT posits that people may give unequal weight to the attributes describing an alternative, and that these weights may vary across product domains. Similarly, the impact of random noise—for example, caused by attending to irrelevant information—is also likely domain dependent. We therefore follow Berkowitsch et al. (2014) in estimating separate weighting parameters, w, and noise parameters, σ, for each product type. This yields a version of MDFT with 15 free parameters: one value each for φ_1, φ_2, and b, and one value of w and σ for each of the six product domains. Although this may appear to be a large number of parameters, please note that for each product domain only five parameters are needed.

For comparison, we used a version of MLBA that makes similar assumptions and has roughly the same level of complexity. As with MDFT, we allowed for the differential weighting of attributes across product domains. We also allowed for the possibility of different subjective value functions across products in case extremeness aversion was domain-specific. This resulted in 15 free parameters: one value of β and m for each of the six product domains, and one value each for λ_+, λ_-, and I_0. Again, this means that for each product domain only five parameters are necessary. Matlab code for fitting both models can be found in the online supplemental materials.

We fit the models to individual data using a maximum likelihood method. For each individual, we found the model parameters that maximized the likelihood of each response.[2] For both models, we used comprehensive grid search plus gradient descent routines for finding optimal parameters. To avoid local minima problems, for each individual we constructed a large (432 point for MLBA, 324 point for MDFT) grid of parameter values as starting points for a bounded Nelder-Mead optimization routine (Nelder & Mead, 1965). Details regarding parameter bounds, start point values, etc. can be found in Appendix B.

We begin our comparison by evaluating the models' overall fit to the full dataset. Quantitatively, MLBA provides the best overall fit: the mean log-likelihood across individuals was higher for MLBA ($\overline{LL}_{\text{MLBA}} = -57.50$) than MDFT ($\overline{LL}_{\text{MDFT}} = -59.09$). However, when examining this result at the individual level it becomes clear that both models do equally well in describing participants' behavior, with 25 (52%) participants best fit by MLBA, compared

[2] Since participants completed different numbers of trials for each context effect, some effects contributed more than others to the overall performance of the models. In Appendix C, we correct for this influence using a weighted fitness function, and show that it does not significantly affect our results.

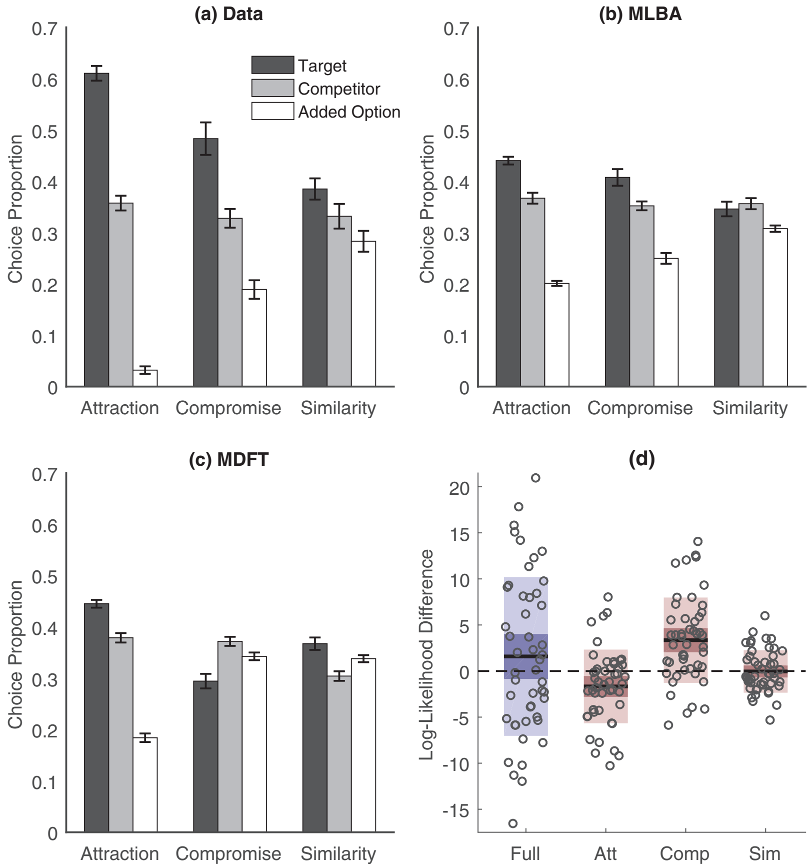

Figure 2. (a) Observed choice shares across the 43 participants in Study 2 of Berkowitsch et al. (2014). Error bars indicate standard errors. (b) Predicted mean choice shares for multiattribute linear ballistic accumulator (MLBA) using parameters estimated from the entire dataset. (c) Predicted mean choice shares for multialternative decision field theory (MDFT) using parameters estimated from the entire dataset. (d) Model comparison of evidence in favor of MLBA based on fits to the entire dataset. Differences were computed by subtracting the maximum log-likelihood according to MDFT from that according to MLBA. Each point represents an individual. Horizontal lines indicate means, dark bands indicate 95% confidence intervals, and light bands indicate standard deviations. Att = attraction effect; Comp = compromise effect; Sim = similarity effect; Full = all effects together. See the online article for the color version of this figure.

with MDFT with 23 (48%) participants. Figure 2d shows the distribution of individual log-likelihood differences produced by the best fitting parameters of each model. Each point represents an individual, and was computed by subtracting the maximum log-likelihood according to MDFT from that according to MLBA. The figure shows that there was considerable variability across individuals with regard to the evidence favoring each model. To better assess the strength of evidence supporting each model, we follow Raftery's (1995) method for using differences in Bayesian information criterion (BIC) values to approximate the relative posterior probability that one of the two models generated the observed data.[3] Probabil-

[3] Since the models have equal numbers of parameters, this essentially reduces to a comparison of maximum likelihoods. However, BIC values provide a means of quantifying the strength of evidence in favor of each model.

ities were classified as weak (.33–.60), positive (.60–.91), strong (.91–.99), and very strong (>.99). The evidence slightly favors MLBA, with 16 (33%) and three (6%) individuals classified as having very strong or strong evidence in favor of MLBA, respectively. However, at the same time, a substantial number of participants yield results favoring MDFT, with 13 (27%) and three (6%) individuals producing very strong or strong evidence favoring MDFT, respectively. These results are consistent with the comparison of Trueblood et al. (2014), who also found evidence in favor of MLBA in comparison to MDFT. However, Trueblood et al. (2014) only examined the models' ability to capture aggregate group-level data, so it is unclear whether their results may have also arisen from roughly equal numbers of participants best fit by each model. One will note that, due to the models having equal numbers of parameters, comparisons of BIC do not reflect any corrections for model complexity. An alternative approach for comparing models while also guarding against overfitting would be to use cross-validation. Unfortunately, for reasons described below, a proper cross-validation analysis was not possible.

Cross-validation analysis. We assessed the models' predictive abilities via a cross-validation analysis (Busemeyer & Wang, 2000). This necessitates the division of data into training and validation sets. For this we used 75% of each individual's responses in each condition for training, with the remaining 25% for validation. Thus, training sets were identical to the full dataset, but with 25% of data removed for later generalization. The same procedure for estimating the models' parameters described above was used to find optimal parameter values for each individual (see Appendix B for additional details). These parameters were then used to calculate predictions for trials in the validation set. Comparing the accuracy of these predictions provides a means of comparing both models' ability to account for context effect data and at the same time implicitly taking the models' complexity into account. An overly complex model might overfit the data of the training set, so that it would perform badly when making out-of-sample predictions. The entire cross-validation process was repeated 10 times, with each making a new random allocation of responses to training and validation sets.

We had hoped to compare the accuracy of each model with respect to the 10 validation sets, but unfortunately, a satisfactory model comparison using cross-validation was not possible. This was due to problems with MLBA's I_0 parameter, stemming from the fact that this parameter is defined in terms of the other parameters. Recall that I_0 is a constant term introduced to avoid situations where all accumulators have negative drift rates and a decision is never made. In practice, this complicates parameter estimation for MLBA because small changes in I_0 can produce large changes in the model's predictions. More importantly, the dependence of I_0 on the rest of the model effectively prevents generalization because parameters estimated from one dataset often produce negative drift rate errors with another dataset. That is, the optimal I_0 estimated from one set of responses is often insufficient to prevent errors across another set of responses. Although a small error rate could be overlooked as noncritical, we observed far too many errors to adequately evaluate MLBA's generalization. Across the 48 individuals and 10 data sets, 29% of MLBA's best-fitting parameter sets yielded uninterpretable predictions, with negative drift rate errors preventing generalization for 43 individuals (90%).

In an effort to prevent negative drift rate errors a priori, we also attempted the cross-validation analysis on a version of MLBA with a fixed value of $I_0 = 5$. This value is much greater than usual, so we hoped that we might prevent the negative drift rate problems. Unfortunately, this too failed because fixing I_0 to a large value during the optimization routine caused the other model parameters to adjust such that the behavior of the model was relatively unchanged, and negative drift rate errors persisted. For these reasons a satisfactory cross-validation analysis was not possible. This issue is substantial considering that this is a relatively small generalization task. All stimuli use the same range of attributes, calibrated to the same individuals, producing responses in the same experiment. MLBA's difficulty generalizing across subsets of this data suggests that comparable problems will likely arise when trying to generalize across less homogeneous data. This presents a significant obstacle for future efforts to generalize MLBA's predictions to genuinely novel, unobserved data.

A More Detailed Comparison of Models

Why are nearly equal numbers of people best fit by each model? Figure 2 shows that, although both models can approximate the observed mean preferences, they have some difficulties. Considering each context effect separately can shed light on the mixed results above. In this section we analyze the data of Berkowitsch et al. (2014) in more detail by examining the three context effects separately. First, we revisit the above comparison—using the same parameter estimates—but consider the model fits separately for each effect. We also use correlations between context effects to examine how each model's mechanisms work in concert to produce each effect. Second, we separately estimated the parameters of the models again by using only the data of one effect at a time. Throughout this section, we highlight key differences between the models, both in terms of goodness of fit and theoretical explanation.

Attraction effect. Figure 2 indicates that the choice pattern of the attraction effect is correctly captured by both models, though the size of the effect is underestimated. The left panel of Figure 3 compares the predicted and observed mean choice proportions for each individual's attraction trials. This analysis demonstrates that both models produce less extreme choice proportions (i.e., smaller effects) at the individual level, though MDFT yields a smaller mean deviation. When only considering choice data of attraction trials 33 people (69%) were best described by MDFT ($\overline{LL}_{\text{MDFT}} = -19.78$), compared with 15 (31%) who were best described by MLBA ($\overline{LL}_{\text{MLBA}} = -21.45$). Again, Figure 2d indicates that the evidence was stronger for some people than others. Using the method described above, 14 people yielded strong or very strong evidence in favor of MDFT, compared with five people with strong or very strong evidence favoring MLBA. This result supports the attraction effect mechanism of MDFT: its process of attention switching and distance-dependent lateral inhibition provided the better account of choice data for the majority of people.

In addition, we estimated the parameters of both models using only choices from attraction trials. Figure 4 shows the distribution of individual log-likelihood differences based on separate fits to each effect. Looking at the leftmost column, it is clear that MDFT is better able to fit choice data from attraction trials when this is its sole focus ($\overline{LL}_{\text{MDFT}} = -14.65$ vs. $\overline{LL}_{\text{MLBA}} = -17.96$). In support of this, 42 people (88%) were best described by MDFT, with 23 having very strong or strong evidence. Only six people were best described by MLBA, with one individual having strong evidence. If we consider the MLBA's parameters (Table 2), we see that some values changed substantially when fitting specifically to attraction data. The median I_0 was smaller for the attraction-only fit, suggesting that drift rates were generally higher. There was also greater curvature to the subjective value functions (i.e., m values increased), indicating greater biases in favor of moderate op-

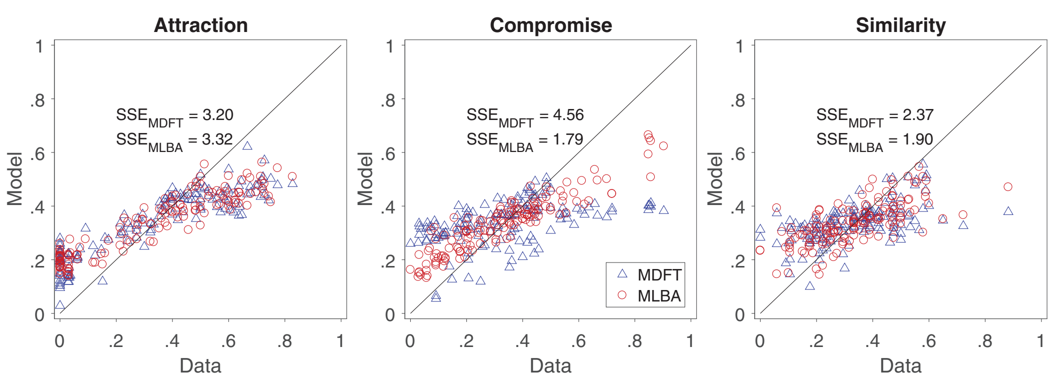

Figure 3. Comparison of observed and predicted individual mean choice proportions, separated by context effect condition. Model predictions are based on parameter estimates using the complete dataset. See the online article for the color version of this figure.

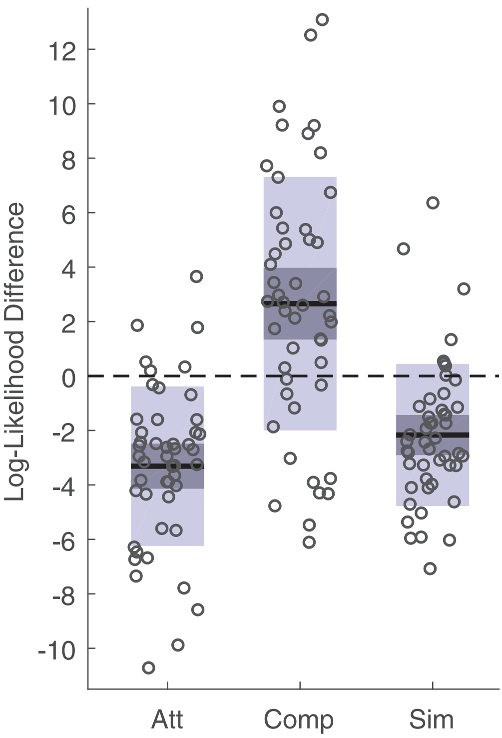

Figure 4. Model comparison of evidence in favor of multiattribute linear ballistic accumulator (MLBA) based on separate fits to each context effect subset. Each point represents an individual. Values were computed by subtracting the maximum log-likelihood according to MDFT from that according to multiattribute linear ballistic accumulator. Att = attraction effect; Comp = compromise effect; Sim = similarity effect. See the online article for the color version of this figure.

tions. If we consider the attraction-specific parameters of MDFT, a different pattern of changes emerges. Table 3 shows increased lateral inhibition (higher φ_1) and a greater importance placed on the dominance, rather than the indifference, dimension (higher *b*). Both of these changes aid the model in predicting the large attraction effect observed in the data. The reliability of this effect in the data may also explain the reduced levels of noise indicated by lower σ values.

Compromise effect. Figure 2 shows that MLBA reproduces the compromise effect, while MDFT predicts that people will, on average, prefer an extreme option over the compromise. The middle panel of Figure 3 shows that both models again produce less extreme choice proportions than were observed at the individual level. Not surprisingly, MLBA also yields a better fit to individual compromise data ($SSE_{\mathrm{MLBA}} = 1.79$ vs. $SSE_{\mathrm{MDFT}} = 4.56$). When only considering compromise trials, 38 people (79%) were best described by MLBA ($\overline{LL}_{\mathrm{MLBA}} = -21.69$), compared with 10 (21%) who were best described by MDFT ($\overline{LL}_{\mathrm{MDFT}} = -25.04$). For 24 individuals there was either strong or very strong evidence in favor of MLBA, while only four individuals produced comparable evidence favoring MDFT.

Estimating model parameters using only compromise trials yields similar results (see Figure 4). Across these data MLBA provides a better fit ($\overline{LL}_{\mathrm{MDFT}} = -21.86$ vs. $\overline{LL}_{\mathrm{MLBA}} = -19.21$), with strong or very strong evidence favoring the model for 20 individuals (42%), compared with eight individuals (17%) with strong or very strong evidence favoring MDFT.

Why is MLBA better able to capture the compromise effect than MDFT? A simple interpretation is that MLBA's compromise effect mechanism is a better reflection of the psychological processes that give rise to compromise effects. Although this is a reasonable conclusion, it is noteworthy that the two models take different approaches to explaining compromise effects. For MDFT, compromise effects emerge slowly, due to the correlations between valances

Table 2
Median Best Fitting Parameters of MLBA

Effect	β_1	β_2	β_3	β_4	β_5	β_6	m_1	m_2	m_3	m_4	m_5	m_6	I_0	λ_+	λ_-
All data	1.56	3.08	.74	.76	4.36	7.20	3.57	1.68	2.35	1.95	1.47	1.42	1.57	.63	.17
Attraction	1.06	2.26	.55	.78	3.69	7.54	14.65	1.63	3.03	3.56	1.49	1.34	.50	.59	.22
Compromise	1.21	3.44	.87	.58	3.37	6.29	2.63	1.49	2.15	2.42	1.34	1.43	2.01	.82	.14
Similarity	3.54	4.15	2.35	2.19	4.82	3.56	17.92	4.35	10.89	13.75	7.01	8.25	1.03	.51	.32

Note. MLBA = multiattribute linear ballistic accumulator.

Table 3
Median Best Fitting Parameters of MDFT

Effect	w_1	w_2	w_3	w_4	w_5	w_6	σ_1	σ_2	σ_3	σ_4	σ_5	σ_6	φ_1	φ_2	b
All data	.66	.61	.31	.45	.57	.69	.31	.18	.29	.14	.13	.10	.05	76.14	1.98
Attraction	.61	.60	.37	.49	.68	.71	.00	.00	.08	.00	.00	.00	.07	333.31	21.08
Compromise	.77	.55	.33	.54	.67	.76	.97	.28	.91	1.55	2.37	.43	.05	38.88	6.91
Similarity	.66	.60	.54	.50	.53	.73	26.63	.00	24.34	.01	.02	22.63	.14	4.22	30.40

Note. MDFT = multialternative decision field theory.

across time. This leads to a "delicate" effect that only appears under particular circumstances. As these results show, MDFT is therefore limited in its ability to predict compromise effects. In contrast, MLBA takes a more direct approach, and uses a subjective value function implying that individuals prefer options with less attribute dispersion. In this sense, compromise options have an inherent advantage because greater preference for them is already built in to their subjective values rather than emerging from the dynamics of MLBA. As noted earlier, Trueblood et al. (2014) allow for $m < 1$, however we find that, across all of the best-fitting parameter set reported here that involve compromise trials, 81% involve $m > 1$. These finding are consistent with those reported by Trueblood et al. (2014) and Tsetsos et al. (2015), so it is clear that, in practice, MLBA employs its subjective value function to account for compromise effects. In the future, it may be interesting to explore ways of incorporating similar subjective value transformation components into other theoretical frameworks, such as MDFT.

If we consider how the models' parameters change when fit to compromise trials only, we also see a divergence. For MDFT, the parameters from the compromise-specific fits show a greater impact of random noise on sampling (i.e., higher σ values). This is sensible, as MDFT produces the compromise effect through sampling of irrelevant information (i.e., noise). Table 3 also shows that competition between alternatives was weaker (i.e., lower φ_1). This may reflect the delicate nature of MDFT's compromise effect, whereby high levels of lateral inhibition produce too much competition and therefore risk swamping preference for the compromise option. In contrast, for MLBA, median parameter values do not change markedly. This is particularly surprising in the case of the m parameters. One might expect to find greater curvature of the subjective attribute space in favor of intermediate options, but this is not the case. Considering that MLBA made accurate compromise effect predictions when fit to the full dataset, this suggests that the model's performance—with respect to the compromise effect—was already "at ceiling." Apparently, increasing m did not improve the model's ability to capture the effect when fit only to compromise trials.

Similarity effect. Figure 2 shows that both models have difficulty accounting for the aggregate similarity effect data. MLBA predicts slightly higher preference for the competitor over the target, thus producing no similarity effect. While MDFT does capture the similarity effect—defined by preference for the target over the competitor—it erroneously predicts that the added option will be chosen more often than the competitor. Using the parameters fit to the full dataset, MDFT gives the better account of responses on similarity trials ($\overline{LL}_{\text{MDFT}} = -14.19$ vs. $\overline{LL}_{\text{MLBA}} = -14.24$). Comparing the models at the individual level also reveals a slight advantage for MDFT. Figure 2d shows that more individuals (27, 56%) were best fit by MDFT, compared with MLBA (21, 44%). Surprisingly, Figure 2b and 2c—which distill each individual's response to one set of choice proportions per condition—misrepresent this result, instead indicating a better fit for MLBA ($SSE_{\text{MLBA}} = 1.90$ vs. $SSE_{\text{MDFT}} = 2.37$). This serves as another reminder of the limitations of analyzing and modeling aggregate context effect choice data.

However, Figure 4 highlights the greater flexibility of MDFT.[4] Refitting the models using only similarity trials yields strong support for MDFT at the individual level ($\overline{LL}_{\text{MDFT}} =$

[4] By flexibility (or model complexity) we refer to the potential number of data patterns that a model can produce.

-9.58 vs. $\overline{LL}_{\text{MLBA}} = -11.75$). Forty-one individuals (85%) were best fit by MDFT, with 17 having very strong or strong evidence, while only seven individuals were best fit by MLBA, with three having very strong or strong evidence.

Table 2 shows some extreme m values for MLBA when estimated with the data of the similarity effect. These large m values imply the most intensely curved subjective value functions that we observe in our modeling. This result was unexpected, as the nonlinear value function is not supposed to be crucial for producing similarity effects. For MDFT, Table 3 shows quite high levels of random noise for Products 1, 3, and 6. This may indicate behavioral inconsistencies across products. Because the model assumes the same psychological process for all products—with only attention-weighting differences—it may attempt to account for inconsistencies using greater attention to irrelevant information for some products. Memory decay rate (φ_2) was also highest for the similarity effect data. A higher rate of memory decay diminishes the long-term impact of competition between alternatives, decreasing the size of context effects. This is sensible when one considers that the similarity effect observed by Berkowitsch et al. (2014) was small.

Relationships between context effect mechanisms. A central goal of the present comparison is to examine the role that various psychological mechanisms play in explaining context effects. We therefore consider the predictions each model makes regarding the relationships between context effects. We quantify each effect for each individual as the relative choice share of the target, RST $= p(T)/[p(T) + p(C)]$, where $p(T)$ is the probability of choosing the target and $p(C)$ is the probability of choosing the competitor. Figure 5 shows the observed and predicted correlations between RST values for each pair of context effects. Participants' responses indicate a positive relationship between attraction and compromise ($r_{\text{Data}} = 0.49$), and negative relationships between compromise and similarity ($r_{\text{Data}} = -0.58$), and between attraction and similarity ($r_{\text{Data}} = -0.53$).

Although both models generally produce smaller context effects than were observed, they both display the correct qualitative pattern of relationships between effects. The left panel of Figure 5 shows that MLBA—which produces the stronger compromise effects—overestimates the relationship between attraction and compromise, while MDFT underestimates the relationship by roughly the same amount. This may indicate that MLBA's distance-dependent comparisons and extremeness aversion tend to work together harmoniously, such that when the former is strong—for the purpose of predicting attraction effects—the latter tends to also be strong—for the purpose of predicting compromise effects. In contrast, it appears that MDFT's distance-dependent inhibition mechanism—the primary source of both attraction and compromise effects—is more limited in its ability to simultaneous produce both effects. The middle

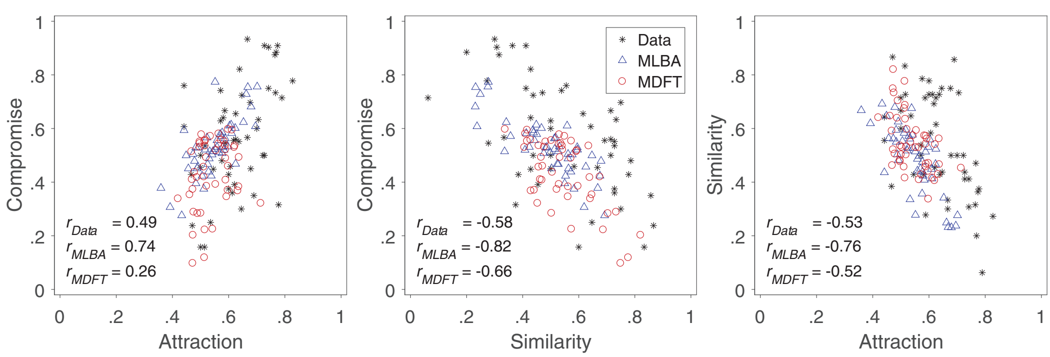

Figure 5. Observed and predicted relationships between individual relative choice share of the target (RST) values across context effect conditions. MLBA = multiattribute linear ballistic accumulator; MDFT = multialternative decision field theory. See the online article for the color version of this figure.

panel of Figure 5 shows that both models overestimate the negative relationship between similarity and compromise effects, though MLBA's prediction is considerably further from the observed correlation. This points to a greater independence between the psychological mechanisms in MDFT, in this case attention switching and distance-dependent inhibition. MLBA's asymmetric attention weighting and extremeness aversion instead appear to work in opposition to one another, with large asymmetries—necessary for predicting similarity effects—associated with weak extremeness aversion—leading to smaller predicted compromise effects. Finally, the right panel of Figure 5 indicates that MLBA again overestimates the negative relationship between attraction and similarity, while MDFT produces a correlation very similar to what is observed. This suggests a high degree of competition between MLBA's asymmetric attention weighting and distance-dependent comparison mechanisms when simultaneously fitting both effects. Of course, the conclusions above do not generalize to all parameter settings of MLBA and MDFT. This analysis does, however, demonstrate how the models' mechanisms interact in practice when fitting preferential multialternative choice data. Additional work will be needed to establish how representative these results are.

Model Generalization Across Context Effects

Another important feature of any psychological model of multialternative choice is its ability to make predictions out of a sample. A potentially illuminating analysis would be to examine how the effect-specific parameters of each model generalize to responses in the other two context effect conditions. For example, studying the effects that fitting MLBA to attraction trials has on its predictions for compromise trials would shed new light on how the model's mechanisms interact and tradeoff to account for different patterns of data. Unfortunately, this comparison is not possible due to the same negative drift rate errors that prevented MLBA's cross-validation. Again, these problems were too common to overlook. When using MLBA's attraction-specific parameters to calculate predictions on compromise and similarity trials, only nine individuals (19%) yielded interpretable results. For the remaining 39 individuals, generalizing attraction-specific parameters to these new trials resulted in trials without positive drift rates. Using the similarity-specific parameters for generalization, only 23 individuals (48%) yielded interpretable results, with the remaining 25 yielding negative drift rate errors. Errors prevented generalization of compromise-specific parameters for 15 people (31%). Again, these problems highlight an important limitation of MLBA to be addressed with future research.

Discussion

The goal of the present work was to compare two prominent models of decision making against each other in predicting three context effects that have received substantial attention in the literature. Contrary to past work we did not only test the theoretical ability of the models to predict the attraction, compromise, and similarity effect, but also used empirical choice data to compare the models. Although both models are, in principle, capable of predicting the three effects simultaneously, our results show a mixture of successes and failures in describing the observed behavior. MDFT is able to capture the attraction and similarity effects accurately, but is less successful in simultaneously producing the compromise effect. In contrast, MLBA is quite good at predicting the compromise effect, but has some difficulties in explaining the attraction and similarity effect at the same time. Across the entire dataset, the models were evenly matched, with roughly half of participants best described by each. When fitting data, MDFT appeared to be the more flexible model as shown by its attraction- and similarity-specific performance. Additionally, it predicts relatively low correlations between context effects, suggesting that its mechanisms are able to operate independently. In contrast, MLBA predicted stronger correlations between effects, indicating that the underlying mechanisms may also interact in important ways. Our findings allow us to examine which psychological mechanisms appear essential to describe the cognitive processes underlying multialternative choice and the context effects arising therein (see below).

The fact that these results contrast with those reported by Trueblood et al. (2014) is noteworthy. One potential explanation revolves around our use of data from a different choice domain. Berkowitsch et al. (2014) examined preferential choices—which is the standard domain of the considered context effects (cf., Heath, Chatterjee, & France, 1995)—while Trueblood et al. (2014) tested their model with data from more perceptual or numerical choice tasks. Our divergent results may reflect differences in the decision process across these domains. Moreover, on the basis of past findings alone, one might expect a priori that a model built on the foundation of LBA—which is typically applied to low-level decision tasks—would perform better in more perceptual domains, and that a model built on the foundation of DFT would excel in the preferential choice domain where it has enjoyed previous success.

Another obstacle to comparing our results to those of Trueblood et al. (2014) stems from the fact that we fit the models separately to each individual rather than aggregate group-level data. As we discuss below, we think with multialternative choice one must be cautious both when drawing conclusions from group data, and when fitting and comparing models using aggregated data. The present results support the idea that insights gleaned from group-level model comparisons can diverge from those originating in individual-level analyses.

Psychological Mechanisms for Explaining the Effects

Three mechanisms stand out as critical for explaining the three context effects. MDFT's attention-switching mechanism is crucial for producing the similarity effect. Initially introduced by Tversky (1972), it explains how correlated activation for the competitor and added options causes the relative preference for the target to increase as attention shifts between attributes. A similar mechanism for shifting attention between distinct cues can also be found in the attention drift-diffusion model (Krajbich, Armel, & Rangel, 2010; Krajbich, Lu, Camerer, & Rangel, 2012; Krajbich & Rangel, 2011).

For the attraction effect, MDFT's distance-dependent inhibition mechanism plays a key role. Because the added option is similar to, but dominated by, the target, its negative activation "boosts" the activation of the target more than the dissimilar competitor option. Similar inhibition mechanisms have also proved valuable to models of learning (Grossberg, 1987) and perception (McClelland & Rumelhart, 1981), as well as other choice models (Usher & McClelland, 2001).

Finally, MLBA's success in accounting for the compromise effect suggests that its subjective value function, with $m > 1$ to create extremeness aversion, constitutes a much more robust mechanism for producing compromise effects. This contrasts sharply with MDFT's compromise mechanism, which is quite delicate and only emerges under certain circumstances. Although MLBA says little about how or why extremeness aversion should occur, the component seems crucial for producing the observed preferences, and requires less fine tuning compared with MDFT's compromise effect mechanisms.

Measurement Models and Cognitive Process Models

One important issue for theories of choice behavior concerns the question of whether a model is proposed as a measurement model or as a cognitive model describing the underlying decision process. For MDFT, it is clear that the model posits an attention-driven process whereby individuals probabilistically sample attribute information. Choice alternatives are compared and compete with one another, causing preferences to evolve dynamically over the course of deliberation until a decision threshold is reached. Thus, the model uses low-level cognitive mechanisms to explain how preferences emerge from a dynamical process.

In contrast, MLBA is presumably perceived less as a model of cognitive processes because its origins lie in the LBA model, which was conceived as an approximation to the drift-diffusion model (DDM). The virtues of the LBA are primarily that it is computationally convenient, compared with the DDM, and captures the most important RT phenomena observed in decision tasks. DDM and LBA are primarily used as measurement models, which allow one to identifying how changes in the decision task affect the processing of information.

LBA is not sufficient to explain context effects, leading Trueblood and colleagues (2014) to propose MLBA, which adds a front-end process. Although this is a sensible move, it is unclear how this front end operates at a process level because it is determined a priori before it "enters" the back-end process of LBA. One possibility is that MLBA's front end also serves as a component for a measurement model without describing a cognitive process. If this is true and MLBA should be considered a pure measurement model, it may be limited in its ability to make predictions for how people behave in novel choice situations.

For this reason, we think of the models as competitors, with different objectives. MDFT aims to explain context effects from the bottom up through the aggregation of many cognitive microprocesses. MLBA is more focused on accurately measuring behavior in terms of psychologically interesting variables, in the spirit of the DDM and LBA. Also, as Tsetsos et al. (2015) point out, MLBA's modular nature—with comparison and competition segregated from accumulation of evidence—allows for careful testing of alternative comparison mechanisms. Future work may exploit this property of MLBA to examine competing assumptions regarding how people judge options against one another, while holding constant the assumptions regarding the accumulation of information, or even replacing the model's back end with a simple-choice rule.

This also highlights the possibility of taking an entirely different approach to testing psychological mechanisms of context effects. A better solution, from the perspective of model selection, might be to develop a grand model that included all relevant mechanisms. Comparing various nested versions of this model would allow researchers to separate the unique influence of each mechanism from the more peripheral characteristics and assumptions that complicate comparisons of models constructed within different theoretical frameworks (see Turner, Schley, Muller, & Tsetsos, 2017 for recent work in this vein).

Time Course and Dynamics of Decision Making: RT Predictions

Models like MDFT and the leaky competing accumulator explain context effects as emerging online over the time course of deliberation. LBA, on the contrary, was designed to ignore this aspect of deliberation. Thus, if we only focus on final, self-terminated choices MLBA might have an advantage over MDFT. However, when also considering the time course for how the effects develop during deliberation MDFT could have an advantage. We did not examine this in the present work as MLBA makes no predictions about preferences during deliberation. In a sense, this issue of whether or not a model is responsible for explaining process data is reminiscent of comparisons between expected utility models and cognitive choice models. Each approach has its own virtues, but the type (and degree) of explanation is quite different. Like expected utility models in economics, MLBA does not aim to explain how different context effects evolve over time. It remains to be seen how important a limitation this is to using MLBA as a tool and theory.

Limitations of Model Comparisons Using Aggregate Data

One striking result from the present investigation involves the use of aggregate data. Because it is often difficult to repeat identical trials in many decision-making domains (e.g., consumer choice), researchers typically collapse data across similar conditions that are assumed to be equivalent. Also, in an effort to "average out" the inherent variability of people's decisions, data is commonly analyzed at the group level. Our results suggest that these aggregate data can be misleading (see also Liew et al., 2016, and Spektor, Kellen, & Hotaling, 2018, for demonstrations of the factors influencing the robustness of context effects). Although MLBA typically gave the best account of context effects across the full dataset, this result obscured a complex mixture of results across conditions and individuals. In many cases the evidence in favor of MDFT was substantial, but difficult to detect on the aggregate. This divergence between individual and aggregated data, and between quantitative and qualitative fit is understandable when one considers that the results shown in Figure 2a reduce the entire experiment to just nine data points. Ignoring the richness of one's experimental design to neatly summarize results is surely a problematic approach for anyone interested in making a detailed study of the role of experimental factors on decision making. It is therefore important to pay close attention to these issues, especially when developing and testing cognitive process models of individual decision makers.

Moving forward, there is a need for better data and experimental paradigms for rigorously comparing models of multialternative choice. Unfortunately, this can be difficult within the preferential choice domain because it is hard to collect enough data from each participant. Preferential choices (e.g., between vacuum cleaners) are difficult to repeat within an individual because participants will likely remember their previous decisions. One could avoid the need to repeat by using many unique stimuli, however this is also problematic because each new stimulus risks introducing more irrelevant information to the task via participants' prior beliefs. As a field, we must overcome these limitations if we are to properly test theories of how individuals make preferential choices. Perhaps one promising avenue for future work lies in the use of hierarchical models where individual-level parameters are represented as samples from group-level distributions, rather than independent estimates. Such an approach can leverage regularities across individuals to improve the accuracy and robustness of parameter estimates, despite small amounts of data for each individual (e.g., Nilsson, Rieskamp, & Wagenmakers, 2011).

Considering that MLBA and MDFT were both designed to account for the similarity, attraction, and compromise effects, we should not be surprised that they perform reasonably well when applied to such data. This highlights the need for new data from novel multialternative choice phenomena where the models make divergent predictions. Through systematic investigation of situations in which each model succeeds and fails we can develop our understanding of when and how various psychological mechanisms impact our decisions.

Conclusion

Multialternative context effects are important because these violations of rational choice axioms give insight into how people compare options to one another when making decisions. Computational models can serve as valuable tools for investigating the mental processes that give rise to these decisions. Exactly how the decision process unfolds is an important topic of recent research. The present work builds on this by identifying mechanisms that appear to be essential when constructing a comprehensive description of the cognitive processes underlying decision making.

Although both MDFT and MLBA are capable of simultaneously producing the similarity, attraction, and compromise effects, we find a mixture of evidence in favor of each model's account of empirical data from a within-subject consumer choice experiment. By focusing our investigation on two prominent models we are able to dive more deeply into the inner workings of each model. Examining each effect individually, we find that MDFT gives the better fit to the similarity and attraction effects, while MLBA performs better on the compromise effect. These results support the idea that attention switching, distance-dependent inhibition, and extremeness aversion are important mechanisms to include in models of multialternative decision making.

References

Berkowitsch, N. A., Scheibehenne, B., & Rieskamp, J. (2014). Rigorously testing multialternative decision field theory against random utility models. *Journal of Experimental Psychology: General, 143,* 1331–1348. http://dx.doi.org/10.1037/a0035159

Berkowitsch, N. A., Scheibehenne, B., Rieskamp, J., & Matthäus, M. (2015). A generalized distance function for preferential choices. *British Journal of Mathematical and Statistical Psychology, 68,* 310–325. http://dx.doi.org/10.1111/bmsp.12048

Bhatia, S. (2013). Associations and the accumulation of preference. *Psychological Review, 120,* 522–543. http://dx.doi.org/10.1037/a0032457

Brown, S. D., & Heathcote, A. (2008). The simplest complete model of choice response time: Linear ballistic accumulation. *Cognitive Psychology, 57,* 153–178. http://dx.doi.org/10.1016/j.cogpsych.2007.12.002

Busemeyer, J. R., & Diederich, A. (2002). Survey of decision field theory. *Mathematical Social Sciences, 43,* 345–370. http://dx.doi.org/10.1016/S0165-4896(02)00016-1

Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review, 100,* 432–459. http://dx.doi.org/10.1037/0033-295X.100.3.432

Busemeyer, J. R., & Wang, Y.-M. (2000). Model comparisons and model selections based on generalization criterion methodology. *Journal of Mathematical Psychology, 44,* 171–189. http://dx.doi.org/10.1006/jmps.1999.1282

Chernev, A. (2004). Extremeness aversion and attribute-balance effects in choice. *Journal of Consumer Research, 31,* 249–263. http://dx.doi.org/10.1086/422105

Diederich, A. (1997). Dynamic stochastic models for decision making under time constraints. *Journal of Mathematical Psychology, 41,* 260–274. http://dx.doi.org/10.1006/jmps.1997.1167

Diederich, A. (2003). MDFT account of decision making under time pressure. *Psychonomic Bulletin & Review, 10,* 157–166. http://dx.doi.org/10.3758/BF03196480

Grossberg, S. (1987). Competitive learning: From interactive activation to adaptive resonance. *Cognitive Science, 11,* 23–63. http://dx.doi.org/10.1111/j.1551-6708.1987.tb00862.x

Heath, T. B., & Chatterjee, S. (1995). Asymmetric decoy effects on lower-quality versus higher-quality brands: Meta-analytic and experimental evidence. *Journal of Consumer Research, 22,* 268–284. http://dx.doi.org/10.1086/209449

Heath, T. B., Chatterjee, S., & France, K. R. (1995). Mental accounting and changes in price: The frame dependence of reference dependence. *Journal of Consumer Research, 22,* 90–97. http://dx.doi.org/10.1086/209437

Hotaling, J. M. (2018). *Decision field theory-planning: A cognitive model of planning on the fly*. Manuscript submitted for publication.

Hotaling, J. M., & Busemeyer, J. R. (2012). DFT-D: A cognitive-dynamical model of dynamic decision making. *Synthese, 189,* 67–80. http://dx.doi.org/10.1007/s11229-012-0157-0

Hotaling, J. M., Busemeyer, J. R., & Li, J. (2010). Theoretical developments in decision field theory: Comment on Tsetsos, Usher, and Chater (2010). *Psychological Review, 117,* 1294–1298. http://dx.doi.org/10.1037/a0020401

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. http://dx.doi.org/10.1086/208899

Johnson, J. G., & Busemeyer, J. R. (2005). A dynamic, stochastic, computational model of preference reversal phenomena. *Psychological Review, 112,* 841–861. http://dx.doi.org/10.1037/0033-295X.112.4.841

Jones, M., & Dzhafarov, E. N. (2014). Unfalsifiability and mutual translatability of major modeling schemes for choice reaction time. *Psychological Review, 121,* 1–32. http://dx.doi.org/10.1037/a0034190

Krajbich, I., Armel, C., & Rangel, A. (2010). Visual fixations and the computation and comparison of value in simple choice. *Nature Neuroscience, 13,* 1292–1298. http://dx.doi.org/10.1038/nn.2635

Krajbich, I., Lu, D., Camerer, C., & Rangel, A. (2012). The attentional drift-diffusion model extends to simple purchasing decisions. *Frontiers in Psychology, 3,* 193. http://dx.doi.org/10.3389/fpsyg.2012.00193

Krajbich, I., & Rangel, A. (2011). Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. *Proceedings of the National Academy of Sciences of the United States of America, 108,* 13852–13857. http://dx.doi.org/10.1073/pnas.1101328108

Liew, S. X., Howe, P. D., & Little, D. R. (2016). The appropriacy of averaging in the study of context effects. *Psychonomic Bulletin & Review, 23,* 1639–1646. http://dx.doi.org/10.3758/s13423-016-1032-7

McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: I. An account of basic findings. *Psychological Review, 88,* 375–407. http://dx.doi.org/10.1037/0033-295X.88.5.375

Nelder, J. A., & Mead, R. (1965). A simplex method for function minimization. *The Computer Journal, 7,* 308–313. http://dx.doi.org/10.1093/comjnl/7.4.308

Nilsson, H., Rieskamp, J., & Wagenmakers, E.-J. (2011). Hierarchical Bayesian parameter estimation for cumulative prospect theory. *Journal of Mathematical Psychology, 55,* 84–93. http://dx.doi.org/10.1016/j.jmp.2010.08.006

Noguchi, T., & Stewart, N. (2014). In the attraction, compromise, and similarity effects, alternatives are repeatedly compared in pairs on single dimensions. *Cognition, 132,* 44–56. http://dx.doi.org/10.1016/j.cognition.2014.03.006

Raftery, A. E. (1995). Bayesian model selection in social research. *Sociological Methodology, 25,* 111–164. http://dx.doi.org/10.2307/271063

Rieskamp, J., Busemeyer, J. R., & Mellers, B. A. (2006). Extending the bounds of rationality: Evidence and theories of preferential choice. *Journal of Economic Literature, 44,* 631–661. http://dx.doi.org/10.1257/jel.44.3.631

Roe, R. M., Busemeyer, J. R., & Townsend, J. T. (2001). Multialternative decision field theory: A dynamic connectionist model of decision making. *Psychological Review, 108,* 370–392. http://dx.doi.org/10.1037/0033-295X.108.2.370

Simonson, I. (1989). Choice based on reasons: The case of attraction and compromise effects. *Journal of Consumer Research, 16,* 158–174. http://dx.doi.org/10.1086/209205

Spektor, M. S., Kellen, D., & Hotaling, J. M. (2018). When the good looks bad: An experimental exploration of the repulsion effect. *Psychological Science*. Advance online publication. http://dx.doi.org/10.1177/0956797618779041

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. http://dx.doi.org/10.1037/a0036137

Trueblood, J. S., Brown, S. D., & Heathcote, A. (2015). The fragile nature of contextual preference reversals: Reply to Tsetsos, Chater, and Usher (2015). *Psychological Review, 122,* 848–853. http://dx.doi.org/10.1037/a0039656

Tsetsos, K., Chater, N., & Usher, M. (2015). Examining the mechanisms underlying contextual preference reversal: Comment on Trueblood, Brown, and Heathcote (2014). *Psychological Review, 122,* 838–847.

Turner, B. M., Schley, D. R., Muller, C., & Tsetsos, K. (2017). Competing theories of multialternative, multiattribute preferential choice. *Psychological Review, 125,* 329–362.

Tversky, A. (1972). Elimination by aspects: A theory of choice. *Psychological Review, 79,* 281–299. http://dx.doi.org/10.1037/h0032955

Tversky, A., & Simonson, I. (1993). Context-dependent preferences. *Management Science, 39,* 1179–1189. http://dx.doi.org/10.1287/mnsc.39.10.1179

Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review, 108,* 550–592. http://dx.doi.org/10.1037/0033-295X.108.3.550

Usher, M., & McClelland, J. L. (2004). Loss aversion and inhibition in dynamical models of multialternative choice. *Psychological Review, 111,* 757–769. http://dx.doi.org/10.1037/0033-295X.111.3.757

Appendix A

Study 2 From Berkowitsch et al. (2014)

Forty-eight participants made repeated choices between products described by two attributes across two experimental sessions. The first session was designed to find pairs of options for which the participant had equal preference. The second session systematically introduced new items to these choice sets in order to elicit context effects.

In the calibration session, participants were presented with two items. One of these was missing an attribute value (e.g. price), and the participant was asked to fill in the value that would make them indifferent between the two items. This resulted in 108 indifference pairs (eighteen per product) for each individual.

A few days later, in the second session, matched indifference pairs were combined with new options carefully designed to produce context effects (see Berkowitsch et al., 2014, Figure 3). On each trial one new option was added. The locations of these additional options were counterbalanced to produce six attraction, compromise, and similarity trials for each matched pair. Unfortunately, due to rounding, some trials no longer constituted a proper context effect triplet, and were not analyzed. Also, approximately half of the intended similarity effect choice triplets became attraction effect triplets because the added option no long laid on the line of difference. Counterbalancing assured that each item in each matched pair was the target for half of the trials. This resulted in a mean of 32.17 (*SD* = 2.91), 32.63 (*SD* = 3.04), and 18.19 (*SD* = 2.17) trials for attraction, compromise, and similarity effects, respectively.

(*Appendices continue*)

Appendix B

Model Fitting Procedure

Equivalent procedures were used for both models, across all datasets. Parameters were optimized using the maximum likelihood method with a Nelder-Mead algorithm implemented in MATLAB's fminsearch function. For each individual a large grid of parameter values was used to define starting points for the optimization procedure. Since each model used 15 free parameters the optimization procedure began with a 15-dimensional grid of start points. From each start point the algorithm used gradient-decent to find parameter values that minimized the negative log-likelihood of the data given the model. Each search was limited to 2,000 iterations, though searches typically stabilized much earlier.

We selected values for the start point grids with the aim of comprehensively searching the parameter space and avoiding local minima. Our goal was to begin the optimization procedure at many locations so that the final result was unaffected by the choice of start points. Beyond this consideration specific parameter start point values were chosen arbitrarily. In choosing the upper and lower bounds of parameters we had the dual goals of avoiding impossible values (e.g. negative attention weights) and allowing a full search of each parameter's allowable range. To further avoid problems caused by local minima, start points were jittered for MLBA's β and m, and MDFT's w and σ. For each parameter, we added uniformly distributed noise to each start point value, according to: β, $U(-0.1, 0.1)$; m, $U(-1, 1)$; w, $U(-0.075, 0.075)$; w, $U(-7.5, 7.5)$. Tables A1 and A2 show the start point values, and upper and lower bounds for the parameter search.

To calculate the likelihood of the data according to each model, the natural logarithm of the predicted probability of each observed response was summed to give each individual's log-likelihood. The optimization algorithm sought parameter values that minimized the negative log-likelihood. Thus, both models were fit to the entire array of responses within a given dataset, without any aggregation. For fits to the full dataset, these were all of the responses a participant made in the experiment. For the subset fits, these were all of the responses made on trials corresponding to one context effect.

Table A1
Fitting MLBA

	β	m	I_0	λ_+	λ_-
Upper bound	10	50	20	.999	.999
Lower bound	.10	.50	.30	.001	.001
Start points	.75, 1, 1.33	1, 10, 20	1, 8, 15	.2, .4, .6, .8	.2, .4, .6, .8

Note. MLBA = multiattribute linear ballistic accumulator.

Table A2
Fitting MDFT

	w	σ	φ_1	φ_2	b
Upper bound	.9	100	.5	1000	100
Lower bound	.1	.001	.001	.1	1
Start points	.4, .5, .6	20, 45, 70	.03, .06, .1	100, 200, 300, 400	10, 20, 30

Note. MDFT = multialternative decision field theory.

(*Appendices continue*)

Appendix C

Model Comparison Using Weighted Fitness Function

Since participants completed different numbers of trials for each context effect, some effects contributed more than others to the overall performance of each model. To test the influence of this on our model comparison we recomputed the fit of each model using a weighted fitness function. Using the same parameters previously estimated from the full dataset, we computed weighted log-likelihoods for each individual. This measure assured that each context effect contributed equally to each individual's overall fitness measure. The weighted log-likelihood of a response was defined as

$$\psi \times \ln(p) \tag{C1}$$

where p is the probability of the observed response according to the best fitting parameters estimated from the full dataset and ψ is a weight corresponding to the condition in which the response was observed. The weight for each context effect was defined as

$$\begin{aligned}\psi_{\text{Att}} &= \frac{N_{\text{Total}}/N_{\text{Att}}}{3}\\ \psi_{\text{Comp}} &= \frac{N_{\text{Total}}/N_{\text{Comp}}}{3}\\ \psi_{\text{Sim}} &= \frac{N_{\text{Total}}/N_{\text{Sim}}}{3}\end{aligned} \tag{C2}$$

where N_{Total} is the total number of trials in the experiment and N_{Att}, N_{Comp}, and N_{Sim} are the numbers of attraction, compromise, and similarity trials, respectively.

The results of our reanalysis closely match those based on the unweighted fitness function. MLBA had the best average fit across individuals ($\overline{wLL}_{\text{MLBA}} = -58.55$ vs. $\overline{wLL}_{\text{MDFT}} = -59.91$). At the individual level both models do equally well in describing participants' behavior, with 24 (50%) participants best fit by each. Figure C1 shows that the distribution of individual weighted log-likelihood differences is quite similar to what is shown in Figure 2d. Using BIC values to approximate relative posterior probabilities, sixteen (33%) and four (8%) individuals were classified as having very strong or strong evidence in favor of MLBA, respectively. Eleven (23%) and five (10%) individuals were classified as having very strong or strong evidence favoring MDFT, respectively.

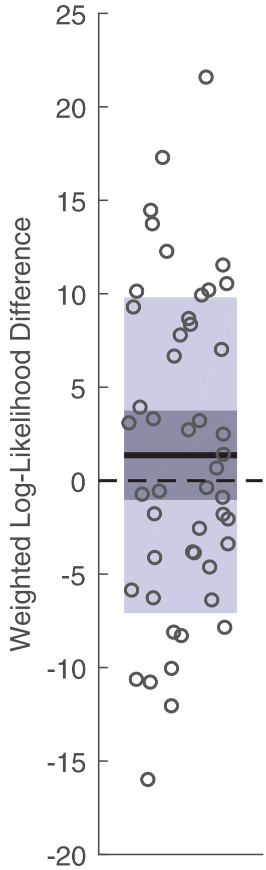

Figure C1. Model comparison of evidence in favor of multiattribute linear ballistic accumulator (MLBA) based on weighted log-likelihoods. Differences were computed by subtracting the weighted log-likelihood according to multialternative decision field theory from that according to MLBA. See the online article for the color version of this figure.

Received September 4, 2017
Revision received January 2, 2018
Accepted January 18, 2018 ■