Modeling accuracy, response time, and bias in continuous orientation judgments

Peter D. Kvam

The Ohio State University

Despite the prevalence of real-world and laboratory tasks where people select among many options, cognitive models have traditionally focused on choices among small sets of alternatives. This has resulted in theoretical and empirical gaps in understanding the decision processes that go into selections among many alternatives or responses that fall along a continuum. In this paper, I address these issues by modeling decisions in a perceptual study where participants produce continuous orientation judgments. I show that manipulations of stimulus difficulty and time pressure have parallel effects to binary choice, with greater stimulus difficulty yielding slower and less accurate responses and time pressure resulting in faster responses at the expense of accuracy. These effects are well accounted for by the circular diffusion model, with drift magnitude parameters shifting as a function of difficulty and threshold parameters shifting with time pressure. However, a manipulation of bias using a predecision cue resulted in bimodal distributions of responses that cannot be explained by the model in its original formulation. To account for this result, I develop a theory of bias based on split attention and racing 2-dimensional diffusion processes. This model suggests that responses are determined by both cue-driven and stimulus-driven evidence accumulation processes, where the winning process determines responses and response times. As a result, it predicts critical features of responses and response times in the conditions with predecision cues, including bimodal distributions of responses and the longer response times observed in conditions where there was a discrepancy between cue and stimulus orientations.

Keywords: decision-making; continuous response; computational modeling; orientation; bias

Significance statement

- This study establishes how accuracy and response times in continuous report change with time pressure, manipulations of stimulus difficulty, and predecision cues.
- Predecision cues that conflict with the stimulus yield bimodal distributions of responses, which cannot be accounted for with traditional theories of bias.
- A new theory of bias in continuous report is developed, where stimulus- and cue-driven accumulators compete to generate responses.

• This model predicts patterns of responses and response times across conditions, providing a more complete account of bias and evidence accumulation in continuous report tasks.

Introduction

Many of the tasks people accomplish require them to make selections among many alternatives or along a continuum. Whether we are deciding how much time or money to invest, figuring out what direction to walk or drive, reproducing the orientation of a stimulus, coming up with the next word in a sentence, or even producing a musical note, there is a plethora of response options available. However, our understanding of the decision processes underlying these sorts of selections is sparse. This is largely due to a traditional focus on developing theories of binary decisions, where a person only has two alternatives from which to choose (Townsend, 2008). Even models of multiplealternative decisions deal primarily with selections among a relatively small set of potential responses (see, e.g., Brown & Heathcote, 2008; Roe et al., 2001; Trueblood et al., 2014; Usher & McClelland, 2001). Fully accounting for responses and response times among larger sets of alternatives requires

Peter D. Kvam, Department of Psychology, The Ohio State University. Correspondence concerning this article should be addressed to Peter Kvam, Department of Psychology, 1827 Neil Ave, Columbus, OH 43210; or at kvam.peter@gmail.com.

^{© 2018,} American Psychological Association. This paper is not the copy of record and may not exactly replicate the final, authoritative version of the article. Please do not copy or cite without authors' permission. The final article will be available, upon publication, via its DOI: 10.1037/xhp0000606

a computationally tractable theory of the decision processes that generate them.

Recent work in visual working memory looking at the allocation of cognitive resources has begun to address this issue by modeling distributions of responses that fall along a circle (Fougnie et al., 2012; Nosofsky & Gold, 2017; Van Den Berg et al., 2012; Van den Berg et al., 2014; Zhang & Luck, 2008). These approaches use circular or wrapped distributions to generated predictions about the density of responses from 0 to 360 degrees (or 0 to 2π radians). However, the models used in these tasks are typically static theories that predict a final distribution of responses without incorporating response dynamics, meaning that they do not account for differences in reaction times [RTs] or temporal shifts in response distributions.

The 2-dimensional circular diffusion model proposed by Smith (2016) adds a dynamic component to these models, generating a process-level explanation of how people come to make responses on a circle. It predicts the complete joint distribution of responses and response times by quantifying the decision process in terms of a decision variable (state), how this state changes over time (sampling), and what conditions result in the generation of a response (choice criterion). While promising, this model has not yet been applied to real data, in part because continuous response tasks are rare and ones considering response times even more so. This paper seeks to remedy both issues by developing a simple continuous-response paradigm and examining how (if) the model is capable of accounting for observed behavior on the task.

Late-breaking work of Ratcliff (in press) has also examined the joint distributions of responses and response times using a continuous response model, though this has primarily looked at responses that can be arranged on a line or plane rather than a circle. Despite this, it has several analogous parameters to the model of Smith (2016) and provided some hints that they shift in sensible ways with manipulations of the stimulus and instructions.

Benchmark phenomena

The development of cognitive models of binary choice has benefited tremendously from a set of benchmark accuracy and response time phenomena that can be observed across task domains. Each phenomenon is triggered by a specific task manipulation, and the resulting changes in behavior correspond to shifts in a particular parameter of the cognitive model. This allows for inferences based on *selective influence*: each parameter of the model can be affected independently by a particular manipulation, so we infer that the parameters describe distinct parts of the cognitive process. Ideally, we would like to be able to do the same thing with manipulations in continuous tasks. Ratcliff (in press) suggested that this might be possible by showing that parameters describing the rate of evidence accumulation and response caution could be uniquely influenced by manipulations of stimulus quality and rewards for fast vs accurate responses, respectively. One of the goals of this paper is to replicate and extend these findings by examining how task manipulations analogous to those in binary choice affect parameters of continuous response models.

Typically, binary choice models will describe the decision process in terms of evidence accumulation, where a person gathers samples of information from the stimulus over time to form their beliefs about the relative support for their choice options. The accumulation process is broken down into an initial belief state (what a person believes before gathering stimulus information), evidence dynamics (how their beliefs change over time as they gather stimulus information), and decision rules (when to stop gathering evidence and select one of the alternatives). The initial state models bias, where a tendency to favor one alternative due to predecision information is represented as a state that starts closer to one alternative (A) than another (B). Evidence dynamics are described in terms of the *drift*, or the average rate of evidence change toward A relative to B, and the noise or random change in evidence over time. Decision rules are specified as a threshold value, where beliefs stronger than θ in favor of an alternative result in a response favoring that option (Stone, 1960; Ratcliff, 1978a). Ratcliff & McKoon (2008) point to three important manipulations that typically impact these different components of the decision process: speed-accuracy manipulations (threshold), coherence or difficulty of stimuli (drift), and bias manipulations (initial state). The effects of each manipulation have been well-established across domains in binary choice tasks, and can reliably produce shifts in one or more parameters when a diffusion or accumulator model is fit to participants' data.

A *speed-accuracy trade-off* results from manipulations that incentivize participants to achieve either faster responses or a higher proportion of correct responses (Bogacz et al., 2010; Heitz & Schall, 2012; Vickers & Packer, 1982; Wick-elgren, 1977). In conditions where they are encouraged to focus on accuracy, participants typically gather more information from the stimulus, resulting in longer response times. Conversely, when they are encouraged to respond faster, participants typically gather less information, resulting in less accurate responses based on inferior evidence. This trade-off suggests that there is an internally adjustable criterion that specifies the amount of information to gather before triggering a decision. In line with this, it ordinarily corresponds to shifts in threshold parameters.

A *difficulty effect* arises when the difficulty or coherence of the stimulus is manipulated across trials. When the information that decision-makers receive is better, they receive information favoring the correct option more quickly relative to noise. As a result, one can expect their accuracy to be higher and/or their responses to be faster. Conversely, when the quality of information the decision-maker receives is lower, their responses will be less accurate or slower. The quality of decision information is indexed by the drift rate, which can often be directly estimated from the evidence a decision-maker receives (see Busemeyer & Townsend, 1993; Krajbich et al., 2012; Link & Heath, 1975; Palmer et al., 2005; Nosofsky & Palmeri, 1997; Ratcliff, 2014).

Finally, decision bias manipulations asymmetrically affect responses and response times to one alternative relative to another.¹ The alternative toward which a person is biased will exhibit faster response times, and the one against which they are biased will exhibit slower ones. When a person is biased toward the alternative which happens to be correct (congruent bias), one can expect that they will be more accurate, and less accurate when biased toward an alternative which is incorrect (incongruent bias). Bias may arise due to unbalanced rewards for hits versus correct rejections or unbalanced penalties for false alarms versus misses (Diederich & Busemeyer, 2006; Pleskac & Busemeyer, 2010), different base rates of stimuli (Wagenmakers et al., 2008; Wolfe et al., 2007), variation in subjective beliefs, or signals of true differences in prior plausibility of various hypotheses (Mulder et al., 2012).

Manipulations of bias are usually connected to shifts of the starting point of the process (equivalent to setting the threshold asymmetrically; Edwards, 1965; Ratcliff, 1985; Mulder et al., 2012), but it can alternatively be modeled by shifts in the drift rate as well (Krajbich et al., 2010; Ratcliff et al., 1999). Diederich & Busemeyer (2006) examined another possibility as well, where pre-decision information was modeled as a distinct process such that evidence was accumulated first by accumulating pre-decision information (stage 1), then accumulating stimulus information (stage 2). In effect, results in a distribution of initial states shifted toward the biasfavored response rather than a single starting point.

The 2-dimensional diffusion process used to model continuous responses (Smith, 2016), described in the next section, has a parameter structure that parallels the evidence accumulation models described above. It uses a drift magnitude to quantify the rate of evidence accumulation toward the best alternative, a variability parameter to describe noise in the accumulation process, a threshold to represent the strictness of the decision rule, and an initial state that can reflect predecision information and bias. Therefore, it should be expected manipulations of time pressure, difficulty, and bias should affect the analogous parameters of the 2-dimensional model – drift, threshold, and start point, respectively.

Although these three manipulations and their resulting phenomena certainly do not constitute an exhaustive list, they represent three of the most important empirical phenomena that have guided construction of binary choice models. They are particularly interesting because each effect corresponds theoretically to a different component of the decision-making process. If continuous responses are generated from an analogous process, then the empirical results should mirror those of binary choice when manipulating stimulus quality, time pressure, and bias. And in turn, these results be predicted by shifts in the corresponding parameters of the continuous response model.

Preliminary support for time pressure and difficulty manipulations in continuous report suggest that they do indeed affect response caution and drift, respectively (Ratcliff, in press). However, it is unknown what exactly the effect of traditional bias manipulations may be, or whether these findings translate to responses that can 'wrap' and fall around a circle (as is typical of studies in visual working memory, including those looking at orientation or motion direction). The experiments and model presented here seek to fill these gaps in our empirical and theoretical understanding.

Model description

As it predicts the joint distribution of responses on a circle and corresponding response times as required for our present purposes, the models presented in this paper use the 2-dimensional circular diffusion model of Smith (2016) as a starting point.² In this model, a decision maker's evidence state is represented as a position in a 2-dimensional space that quantifies how much they have for all of the available hypotheses. Typically, we think of 2-dimensional positions in terms of x and y coordinates, but it is often more convenient to describe the evidence state in terms of polar coordinates r and ϕ . Respectively, these correspond to the distance of the state to the origin on $[0,\infty)$ and its orientation relative to the positive x direction on $[0,2\pi)$ (radians) or $[0^{\circ},360^{\circ})$ (degrees). The angle coordinate ϕ can be understood as the alternative a person favors or believes most, while the radius coordinate r describes the strength of evidence in favor of that alternative. For example, a person judging the direction of motion of a stimulus from 0 to 360 degrees might have a state of $r = 0.1, \phi = 45^{\circ}$, which would indicate (assuming r varies from say, 0 to 1) a weak belief that the stimulus is moving toward the upper-right corner of the screen. Put together, these two coordinates describe the evidence at any

¹In this paper, bias refers to a tendency to respond in favor of one set of alternatives over another by virtue of some information external to the target stimulus. It is not intended to refer to patterns of non-coherent behavior, as the term is used in the judgment and decision-making literature on heuristics and biases (Tversky & Kahneman, 1974; Gilovich et al., 2002). In fact, predecision biases in our framework may actually be adaptive and even correspond to an optimal prior for a given decision (Bogacz et al., 2006).

²Although the model of Ratcliff (in press) could probably be extended to model responses on a circle, the increased number of parameters and lack of analytic likelihoods make it unnecessarily difficult to work with for the purposes of this paper.

point in time, including the *initial state* of the process.

The evidence state shifts as a person gathers new information from the stimulus, and these dynamics are determined by a 2+ parameter distribution that describes how r and ϕ change over time. The model of Smith (2016) describes it in terms of a pair of drift components and a diffusion rate. The drift components describe what direction the state moves on average in terms of the ϕ coordinate [drift *direction*] and how quickly it moves in that direction in terms of the r coordinate [drift magnitude]. Diffusion describes the noise in this process - the random information generated during the evidence accumulation process from the stimulus or from noisy neural processes. Finally, the evidence accumulation process terminates when a person is certain enough of their current leaning - when the r coordinate exceeds a certain value. This value is called the *threshold* of the decision process, and creates a circular boundary surrounding the origin.



Figure 1. State representation, sampling, threshold, and evidence trajectories for an orientation matching task.

A diagram of this model is shown in Figure 1. I test manipulations and parameters of the model using an orientation task, described in the next section. In this task, participants can give orientation responses anywhere from 0 to 180 degrees, which is mapped onto 0 to 360 degrees on the response circle by simply multiplying orientations by 2.³ As participants gather evidence, their support for the various orientation responses shifts over time, moving their state around inside the response circle until they hit its edge. The response time for a particular trial is given by the time it takes to reach this response threshold (gray responses at the upper right of Figure 1). Thus, both response and response time are predicted as a function of the initial state, drift direction, drift magnitude, and threshold, plus an additional parameter quantifying the amount of time it takes for non-decision processes (such as stimulus encoding and motor actions).

Methods

This study seeks to establish empirical results in continuous choice by exploring continuous analogues of the three classic binary choice phenomena. The first priority is to characterize these effects descriptively, and the second is to investigate if and how a continuous response model can account for them. In doing so, I examine what parametric shifts allow the circular diffusion model to account for the results or what modifications are needed in order to fully capture them. If the structure of the phenomena and models line up with binary choice, one should expect that speed-accuracy trade-offs correspond to shifts in threshold, difficulty corresponds to shifts in drift, and predecision information corresponds shifts in either initial states or drift.

Each of these predictions is tested using an orientation judgment task where participants see a rapid series of Gabor patches and must judge the mean orientation of the patches by giving a response from 0° to 180° . In this continuous task, assessing accuracy in terms of proportion correct would require us to discretize responses and throw away informative data about the continuum of responses. Instead, I will refer to the *deviation* of a response from the true answer or a cue. Because the stimuli in this study all fall on 0 to 180 degree orientations, any response deviations are bounded between 0 and 90 degrees. Note that being 120 degrees clockwise is the same as being 60 degrees counterclockwise from the true mean orientation. Because the stimuli fall only on a half-circle, responses at ± 180 degrees from one another are equivalent. The smallest distance to the target / true response is taken as the actual response deviation. Hence, perfect accuracy will result in response deviations of 0, uniform random guesses will result in average response deviations of 45 degrees, and perfect inaccuracy will produce response deviations of 90 degrees.

Experiment

As in studies of binary choice, a speed-accuracy trade-off was implemented by changing the incentive structure of the task to reward either response speed or response accuracy. Difficulty was manipulated by changing the coherence of the stimulus, and predecision bias was manipulated by providing informative cues prior to the stimulus appearing on-screen.

Task. The task is shown in Figure 2. This was primarily an orientation detection task, where participants saw a

³Motion direction tasks commonly suffer from an issue where motion in one direction will be confused for motion in the opposite direction along the same axis. For example, motion at 90 degrees is frequently confused for motion at 270 degrees, apparently because the axis of motion is sampled in combination with its motion along that axis. To avoid these types of confusions, I examined orientation, which does suffer from these complications due to symmetry.

rapid series of Gabor patches and had to select the mean orientation of the patches they saw. Each Gabor stimulus was generated by first choosing a random mean orientation from a uniform distribution on [0,180) degrees. On each screen refresh (drawn every 16.7 ms / at 60 Hz), a new Gabor patch would be drawn from a wrapped normal distribution centered on this mean orientation. The standard deviation of this wrapped normal depended on the difficulty manipulation: the standard deviations were 15, 30, or 45 degrees, corresponding to easy, medium, and hard difficulty (larger standard deviations result in noisier stimulus information). The difficulty was drawn randomly between these levels on each trial.

To begin a trial, participants clicked on a small white circle in the middle of the screen. Their cursor was then moved to the exact center of the screen, which involved only minor adjustments as the white circle closely circumscribed this point. In the cued condition (top left of Figure 2), constituting half the trials, they then saw a green line on the screen. This line corresponded to the true mean orientation of the stimulus in 50% of cued trials, and on the other 50% of cued trials was randomly set to 20, 50, or 70 degrees away from the true orientation (either clockwise or counterclockwise with a 50-50 chance). This allowed the cue to be informative but not overwhelm stimulus information in terms of usefulness. For uncued trials, they simply saw an uninformative green circle appear (bottom left of Figure 2) before the stimulus.

The green line cue or circle disappeared after 1000ms, the screen was left blank for 250ms, and then the rapid Gabor stimulus was presented. Participants could respond at any time by moving their mouse across the edge of the white response circle at their desired orientation (outer white circle, Figure 2). Their response could be made on the lower or upper half of the circle – a response at 30 degrees was equivalent to one at 210 degrees – though nearly all participants preferred to use the upper half.

Participants received points in the task for both the speed and accuracy of their responses, contingent on whether a block of trials was a speed-emphasis block or accuracyemphasis one. In the speed condition, they received 100 points for responding within 800ms of the stimulus onset and up to 100 points based on how close they were to the true mean orientation of the stimulus: each degree away decreased their reward by 1.1 points. In the accuracy condition, they received up to 200 points based solely on how close they were to the true mean orientation - in this case, each degree away decreased their reward by 2.2 points. They were encouraged to have ballistic mouse movement by penalizing them based on the amount of time the mouse pointer spent between the center (circle that they clicked to start the trial) and the edge of the response circle. This was done both to control the non-decision time (time required for non-decision

response processes like moving the mouse) across trials as well as avoid having participants externalize their beliefs or use the mouse position to test orientations before responding. Participants who accumulated a sufficient number of points during the task received additional research credit for their participation.

Participants. The study materials and participants were approved by the Michigan State University Institutional Review Board. A total of 12 Michigan State University students (8 female, 4 male) each completed 960 trials of the experiment for class credit. Participants were primarily 18-26 years old. One additional participant completed the task, but was removed from further analyses for having a median response time well outside the normal range for participants (> 5000ms, compared to 500-1500ms for other participants). The full task, including introduction and practice trials, took approximately 1 to 1.5 hours to complete.

The sample size of 12 participants was selected so that the group-level posterior distributions of parameter estimates were both sufficiently constrained and suitably wellpopulated within each person. Because each participant provided a large number of data points and all statistics were estimated in a hierarchical Bayesian way, this sample was sufficient to obtain precise estimates on both individual and group-level parameters. For the cognitive model (circular diffusion) analyses, each individual's data was fit separately.

Materials. All stimuli were generated and responses recorded using MATLAB and Psychtoolbox-3 (Brainard, 1997; Kleiner et al., 2007). Responses were recorded on the mouse.

Procedure. Upon entering the laboratory, participants were briefed on the content of the task and completed informed consent. They were then seated in a dark, windowless office to complete the task. They completed approximately 60 practice trials of the task (more trials were generated when responses were too inaccurate or slow), with immediate feedback on how quick and how accurate their responses were on each practice trial.

Upon finishing the practice trials, participants completed 960 trials of the main task as described above. These were organized into 12 blocks of 80 trials. The blocks were organized so that participants would see three blocks of accuracy with a predecision orientation cue, three blocks of accuracy trials with no orientation cue, three blocks of speed with the orientation cue, and three blocks of speed trials with no predecision orientation cue. Difficulty was randomly determined on each trial, so it varied within blocks. The order of the blocks of trials was randomized across participants.

Before each block, participants received instructions for the upcoming block's condition, indicating both how they would earn points (for the speed / accuracy manipulation) as well as whether there would be a predecision cue. They were also reminded before each trial whether they were in a



Figure 2. Time course of a trial of the orientation detection task. Pre-decision cues were given as a green line (left panels), then the series of Gabor patches was presented (center). They entered their response at any time by moving the mouse across the white response circle (right).

speed or accuracy block with a small piece of text above the white pre-trial center circle, which read 'SPEED' or 'ACCU-RACY'.

Upon completing the experiment, participants were debriefed on the purpose of the study and told how much credit they would receive based on the number of points they accumulated during the task.

Results

The analyses focused on two main outcomes. The first was simple response times – the number of seconds it took for a person to enter their response on a trial. The second was error magnitude, or the distance in degrees between a response on a particular trial and the true mean orientation of the stimulus for that trial.

All analyses used a hierarchical Bayesian linear model, where individual-level parameters (coefficients for the size of each standardized effect) were constrained by a grouplevel distribution. Unless otherwise specified, the Bayesian models used diffuse priors so that the data had maximal influence over parameter estimates. In reporting, I provide the mean posterior parameter estimates as well as the best estimate of the range that includes the 95% most credible values from this posterior distribution of parameter values. This is referred to as the 95% highest density interval [95% HDI] (see Kruschke, 2014). Here, I report only the group-level parameter estimates for each effect for the sake of brevity.

The data as well as the code for each of the JAGS models is available on the Open Science Framework at osf.io/q3ytj (Kvam, n.d.).

Response times. Mean response times increased with greater stimulus difficulty, increased in the accuracy condition relative to the speed condition, and increased with the degree of mismatch between the predecision cue and the stimulus orientation. These results are shown in Figures 3 and 4.

Before analysis, response times were log transformed in order to make them approximately normally distributed.



Figure 3. Response times (left) and accuracy in terms of degrees deviation from the correct response (right) across difficulty levels and speed-accuracy manipulation. Note that higher values indicate less accurate responses. Error bars indicate pooled standard error across participants.

In the descriptive model, response times were predicted as a function of four main factors. The first was stimulus difficulty, operationally defined as the standard deviation of the orientation of the jittering Gabor patches (15 / 30 / 45 degrees). The second was the time pressure manipulation, coded in the models as 0 for accuracy and 1 for speed. Third was the presence of the cue; the default cue was assumed to be the orientation that matched the true stimulus orientation, so this factor describes the benefit conferred by an informative cue. The final factor was the orientation of the cue, which corresponded to the number of degrees that the cue deviated from the true stimulus orientation (i.e. 0/20/ 50 / 70 degrees). The inclusion of this final factor was conditioned on the third factor, so that cue orientation was not considered when there was no cue. Before analysis, response times were log transformed in order to make them approximately normally distributed. The stimulus difficulty and the cue orientation were standardized before using them to predict responses or response times. Put together, the hierarchical Bayesian model had an intercept, a linear main

effect of the difficulty of the stimulus, a binary main effect of whether each trial was a speed-emphasis trial, a binary main effect of whether a trial included a cue, a linear main effect for the orientation of the cue (which was conditioned on the binary indicator for cued trials so it did not come into play on uncued trials), and all 2-, 3-, and 4-way interactions between these factors.

The mean estimates and 95% Highest Density Intervals for each coefficient are shown in Table 1. As shown, the only effects with substantial contributions to differences in response times were the main effects, including the main effect of cue orientation that was contingent on the cue being present (and is hence presented as an interaction). Response times increased with difficulty as well as with less accurate cue orientations: the further the cue was from the true mean stimulus orientation, the longer participants took to respond. By contrast, response times decreased with speed instructions as well as with the presence of an informative cue.



Figure 4. Mean response times (left) and response deviation / error magnitudes (right) as a function of the orientation of the cue (x). Note that higher response deviations indicate less accurate responses. Error bars indicate pooled standard error across participants. The black dotted line with gray error box corresponds to mean response time and accuracy from the uncued conditions.

Each of the interactions (except of course the cue interaction, which is a conditional main effect) in the model included zero as a credible value, suggesting that the various manipulations did not interact strongly enough to produce substantial changes in mean response times.

This model confirms that difficulty, speed-accuracy, and cue / cue orientation manipulations had their anticipated effects. The raw data for difficulty and speed-accuracy conditions can be seen in Figure 3 (left panel), and data for cue orientation conditions can be seen in Figure 4 (left panel).

Accuracy / response deviation. The second characteristic of responses examined was their response deviations from the true orientation. For this metric, I took the true mean orientations of the Gabor stimuli and compared it to participants' responses on the circle. A person's response deviation for a particular trial was the number of degrees clockwise that their response fell relative to the true orientation of the stimulus. For example, if their response was 2 degrees counterclockwise from the actual mean orientation, their response deviation would be -2 degrees. All response deviations therefore fell between -90 and +90 degrees. On average, response deviations were approximately zero, indicating no systemic clockwise or counterclockwise bias in responses.

In order to gauge accuracy, it is convenient to refer to the absolute value of the response deviation, henceforth called the error magnitude, which is simply the the distance between stimulus and response in degrees without regard to its direction. The error magnitudes across conditions are shown in red in Figures 3 and 4. As expected, error magnitudes increased in the speed relative to the accuracy condition, increased with higher stimulus difficulty, and increased as the predecision cue conflicted more with the stimulus.

To formally describe how participants' responses shifted across conditions, I used a model that is highly similar to the one for response times. However, rather than the mean location of responses – which is essentially zero, as they are generally centered on the correct response – each factor predicted the *variance* of the distribution of responses. ⁴

As before, these outcomes were allowed to change as a function of main effects of difficulty (linear effect), speed manipulation (binary factor), cue presence (binary factor), and cue orientation (conditional linear effect), as well as the interactions between all combinations of these factors. For simplicity, I discuss the results in terms of accuracy. This means that lower coefficients for effects on response deviation correspond to higher accuracy and higher coefficients correspond to lower accuracy.

Estimates for all coefficients are shown in Table 2. The same main effects of manipulations appeared as in response times. High difficulty, speed emphasis, and inaccurate orientation cues all increased error magnitudes while low difficulty, accuracy emphasis, and the presence of an accurate cue all decreased error magnitudes.

However, unlike the response time results, the response results were somewhat complicated by interactions between manipulations. The 2-way interaction between difficulty and speed emphasis as well as the 3-way interaction between difficulty / cue / orientation and 4-way interaction between difficulty / speed / cue / orientation, each of which increased accuracy, can be at least partially attributed to a floor effect. These manipulations put together would result in the lowest accuracy conditions – note the large main effects of difficulty, speed, and cue orientation on accuracy – but average response deviations have an upper bound. Guesses will yield responses that are on average ± 45 degrees away from the

⁴Formally, they predicted the log of the variance of this distribution in order to allow the sum of all factors to take any value and to make the predicted variances log-normally distributed (for a similar approach, see Kvam & Pleskac, 2016).

PETER D. KVAM

Table 1

Mean estimates of coefficients for main effects and interactions of stimulus difficulty, time pressure, cue presence, and cue orientation on response times. The ranges containing the 95% Highest Density Interval (HDI) are also provided. Intervals excluding zero are starred for quick reference.

| Coefficient | Mean Estimate | 95% HDI |
|---|---------------|-----------------|
| Difficulty | 0.06 | [0.04, 0.08]* |
| Speed | -0.13 | [-0.19, -0.06]* |
| Cue | -0.09 | [-0.14, -0.03]* |
| $Cue \times orientation$ | 0.04 | [0.02, 0.05]* |
| Difficulty \times speed | -0.02 | [-0.03, 0.00] |
| Difficulty \times cue | -0.02 | [-0.04, 0.00] |
| Difficulty \times speed \times cue | 0.00 | [-0.02, 0.02] |
| Difficulty \times cue \times orientation | -0.00 | [-0.02, 0.01] |
| Speed \times cue | -0.02 | [-0.06, 0.02] |
| Speed \times cue \times orientation | -0.01 | [-0.01, 0.00] |
| Difficulty \times speed \times cue \times orientation | 0.01 | [-0.01, 0.02] |

Table 2

Mean estimates of coefficients for main effects and interactions of stimulus difficulty, speed manipulation, cue presence, and cue orientation on accuracy of responses. The ranges containing the 95% Highest Density Interval (HDI) are also provided. Intervals excluding zero are starred.

| Coefficient | Mean Estimate | 95% HDI |
|---|---------------|----------------------|
| Difficulty | 0.79 | [0.70, 0.89]* |
| Speed | 0.16 | [0.07, 0.26]* |
| Cue | -0.31 | $[-0.49, -0.14]^*$ |
| $Cue \times orientation$ | 0.52 | [0.40, 0.65]* |
| Difficulty \times speed | -0.09 | $[-0.15, -0.04]^*$ |
| Difficulty \times cue | -0.24 | $[-0.46, -0.02]^*$ |
| Difficulty \times speed \times cue | 0.03 | [-0.06, 0.13] |
| Difficulty \times cue \times orientation | -0.09 | $[-0.15, -0.03]^{3}$ |
| Speed \times cue | -0.05 | [-0.21, 0.09] |
| Speed \times cue \times orientation | 0.06 | [0.00, 0.12] |
| Difficulty \times speed \times cue \times orientation | -0.05 | [-0.09, 0.00] |
| | | |

true orientation, so mean response deviations will generally not exceed these values in any conditions.

The interaction between difficulty and cue presence was the strongest effect of these and the only one to rival any of the main effects. This seems to be related to the effect of the cue on response distributions. As I show in the next section, it appears that participants often responded consistent with the cue without merging it with stimulus information. I attribute this to a competition process between the cue and stimulus, where a cue-driven and stimulus-driven process race to threshold. Because the stimulus-driven process is more likely to win the race. As a result, accuracy is particularly strongly affected by the cue in the more difficult conditions, yielding the cue-difficulty interaction.

Response distributions

One of the most striking results of the study comes from the distributions of responses observed in the cued conditions, shown in Figure 5. When the cue indicates the correct stimulus orientation (top panel), responses form a typical von Mises distribution centered on the true orientation of the stimulus. However, when the cue substantially differs from the mean orientation of the stimulus, the pre-decision cue resulted in bimodal distributions of responses (bottom panels). In these cases, one mode is centered on the stimulus orientation and the other is centered on the cue orientation.

This bimodal distribution of responses does not make sense from the point of view of models that propose an integration of the evidence from the two sources of information (cue and stimulus). Models that posit a shift in starting point (including the 2-stage model of Diederich & Busemeyer, 2006) or drift distributions would suggest that a compromise between cue- and stimulus-consistent information could be reached, resulting in responses that are concentrated partway between the stimulus mean and the cue. Instead, decision-makers appear to be using either the cue or the stimulus to generate their responses rather than merging informa-

8



Figure 5. Distribution of responses in cued conditions. Dotted black lines indicate the stimulus mean orientation, and vertical gray lines indicate the orientation of the cue relative to the stimulus.

tion from the two sources.

Given the commonplace use of integration-type bias models in binary choice, this result is somewhat surprising. At the same time, it illustrates one of the benefits of using continuous tasks. In binary choice, the effects of drift bias and starting point bias both swing decisions toward the bias-favored alternative in similar ways, so the exact nature of this bias (integrative vs competitive) would not be observable from binary choice proportions. The continuous task – and in particular the distribution of responses it generates – therefore grants us a richer picture of the biasing effect of cues.

The bimodal distribution of responses cannot be handled by the lone 2-dimensional diffusion process of Smith (2016), as it predicts unimodal distributions of responses on the circle. As a result, it is necessary to develop an extended theory of how responses are generated in the presence of a cue. Bimodality is accounted for in the model of Ratcliff (in press) by virtue of competing accumulation processes; a similar approach can be gainfully applied here. I supplement the circular diffusion model with an analogous mechanism by constructing a race between stimulus-driven and cue-driven response processes for the cued trials. To reach a point where this can be done, it is helpful to focus first on understanding the stimulus-driven component of the response process. Therefore, I present a model for the uncued conditions next, which can in turn be used to model the stimulus-driven component of responses in the cued conditions.

Modeling

Although it may not work well for modeling cued conditions, the circular diffusion model can still be applied to modeling orientation responses in the uncued condition. To do so, minor transformations to the response space are required to map the stimuli – which are rotationally symmetric at 180 degrees – to responses on the circle that are dispersed across 0 to 360 degrees. The simplest transformation here is to multiply all orientation responses by two. Doing so constructs a response space where stimuli at orthogonal rotations (i.e. 45 / 135 degrees or horizontal / vertical) provide evidence against one another. The response options can then be arranged in a circle as shown in Figure 1: vertical responses are on the upper part of the circle opposite horizontal responses, and rightward learning responses are on the right side of the circle opposite leftward leaning responses.

A person's beliefs at any given time are represented as a point within this circle centered on the origin. Beliefs in favor of a vertical orientation are represented as a point further upward, those in favor of a leftward orientation further left, and so on. A person's initial beliefs can be described as one of these points – in the case where they believe the relative likelihoods of the different stimulus states are even (e.g., if they have not seen any cue or stimulus information), this point would be located at the origin or the center of the circle shown in Figure 1. The initial state is then updated with new information as a person views the stimulus: as they sample left-leaning orientations, the state moves leftward; with horizontal orientations, it moves downward; with vertical orientations, upward; and so on.

As this process unfolds over time, the samples add up together to form the trajectory of a random walk, shown as the gray trajectories in Figure 1. As the steps of the random walk become very small, this motion approximates Brownian motion with drift (Stroock & Varadhan, 2007; Itô, 1974). As I described in the introduction, this can be characterized in terms of drift direction ρ , the drift magnitude $|\mu|$, and diffusion parameter σ^2 .

The evidence accumulation process halts when the state crosses a threshold θ , which indicates how much total information a person wishes to have before making their selection. It specifically adjusts the radius of the circle shown in Figure 1: larger θ will mean that they wait for more information before selecting, and a smaller θ will mean they wait for less. When a person's state crosses the edge of this circle, they make a response corresponding to the location at which they crossed the circle. For example, crossing at 90 degrees on the circle would lead to a "vertical" response.

This gives a distribution of responses for the model, but does not yet predict the amount of time it takes to complete the response. The amount of time it takes the random walk (drift + diffusion process) to reach the threshold determines decision component of response time. The remaining time components of the response process that are not decision-related, such as perceiving the stimulus and executing the motor response corresponding to a selection, are quantified in the non-decision time ndt.

Put together, this yields 6 main components of the model: initial state, drift direction, drift magnitude, diffusion, threshold, and non-decision time. One of these must be fixed in order to set the scale of the model; otherwise, one could inflate threshold, drift magnitude, and diffusion and obtain the same accuracy and response time predictions. To ensure that all parameters are identifiable, the diffusion coefficient is set to 1 across all models reported in this paper. The starting point can also be set to [0,0], as participants have no information before the stimulus appears in the uncued conditions.

The joint likelihood for the responses *Resp* and response times *RT* are given as a function of the remaining parameters, $Pr(Resp = x, RT = t | \mu, \theta, \rho, ndt)$.

$$Pr(Resp = x, RT = t | \mu, \theta, \rho, ndt) =$$

$$exp\left(\theta | \mu | (cos(x)cos(\rho) + sin(x)sin(\rho)) - \frac{|\mu|^2(t - ndt)}{2}\right)$$

$$\cdot \theta^2 \sum_{k=1}^n \frac{j_{0,k}}{J_1(j_{0,k})} exp\left(\frac{-j_{0,k}^2(t - ndt)}{2\theta^2}\right)$$

The function $J_1()$ is a first-order Bessel function of the first kind, and the elements $j_{0,k}$ are the zeros of the zeroorder Bessel function of the first kind. The series $\sum_{k=1}^{n}$ computes and evaluates the function for the first *n* zeros. The true likelihood is given by an infinite series, $n = \infty$. However, for the practical purposes of this study, I found that n = 50 to 150 yields sufficiently precise approximations. This value can be decreased for computational efficiency or increased for precision; higher values of *n* are particularly important to use for accurately computing likelihoods at the leading edge of the RT distribution.

This model was implemented in MATLAB and estimated using an Markov chain Monte Carlo sampler with a standard Metropolis-Hastings algorithm (van Ravenzwaaij et al., 2016). This is one of the first applications of this model to real data, so a custom-built sampler was used in order to allow for better troubleshooting and debugging in MATLAB. The model likelihoods and simulation code can also be implemented in R using the CircularDDM package available at cran.r-project.org/web/packages/CircularDDM (Lin et al., 2017). This code can also be put together with other types of samplers for greater sampling efficiency and convergence between chains.

Uncued conditions model

The first cognitive model examines only uncued conditions in order to establish relationships between the speedaccuracy and difficulty manipulations and the drift magnitude, drift direction, threshold, and non-decision time parameters. Each of these parameters was permitted to vary freely across the six speed / accuracy and difficulty level conditions, and they were estimated separately for each participant and condition. The priors were set as drift direction $\rho \sim Uniform(0, 360)$ (in degrees), drift magnitude $\mu \sim$ Gamma(2.0, 1.5), threshold $\theta \sim Gamma(2.0, 1.5)$, and nondecision time $ndt \sim Gamma(2.0, 0.2)$ (Gamma distributions used the shape-scale parameterization). These priors were designed to constrain the parameters to reasonable ranges while not heavily influencing the posterior distributions. Additionally, the priors were the same for each condition and participant, so any differences between posterior distributions can be attributed purely to differences in the data rather than to priors.

Each participant and condition was run for 5 chains of length 10,000 samples. A few failures to converge occurred on initial runs, primarily due to errors in estimating likelihoods at the leading edge of the distribution that led to extremely high likelihoods for extremely unreasonable parameter values. Chains that failed to converge were re-run with a higher value for n in the likelihood computation (yielding greater precision in approximating the infinite series), which successfully addressed this issue. The final estimates of these parameters across participants are shown in Figure 6 and provided online at the Open Science Framework, osf.io/q3ytj/.

Drift direction estimates (top left of Figure 6) were consistently centered at zero across all uncued conditions, confirming no clockwise or counterclockwise bias to participants' responses. Drift magnitude estimates (bottom left of Figure 6) decreased with stimulus coherence, indicating that stimuli with higher variance in orientation were indeed providing weaker evidence. Drift magnitude estimates were similar between speed and accuracy conditions as well, suggesting that participants were gathering information at roughly the same rate and quality regardless of time pressure.

Threshold estimates (top right of Figure 6) were consistently higher in the accuracy relative to the speed condition, suggesting that participants were applying stricter criteria for their decisions to the evidence in the accuracy condition relative to the speed condition. This shift explains the longer response times and higher accuracy in the accuracy condition as shown in Figure 3. There is a small suggestion that threshold estimates decreased with increasing difficulty, though this effect is quite small in comparison to that of the speed-accuracy manipulation. Similar results were found by Kvam & Pleskac (2016) – the authors suggested that such an effect may be the result of increasing thresholds when high-quality information is available, suggesting that



Figure 6. Mean estimates and 95% highest density intervals for for drift direction (top left), drift magnitude (bottom left), threshold (top right), and non-decision time (bottom right) for each participant (each participant is a different color) and uncued condition. The mean parameter values across participants for each condition are shown in black.

participants are adjusting their choice criteria on-line, or it could be the result of decision boundaries that collapse over time. The latter possibility would indicate that participants are trading off accuracy for the amount of time it takes to finish a trial (Bowman et al., 2012; Drugowitsch et al., 2012; Ratcliff & Frank, 2012; Palestro et al., 2018). Fortunately, fixed-boundary versus collapsing-boundary models in binary choice predict highly similar distributions of response times (Voskuilen et al., 2016), suggesting that collapsing boundaries are unlikely to be a necessary component of the model at present. A simpler explanation is that this is simply a side effect of fixing the diffusion rate across conditions, as the noise sampled during accumulation is actually likely to increase with greater difficulty, causing the *relative* threshold estimates to appear to decrease as stimulus signal gives way to noise. Functionally, each of the parameters estimates is relative to the set diffusion rate, so they are not each uniquely identified. In this sense, a diffusion rate shift is mathematically equivalent to particular shifts in drift and threshold, so one will naturally be expressed in terms of the other when a parameter is fixed to set the scale of the model.

Finally, non-decision time (bottom right of Figure 6), like drift direction, appeared not to change substantially between speed and accuracy conditions or across difficulty levels. With the exception of a single participant (red participant), within-subject shifts in non-decision time were typically within approximately 20-30 milliseconds.

Overall, the model was able to capture the distributions of responses and response times well across conditions. The predictions generated from the mean posterior parameter estimates for each participant are shown in Figures 7 and 8. Most notably in responses, the model tracked the decreasing precision of responses with greater stimulus difficulty (left / middle / right panels of each subplot in Figure 7) and the slightly lower precision in the speed relative to the accuracy condition (pink and blue plots in Figure 7, respectively). It also tracked the slowing response times with increasing stimulus difficulty (left / middle / right panels of each subplot) and the slower response times in the accuracy (blue) relative to the speed (pink) condition as shown in the response time plots in Figure 8.

Cued conditions model. As I covered in the results, the cued data frequently exhibit a bimodal distribution of responses, where one mode is centered on the mean stimulus orientation (as in the uncued trials) and the second mode is centered on the location of the cue (Figure 5). This suggests that on many trials, the cue information was not integrated together with conflicting stimulus information, but rather it was used instead of stimulus information. The resulted in either responses consistent with the stimulus *or* responses consistent with the cue, rather than responses that were a compromise between stimulus and cue.

There are a few possibilities here that can be ruled out on the basis of the bimodal responses alone. For example, starting point bias and drift bias models predict a single mode to the response distributions. This occurs because drift and starting point bias models assume that participants integrate cue and stimulus information into a unified evidence representation, which winds up producing responses that are a weighted average of the stimulus and cue orientations. This



Figure 7. Observed (histograms) and predicted (lines, generated from the mean posterior parameter values) distributions of responses for each participant and condition. Each subplot contains the data for a separate participant. Within each subplot is shown the distributions for different levels of stimulus difficulty, given as the standard deviation of the distribution of stimulus orientations (15/30/45). The speed emphasis condition is shown in pink, and the accuracy emphasis condition is shown in blue.

would result in a single peak in the response distribution, located partway between the stimulus mean and the cue direction.

A two-stage model, where the cue drives an initial evidence accumulation stage and the stimulus drives a subsequent one (Diederich & Busemeyer, 2006), typically makes similar predictions to the starting point model and can be ruled out for similar reasons. If the evidence accumulated from the cue-processing stage does not trigger a response, it can instead be represented as a distribution of starting points for a subsequent stimulus-driven process. This would again result in a unimodal distribution of responses. However, in such a model it would also technically be possible to cross the threshold during the cue-processing stage, before reaching the stimulus-processing stage. However, this would result in cue-consistent responses concentrated near the nondecision time (as there would be no additional decision time required to make a selection), which does not pan out in the empirical RT distributions. Instead, it seems that cueconsistent responses and stimulus-consistent responses have similar response time distributions, with the cue-consistent responses often being slower than the stimulus-consistent ones – something which is directly in conflict with a twostage account.

Another possible explanation is that the cue and stimulus compete with one another to generate responses. According to this hypothesis, the cue would drive its own evidence accumulation process, where support for a cueconsistent response grows over time in parallel with support for a stimulus-consistent response. The process that hits a response boundary first "wins" the race and generates a response consistent with its representation.

An explanation based on competition between cue and stimulus provides a clear justification for the appearance of bimodal distributions of responses. But there are still outstanding phenomena to explain in the cued conditions. In addition to the bimodal distribution of responses, it is important to return take note of the pattern of results in Figure 4: responses are both faster and more accurate when the cue and stimulus agree with one another, and slower when they are in conflict. This suggests some interaction between the cue and stimulus evidence accumulation.

Accounting for the results shown in Figures 4 and 5 is far from straightforward. It suggests a number of additional



Figure 8. Observed (histograms) and predicted (lines, generated from the mean posterior parameter values) distributions of responses for each participant and condition. Each subplot contains the data for a separate participant. Within each subplot is shown the distributions for different levels of stimulus difficulty, given as the standard deviation of the distribution of stimulus orientations (15/30/45). The speed emphasis condition is shown in pink, and the accuracy emphasis condition is shown in blue.

assumptions must be made to account for behavior on cued trials. The first assumption has already been covered: I propose that the cue drives a separate accumulation process that can trigger a response consistent with the cue's orientation. The second assumption is one regarding attention: I propose that the cue draws attention to orientations similar to the one shown on the screen, making a person more likely to sample orientations consistent with the cue when gathering information from the stimulus.

The practical implementation of the second assumption is to scale the drift magnitude by the position of cue relative to the stimulus. When the cue and stimulus are close to one another, drift magnitude for the stimulus should be higher. And when they are further apart, drift magnitude should be lower because attention derived from the cue does not facilitate the perception of orientations coming from the stimulus. To account for this effect, I simply add an excitation / inhibition factor to the drift magnitude of the stimulus-driven accumulation process. The size of this effect is controlled by a free parameter γ , which determines how much is added or subtracted from the value of μ as a function of the difference between stimulus and cue orientations, $\rho_s - \rho_c$. Formally, the drift magnitude for the stimulus μ_s is re-specified as

$$\mu'_{s} = \mu_{s} + \gamma \cdot \cos(\rho_{s} - \rho_{c}) \tag{1}$$

Note that the cosine function will result in excitatory interactions (positive values) between stimulus and cue when they are within 45 degrees of one another, and inhibitory interactions (negative values) when they are 45 to 90 degrees removed from one another. The scalar γ simply determines the magnitude of the effect. This simple transformation agrees well with the empirical result shown in Figure 4, where the crossover from improving to diminishing accuracy and RT happens between cue orientations of 20 and 50 degrees.

In the model below, the difference in drift directions for the cue and stimulus are simply set to the orientation of the cue and true mean orientation of the stimulus, respectively. As shown, this allows the model to account for the shift in response times across cue orientations with no trouble.

Model predictions. With these two assumptions, it is possible to account for the important effects of the cue. Formally, the stimulus-driven process is identical to the one presented in the previous section except for the attention effect

added to the drift magnitude (μ_s replaced by μ'_s from Equation 1). It is defined in terms of a drift direction ρ_s , drift magnitude μ'_s , noise (set to 1), threshold θ_s , and non-decision time ndt_s . The cue-driven process is characterized in a similar way, by setting a drift direction ρ_c , drift magnitude μ_c , diffusion rate σ_c^2 , threshold θ_c , and non-decision time component ndt_c .

As I show below, it is only necessary to allow 2 of the cuedriven process parameters to vary across conditions. One has to be fixed for scale (threshold was set because it seemed the least interpretable aspect of the cue-driven process), drift direction for the cue-driven process can simply be specified as the orientation of the cue, and non-decision time can be fixed at the same value as the uncued conditions. This leaves only drift magnitude for the cue μ_c , diffusion rate for the cue σ_c^2 , and the excitation/inhibition factor γ as free parameters in the model.

Because the two processes are racing, the likelihoods of finishing at a particular time are conditioned on the likelihood of the opposing process finishing after that time. This gives the likelihood for the stimulus-driven responses:

$$Pr(Resp = x, RT = t | \text{stimulus}) =$$

$$Pr(Resp = x, RT = t | \mu_s, \theta_s, \rho_s, ndt_s)$$

$$\cdot \left(1 - \int_0^t Pr(Resp = x, RT = t | \mu_c, \theta_c, \rho_c, ndt_c)\right)$$
(2)

Here, $Pr(Resp = x, RT = t | \mu_s, \theta_s, \rho_s, ndt_s)$ is given from Equation. Conversely, the cue-driven responses will be conditioned on the stimulus-drive response not finishing by time *t*.

$$Pr(Resp = x, RT = t | cue) =$$

$$Pr(Resp = x, RT = t | \mu_c, \theta_c, \rho_c, ndt_c)$$

$$\cdot \left(1 - \int_0^t Pr(Resp = x, RT = t | \mu_s, \theta_s, \rho_s, ndt_s)\right)$$
(3)

The overall likelihood of responses is then simply given by adding the two conditional distributions from Equations 2 and 3 together.

$$Pr(Resp = x, RT = t | \mu_s, \mu_c, \rho_s, \rho_c, \theta_s, \theta_c, ndt_s, ndt_c) =$$

$$Pr(Resp = x, RT = t | \text{stimulus})$$

$$+ Pr(Resp = x, RT = t | \text{cue})$$
(4)

Unfortunately, computing the complete joint probability density function requires the use of numerical integration to evaluate the cumulative distributions in Equations 2 and 3. Alternatively, simulated trials can be generated from the model and used to generate an approximate likelihood function by passing a kernel density estimator over the simulated data (Holmes, 2015; Turner & Sederberg, 2012). Either approach will be more computationally strenuous than the analytic likelihoods for the uncued model.

For the present purposes, I approximated the conditional distributions from Equation 4 by calculating the cumulative response time distributions for the cue- and stimulus-driven processes separately at 10-ms intervals, then conditioned the cue RT distribution on the cumulative stimulus RT distribution and conditioned the stimulus RT distribution on the cumulative cue RT distribution. In general, this conditioning in the model should result in a shorter tail to the response time distributions in cued conditions, which did seem to appear in the empirical RT distributions (compare Figure 8 against Figure 9). The diffusion coefficient for the cuebased process seemed more informative than the threshold, as it directly gives information about the variability of cuebased responses and the rate at which the cue-driven process moves, so the threshold for the cue-driven process was fixed at $\theta_c = 2.0$ and the drift magnitude μ_c and diffusion rate σ_c^2 for the cue-driven process were estimated instead. For those readers interested in looking at threshold rather than diffusion rate estimates, these can be computed by relating μ , σ^2 , and θ from Equation (or Equations 22 and 23 in Smith, 2016). The model did not seem to require a shift in non-decision time for the cue-driven process, so cue nondecision time was fixed to the value of stimulus non-decision time, $ndt_c = ndt_s$. The stimulus-driven process was restricted to the parameter values estimated from the non-cued trials. Finally, data was collapsed between symmetric cue orientations (e.g. +70 and -70 degrees) to ensure there was sufficient data in each condition. In the end, the cued condition model was quite simple but accounted well for the data.

The model predictions for the cued conditions are shown in Figure 9 overlaid on the response and response time data. Critically, the inclusion of the parameter γ allowed the model to account for the mean shifts in RT across orientations of the cue (bottom panels). The model suggests that as the cue moves further away from the stimulus, attention is divided between orientations, resulting in a lower drift magnitude for the stimulus and consequently a longer distribution of response times. As a result, it predicts well the pattern of mean response times shown in Figure 4 (left panel). This effect apparently cannot be captured without some sort of interaction between the cue and stimulus: a race by itself without the effect from Equation 1 will predict identical distributions of response times, as will a simple probability mixture model of cue- and stimulus-driven responses.

Of course, the model also captures the bimodal distributions of responses in the cued conditions (top panels of Figure 9). As a consequence, it will have no trouble capturing the shifts in accuracy based on cue orientations shown in the left panel of Figure 4. The model allows accuracy to improve when the cue and stimulus agree; this happens both be-



Figure 9. Histograms of data (bars) and model prediction (lines) for the cue-inclusive conditions, including response distributions (top) and response time distributions (bottom) for 0, 20, 50, and 70 degree cue deviations (left to right).

cause the cue will sometimes yield the exact correct answer (naturally improving accuracy) and because of the excitatory effects of γ yielding higher drift magnitudes when stimulus and cue agree. Although the model tends to very slightly over-predict the variability of responses in the 0 degree cue condition, this effect is still quite clearly captured in general.

Of course, the results shown in Figure 9 might not hold for every individual participant. As one reviewer rightly suggested, it could be the case that some participants use the cue and others do not, resulting in a mixture of strategies across people. This would create different distributions of responses at the individual level that appeared as bimodal on the aggregate. As shown in Figure 10, this seems not to be the case. Though they do so to varying degrees, essentially every participant shows some amount of bimodal responding, which is most evident in the 70-degree conditions.

Fits to each individual's data, which also shows these response distributions in finer-grained detail, can be found in the online materials on the Open Science Framework at osf.io/q3ytj. Table 3 shows the mean posterior parameter values for each participant, giving some insight into how much of an influence the cue had on different people's responses. For example, participant 5 had a large drift magnitude for the cue-driven response process (high μ_c), but very little benefit from the cue on the stimulus-driven process (low γ). This may indicate a participant who was inclined to trust the cue completely rather than using it to inform their beliefs about the stimulus. Conversely, participant 1 had a relative low drift for the cue-driven process but a high value for γ , indicating that they allowed the cue to redirect their attention toward corresponding orientations but not to override the information they sampled from the stimulus.

Put together, the two assumptions of competing accumulators and excitatory / inhibitory relations between cue and

Table 3

Mean estimates of the drift magnitude of the cue (μ_c) , diffusion coefficient of the cue (σ_c) , and excitation / inhibition due to the relative orientation of the cue and stimulus (γ) for each participant. These were used in combination with the estimates from the original uncued condition for the stimulus-driven process. These estimates and their HDIs are provided in online material on the Open Science Framework (osf.io/q3ytj).

| Participant | μ_c | σ_c | γ |
|-------------|---------|------------|------|
| 1 | 1.20 | 1.06 | 1.31 |
| 2 | 1.51 | 0.15 | 1.32 |
| 3 | 1.30 | 0.67 | 0.81 |
| 4 | 1.29 | 0.36 | 0.61 |
| 5 | 2.91 | 0.28 | 0.00 |
| 6 | 1.64 | 0.35 | 0.74 |
| 7 | 2.57 | 1.06 | 0.15 |
| 8 | 2.47 | 0.36 | 0.35 |
| 9 | 0.89 | 0.23 | 0.62 |
| 10 | 1.54 | 0.59 | 0.73 |
| 11 | 1.31 | 0.97 | 0.78 |
| 12 | 1.19 | 1.13 | 0.84 |
| | | | |

stimulus seem to account for the major findings associated with the cued conditions. The two primary empirical challenges come from Figures 4 and 5: responses in cued conditions are bimodal, and greater discrepancy between stimulus and cue results in slower response times. These two findings put together pose a major stumbling block for classical theories of bias (which predict unimodal response distributions) as well as theories based on probability mixtures or fully independent cue and stimulus processes (which predict uniform response times across cue orientations). By includ-

PETER D. KVAM



Figure 10. Distribution of responses in cued conditions for each participant, including cue orientations of 0, 20, 50, and 70 degrees relative to the stimulus (left to right).

ing both cue-driven and stimulus-driven responses that interacted based on the discrepancy between their orientations, the model developed here was able to cover both phenomena. As a result, the proposed model provides the most complete theory of bias in tasks where responses fall along a scale or continuum.

A note on parameter variability

Diffusion and accumulator models of binary choice typically include trial-to-trial variability in the parameters of the model, particularly for drift rates and starting points (Brown & Heathcote, 2008; Ratcliff, 1978b; Ratcliff & Smith, 2004). This is usually done because an absence of variability in these parameters will result the same distribution of response times for error and correct responses. Drift rate variability, which in our case would be reflected by variability in drift magnitude, is usually introduced to account for slow errors. Conversely, starting point variability is frequently introduced to account for fast errors.

Because responses can fall along a continuum in this task, the analogous phenomenon would be covariance between error magnitude (response deviation) and response time. Drift variability would create a positive correlation between error magnitude and response time, while starting point variability could create a negative correlation between them. This occurred in the cued conditions in our task, where it can be explained as a natural consequence of multiple response processes arriving at different times. However, the model of the uncued condition – with no variability in parameters or competing processes – should generate response times that have the same distribution regardless of the orientation response a person gives.

Of course, this can be violated across conditions such as the speed-accuracy manipulation, as drift rates and thresholds shift systematically across them. However, within an individual uncued condition there should be no substantial relationship between response time and response deviation. This prediction was examined using a hierarchical Bayesian linear model, with each combination of participant and condition having its own correlation between RT and response deviation. Each of the condition-level correlations was constrained by a group-level distribution of RT-response relationships to obtain an estimate for the relationship between RT and response for each condition. Looking at these estimates for each condition garners an overall sense of how response deviation and response time are related across the uncued trials.

A plot of the relationship between the magnitude of response deviations (x) and response times (y) is shown in Figure 11. While occasional participants and conditions seemed to show a positive or negative relationship, this was not reflected in the overall results. The group-level mean estimate for the standardized correlation between response time and response deviation, across all conditions and participants, was centered at 0.01, and the 95% highest density interval

MODELING ORIENTATION JUDGMENTS



Figure 11. Scatter plot and linear estimate of the relationship between response time and response deviation for each participant (shown in separate colors).

suggested that the correlation was between -.03 and .06. Therefore, it seems that the relationship between RT and response deviation is small enough (if it exists) that the model does not need to be modified to account for the result.

The lack of relationship indicates that drift or start point variability is not required to account for qualitative patterns in the results. This is particularly relevant to the model of Ratcliff (in press), which posits trial to trial variability in drift in a continuous process. Both the present paper and the experiments in Ratcliff's work seem to find no strong connection between error magnitude and response times (within conditions), which for the time being suggests that a simpler model without parameter variability is sufficient. Of course, this may not hold for all continuous response tasks. As more work is done on the topic, it should become evident which tasks require parameter variability and which ones do not.

Discussion

The results provide strong support for the presence of continuous analogues of the three binary choice effects: a speedaccuracy trade-off, a stimulus difficulty effect, and a biasing effect from predecision cues. As a result, a model with similar structure to diffusion or accumulator models should be well-suited to describing behavior in continuous-response tasks. This is reflected in how the parameters of the model change with manipulations. As in binary choice and the continuous work of Ratcliff (in press), speed-accuracy manipulations yield shifts in thresholds and difficulty manipulations yield changes in drift.

However, the effects associated with predecision cues seem to be more complex than previous theories of bias would suggest. Not only did cues improve or hinder response accuracy and response time based on the match between cue and stimulus, but cues that conflicted with the stimulus generated bimodal response distributions. Accounting for these results required additional assumptions: namely, that the cue and stimulus drive separate accumulation processes and that the cue brings attention to those orientations consistent its indicated direction.

The exact nature of the attentional effect of cue orientation, reflected in drift magnitude, is likely to be connected to the tuning functions of orientation columns in the brain. For example, vertical columns will be most highly activated by vertical orientations of a cue, but they may also be activated more weakly by orientations of 0-20 degrees away from vertical. Activity of the true stimulus direction may be facilitated by off-mean orientations simply because the tuning function of orientation columns overlaps with nearby cue orientations (Hubel & Wiesel, 1974; Yacoub et al., 2008). Conversely, they may be inhibited by orientations close to horizontal, whether this is due to direct inhibition between orientations or simply a reallocation of attentional resources resulting from the cue.

The interaction between competing cue-driven and stimulus-driven processes is likely to have neural correlates as well. Strong evidence from conflict processing suggests that anterior cingulate cortex should be active when monitoring conflicts between multiple information streams (Gehring et al., 1993). Computational models of this process, using tasks such as the Flanker task (Eriksen & Schultz, 1979), have suggested that similar interactive processes may be responsible for the delayed response times when different sources of information indicate conflicting responses (Botvinick et al., 2001; Lu & Proctor, 1995), even when the cue is irrelevant (Simon & Rudell, 1967). It is entirely possible that the same conflict mediation processes take place when the cue and stimulus disagree in experiments like the one here, and it seems likely that ACC activation might correspond to the output of Equation 1 and the values of γ estimated in the model. Certainly these possibilities pose interesting directions for future neuroimaging work using the task

and/or model developed here.

Finally, one caveat to the model presented here is that trial to trial variability is not introduced for any of the parameters. In binary choice models, drift rate or starting point variability are important to include because they generate the asymmetric distributions of correct and error response times observed in empirical data. However, such variability appears not to be required for the present results, where the magnitude of errors in uncued conditions was unrelated to mean response times. This allows us to use a simpler model that does not include trial to trial variability for the time being. Of course, the cued conditions drive a pair of accumulators that can result in response time asymmetries for responses generated by the cue and those generated by the stimulus. This seems to suggest that sufficiently different or variable stimuli might drive different response processes. How much variability between stimuli is required to result in variability in drift parameters remains an open question.

Put together, the results presented here underscore the benefits of using continuous response tasks and models, which provide rich information for making inferences about cognitive processes. This study not only established benchmark results for responses and response times in continuous report, but showed that using a continuous response task could grant insights that are not permitted by binary response paradigms. Most notably, bimodal response distributions reveal that competing cue-drive and stimulus-driven processes seem to drive bias in responses. This allowed us to extend our theoretical understanding of bias in addition to providing steps toward understanding how decision makers cope with continuous selection tasks.

References

- Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (2006). The physics of optimal decision making: A formal analysis of models of performance in twoalternative forced-choice tasks. *Psychological Review*, *113*(4), 700-765. doi: 10.1037/0033-295X.113.4.700
- Bogacz, R., Wagenmakers, E.-J., Forstmann, B. U., & Nieuwenhuis, S. (2010). The neural basis of the speed-accuracy tradeoff. *Trends in Neurosciences*, *33*(1), 10–16. doi: 10.1016/j.tins.2009.09.002
- Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological review*, 108(3), 624.
- Bowman, N. E., Kording, K. P., & Gottfried, J. A. (2012). Temporal integration of olfactory perceptual evidence in human orbitofrontal cortex. *Neuron*, 75(5), 916–927. doi: 10.1016/j.neuron.2012.06.035
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, *10*, 433–436.

- Brown, S. D., & Heathcote, A. (2008). The simplest complete model of choice response time: Linear ballistic accumulation. *Cognitive Psychology*, 57(3), 153–178. doi: 10.1016/j.cogpsych.2007.12.002
- Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, 100(3), 432–459. doi: 10.1037/0033-295X.100.3.432
- Diederich, A., & Busemeyer, J. R. (2006). Modeling the effects of payoff on response bias in a perceptual discrimination task: Bound-change, drift-rate-change, or twostage-processing hypothesis. *Perception & Psychophysics*, 68(2), 194–207. doi: 10.3758/BF03193669
- Drugowitsch, J., Moreno-Bote, R., Churchland, A. K., Shadlen, M. N., & Pouget, A. (2012). The cost of accumulating evidence in perceptual decision making. *The Journal of Neuroscience*, 32(11), 3612–3628. doi: 10.1523/JNEUROSCI.4010-11.2012
- Edwards, W. (1965). Optimal strategies for seeking information: Models for statistics, choice reaction times, and human information processing. *Journal of Mathematical Psychology*, 2(2), 312–329. doi: 10.1016/0022-2496(65)90007-6
- Eriksen, C. W., & Schultz, D. W. (1979). Information processing in visual search: A continuous flow conception and experimental results. *Perception & Psychophysics*, 25(4), 249–263.
- Fougnie, D., Suchow, J. W., & Alvarez, G. A. (2012). Variability in the quality of visual working memory. *Nature communications*, *3*, 1229.
- Gehring, W. J., Goss, B., Coles, M. G., Meyer, D. E., & Donchin, E. (1993). A neural system for error detection and compensation. *Psychological science*, 4(6), 385–390.
- Gilovich, T., Griffin, D., & Kahneman, D. (2002). *Heuristics* and biases: The psychology of intuitive judgment. Cambridge University Press.
- Heitz, R. P., & Schall, J. D. (2012). Neural mechanisms of speed-accuracy tradeoff. *Neuron*, 76(3), 616–628. doi: 10.1016/j.neuron.2012.08.030
- 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.
- Hubel, D. H., & Wiesel, T. N. (1974). Sequence regularity and geometry of orientation columns in the monkey striate cortex. *Journal of Comparative Neurology*, 158(3), 267– 293.

Itô, K. (1974). Diffusion processes. Wiley Online Library.

- Kleiner, M., Brainard, D., Pelli, D., Ingling, A., Murray, R., & Broussard, C. (2007). What's new in psychtoolbox-3. Perception, 36(14), 1.
- Krajbich, I., Armel, C., & Rangel, A. (2010). Visual fixations and the computation and comparison of value in simple choice. *Nature neuroscience*, 13(10), 1292–1298.
- Krajbich, I., Lu, D., Camerer, C., & Rangel, A. (2012). The attentional drift-diffusion model extends to simple purchasing decisions. Frontiers in Psychology, 3. doi: 10.3389/fpsyg.2012.00193
- Kruschke, J. (2014). Doing bayesian data analysis: A tutorial with R, JAGS, and STAN. Academic Press.
- Kvam, P. D. (n.d.). Materials for "modeling accuracy, response time, and bias in continuous orientation judgments". Retrieved from osf.io/q3ytj.
- Kvam, P. D., & Pleskac, T. J. (2016). Strength and weight: The determinants of choice and confidence. Cognition. 152, 170-380. doi: 10.1016/j.cognition.2016.04.008
- Lin, Y.-S., Kvam, P. D., & Heathcote, A. (2017). Circularddm: Circular drift-diffusion model. Retrieved from https://cran.r-project.org/web/packages/CircularDDM/ Human Perception and Performance, 40(2), 870.
- Link, S., & Heath, R. (1975). A sequential theory of psychological discrimination. Psychometrika, 40(1), 77-105. doi: 10.1007/BF02291481
- Lu, C.-H., & Proctor, R. W. (1995). The influence of irrelevant location information on performance: A review of the simon and spatial stroop effects. Psychonomic bulletin & review, 2(2), 174–207.
- Mulder, M. J., Wagenmakers, E.-J., Ratcliff, R., Boekel, W., & Forstmann, B. U. (2012). Bias in the brain: a diffusion model analysis of prior probability and potential payoff. Journal of Neuroscience, 32(7), 2335–2343.
- Nosofsky, R. M., & Gold, J. M. (2017). Biased guessing in a complete-identification visual-working-memory task: Further evidence for mixed-state models. Journal of Experimental Psychology: Human Perception and Performance.
- Nosofsky, R. M., & Palmeri, T. J. (1997). An exemplar-based random walk model of speeded classification. Psychological Review, 104(2), 266-300. doi: http://dx.doi.org/10.1037/0033-295X.104.2.266

- Palestro, J. J., Weichart, E., Sederberg, P. B., & Turner, B. M. (2018). Some task demands induce collapsing bounds: Evidence from a behavioral analysis. Psychonomic bulletin & review, 1–24.
- Palmer, J., Huk, A. C., & Shadlen, M. N. (2005). The effect of stimulus strength on the speed and accuracy of a perceptual decision. Journal of Vision, 5(5), 1. doi: 10.1167/5.5.1
- Pleskac, T. J., & Busemeyer, J. R. (2010). Two-stage dynamic signal detection: a theory of choice, decision time, and confidence. Psychological Review, 117(3), 864. doi: 10.1037/A0019737
- Ratcliff, R. (1978a). A theory of memory retrieval. Psychological Review, 85(2), 59-108. doi: 10.1037/0033-295X.85.2.59
- Ratcliff, R. (1978b). A theory of memory retrieval. Psychological Review, 85(2), 59-108.
- Ratcliff, R. (1985). Theoretical interpretations of the speed and accuracy of positive and negative responses. Psychological Review, 92(2), 212-225. doi: 10.1037/0033-295X.92.2.212
- Ratcliff, R. (2014). Measuring psychometric functions with the diffusion model. Journal of Experimental Psychology:
- Ratcliff, R. (in press). Decision making on spatially continuous scales. Psychological Review.
- Ratcliff, R., & Frank, M. J. (2012). Reinforcementbased decision making in corticostriatal circuits: mutual constraints by neurocomputational and diffusion models. Neural Computation, 24(5), 1186–1229. doi: 10.1162/NECOa00270
- Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision Neural Computation, 20(4), 873–922. tasks. doi: 10.1162/neco.2008.12-06-420
- Ratcliff, R., & Smith, P. L. (2004). A comparison of sequential sampling models for two-choice reaction time. Psychological Review, 111(2), 333-367. doi: 10.1037/0033-295X.111.2.333
- Ratcliff, R., Van Zandt, T., & McKoon, G. (1999). Connectionist and diffusion models of reaction time. Psychological Review, 106, 261-300.
- Roe, R. M., Busemeyer, J. R., & Townsend, J. T. (2001). Multialternative decision field theory: A dynamic connectionst model of decision making. Psychological Review, 108(2), 370-392. doi: 10.1037/0033-295X.108.2.370

- Simon, J. R., & Rudell, A. P. (1967). Auditory sr compatibility: the effect of an irrelevant cue on information processing. *Journal of applied psychology*, 51(3), 300.
- Smith, P. L. (2016). Diffusion theory of decision making in continuous report. *Psychological Review*, 123(4), 425– 451. doi: 10.1037/rev0000023
- Stone, M. (1960). Models for choice-reaction time. *Psy-chometrika*, 25(3), 251–260.
- Stroock, D. W., & Varadhan, S. S. (2007). Multidimensional diffusion processes. Springer.
- Townsend, J. T. (2008). Mathematical psychology: Prospects for the 21st century: A guest editorial. *Journal of mathematical psychology*, *52*(5), 269–280. doi: 10.1016/j.jmp.2008.05.001
- Trueblood, J. S., Brown, S. D., & Heathcote, A. (2014). The multiattribute linear ballistic accumulator model of context effects in multialternative choice. *Psychological Review*, 121(2), 179–205. doi: 10.1037/a0036137
- Turner, B. M., & Sederberg, P. B. (2012). Approximate bayesian computation with differential evolution. *Journal* of Mathematical Psychology, 56(5), 375–385.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157), 1124–1131. doi: 10.1126/science.185.4157.1124
- Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: the leaky, competing accumulator model. *Psychological Review*, *108*(3), 550–592. doi: 10.1037/0033-295X.108.3.550
- Van den Berg, R., Awh, E., & Ma, W. J. (2014). Factorial comparison of working memory models. *Psychological review*, 121(1), 124.
- Van Den Berg, R., Shin, H., Chou, W.-C., George, R., & Ma, W. J. (2012). Variability in encoding precision accounts

for visual short-term memory limitations. *Proceedings of the National Academy of Sciences*, *109*(22), 8780–8785. doi: 10.1073/pnas.1117465109

- van Ravenzwaaij, D., Cassey, P., & Brown, S. D. (2016). A simple introduction to markov chain monte–carlo sampling. *Psychonomic bulletin & review*, 1–12.
- Vickers, D., & Packer, J. (1982). Effects of alternating set for speed or accuracy on response time, accuracy and confidence in a unidimensional discrimination task. *Acta Psychologica*, 50(2), 179–197. doi: 10.1016/0001-6918(82)90006-3
- Voskuilen, C., Ratcliff, R., & Smith, P. L. (2016). Comparing fixed and collapsing boundary versions of the diffusion model. *Journal of Mathematical Psychology*, 73, 59–79. doi: 10.1016/j.jmp.2016.04.008
- Wagenmakers, E.-J., Ratcliff, R., Gomez, P., & McKoon, G. (2008). A diffusion model account of criterion shifts in the lexical decision task. *Journal of memory and language*, 58(1), 140–159.
- Wickelgren, W. A. (1977). Speed-accuracy tradeoff and information processing dynamics. *Acta Psychologica*, 41(1), 67–85. doi: 10.1016/0001-6918(77)90012-9
- Wolfe, J. M., Horowitz, T. S., Van Wert, M. J., Kenner, N. M., Place, S. S., & Kibbi, N. (2007). Low target prevalence is a stubborn source of errors in visual search tasks. *Journal of Experimental Psychology: General*, 136(4), 623–638. doi: 10.1037/0096-3445.136.4.623
- Yacoub, E., Harel, N., & Uğurbil, K. (2008). High-field fmri unveils orientation columns in humans. *Proceedings of the National Academy of Sciences*.
- Zhang, W., & Luck, S. J. (2008). Discrete fixedresolution representations in visual working memory. *Nature*, 453(7192), 233–235.