Supplementary materials for "Decisions among shifting choice alternatives reveal option-general representations of evidence"

Peter D. Kvam*, Konstantina Sokratous, and Anderson Fitch
University of Florida

In this supplementary material, we present additional modeling details as well as tests of several more versions of the DDM and RDM models. None of these models were ultimately favored against their counterparts presented in the main text; however, we believe they may be informative for readers seeking alternative explanations for our results. In general, there does not appear to be a good way to save alternative-specific models even with drastic changes to their structure to increase complexity and flexibility.

For those interested in additional modeling details, we refer to the Open Science Framework page attached to this project (osf.io/gwz3j). The code for each of the models detailing their exact structure, as well as for the comparisons and network training, are all provided in the model code folder there.

Flexible re-mapping

As we described briefly in the main text, it may be possible for alternative-specific accounts of decision-making to allow for "imputing" the degree of support one should have for a new option, even when the stimulus information for the new option has not been accumulated. In such a view, the degree of support for a new option might depend on its similarity to the original options, rather than just the residual evidence for the original two options from the first stage. For example, if a decision-maker's options changed from A, B to A, C, then the support for C could potentially be inferred from the degree of support from A and B.

This type of proposal would not produce the pattern of results we observed in Experiment 2 – where differences in accuracy can only be attributed to stimulus evidence that was accumulated before the options changed. However, such an explanation could in theory describe some of the patterns of behavior present in Experiment 1.

To test this proposal, we created and tested new versions of the racing diffusion model (RDM) and diffusion decision model (DDM). In these models, the degree of carryover was allowed to vary across the different switch conditions. This allowed us to examine how evidence might be transferred from the old set of options to the new set of options. To avoid making any overly strong assumptions about how evidence might be transferred to new options, we simply fit the degree of transfer from old to new options as a free parameter. In total, there were 10 carryover parameters for 10 switch conditions, allowing the new models a tremendous degree of flexibility in how they mapped support for the initial set of options onto support for the final options.

There is a fair argument to be made that models with such flexible carryover functions requires some knowledge of the relationships among the choice options, and thus that they are in a way alternative-general. We test these alternative theories more for completeness and at the suggestion of a reviewer, rather than because we believe that this is likely how people maintained performance on the task.

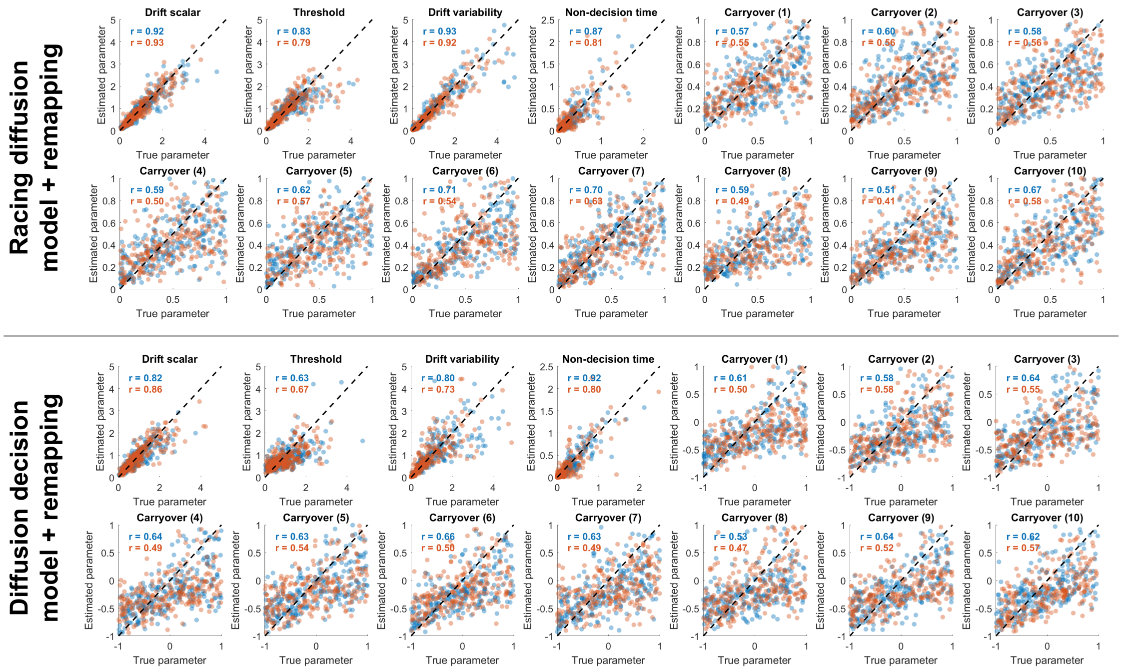

Figure S1

Parameter recovery for the RDM+remapping (top) and DDM+remapping (bottom) models.

Doing so yielded two new models: the DDM+remapping and the RDM+remapping model. In the RDM-remapping, the carryover parameter was simulated as a draw from a uniform distribution on [0,1] – the same as the RMD presented in the main text, but with 9 additional carryover parameters permitting this remapping to be flexibly implemented across conditions. At the suggestion of a reviewer, the DDM additionally allowed the carryover parameter to take negative values. The implication of this is that, depending on the way in which the options changed, the degree of support between blue and orange options could invert. This could potentially allow it to better account for cases like the swap condition, where the switch actually inverted the two options, by inverting the degree of support between blue and orange options. It afforded the DDM more flexibility, but ultimately did not help it fit the data.

As before, we trained a neural network to estimate the parameters of each of the models. The number of simulations, simulation method, layer structure, training method, and hyperparameters were identical to those of neural networks used to train the original models. The only exception was the output layer, which had 14 rather than just 5 output nodes to accommodate the additional 9 parameters in the remapping models.

The results of the parameter estimation networks are shown in Figure S1. The drift scalar, threshold, drift variability, and non-decision time showed the same patterns of recovery as in the original DDM and RDM models, which were generally good to excellent. The carryover parameters, as in the original models, were somewhat more difficult to estimate, doubly so here because each one accounted for behavior on a smaller amount of data (1 condition rather than 10). However, the recover was good enough that we should expect reasonable accuracy when the networks are applied to real data.

The neural networks trained on simulated data were then applied to the real data, estimating 14 parameters each for the DDM+remapping and RDM+remapping on each participant. As before, the parameter estimates were then used to generate a synthetic data set from each model, allowing us to compare the model predictions against the real data.

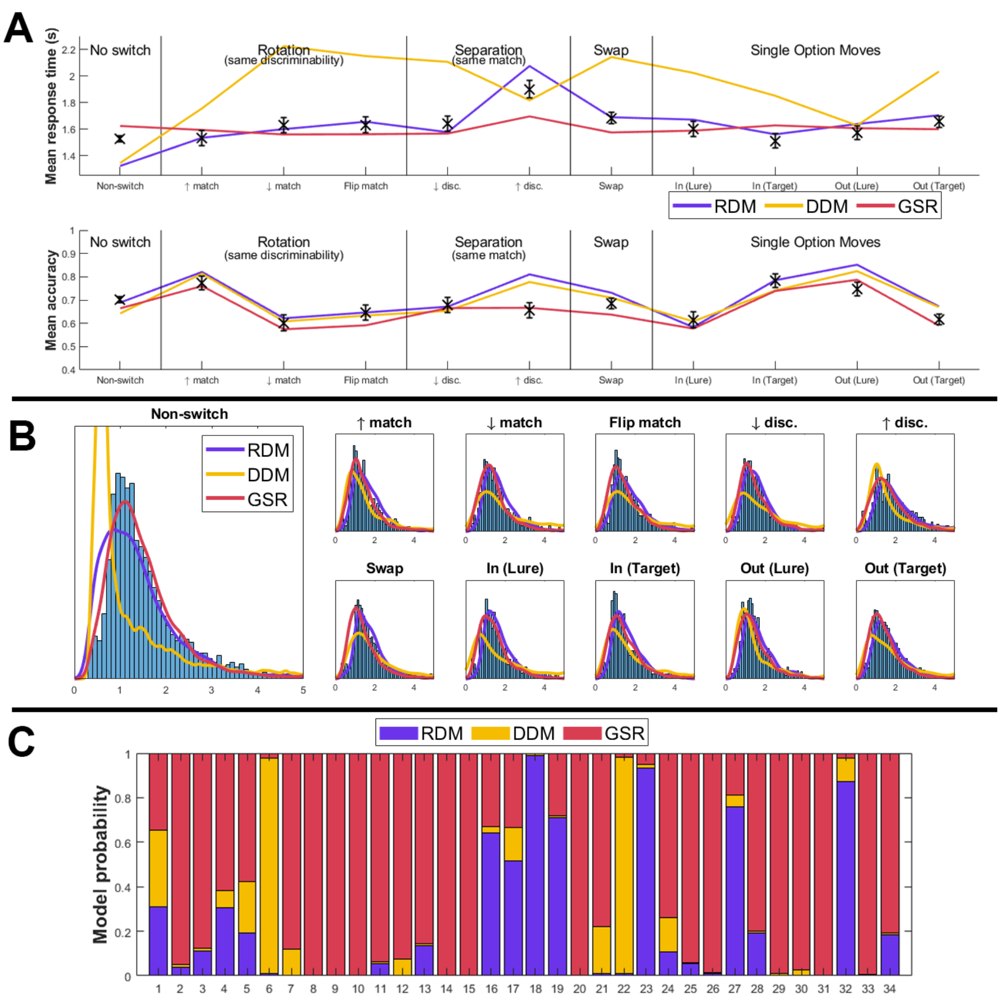

Figure S2

Fits of the flexible-remapping versions of the RDM and DDM, compared to the original GSR. (A) Data (black x) and model predictions from the flexible-remapping RDM (purple), flexible-remapping DDM (yellow, and the GSR model. (B) Histograms of response time distributions for each condition, and the predicted probability density from each model. (C) The posterior probability of each model given the data from each participant (x axis). Bar heights correspond to the probabilities of each model.

The results of the model fitting procedure, and the posterior model predictions from the RDM+remapping and DDM+remapping as well as the original GSR model, are shown in Figure S2. For the RDM+remapping model, the fits are almost identical to those of the original RDM presented in the main text – there appeared to be no major advantage to allowing the carryover parameter to vary from condition to condition. Of course, the fits of the original RDM were generally quite good, so it is possible that there was just not much room for improvement on the original fits. All three models did fairly well at accounting for the accuracy data in particular, although the remapping models showed the same issue as before in accounting for the increased versus decreased discriminability conditions.

Although the RDM+remapping model performed similarly, the DDM+remapping model showed generally poorer fits to the response time data than the original model. The most likely reason for this is that the additional flexibility in the carryover parameter – permitting it to vary between -1 and +1 rather than between 0 and 1 – made it more difficult to estimate. The DDM+remapping model tended to predict negative values for carryover quite frequently, between 17 and 88% across conditions. This seems to have resulted in a long tail to response time distributions, as shown in the yellow distributions in Figure S2B. This tail most likely resulted from negative carryover, such that the evidence during the second stage had to "catch up" to a starting point bias that favored the wrong option in some cases. More generally, the failure of the DDM+remapping model illustrates the potential pitfalls of introducing additional flexibility in the model – although it allows for the model to potentially make more right predictions, it also allows it to make more wrong predictions. As a result, greater model flexibility can also generate variability in model predictions that creates mismatches with the empirical the results (Myung, 2000).

In addition to the parameter estimation and posterior predictions shown in Figure S2A-B, we also carried out a model comparison between the DDM+remapping, RDM+remapping, and GSR models. To do so, we simulated 100,000 simulated data sets from each model, using the same priors as before: all carryover parameters were simulated form a uniform prior on [0,1] for the RDM+carryover and [-1,1] for the DDM+carryover. These simulated data were then used to train a deep neural network to map data onto which model was used to generate it. The layer structure, training algorithm, hyper-parameters, inputs, outputs, and train-test split were identical to the model classification network shown in the main text.

This neural network produced excellent discrimination accuracy, with a correct classification rate of 96.90% on the training set and 94.44% on the validation set.

The trained neural network was then applied to the real data to obtain posterior probabilities of each model given the data. The results are shown in Figure S2C, and are quite similar to the results of the model comparison presented in the main text. The DDM+remapping provided the best account of the data from 3 participants (overall posterior probability = .11), the RDM+remapping provided the best account for 7 participants (overall posterior probability = .21), and the GSR accounted best for the remaining 24 participants (overall posterior probability = .68).

Surprisingly, the proportion of participants best fit by the RDM+remapping was lower than the proportion best fit by the basic RDM in the main text. We suspect that this is due to the increase in model complexity that accompanied the additional 9 parameters. Ultimately, the additional flexibility conferred by adding flexible evidence re-mapping to the DDM and the RDM was not warranted, and did not confer enough advantage over the GSR to substantively change the outcome of the model comparison.

Even more flexibility

At the request of a reviewer, we also tested a version of the DDM where the drift rates could vary across all of the stimulus manipulations and across all of the switch conditions. In the OSF files, this is referred to as the “DDM - extreme flex” model to indicate that it has an extremely high degree of flexibility to capture many patterns of results that result from combinations of stimuli and changes in the choice options. This model included all of the start point / carryover changes that were included in the DDM+remapping model described above for a total of 11 carryover parameters, and estimated a total of 12 drift rates, one for each stimulus level and set of options: 2 levels of discriminability $\times$ 2 levels of coherence $\times$ 3 levels of match. Additionally, it allowed the drift rates to change from the first to the second (post-switch) stage by way of 11 additional drift scalar manipulations. This allowed for the degree of evidence accumulation during the second stage to depend on the switch type, and potentially to include a degree of decay as indicated by other two-stage models of evidence accumulation (Yu et al., 2015; Kvam et al., 2015).

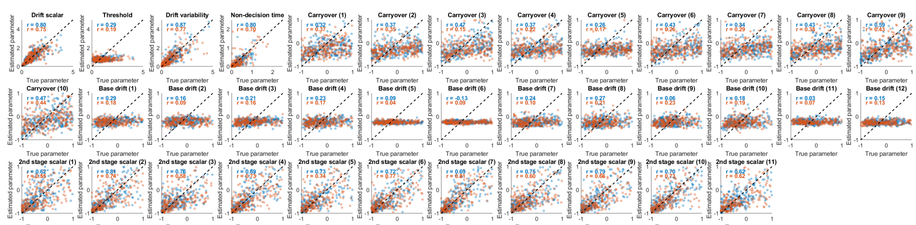

Figure S3

Plot of the true (x) and predicted (y) parameter values from a model recovery study, for the DDM with drift rates that could vary across all conditions. Each subplot corresponds to a different parameter in this model. Drift rates, carryover, and second-stage drift scalar parameters are numbered according to which condition they described.

The results of model recovery for this DDM are shown in Figure S3. Overall, it was substantially more difficult to recover the parameters of this model than to recover the parameters of the models presented in the main text, and even more difficult than the remapping models presented above. In particular, all the parameters related to drift rates were difficult to recover, as they tended to describe behavior across only a few trials at most. However, there was still some correspondence between the true values used to generate simulated data and the estimates of the model, meaning that we should be able to extract some meaningful insights from its estimates.

We attempted to construct an RDM model with similar flexibility, by allowing both of its drift rates to vary across stimulus and option conditions. However, because of the number of parameters of this model, its drift rates were ultimately impossible to recover with the data we had available. We therefore do not include the flexible-drift RDM in our model comparisons, and instead compare the flexible-drift DDM against the base RDM and GSR models.

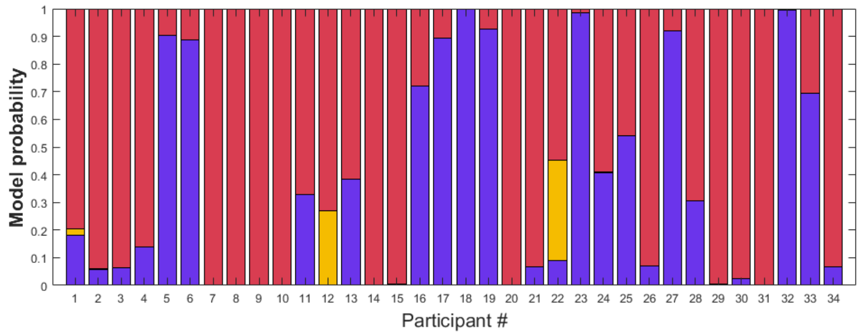

Figure S4

Posterior probabilities of the RDM (purple), flexible-drift DDM (yellow), and GSR (red) models for each of the participants in our study. The proportion of each bar occupied by each model corresponds to the relative posterior probability of that model for the corresponding participant (x).

The result of this model comparison is shown in Figure S4. As shown, the flexible-drift DDM did not fit best for any of the participants in our study. It never fit better than the GSR, and in fact it only ever fit better than the RDM for two participants. We expect that this model is simply too flexible – it makes such wide-ranging predictions that any individual prediction it makes is very unlikely to be borne out in the data (Villarreal et al., 2023). This continues the trend of poorer performance with increased model complexity that was set by the RDM+remapping and DDM+remapping models, suggesting that more complex alternative-specific models will not be able to explain our results and that their additional complexity is unwarranted. In each of the cases, they are still outperformed by the relatively simple GSR.

Implications for continuous models

Although the comparisons presented in the main text are the ones most critical to discriminating between alternative-specific and alternative-general models, some of the condition manipulations from Study 1 can be informative in understanding what kinds of models might be plausible within the class of alternative-general models. Specifically, we look at the geometric similarity representation [GSR] model (Kvam, 2019a), the circular diffusion model [CDM] (Smith, 2016), and two implementations of the spatially-continuous diffusion model [SCDM] (Ratcliff, 2018). These models are shown in Figure S5.

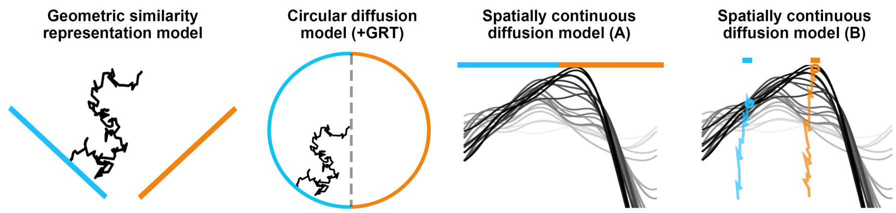

Figure S5

Diagram of the structure of four alternative-general models that we consider in light of the results.

The GSR and CDM are formulated the same as in prior work (Kvam, 2019a,b; Kvam & Turner, 2021; Smith, 2016, 2019). However, we consider two versions of the SCDM for binary choice: one where an option is described by a range of values along the continuum (A), and one where each option is described by a single point (B). In the SCDM-A, evidence accumulates according to a Gaussian process and halts whenever the height of the Gaussian process at any point along the continuum reaches the threshold. In the SCDM-B, evidence accumulates according to a Gaussian process, but thresholds are only applied to the options currently under consideration. In the SCDM-B, a decision will only be made if the accumulators corresponding to the orange or blue option hit a threshold level of activation – not if some other option along the continuum hits the threshold first.

The CDM and SCDM-A both predict that manipulating the discriminability of the response options during a switch should have no effect. Because there is a single threshold divided into regions corresponding to each option, the options moving away from one another or closer to one another will not shift the point at which a decision is made. This means that response times should be identical in switch conditions of Experiment 1 where discriminability is manipulated (separation / same-match conditions), which is not the case in the data. Instead, response times in the condition where discriminability decreases as the options converge (M = 1.64, 95% HDI = [1.85, 1.70]) are much faster than those in the condition where discriminability increases as the options diverge (M = 1.90, 95% HDI = [1.83, 1.97]; BF > 10,000). Accuracy was also effectively the same in the decreased-discriminability condition (M = .68, 95% HDI = [.65, .71]) as in the increased-discriminability condition (M = .66, 95% HDI = [.62, .69], BF_0 = 87.23) indicating that this is not merely an effect of higher thresholds.

Ironically, this effect is the opposite of what we might expect if participants were using some metacognitive knowledge about the difficulty of the choice options to adjust their thresholds. If participants were somehow able to adjust their thresholds from trial to trial (which seems unlikely a priori), then they ought to use a lower threshold in higher-discriminability conditions because the options should be easier to tell apart (Kvam et al., 2023). Such a change would result in slower responses in the high-discriminability condition, rather than the low-discriminability condition. By contrast, the SCDM-B and GSR can capture the observed pattern of data: when the options are far apart (more discriminable), the stimulus matches up poorly with both available response options, resulting in low drifts relative to both choice options and thus slower response times and lower accuracy.

Even if we assume that participants could adjust their thresholds trial-to-trial and make them lower when the options are closer together, this would make the CDM and SCDM-A unable to handle other conditions. In particular, the options get less similar to one another in the lure-out condition – meaning if participants increase their thresholds in the diverge condition, they should also increase their thresholds in the lure-out condition. Again, this is contradicted by the data, where the lure-out condition's response times are among the fastest (M = 1.51, 95% HDI = [1.46, 1.57]) and its accuracy is the highest (M = .75, 95% HDI = [.72, .78]) compared to all other conditions. Put together, these findings strongly indicate that the GSR and SCDM-B should be favored over the CDM and SCDM-A.

Among the GSR and SCDM-B, there are few phenomena that would allow us to clearly draw conclusions in favor of one or the other. In part, this is because the GSR can mimic nearly perfectly the performance of the SCDM as the number of dimensions in the GSR model increases (Kvam & Turner, 2021). For now, these two models both seem to be plausible accounts of how people make decisions among shifting sets of options, which sets them a step above any other dynamic models of decision-making.

References

Kvam, P. D. (2019a). A geometric framework for modeling dynamic decisions among arbitrarily many alternatives. *Journal of Mathematical Psychology*, *91*, 14–37.

Kvam, P. D. (2019b). Modeling accuracy, response time, and bias in continuous orientation judgments. *Journal of Experimental Psychology: Human Perception and Performance*, *45*(3), 301–318.

Kvam, P. D., Marley, A., & Heathcote, A. (2023). A unified theory of discrete and continuous responding. *Psychological Review*, *130*(2), 368–400.

Kvam, P. D., Pleskac, T. J., Yu, S., & Busemeyer, J. R. (2015). Interference effects of choice on confidence: Quantum characteristics of evidence accumulation. *Proceedings of the National Academy of Sciences*, *112*(34), 10645–10650. doi: 10.1073/pnas.1500688112

Kvam, P. D., & Turner, B. M. (2021). Reconciling similarity across models of continuous selections. *Psychological Review*, *128*(4), 766—786.

Myung, I. J. (2000). The importance of complexity in model selection. *Journal of Mathematical Psychology*, *44*(1), 190–204. doi: 10.1006/jmps.1999.1283

Ratcliff, R. (2018). Decision making on spatially continuous scales. *Psychological Review*, *125*(6), 888–935.

Smith, P. L. (2016). Diffusion theory of decision making in continuous report. *Psychological Review*, *123*(4), 425–451. doi: 10.1037/rev0000023

Smith, P. L. (2019). Linking the diffusion model and general recognition theory: Circular diffusion with bivariate-normally distributed drift rates. *Journal of Mathematical Psychology*, *91*, 145–158.

Villarreal, M., Etz, A., & Lee, M. D. (2023). Evaluating the complexity and falsifiability of psychological models. *Psychological Review*, *130*, 853–872.

Yu, S., Pleskac, T. J., & Zeigenfuse, M. D. (2015). Dynamics of postdecisional processing of confidence. *Journal of Experimental Psychology: General*, *144*(2), 489.