Supplementary Materials: Competing theories of multi-alternative, multi-attribute preferential choice

Brandon M. Turner[a,*], Dan R. Schley[b], Carly Muller[a], Konstantinos Tsetsos[c]

[a] *Department of Psychology, The Ohio State University*
[b] *Rotterdam School of Management, Erasmus University*
[c] *Department of Neurophysiology and Pathophysiology, University Medical Center Hamburg*

Key words: Multiattribute Linear Ballistic Accumulator model, Leaky Competing Accumulator model, Multialternative Decision Field Theory model, Associations and Accumulation model

---

*Corresponding Author
Email addresses: turner.826@gmail.com (Brandon M. Turner)
This research was supported by National Science Foundation grant 1358507.

Contents

The supplementary materials below consist of two sections. First, we provide additional details from our hierarchical analysis of Experiment 1, as it offers finer granularity that what was reported in the main text. Second, we provide additional results from our analyses of Experiment 2. Specifically, we provide plots of the best-fitting model predictions – both the aggregated and hierarchical models – against the observed data, as well as fits to individual subjects.

1. Study 1: Additional Results

In this section, we provide three additional details that were not reported in the paper. First, we present the technical specification of the models and the details of our sampling algorithms. Second, we show estimated posterior distributions for each of the nonhierarchical models. Third, we show fit statistics for the hierarchical models for each individual subject.

1.1. Modeling Analysis 1: No Individual Differences

The first analysis we performed collapsed across subjects, resulting in a total of nine probabilities: three probabilities corresponding to the three choice options, times the three context effects. This is a sparse amount of data, especially given that the three probabilities within a context effect must sum to one. However, because we collapsed across subjects for this analysis, each observation consists of 48×3 judgments, where the three effects were collapsed across the various items (see Berkowitsch et al., 2014).

We had many choices in fitting the models to data. Each model has a custom set of mechanisms available, and it is easy to imagine how mechanisms from other models can be borrowed to obtain better fits to the data. To limit the number of possibilities, we investigated at least two versions of each model: the core version put forth in the original publication, and the core version adapted to incorporate biases in the attribute space. Because the products offered in the consumer goods experiment of Berkowitsch et al. (2014) varied along different dimensions, the additional attribute bias parameters allowed the models to capture effects that may have been an artifact of the data aggregation process. Furthermore, we tested many simpler (i.e., having fewer parameters) nested variants of the "base" models described below, such as removing the leakage parameters of the MLCA model. However, these simpler models usually performed worse, and so we do not report them here. For the AAM, we found that adding a lateral inhibition mechanism

greatly improved the model fit, and so we report it below for sake of discussion. The technical details of each model can be found in the supplementary materials.

1.1.1. Model Specification: No Individual Differences

A total of four stock models were fit to the data: MDFT, MLCA, AAM, and MLBA. Within each stock model, we investigated at least two variants of each model, one variant having attribute bias parameters, and one that did not. For the AAM, although not used in Bhatia (2013), we investigated whether having a lateral inhibition term improved the model's fit to data. In parallel with the introduction, we now discuss each of the four models chronologically.

MDFT. We investigated two variants of the MDFT model. The first variant, which we call MDFT 1.0, is the base version of the model following Roe et al. (2001) with the additional insights from Hotaling et al. (2010). This version of the model has a total of four parameters: the within-trial variability term Σ, the sensitivity parameter ϕ_1, the decay parameter ϕ_2, and the bias term β for gains. MDFT 1.0 assumes that each attribute dimension is sampled with equal probability (i.e., $\omega = 0.5$), but the bias term β does allow for the attributes to have larger magnitudes in affecting the preference states when a particular dimension is sampled (see the Supplementary Materials for derivations). Considering this, we also investigated a version of the model where the attribute dimensions could be sampled with unequal probability. This model, which we refer to as MDFT 2.0, freely estimates the parameter ω. Hence, MDFT 2.0 allows for a bias in the input of a particular attribute dimension to the accumulators, and also a bias in the saliency of each attribute.

The MDFT – like the MLCA and AAM models below – has two forms of accumulation termination policies. The first version is that the preference states stochastically grow until a predefined threshold θ is reached. The second version simply evaluates the preference states up to a pre-specified number of iterations T. To maintain consistency with all models, we used the second termination policy by setting $T = 100$.

The final consideration is the constraints on the model parameters. The parameters Σ, ϕ_1, ϕ_2, and β are all constrained to be positive values. As such, we modeled these parameters on the log scale to facilitate posterior sampling. Specifically, when the parameters are near their boundary, the

efficiency of our particular sampling algorithm decreases. Transformations to log scale help to avoid this problem altogether (see Turner and Sederberg, 2014, for a discussion). A similar issue could arise for the attribute dimension bias parameter ω which is bound by $[0, 1]$. As such, we applied a logit transformation to ω so that it also had continuous, uniform support. Considering the scale of the parameters Σ, ϕ_1, ϕ_2, and ω, we specified a uniform prior on the interval $(-1000, 1000)$. However, for β, we followed Hotaling et al. (2010) and constrained $\beta > 1$. The motivation for this choice is that the distance function (see the Supplementary Materials) should weigh changes in the dominance direction more than the indifference direction, and so the weight β should only increase the value of the dominance distance. Because we modeled β on the log scale, we incorporated this constraint by specifying a uniform prior on the interval $(0, 1000)$.

MLCA. We investigated two versions of the MLCA model. The first version of the model had a total of four parameters: the lateral inhibition parameter L, the leakage parameter k, the within-trial noise parameter η, and the baseline level of input I_0. We refer to this model as MLCA 1.0 because it is the most basic model and assumes each attribute dimension is equally likely to be sampled in the accumulation process (i.e., $\omega = 0.5$). To extend the basic model to have a biased attribute-sampling process, we freely estimated one additional parameter ω. We refer to this extended variant of the model as MLCA 2.0.

As discussed above (also see the Supplementary Materials), we chose a loss aversion function that did not require any parameters to be estimated, consistent with Usher and McClelland (2004). Unfortunately, this particular loss aversion function is not well suited for raw stimulus inputs on the range $(0, 10)$. Because the values of the stimulus dimensions are in arbitrary units, we scaled the stimulus values by first dividing each input by 10.

The MLCA is unique in the way the time steps are specified to approximate the stochastic differential equation that governs its accumulation process. For our purposes, we set $dt = 0.1$ and $\tau = 0.1$. Similar to the MDFT model above, we must also specify a termination policy for the accumulation process. Here, given the settings of dt and τ, we required longer integration times to be on the same scale as the AAM and MDFT model, and so we set $T = 500$.

The final specification is how the model parameters are constrained. The parameters η and I_0 are constrained to be positive, and we estimated them

on the log scale. The parameters k, L, and ω are assumed to be between zero and one, so we estimated them on the logit scale. Given the transformations for each model parameter, we specified completely uniform priors on the interval $(-1000, 1000)$.

AAM. In total, we report four different variants of the AAM. The first variant, AAM 1.0, is the stock version of the model that assumes each attribute dimension is preferred equally and has no lateral inhibition component. In total, AAM 1.0 has four parameters, the baseline activation input a_0, the attribute association α, decay or leakage d, and the within-trial noise term e. Although not explored in Bhatia (2013), we investigated the effects of lateral inhibition on the model's ability to capture data. To this end, we added the parameter l to enforce inhibition across the response alternatives in an analogous way to the MLCA and MDFT models. We refer to this extended model as AAM 2.0. Consistent with the other model variants, we also investigated AAM 3.0 which instantiates an attribute bias. In AAM 1.0, the baseline activation input parameter a_0 is assumed to be constant across time and across dimensions. To instantiate a bias in the baseline activation across dimensions, we replaced a_0 with the parameter a_0^1 which only operates on the first attribute dimension, and the parameter a_0^2 which only operates on the second attribute dimension. This change added a single parameter, leaving AAM 3.0 with five total parameters. The final model, AAM 4.0, investigates the interaction of attribute bias and lateral inhibition by adding a lateral inhibition term to AAM 3.0. In counting the number of parameters, AAM 4.0 has six total parameters, which is the most of any of the models we investigated.

Similar to the MDFT model above, the AAM allows two different stopping policies. To ensure consistency with the other models we investigated here, we chose the second version of stopping policy by running the model for $T = 100$ iterations for each parameter proposal in our fitting algorithm.

The final specification for these models is the choice of prior. For each of the model parameters, the only explicit constraint is that the parameters be greater than zero. Because the AAM has not been fit to data prior to this article, we wished to avoid any undue constraint on the parameters, and so we specified uniform priors on the interval $(0, 1000)$.

MLBA. We also investigated two versions of the MLBA model. The first version is the basic form of the model from (Trueblood et al., 2014), which

we call MLBA 1.0. For this version of the model, a total of four parameters were freely estimated: the base input parameter I_0, the curvature exponent parameter m, and two attention weight parameters λ_1 and λ_2. Because MLBA 1.0 assumes that the attribute weights are equal, we also investigated a version of the model that allows for an asymmetry in the attribute weights. This model, which we call MDFT 2.0, introduces a bias parameter β that mediates the attention weights across the attribute dimensions (cf. Trueblood et al., 2014).

Unlike the other three models, the MLBA model has a specific stopping rule in both free response and interrogation paradigms. The decision mechanism is equivalent to that of the LBA model where three accumulators race to a threshold amount of evidence χ. There are also between-trial variability parameters in the starting point A, and drift rate s. Following Trueblood et al. (2014), we set these parameters to fixed values across all model fits by setting $\chi = 2$, $A = 1$, and $s = 1$.

Following Trueblood et al. (2014), the parameters I_0, m, λ_1, and λ_2, were all constrained to be positive. For sampling purposes, these parameters were estimated on the log scale. The parameter β does not have a lower bound constraint, and so we estimated it on the regular scale. Given these constraints, we specified uniform priors for each parameter on the interval $(-1000, 1000)$.

1.1.2. Fitting the Models

Of the models we investigated here, only the MLBA model has a tractable likelihood function. To fit the other models to data in a Bayesian context, we must rely on algorithms to approximate the likelihood function. We chose a likelihood-approximation technique called the probability density approximation (PDA; Turner and Sederberg, 2014) method. The PDA method works by generating a synthetic set of data for each parameter proposal by simulating the model many times. Once the synthetic data have been created, the distribution of data can be smoothed using a kernel density estimate (Silverman, 1986) to provide better resolution in the approximation of the likelihood function. For each parameter proposal, we simulated the model 30,000 times for each of the three conditions.

Because the MLBA model has a tractable likelihood function, it provides us an opportunity to study the accuracy of the PDA method for this particular paradigm. To this end, we used both likelihood-free and likelihood-informed methods to estimate the posterior distribution of the MLBA model

parameters. By comparing the estimated posterior distributions obtained from each method, we can evaluate the quality of the posterior estimates obtained using our PDA method.

To generate good parameter proposals, we used differential evolution with Markov chain Monte Carlo (DE-MCMC; ter Braak, 2006; Turner et al., 2013b) to estimate the shape of the joint posterior distribution. We used 24 chains and obtained 3,000 samples after a burn-in period of 1,000 samples. The burn-in period allowed us to converge quickly to the high-density regions of the posterior distribution, while the rest of the samples allowed us to improve the reliability of the estimates. For the first 500 iterations, a migration step was used (see Turner and Sederberg, 2012, 2014, for details) with probability 0.1, but after this period, the migration probability was set to 0. Thus, our estimates of the joint posterior distribution for each model are based on 48,000 samples.

Finally, one other step was necessary to arrive at quality estimates of the posterior. Sometimes in estimating posteriors for simulation-based models, proposals can generate data that have a superficially close fit to data. When this occurs, a practical solution is to periodically reevaluate the likelihood approximation for that proposal (see Holmes, 2015). This step, which we refer to as a "purification" step, was performed for all chains every 10 iterations. Convergence of the chains was assessed through visual inspection.

1.2. Modeling Analysis 2: Hierarchical Models

1.2.1. Model Specification: Hierarchical Models

Two versions of each model were fit to the data, closely following the models 1.0 and 2.0 presented in our Analysis 1 above. Specifically, the first version contained no bias terms to account for attribute asymmetry, whereas the second version possessed additional bias terms to allow for unequal weighting of the attribute dimensions.

MDFT. The hierarchal MDFT (HMDFT) models we used closely followed the basic structure of the model presented in our Analysis 1 above. Except where stated, all details of model implementation were identical to that of Analysis 1. The first hierarchical model contained four parameters per subject: a within-trial variability term Σ_j, a sensitivity parameter $\phi_{1,j}$, a decay parameter $\phi_{2,j}$, and a bias term β_j for gains. We maintained that all four of these model parameters should be modeled on the log scale. This transformation provided two benefits. First, it enforced that all model parameters

should be positive once they were exponentiated. Second, it facilitated the development of the hierarchical models here. Specifically, with appropriate choices for priors on the subject-specific parameters, we could establish a conjugate relationship between the prior and the posterior. A conjugate relationship makes posterior sampling more efficient, enabling us to gather more samples at a faster rate.

To build the hierarchical model, we assumed

$$\begin{aligned}
\log(\Sigma_j) &\sim \mathcal{N}(\Sigma_\mu, \Sigma_\sigma), \\
\log(\phi_{1,j}) &\sim \mathcal{N}(\phi_\mu^{(1)}, \phi_\sigma^{(1)}), \\
\log(\phi_{2,j}) &\sim \mathcal{N}(\phi_\mu^{(2)}, \phi_\sigma^{(2)}), \text{ and} \\
\log(\beta_j) &\sim \mathcal{N}(\beta_\mu, \beta_\sigma),
\end{aligned}$$

where $\mathcal{N}(a, b)$ denotes a normal distribution with mean a and standard deviation b. For the group-level mean parameters, we specified mildly informative priors, hoping to remain agnostic so as to avoid any undue constraint:

$$\begin{aligned}
\Sigma_\mu &\sim \mathcal{N}(\log(10), 100), \\
\phi_\mu^{(1)} &\sim \mathcal{N}(\log(0.5), 100), \\
\phi_\mu^{(2)} &\sim \mathcal{N}(\log(0.5), 100), \text{ and} \\
\beta_\mu &\sim \mathcal{N}(\log(12), 100).
\end{aligned}$$

For the group-level standard deviation parameters, we supplied more informative priors based on previous research for the spread of subject-to-subject parameters for other cognitive models (Turner et al., 2013b):

$$\Sigma_\sigma, \phi_\sigma^{(1)}, \phi_\sigma^{(2)}, \beta_\sigma \sim \Gamma(4, 10),$$

where $\Gamma(a, b)$ denotes the Gamma distribution with shape parameter a, and rate parameter b.

For the first hierarchical MDFT model, HMDFT 1.0, the attribute bias parameter ω was set to 0.5 for each subject. However, the second hierarchical model HMDFT 2.0 added the possibility for the attribute dimensions to be sampled asymmetrically. As in MDFT 2.0 above, in HMDFT 2.0 we maintained that ω_j should be modeled on the logit scale so that group level distributions could be easily specified. Here, ω_j was freely estimated for each subject, with similar prior specifications as the other model parameters from

HMDFT 1.0. Namely,

$$\begin{aligned} \text{logit}(\omega_j) &\sim \mathcal{N}(\omega_\mu, \omega_\sigma), \\ \omega_\mu &\sim \mathcal{N}(0, 100), \text{ and} \\ \omega_\sigma &\sim \Gamma(4, 10). \end{aligned} \tag{1}$$

MLCA. As with HMDFT, we investigated two versions of the hierarchical MLCA (HMLCA) model. The first version of the model, HMLCA 1.0, had a total of four subject-specific parameters: a lateral inhibition parameter L_j, a leakage parameter k_j, a within-trial noise parameter η_j, and a baseline level of input $I_{0,j}$. To build the hierarchal model, we assumed

$$\begin{aligned} \text{logit}(L_j) &\sim \mathcal{N}(L_\mu, L_\sigma), \\ \text{logit}(k_j) &\sim \mathcal{N}(k_\mu, k_\sigma), \\ \log(\eta_j) &\sim \mathcal{N}(\eta_\mu, \eta_\sigma), \text{ and} \\ \log(I_{0,j}) &\sim \mathcal{N}(I_\mu, I_\sigma). \end{aligned}$$

Maintaining consistency with the HMDFT model above, we assumed the following priors for the HMLCA model:

$$\begin{aligned} L_\mu &\sim \mathcal{N}(\log(0.5), 100), \\ k_\mu &\sim \mathcal{N}(\log(0.5), 100), \\ \eta_\mu &\sim \mathcal{N}(\log(0.25), 100), \\ I_\mu &\sim \mathcal{N}(\log(1), 100), \text{ and} \\ L_\sigma, k_\sigma, \eta_\sigma, I_\sigma &\sim \Gamma(4, 10). \end{aligned}$$

In the HMLCA 2.0 model, we investigated the role of the attribute bias parameter ω. As in HMDFT 2.0, we freely estimated ω for each subject, and specified the same hierarchical structure as in Equation 1.

AAM. Because of the results from our Analysis 1 above, we only investigated two hierarchical versions of the AAM (HAAM). The first version, HAAM 1.0, included five subject-specific parameters: a baseline activation input $a_{0,j}$, an attribute association α_j, decay d_j, lateral inhibition l_j, and the within-trial

noise term e_j. To build the hierarchical model, we assumed:

$$\begin{aligned}
\log(a_{0,j}) &\sim \mathcal{N}(a_\mu, a_\sigma), \\
\log(\alpha_j) &\sim \mathcal{N}(\alpha_\mu, \alpha_\sigma), \\
\log(d_j) &\sim \mathcal{N}(d_\mu, d_\sigma), \\
\log(l_j) &\sim \mathcal{N}(l_\mu, l_\sigma), \text{ and} \\
\log(e_j) &\sim \mathcal{N}(e_\mu, e_\sigma).
\end{aligned}$$

To be consistent with our other hierarchical models, we assumed the following priors for the group-level parameters:

$$\begin{aligned}
a_\mu &\sim \mathcal{N}(\log(45), 100), \\
\alpha_\mu &\sim \mathcal{N}(\log(0.65), 100), \\
d_\mu &\sim \mathcal{N}(\log(0.9), 100), \\
l_\mu &\sim \mathcal{N}(\log(0.1), 100), \\
e_\mu &\sim \mathcal{N}(\log(2.5), 100), \text{ and} \\
a_\sigma, \alpha_\sigma, d_\sigma, l_\sigma, e_\sigma &\sim \Gamma(4, 10).
\end{aligned}$$

To instantiate an attribute bias in the model, we incorporated a second baseline activation parameter a_0 to correspond to the second attribute dimension, as in the MDFT 3.0 and MDFT 4.0 models from our Analysis 1. This new parameter $a^2_{0,j}$ was treated equivalently to the $a_{0,j}$ parameter above by enforcing the same hierarchal structure (i.e., $\log(a^2_{0,j}) \sim \mathcal{N}(a^2_\mu, a^2_\sigma)$) with identical priors for the group level parameters a^2_μ and a^2_σ. We refer to this hierarchical AAM models with attribute bias as HAAM 2.0.

MLBA. We fit two hierarchical versions of the MLBA model to data. The first version, HMLBA 1.0, included four subject-specific parameters: a base input parameter $I_{0,j}$, a curvature exponent parameter m_j, and two attention weight parameters $\lambda_{1,j}$ and $\lambda_{2,j}$. To build the hierarchical model, we assumed:

$$\begin{aligned}
\log(I_{0,j}) &\sim \mathcal{N}(I_\mu, I_\sigma), \\
\log(m_j) &\sim \mathcal{N}(m_\mu, m_\sigma), \\
\log(\lambda_{1,j}) &\sim \mathcal{N}(\lambda^1_\mu, \lambda^1_\sigma), \text{ and} \\
\log(\lambda_{2,j}) &\sim \mathcal{N}(\lambda^2_\mu, \lambda^2_\sigma).
\end{aligned}$$

To be consistent with our other hierarchical models, we assumed the following priors for the group-level parameters:

$$\begin{aligned} I_\mu &\sim \mathcal{N}(\log(13), 100), \\ m_\mu &\sim \mathcal{N}(\log(22), 100), \\ \lambda_\mu^1 &\sim \mathcal{N}(\log(0.1), 100), \\ \lambda_\mu^2 &\sim \mathcal{N}(\log(0.1), 100), \text{ and} \\ I_\sigma, m_\sigma, \lambda_\sigma^1, \lambda_\sigma^2 &\sim \Gamma(4, 10). \end{aligned}$$

Following MLBA 2.0 from our Analysis 1, the second hierarchical MLBA model – HMLBA 2.0 – added a bias parameter β_j for each subject to mediate the attention weights across the attribute dimensions. For this parameter, we specified a similar hierarchical structure as the other model parameters, such that

$$\begin{aligned} \log(\beta_j) &\sim \mathcal{N}(\beta_\mu, \beta_\sigma), \\ \beta_\mu &\sim \mathcal{N}(0, 100), \text{ and} \\ \beta_\sigma &\sim \Gamma(4, 10). \end{aligned} \tag{2}$$

1.2.2. Fitting the Models

At the level of the subject-specific effects, our strategy for fitting the hierarchical models to data was similar to that applied in Analysis 1: we used the PDA (Turner and Sederberg, 2014) method to approximate the joint posterior distribution for each subject. For each parameter proposal, a synthetic data set of the same size as the observed data was generated, and probabilities for each response were calculated for each of the three context effects. These probabilities were used as the long term predictions of the model, and the multinomial likelihood was used to characterize the mismatch between model predictions and observed data (see Turner and Van Zandt, 2012; Turner et al., 2013a). Fewer model simulations per proposal were necessary compared to our Analysis 1 because we did not need to integrate over the subject-to-subject variability as we do when aggregating across subjects.

To estimate the group level effects, we relied on the Gibbs ABC algorithm (Turner and Van Zandt, 2014). Gibbs ABC is a likelihood-free algorithm that is well suited for fitting hierarchical models to data because it efficiently partitions the parameter space into group and subject-specific levels. The Gibbs ABC algorithm, combined with the conjugate relationship for the group-level

parameter, allowed us to estimate the group-level parameters simultaneously with the subject-level parameters using the PDA method.

We again used DE-MCMC to generate good parameter proposals, and a purification step every 10 iterations. We used 24 chains and obtained 9,000 samples after a burn-in period of 5,000 samples. For these models, we did not use a migration step and thinned the samples by retaining every other iteration. Thus, our estimates of the joint posterior distribution for each model are based on 48,000 samples.

1.3. Estimated Posterior Distributions: No Individual Differences

Figure 1 shows the estimated posterior distributions for the parameters in each model. Starting with the MDFT model (i.e., Row 1), we felt that two parameters estimates are worth mentioning. First, the parameter β is much larger than 1 (i.e., the mean of β is 66.69), which implies that the dominance dimension greatly outweighs the indifference dimension. Second, the parameter ω is slightly less than 0.5 (i.e., its mean is 0.48). This setting of ω means that the second attribute dimension is sampled more frequently moment-by-moment, a feature of the model that is consistent with the AAM.

The MLCA model (i.e., Row 2) has four parameters of interest. The first parameter is the baseline input parameter I_0, which has a mean of 0.14. It appears that for the MLCA model, some baseline input is required to obtain good fits to data, whereas in the MLBA model (i.e., Row 4), the baseline input was completely unnecessary. We suspect this has to do with the stochastic nature of the MLCA model interacting with the loss aversion function. Because the attribute values are relatively small, some baseline input is necessary to allow for the response competition that ensues in the model. This is not the case for the MLBA model, whose evidence accumulation process in completely independent. Second, the attribute dimension weight ω has a mean of 0.506, suggesting a small preference for the first attribute dimension. This result is consistent with the MLBA model, but inconsistent with the AAM and MDFT models. Finally, the estimates of leakage k and lateral inhibition L suggest that the accumulation process is leak dominant (e.g., Gao et al., 2011). Also note that the parameter estimate for the parameter L on the logit scale spans the negative axis of the parameter space. This simply implies that that the value of L is very near zero (i.e., its mean is 0.001).

For the AAM (i.e., Row 3), we first note that while the a_0^1 and a_0^2 parameters are quite similar, the a_0^2 parameter is slightly larger which indicates

a preference for the second attribute dimension for these data. Another interesting aspect of the model parameters is that to fit these data, the AAM is lateral inhibition dominant, meaning that the parameter l is larger than the decay parameter d. When examining the joint posterior distribution of (d, l), we observed a strong negative correlation ($r = -0.93$) between the parameters.

Finally, three parameters are of interest for the MLBA model (i.e., Row 4). First, the estimates for the curvature parameter m suggest that the mapping function is convex (i.e., the mean for m is 1.22). Having a convex mapping function is what produces the compromise effect in the model, something that allows a preference advantage for the target option (see the main text). However, as Figure 1 reveals, the curvature function is not convex enough as the MLBA model undershoots the strength of the compromise effect for these data. However, simply increasing the estimate for the m parameter would not cause the model to fit the data better because it could cause the model to miss other effects. A second parameter of interest is the β parameter, whose mean is 0.89. When $\beta < 1$, a bias toward the first attribute dimension is observed, something that is inconsistent with AAM or MDFT models. Finally, the mean for the baseline input I_0 parameter is 0.0001 on the normal scale, suggesting that no input is necessary for these data.

Recall that for the MLBA model we estimated the parameters using both likelihood-free and likelihood-informed methods. The black densities in the third row of Figure 1 are the estimates obtained when using the analytic expressions for the likelihood function. Comparing across the panels, Figure 1 shows that the estimates using both methods are similar, with an exception being the baseline input parameter. We suspect that this misestimation occurred due to some numerical instabilities in calculating the likelihood function. We conclude this on the basis of several posterior recovery studies we performed on synthetic data when the baseline input parameter I_0 was larger than zero (e.g., $I_0 = 5$, following Trueblood et al., 2014). In these studies, the PDA-estimated posteriors were extremely accurate, and in particular, no misestimation for I_0 was observed. Although there is no way to guarantee that the posterior distributions for the other parameters are being correctly estimated, we find the convergence of the two estimated posterior distributions in the MLBA model to be reassuring when using PDA to approximate the likelihood function.

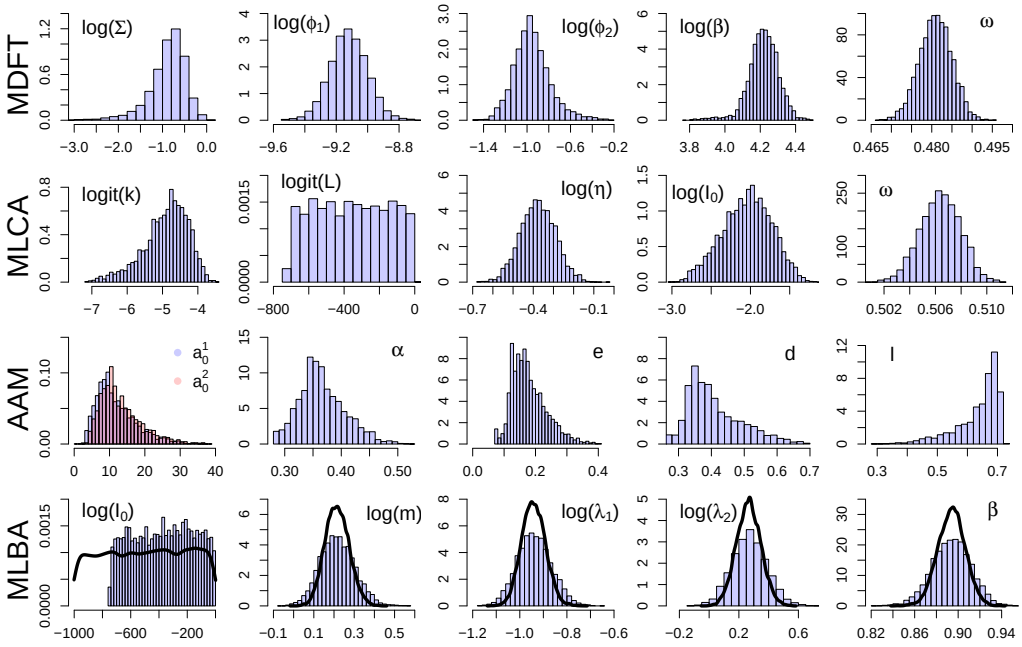

Figure 1: Estimated posterior distributions for each parameter (columns) of each of the best-fitting model variants: MDFT (first), MLCA (second), AAM (third), and MLBA (fourth). For the MLBA model, we estimated the model parameters using both likelihood-informed (black densities) and likelihood-free (i.e., the probability density approximation) method (histograms).

1.4. Hierarchical Fits to Data by Subject

When fitting a hierarchical model to data, we have a few choices about how to assess the model's performance on the levels of the hierarchy. One measure is to collapse across the subject level and provide a single measure for model fit and complexity – measures that we reported in Table 2 of the main text. Another measure though, is to break this statistic apart and assess how well the model fits to each individual subject. Here, we simply assessed the deviance information criterion (DIC) for each subject by model combination separately, and placed them in Table 1. We could then evaluate these DIC values relative to each of the best-fitting models (i.e., the symmetric models) from Experiment 1. The final line in Table 1 shows the total number of wins obtained by each model. In this analysis, the HAAM and HMLCA models perform very similarly, whereas HMLBA and HMDFT perform substantially worse relative to these models. Although the ordering of the model performance is identical to the results at the aggregate level (i.e., Table 2 in the main text), this is not always necessarily the case.

Subject	HMDFT 1.0	HMLCA 1.0	HAAM 1.0	HMLBA 1.0
1	63.18	**51.71**	54.89	53.46
2	40.99	34.49	38.76	**32.23**
3	62.38	51.11	45.40	**44.76**
4	106.91	72.38	**34.23**	36.40
5	32.68	**26.67**	34.78	42.58
6	**29.18**	34.33	36.42	42.63
7	66.93	51.06	**32.88**	42.07
8	39.70	43.74	**39.08**	43.68
9	35.80	**33.79**	37.96	37.94
10	56.61	50.66	51.07	**37.56**
11	40.91	44.29	48.79	**39.19**
12	31.77	**30.73**	34.02	45.07
13	37.76	33.62	**31.27**	38.32
14	81.08	60.21	**59.55**	77.07
15	54.71	43.69	**38.02**	47.21
16	54.20	44.65	**36.89**	43.49
17	89.21	61.78	**42.02**	73.50
18	48.49	**34.14**	41.98	43.30
19	96.92	63.73	**34.34**	60.94
20	**28.86**	30.76	30.62	45.88

21	40.57	37.77	**36.37**	43.00
22	33.89	**30.31**	41.79	41.49
23	82.08	**36.18**	37.86	63.72
24	48.26	**42.84**	46.45	54.07
25	88.06	76.63	**64.66**	75.92
26	109.50	78.87	36.76	**30.06**
27	**29.65**	32.45	37.55	38.20
28	47.66	**36.18**	41.43	41.61
29	**36.17**	36.98	40.99	36.56
30	52.74	**51.36**	58.49	63.71
31	**42.73**	43.60	46.04	60.06
32	30.44	**28.03**	32.95	50.63
33	**35.76**	38.53	44.09	36.57
34	**30.08**	30.73	29.30	45.40
35	73.22	50.98	**39.10**	40.98
36	107.62	**39.45**	38.90	62.87
37	57.66	47.04	**41.31**	44.00
38	37.39	36.67	**34.70**	34.94
39	81.86	72.40	59.28	**55.86**
40	75.02	60.48	49.22	**47.66**
41	33.78	**33.64**	38.09	39.17
42	35.86	**31.52**	36.32	46.83
43	48.98	41.84	**32.56**	39.08
44	102.66	74.81	**42.33**	61.26
45	31.63	**30.13**	34.91	45.92
46	60.06	38.09	**30.94**	36.90
47	67.17	51.63	**46.20**	52.66
48	36.85	**35.66**	38.88	49.65
Total	7	16	18	7

Table 1: Deviance information criterion (DIC) values for each subject (rows) and each best-fitting hierarchical model (columns) applied to Experiment 1. Each value represents the mean statistic obtained across all chains in the sampling algorithm. The best performing model for each subject is shown in bold-face type.

2. Study 2: Additional Results

For the analyses of Experiment 2, we provide four elaborated results. First, we provide the technical specifications of each model and the details of the sampling algorithms used to estimate the model parameters. Second, plots of the best-fitting single-level model fits against the aggregated data from each condition of our experiment. Third, we provide analogous plots of the best-fitting hierarchical model fits against the aggregated data from each condition of our experiment. These plots essentially separate out the points shown in Figures 5 and 6 so that fits to individual conditions can be appreciated. These plots reveal interesting weaknesses and strengths of each model, which we discussed in the main text. Finally, we provide a table of fit statistics comparing the best-fitting hierarchical models on a subject-by-subject basis. The results from this analysis reveal that while at the aggregated level, the hierarchical AAM fit best, the competition was much closer at the individual subject level.

2.1. Modeling Analysis 1: No Individual Differences

2.1.1. Model Specification: No Individual Differences

MDFT. For the MDFT model, we investigated both MDFT 1.0 and MDFT 2.0 from the Study 1 above. Recall that the MDFT 1.0 model assumes a within-trial variability term Σ, a sensitivity parameter ϕ_1, a decay parameter ϕ_2, and a bias term for gains β. MDFT 2.0 extends MDFT 1.0 by allowing one additional attribute bias parameter ω to be freely estimated, whereas MDFT 1.0 sets $\omega = 0.5$. Similar to the MDFT model variants above, we instantiated a termination policy by setting $T = 100$. We again modeled the attribute bias parameters (of MDFT 2.0) on the logit scale, and all of the other parameters on the log scale. Given these transformations, we specified a continuous uniform prior on the interval $(-1000, 1000)$ for all model parameters except the bias for gains parameters β. For these parameters, we constrained $\beta > 1$ by specifying a uniform prior on the interval $(0, 1000)$ on the log scale. We chose against using the priors from the previous study to constrain the parameter estimates because the previous study was based completely on consumer goods, which may have adversely affected the models ability to fit data from a perceptual experiment.

MLCA. We investigated both the MLCA 1.0 and MLCA 2.0 models from Study 1 above. Recall that base MLCA 1.0 model include a base input parameter I_0, a lateral inhibition parameter L, a leakage parameter k, and a

within-trial noise parameter η. MLCA 2.0 includes an attribute bias parameter ω to allow for asymmetric attribute dimension sampling, whereas MLCA 1.0 sets $\omega = 0.5$. To implement the model with our specific loss aversion function, we again scaled the stimulus values by first dividing each input by 10. As in the MLCA variants in Study 1, we approximated the stochastic differential equation by setting $dt = 0.1$ and $\tau = 0.1$ (Brown et al., 2006), and instantiated a stopping rule by running the model for $T = 500$ iterations. The parameters η and I_0 were log transformed, whereas the parameters L, k, and ω were all logit transformed. Following these transformations, we again specified a continuous uniform prior for all model parameters on the interval $(-1000, 1000)$.

AAM. For the AAM, we used the same complete set of model parameters from AAM 2.0 for the symmetric version of the model, and AAM 4.0 for the asymmetric version. Recall that for the AAM 2.0, a total of four parameters are allowed to vary: an attribute association parameter α, a decay parameter d, a lateral inhibition parameter l, a within-trial noise term e, and an attribute dimension parameter a_0. For the AAM 4.0, we replaced the attribute dimension parameter a_0 with two parameters corresponding to each attribute dimension: a_0^1 (i.e., for width) and a_0^2 (for height). We again instantiated a termination policy by setting $T = 100$. Consistent with our prior specifications from above, we specified completely uniformed priors on the interval $(0, 1000)$ for each of the model parameters, which were modeled on their regular scale.

MLBA. We investigated both the MLBA 1.0 and MLBA 2.0 models from Study 1. Recall that the MLBA 1.0 model assumes a base input parameter I_0, a curvature exponent parameters m, and a pair of attention weight parameters λ_1 and λ_2. To capture potential attribute dimension biases, the MLBA 2.0 model allowed β to freely vary, whereas in MLBA 1.0, $\beta = 1$. To instantiate a stopping rule in the MLBA model, we maintained the same specifications for as the MLBA model above (as well as Trueblood et al., 2014): $A = 1$, $\chi = 2$, and $s = 1$. Consistent with the variants of the MLBA model investigated in Study 1, we applied a log transformation for each model parameter, and specified a continuous uniform prior on the interval $(-1000, 1000)$ on these transformed parameters.

2.1.2. Fitting the Models

We again used the probability density approximation (PDA; Turner and Sederberg, 2014) method to fit the models to data. This time, because there were many more cells to the experimental design (i.e., 15 cells in total requiring model simulations), we used a smaller number of model simulations to approximate the likelihood function. Specifically, we used 5,000 simulations per cell, totaling 300,000 model simulations per parameter proposal. We again used a purification step for all chains every 10 iterations (Holmes, 2015).

We again used differential evolution with Markov chain Monte Carlo (DE-MCMC; ter Braak, 2006; Turner et al., 2013b) to estimate the shape of the joint posterior distribution. We used 24 chains and obtained 8,000 samples after a burn-in period of 5,000 samples. For the first 2,500 iterations, a migration step was used (see Turner and Sederberg, 2012, 2014, for details) with probability 0.1, but after this period, the migration probability was set to zero. Once the samples had been obtained, we thinned the chains by retaining every other sample, which helps to reduce autocorrelation and improve the quality of the estimates. Hence, our estimates of the joint posterior distribution for each model are based on 96,000 samples. Convergence of the chains was assessed through visual inspection.

2.2. Modeling Analysis 2: Hierarchical Models

2.2.1. Fitting the Models

Unless otherwise stated, we used an identical process to fit the hierarchical models in this study as was reported in our Study 1. We again used a mixture of the Gibbs ABC algorithm for the group-level effects Turner and Van Zandt (2014), and the PDA method Turner and Sederberg (2014) nested within Gibbs ABC to estimate the subject-specific effects (i.e., to form an approximation of the likelihood). We again used DE-MCMC to generate good parameter proposals, and a purification step every 10 iterations. We used 24 chains and obtained 7,000 samples after a burn-in period of 3,000 samples. For these models, we did used a migration step probability of 0.1 for the first 900 iterations, and then set the migration probability to 0. We thinned our chains following the burn-in period by retaining the samples from every other iteration. Thus, our estimates of the joint posterior distribution for each model are based on 48,000 samples.

2.3. Aggregated Fits to Data by Condition

Table 3 in the main text describes the set of stimuli we used in our Experiment 2 for both ternary and binary trials. Examination of this table reveals that only seven unique stimuli were used, as it is the comparison of these stimuli in the context of other items that distorts the choice share in interesting ways. In conjunction with Table 3 of the main text, we also provide a map of the unique stimuli used in Experiment 2. Figure 2 shows the seven stimuli in their attribute space, where the dots indicate the physical value used in the experiment, and the corresponding text labels the individual items. For example, the three primary comparison items are labeled "X1", "X2", and "X3", respectively. As the sum of the attribute dimensions of these items is equal to 10, subjects should be indifferent to these choice options, assuming no bias is present in the relative weighting of the attribute dimensions. The stimulus set also contains two types of decoys, two that are dominated by the other options (i.e., the sum of their attribute dimensions is less than 10), and two that should be of equal valuation relative to the primary comparison items (i.e., the sum of their attribute dimension is equal to 10). The two dominated decoys are labeled "D1" and "D2", whereas the undominated decoys are labeled "S1" and "S2". Each item in the stimulus set is color coded to facilitate a comparison across conditions in the model fit plots below. For example, Item X3, which sits between X1 and X2 in the attribute space, is colored purple as it is the average of blue (i.e., X1) and red (i.e., X2). Dominated options are labeled with colors expressing a dilution of the primary option, such that D1 is light blue as it is dominated by X1 (i.e., blue), and D2 is pink as it is dominated by X2 (i.e., red). The undominated decoys are colored with secondary colors that are closest in hue to their primary-colored primary comparison item such that S1 is green (i.e., closest to the blue X1), and S2 is orange (i.e., closest to the red X2).

Figures 3, 4, 5, and 6 show bar plots of the observed data against the best-fitting model predictions for the MDFT 1.0, MLCA 2.0, AAM 2.0, and MLBA 1.0, respectively. Recall that for these fits, no clear pattern of asymmetric or symmetric weighting was revealed across models in our evaluation (see Table 4 in the main text). For MLCA and AAM, asymmetric variants performed best, whereas for MLBA and MDFT, symmetric models performed best. Regarding the model predictions, we first obtained the best fitting parameter set for each model, and then simulated data from model with identical numbers of observations relative to the observed data. Similarly, the data in each of the bar plots below are the average preference for each item, averaged

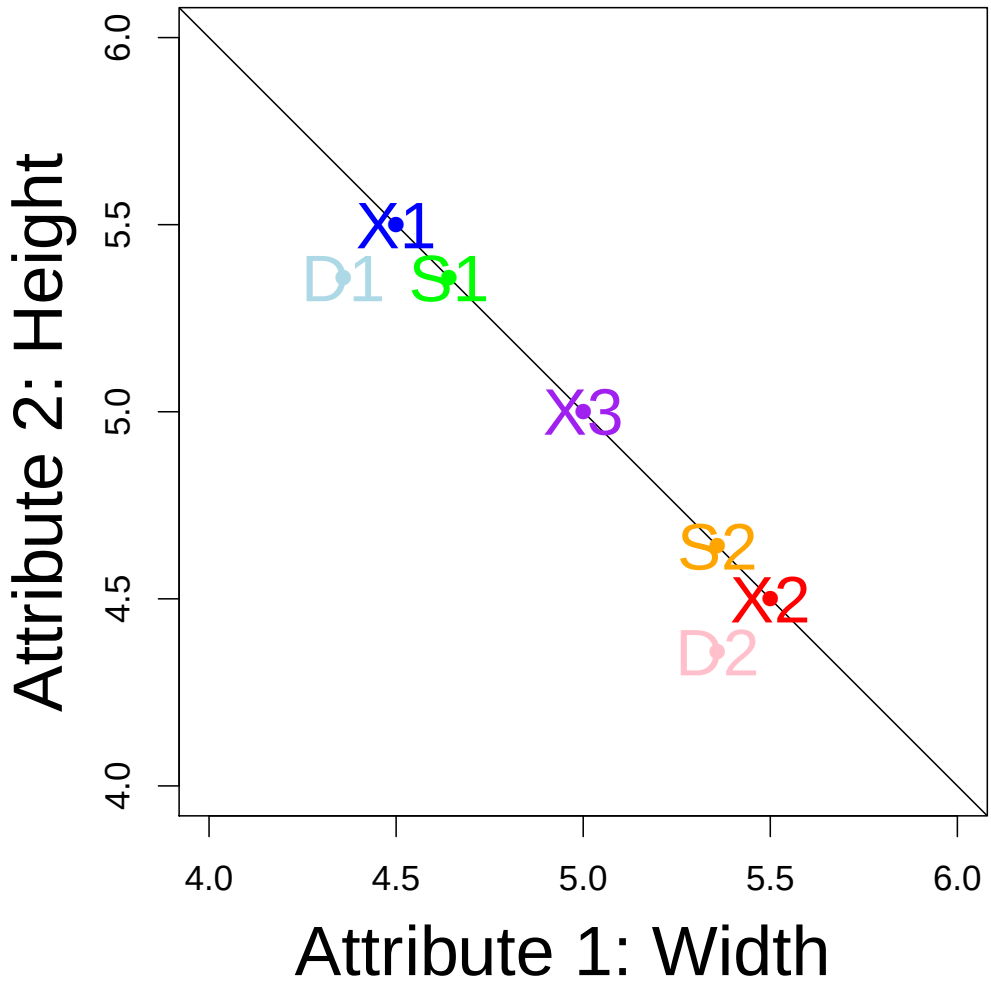

Figure 2: Illustration of the stimulus set used in Experiment 2. The stimuli are spread across the attribute space in a symmetric manner, and they consist of three "targets" (i.e., Xs), two dominated "decoys" (i.e., Ds), and two undominated decoys (i.e., Ss). In all, we used 15 different combinations of these stimuli, with 7 binary trials and 8 ternary trials (i.e., see Table 3 in the main text). The color of each stimulus value in the set corresponds to the bar plots in the model fits below.

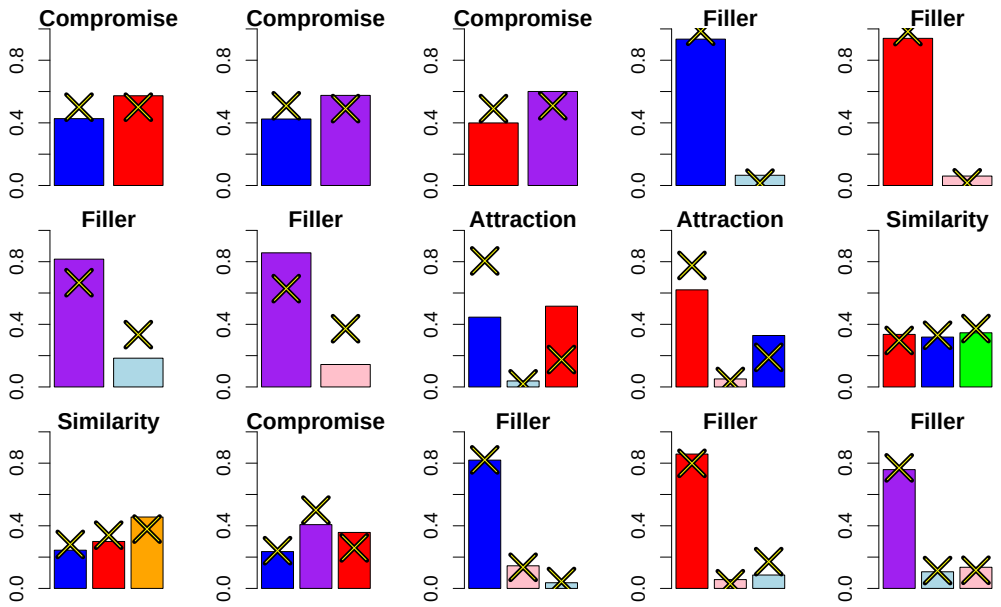

Figure 3: Fits of the MDFT 1.0 ("X" symbols) against the data from Experiment 2. Each panel shows the data and best-fitting model predictions for a specific condition, where the bars are color coded according to the stimuli presented in that condition, corresponding to Figure 2.

over subjects and displayed for each condition in Experiment 2. The color of each bar corresponds to the stimulus used in each condition according to the stimulus map in Figure 2.

2.4. Hierarchical Fits to Data by Condition

Figures 7, 8, 9, and 10 show bar plots of the observed data against the best-fitting model predictions for the HMDFT 2.0, HMLCA 2.0, HAAM 2.0, and HMLBA 2.0, respectively. Recall that when applying hierarchical models to the data from Experiment 2, asymmetric models were preferred (see Table 5 in the main text), suggesting that some distortion of the preference space is observed when averaging across subjects, as we did in the aggregate analyses above. The model predictions were performed slightly differently that in the preceding section. Here, we looped through each subject in the experiment, found the best fitting parameter values for that subject, and then generated data with the same number of trials that the corresponding

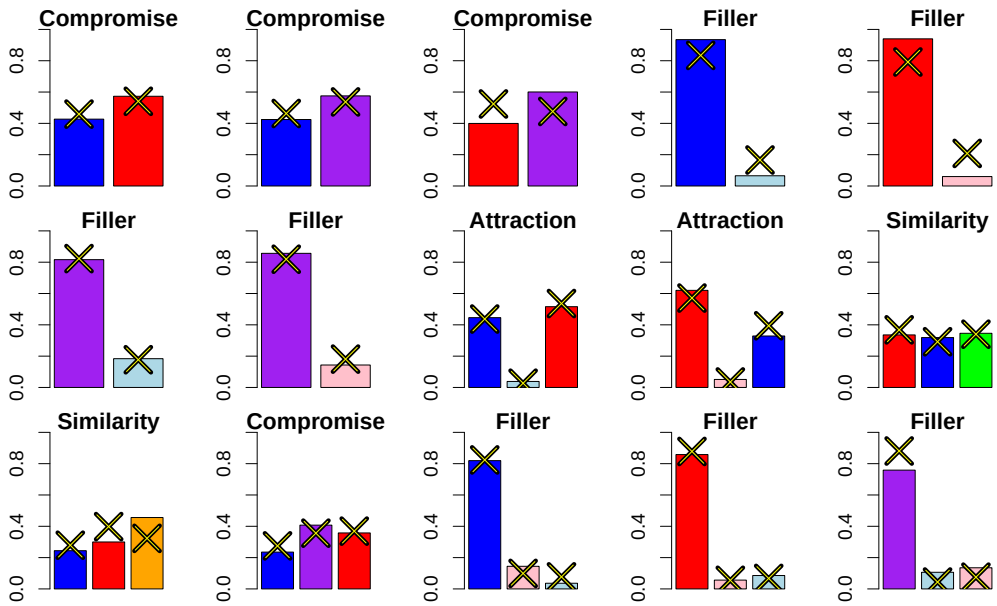

Figure 4: Fits of the MLCA 2.0 ("X" symbols) against the data from Experiment 2. Each panel shows the data and best-fitting model predictions for a specific condition, where the bars are color coded according to the stimuli presented in that condition, corresponding to Figure 2.

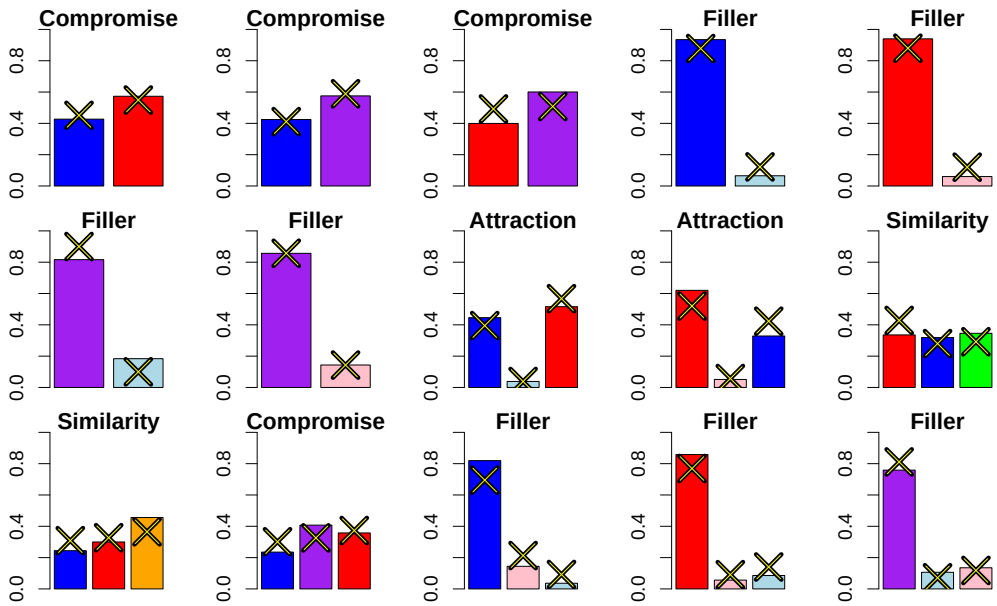

Figure 5: Fits of the AAM 2.0 ("X" symbols) against the data from Experiment 2. Each panel shows the data and best-fitting model predictions for a specific condition, where the bars are color coded according to the stimuli presented in that condition, corresponding to Figure 2.

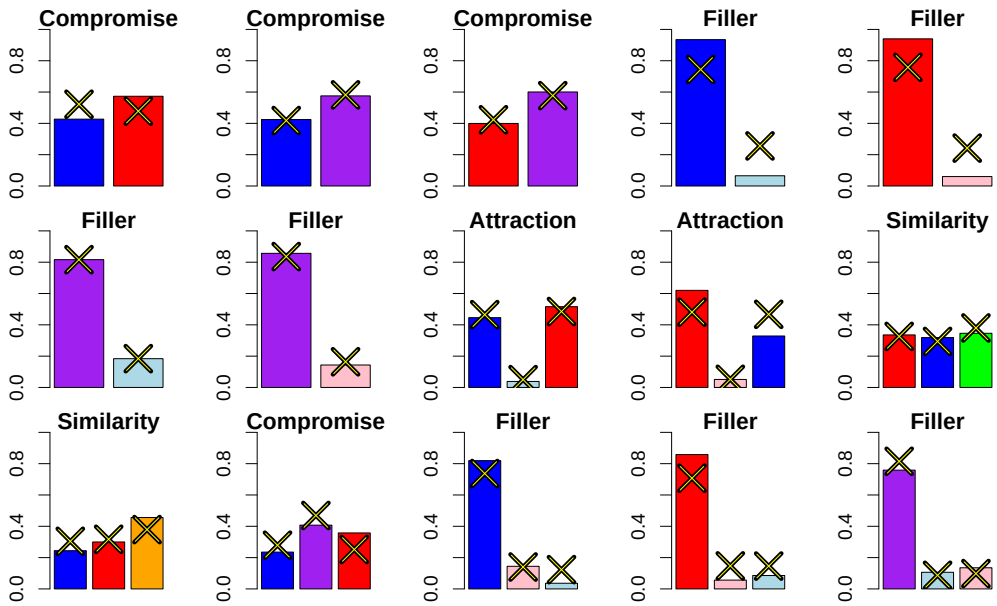

Figure 6: Fits of the MLBA 1.0 ("X" symbols) against the data from Experiment 2. Each panel shows the data and best-fitting model predictions for a specific condition, where the bars are color coded according to the stimuli presented in that condition, corresponding to Figure 2.

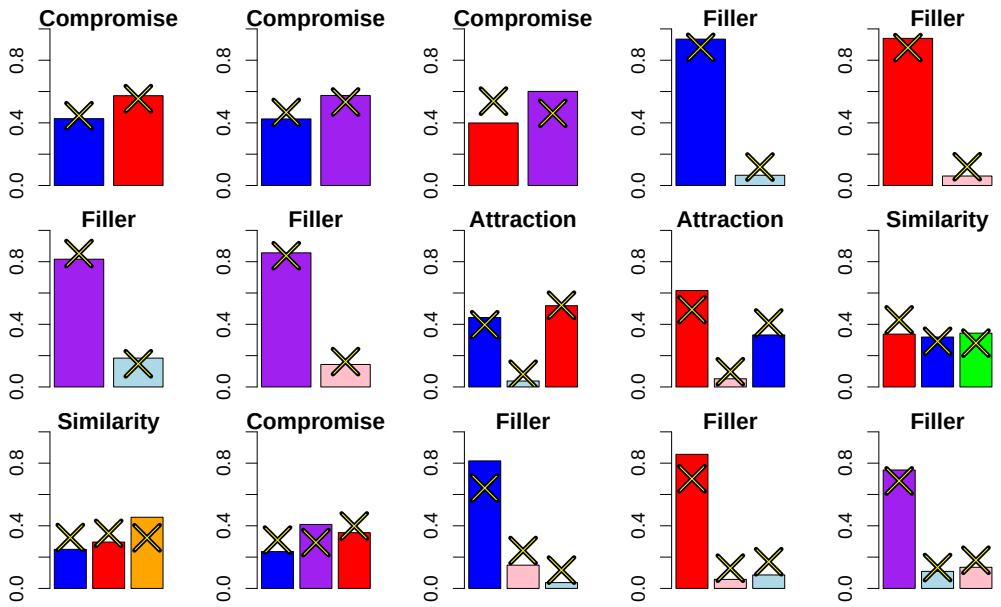

Figure 7: Fits of the HMDFT 2.0 ("X" symbols) against the data from Experiment 2. Each panel shows the data and best-fitting model predictions for a specific condition, where the bars are color coded according to the stimuli presented in that condition, corresponding to Figure 2.

subject provided. With these subject-by-subject predictions in hand, we averaged the predictions so that they appeared on the same level as that of the observed data. To gain some insight into the variability inherent in these predictions, the reader is directed to Figure 6 in the main text. As in the bar plots in the preceding section, the color of each bar corresponds to the stimulus used in each condition according to the stimulus map in Figure 2.

2.5. Hierarchical Fits to Data by Subject

Similar to the analyses presented in Table 1, we computed the DIC values for each subject in our Experiment 2, for each of the best-fitting hierarchical models. Recall that the best fitting hierarchical models in Experiment 2 were asymmetric, meaning they required attribute bias parameters enabling differential weighting of height relative to width. Once the DIC values were computed for each subject by model combination, we could tabulate the number of subjects for which a given model provided the best fit. The last line of Table 2 shows these results. While in the main text we reported

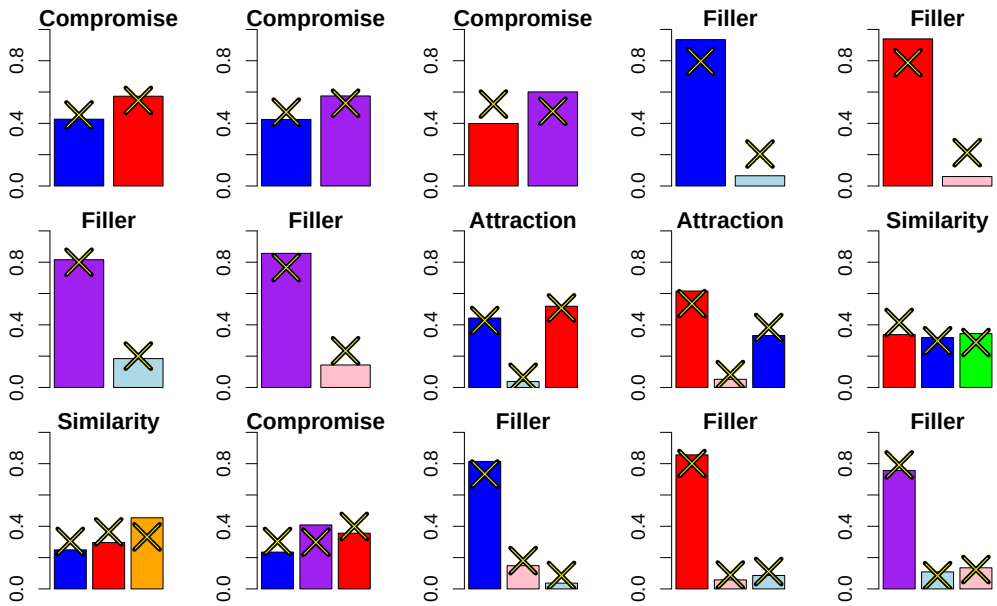

Figure 8: Fits of the HMLCA 2.0 ("X" symbols) against the data from Experiment 2. Each panel shows the data and best-fitting model predictions for a specific condition, where the bars are color coded according to the stimuli presented in that condition, corresponding to Figure 2.

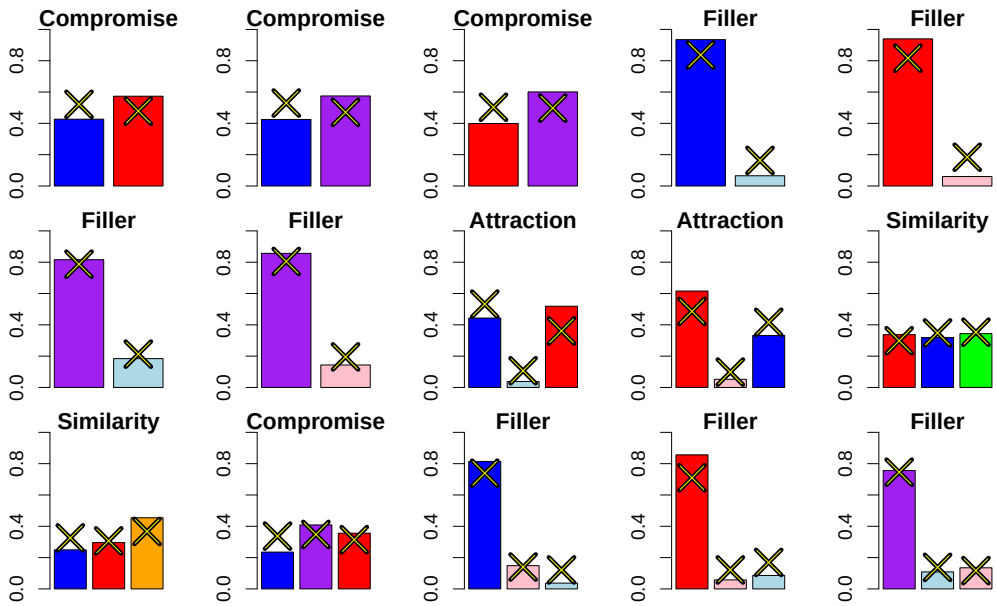

Figure 9: Fits of the HAAM 2.0 ("X" symbols) against the data from Experiment 2. Each panel shows the data and best-fitting model predictions for a specific condition, where the bars are color coded according to the stimuli presented in that condition, corresponding to Figure 2.

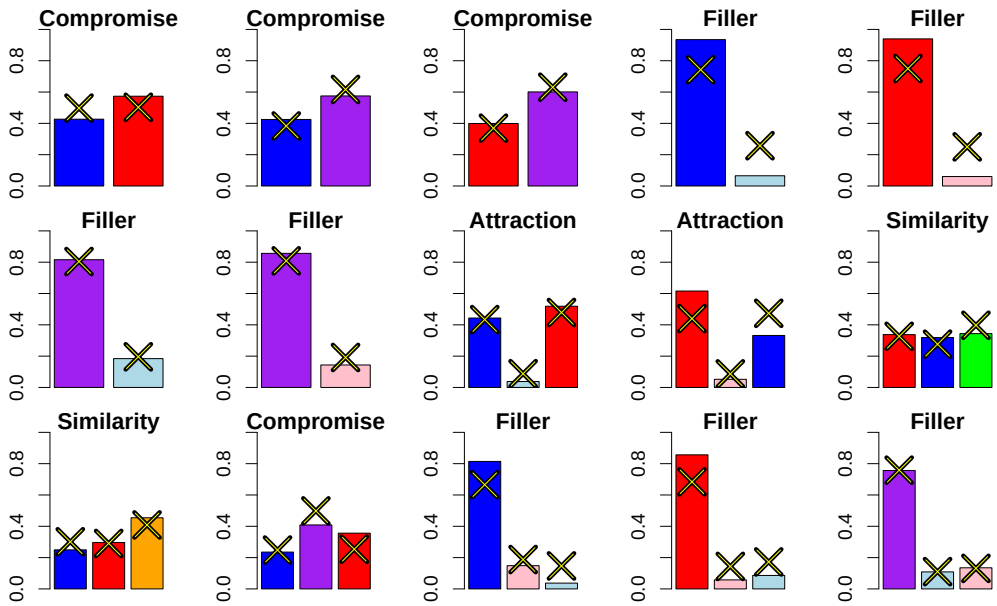

Figure 10: Fits of the HMLBA 2.0 ("X" symbols) against the data from Experiment 2. Each panel shows the data and best-fitting model predictions for a specific condition, where the bars are color coded according to the stimuli presented in that condition, corresponding to Figure 2.

that the order of the best models was (1) HAAM 2.0, (2) HMLCA 2.0, (3) HMLBA 2.0, and (4) HMDFT 2.0, the results in Table 2 show that the competition among models was more close, suggesting a tie for first between HAAM, HMLCA, and HMLBA, and a distant second for HMDFT. What this suggests at the aggregate level is that HAAM 2.0 and HMLCA 2.0 fit the data relatively consistently, providing good fits for each subject even when it did not necessarily provide the best fit overall. However, models like HMLCA 2.0 had more variance in the quality of their fits, providing fits on par with the best fitting models for some subjects, but also providing poor fits to other subjects. When averaging the model fits across subjects, this leads to MLBA 2.0 performing worse relative to HAAM 2.0 and HMLCA 2.0.

Subject	HMDFT 2.0	HMLCA 2.0	HAAM 2.0	HMLBA 2.0
1	384.01	479.19	**322.15**	389.54
2	715.87	**177.11**	665.80	194.84
3	189.63	303.82	**164.59**	183.03
4	220.53	**209.14**	243.84	248.00
5	464.92	738.74	**406.90**	449.99
6	**251.31**	274.26	263.96	366.04
7	251.36	395.54	221.05	**183.00**
8	**292.30**	324.68	295.75	463.41
9	312.88	462.97	**269.21**	371.46
10	298.93	**262.74**	314.79	293.56
11	632.38	396.51	571.87	**247.60**
12	**223.20**	363.28	226.71	245.77
13	223.27	282.69	**208.45**	287.59
14	242.78	246.01	**219.76**	264.20
15	738.19	454.31	645.46	**228.01**
16	314.36	**233.38**	286.83	248.21
17	311.57	302.70	**270.43**	369.43
18	184.43	480.53	207.80	**181.63**
19	**356.30**	403.27	362.45	387.41
20	282.74	287.35	284.61	**280.04**
21	295.14	308.31	**246.88**	432.49
22	287.20	302.44	260.26	**213.38**
23	397.46	511.61	343.55	**338.32**
24	318.85	**312.96**	386.19	326.53

25	336.68	357.63	**304.23**	324.04
26	554.29	**208.88**	476.61	242.27
27	211.40	308.92	**197.21**	246.25
28	464.98	365.47	421.94	**283.81**
29	277.11	486.58	**274.35**	344.99
30	518.22	476.95	512.59	**244.38**
31	908.00	**402.34**	868.21	487.77
32	418.86	449.68	373.30	**370.93**
33	402.73	**189.93**	318.95	307.82
34	339.41	474.09	**264.85**	351.88
35	580.26	635.80	512.66	**260.35**
36	**226.37**	345.71	232.74	300.95
37	493.73	**246.58**	420.73	288.19
38	757.37	**181.71**	722.54	228.55
39	386.64	**209.88**	325.98	531.73
40	375.91	686.76	378.29	**222.71**
41	522.11	**445.86**	518.79	494.52
42	**150.24**	313.73	152.98	176.18
Total	6	12	12	12

Table 2: Deviance information criterion (DIC) values for each subject (rows) and each best-fitting hierarchical model (columns) applied to Study 2. Each value represents the mean statistic obtained across all chains in the sampling algorithm. The best performing model for each subject is shown in bold-face type.

References

Berkowitsch, N. A. J., Scheibehenne, B., Rieskamp, J., 2014. Rigorously testing multialternative decision field theory against random utility models. Journal of Experimental Psychology 143, 1331–1348.

Bhatia, S., 2013. Associations and the accumulation of preference. Psychological Review 120, 522–543.

Brown, S., Ratcliff, R., Smith, P. L., 2006. Evaluating methods for approximating stochastic differential equations. Journal of Mathematical Psychology 50, 402–410.

Gao, J., Tortell, R., McClelland, J. L., 2011. Dynamic integration of reward and stimulus information in perceptual decision-making. PLoS ONE 6, 1–21.

Holmes, W. R., 2015. A practical guide to the probability density approximation (pda) with improved implementation and error characterization. Journal of Mathematical Psychology 68, 13–24.

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.

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.

Silverman, B. W., 1986. Density estimation for statistics and data analysis. London: Chapman & Hall.

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.

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

Turner, B. M., Dennis, S., Van Zandt, T., 2013a. Bayesian analysis of memory models. Psychological Review 120, 667–678.

Turner, B. M., Sederberg, P. B., 2012. Approximate Bayesian computation with Differential Evolution. Journal of Mathematical Psychology 56, 375–385.

Turner, B. M., Sederberg, P. B., 2014. A generalized, likelihood-free method for parameter estimation. Psychonomic Bulletin and Review 21, 227–250.

Turner, B. M., Sederberg, P. B., Brown, S., Steyvers, M., 2013b. A method for efficiently sampling from distributions with correlated dimensions. Psychological Methods.

Turner, B. M., Van Zandt, T., 2012. A tutorial on approximate Bayesian computation. Journal of Mathematical Psychology 56, 69–85.

Turner, B. M., Van Zandt, T., 2014. Hierarchical approximate Bayesian computation. Psychometrika 79, 185–209.

Usher, M., McClelland, J. L., 2004. Loss aversion and inihibition in dynamical models of multialternative choice. Psychological Review 111, 757–769.