Some Tasks' Demands Require Collapsing Bounds: Evidence from a Behavioral Analysis

A Thesis

Presented in Partial Fulfillment of the Requirements for the Degree Master of Arts in the Graduate School of The Ohio State University

By

James Palestro, B.A.

Graduate Program in Psychology

The Ohio State University

2018

Master's Examination Committee:

Brandon Turner, Advisor

Andrew Leber

Per Sederberg

Abstract

Traditional models of choice response time assume that sensory evidence accumulates for choice alternatives until a threshold amount of evidence has been obtained. Although some researchers have characterized the threshold as varying randomly from trial to trial, these investigations have all assumed that the threshold remains fixed across time within a trial. Despite decades of successful applications of these models to a variety of experimental manipulations, this time-invariance assumption has recently been called into question, and a time-variant alternative implementing collapsing decision thresholds has been proposed instead. Here, we investigated the fidelity of the collapsing threshold assumption by assessing relative model fit to data from a highly constrained experimental design that coupled a within-subject mixture of two classic response time paradigms - interrogation and free response - within a random dot motion (RDM) task. Overall, we identified strong evidence in favor of the time-variant model with collapsing decision boundaries, suggesting that subjects may adapt a dynamic decision threshold due to task characteristics, specifically to account for the mixing of response time paradigms and motion strengths across trials in the mixed RDM task. We conclude that time-variant mechanisms may serve as a viable strategy for human

subjects when completing a task that significantly constrains the amount of evidence ostensibly giving rise to perceptual decisions.

Vita

2017-Present	Graduate Administrative Associate, The Ohio State University.
2016-2017	Graduate Research Associate, The Ohio State University.
2016	Summer Research Excellence Fellowship, The Ohio State University
2015	Psychology Department Fellowship, The Ohio State University
2015	B.A. Psychology, Youngstown State University

Publications

Research Publications

M.A. Kandemir, M.M. Jameson, J.J. Palestro "Translating the Children's Anxiety in Math Scale to Turkish for use with Turkish children". *British Journal of Education, Society & Behavioural Science*, 15(1):1–10, 2015.

Fields of Study

Major Field: Psychology

Table of Contents

List of Tables

List of Figures

Chapter 1: Introduction

1.1 Introduction

As human beings, our days are spent making decisions. These decisions can range from simple and trivial, such as remaining stopped at a red light, to complex and vital, such as choosing a retirement plan. As decision making plays an integral part of every day life, understanding how human beings deliberate among choices has been a major area of psychological research for decades. However, most of the decisions we make from day to day are incredibly complex, and often involve different strategies, time courses, or multiple aspects that must be considered. Unfortunately, the complexity of these decisions prevent us from directly studying them in a controlled manner, as we cannot logically rule out the many different factors that might contribute to a particular decision. As a consequence, researchers typically reduce the types of decisions in a laboratory to simple perceptual judgments, where one of the most popular tasks is a choice between two options consisting of unidimensional stimuli: the two-alternative force choice (2AFC) task. The simplicity of the 2AFC task provides an opportunity to control both the complexity of the decision and the environment in which the decision takes place. And while it may lack ecological validity, having

a controlled environment allows for a greater understanding of the underlying choice mechanisms involved in the decision-making process at a finer granularity.

1.1.1 Random Dot Motion Task

One particularly common stimulus in the 2AFC task is a display of dots that move in random directions. In this task, commonly referred to as the random dot motion (RDM) task, subjects are asked to report the net direction of a cloud of moving dots, where some proportion of the dots are moving coherently in one direction while the rest are moving randomly (see Newsome & Pare, 1988). To determine the direction of dot motion, subjects must track the motion of the dots across time. Theoretically, a common way to think of the tracking process is to assume that over time, evidence is collected from one moment to the next and is then aggregated in some way, such as through integration or a running average (cf. Turner, Gao, Koenig, Palfy, & McClelland, 2017).

One common experimental manipulation is to increase the amount of coherence (i.e., the number of coherently moving dots), where the strength of coherent motion is treated as the independent variable, and choice probabilities are often treated as the dependent variable. This manipulation relates to theory because as the coherency increases, the diagnosticity of the evidence at each moment in time is more closely related to the actual strength of coherent motion.

Two Major Paradigms

If one assumes that an observer uses the amount of collected evidence at each moment when making a choice, then one can consider how the amount of evidence relates to the time of the choice. Two canonical ways of understanding this interaction

is to treat time as either an independent or dependent variable. When treating time as an independent variable, the experimenter chooses how long the cloud of dots should be presented on the screen, after which the observer is required to make a decision (Gao, Tortell, & McClelland, 2011; Ratcliff, 2006; Turner, Gao, et al., 2017; Usher & McClelland, 2001; Wickelgren, 1977). In this instance, a cue is presented to prompt the subjects to respond, and subjects are instructed to withhold their response until this cue appears. This particular paradigm is referred to as the "interrogation" or "signal to respond" paradigm, although we will use the former throughout this paper.

Thinking back to how evidence is related to time, the interrogation paradigm allows us to explicitly control the amount of evidence that could have been collected up to some time point. The typical pattern of results is that increases in viewing times are associated with increases in accuracy. At short interrogation times, where the observer has not had an opportunity to gather much evidence, the probabilities of each response are near chance (i.e., near 0.5). However, with longer viewing durations, the probabilities of each response differentiate, such that rightward and leftward responses are more accurate (i.e., away from 0.5). These functions are often called time-accuracy curves, and they can be further modulated by coherency, where larger coherencies produce time-accuracy curves that differentiate quicker with respect to time.

In another paradigm, the "free response" paradigm, the viewing times of the stimuli are not experimentally controlled, but instead both choice and response time are dependent variables. In these paradigms, subjects are allowed to respond whenever they are ready to make a choice. Unlike the interrogation paradigm, the free response paradigm allows researchers to investigate the joint distribution of choice and response time, so that the accuracy of choices can be examined conditional on the

time at which the choice was made. When subjects are left to respond at their own pace, the typical pattern of results is that decisions are more accurate with longer viewing times, producing the weak form of the speed-accuracy tradeoff (SAT; Garrett, 1922; Wickelgren, 1977).

Both paradigms have strengths and weaknesses. The interrogation paradigm is potentially problematic because despite instructions to respond immediately following the cue, subjects sometimes take longer to respond for the shorter viewing durations relative to the longer ones. Additionally, it is possible that subjects continue to integrate information after the go-cue (e.g., Moran, Teogorescu, & Usher, 2015), and this extra integration can complicate the interpretation of the time-accuracy curves. The free response paradigm circumvents this issue by simply conditioning on the time at which a response was made. However, one potential limitation of the free response paradigm is that it only obtains information about the state of gathered evidence at the time of the response. In other words, it gives us no insight into how subjects gather evidence from one moment to the next. As such, it is difficult to construct a concrete timeline for how evidence accumulates in the same way as in the interrogation paradigm.

The paradigms above have played a critical role in providing benchmarks that any newly proposed theory about the decision process must match. Up to this point, we have only mentioned in an abstract sense how observers might go about completing the RDM task. However, there currently exists a number of interesting theories about how observers make these types of decisions, and these theories are successful in that they provide an explanation across the two paradigms. Having described the stimulus

and various task manipulations, we can now turn to a more direct description of the theoretical accounts of the data from each task.

1.1.2 Sequential Sampling Theory

Currently, the most successful attempts to explain and understand perceptual decision making behavior in a 2AFC framework involve sequential sampling theory (Laming, 1968; Ratcliff, 1978; Ratcliff & Rouder, 1998; Stone, 1960). In their most basic form, models that embody sequential sampling theory typically assume that decisions are made by sequentially accumulating sensory evidence from a starting point toward a decision threshold. The implementation of this decision threshold is contingent on the response time paradigm. In the interrogation paradigm, it is unnecessary to terminate the evidence accumulation process prior to the presentation of a go-cue, so it is commonly assumed that observers will continue to integrate evidence until they are prompted to respond via a go-cue (but see Ratcliff, 2006). Here, the go-cue itself serves as an exogenous stopping rule, effectively replacing the endogenous decision threshold. In this way, the lag time of the go-cue can be used to terminate the evidence accumulation process at different times, allowing for more or less evidence to be accumulated.

In a free-response paradigm, because subjects have an unbounded amount of time to respond, assuming the presence of a decision threshold is essential in accounting for self-terminating nature of the task. In addition, the inclusion of a decision threshold can help to explain the relative speeds of correct and error responses typically observed in experiments with a SAT manipulation. By reducing the distance between decision boundaries or the distance between the boundary and the starting point, response

times will decrease at the cost of accuracy. By contrast, increasing the distance from the starting point to the boundary will increase the accuracy at the cost of an increase in the response time.

In sequential sampling models, the distance to the decision threshold can be assumed to fluctuate from one trial to the next, but the total amount of evidence that must be accumulated before a decision is made is traditionally invariant with respect to time. However, recent evidence, mostly originating from neuroscientific studies, suggests that time invariance may be too strong of an assumption. Instead, two major classes of time-variant models have been proposed: urgency signals and collapsing boundaries. In the next section, we briefly review some of the evidence for time-variant mechanisms within sequential sampling theory.

1.1.3 Evidence for Time Variance

The general class of sequential sampling models have enjoyed widespread success as they have continued to provide a successful account of a variety of empirical benchmarks observed in decision-making tasks. One of the most successful instantiations of sequential sampling theory is the diffusion decision model (DDM; Ratcliff, 1978), which assumes that decisions are based on the continuous accumulation of noisy sensory evidence across time, gradually evolving from an initial starting point toward one of two boundaries, each representing a particular choice alternative.

While the DDM has been incredibly successful since its inception in accounting for a variety of choice response time data, it has evolved significantly during that time. The original DDM (Ratcliff, 1978) included three sources of variability: within-trial variability in the accumulation process, and between-trial variability in both the

mean rate of evidence accumulation (i.e., drift rate) and nondecision time. These sources of variability were essential for the model to account for a variety of empirical benchmarks that early sequential sampling models could not produce (Ratcliff & McKoon, 2008; Ratcliff & Tuerlinckx, 2002). The "modern" DDM (Ratcliff & Rouder, 1998) implemented an additional source of variability – between-trial variability in the starting point – that allowed the diffusion model to account for the relative speeds of errors and correct responses in the SAT manipulation (e.g., Ratcliff, Thapar, & McKoon, 2006; Starns & Ratcliff, 2010; White, Ratcliff, Vasey, & McKoon, 2009).

However, despite the model's ability to account for behavioral data, recent literature has noted that it suffers from a lack of neural plausibility in that it does not provide an explicit explanation for how its mechanisms are implemented in the brain. Ditterich (2006a) describes this problem in terms of the black box view of cognitive research, such that there is a clear separation between models of behavior and the actual neural mechanisms that produce said behavior. While choice accuracy and response times can be modeled with the hypothesized parameters of the DDM, the black box view asserts that one's knowledge remains incomplete if the model fits cannot be substantiated by the true neural correlates (but see Turner, Van Maanen, & Forstmann, 2015, for such an analysis). In other words, if we are to assume that the DDM is the ground truth of the decision-making process, then there needs to be a convincing link between the parameters governing the behavioral response and the observed neural activity (Schall, 2004; Turner, Forstmann, Love, Palmeri, & Van Maanen, 2017).

1.1.4 Neural Correlates and the Diffusion Model

While the current state of the literature does not quantitatively link diffusion model parameters to neural correlates, strong evidence from single cell recording suggests that firing rates of neurons may mimic processes in the diffusion model (see Gold & Shadlen, 2007, for review). Groundbreaking work in monkey neurophysiology has linked neuronal firing with the accumulation of sensory evidence for decisions about vibration frequency (primary somatosensory cortex, S1 Hernandex, Zainos, & Romo, 2002; Romo, Hernandez, & Zainos, 2004), direction of motion (middle temporal area, MT Britten, Shadlen, Celebrini, & Movshon, 1996), novelty detection (frontal eye field, FEF; see Schall & Thompson, 1999, for review), and face versus object discrimination (inferotemporal cortex, IT; Afraz, Kiani, & Esteky, 2006).

Neural firing rate in the lateral interpaietal area (LIP) in particular has been studied extensively (Roitman & Shadlen, 2002; Shadlen & Newsome, 1996). The LIP response to sensory stimuli is unique in that the buildup in firing not only scales with the quality and strength of evidence, but it also predicts motor responses (Roitman & Shadlen, 2002), suggesting that the LIP tracks an implicit decision variable rather than sensory evidence only (see Kelly & O'Connell, 2013, for review). Importantly, firing rate of LIP neurons appears to increase up to a bound before a decision is made, regardless of evidence strength (Roitman & Shadlen, 2002). Applying a diffusion model framework to LIP data has further shown that mean firing rate maps onto drift, and the variance in firing patterns maps onto diffusion (Churchland et al., 2011).

Other modeling efforts, however, have revealed potential discrepancies between the standard diffusion model account of decision making and data from neurophysiology. In the LIP literature, neural firing rate only maps onto the decision variable within a diffusion framework after approximately 220 ms of the stimulus being on the screen (Mazurek, Roitman, Ditterich, & Shadlen, 2003). This is troubling, considering evidence from other studies shows that the evidence required to reach a decision threshold is accumulated in a window typically between 100 and 300 ms post-stimulus onset (e.g. Cook & Maunsell, 2002; Luna, Hernndez, Brody, & Romo, 2005). Even in the choice and response time realm, a standard bounded accumulator model can explain the basic properties of the neurophysiological data as well as the mean response time distribution of correct responses, but it cannot account for longer response times for error responses (Ditterich, 2006a).

A similar pattern of results has emerged from human decision-making literature, such that a standard diffusion framework accounts for some, but not all, aspects of the neurophysiological correlates of evidence accumulation and the subsequent choice response time data. Using electroencephalography (EEG) methods, researchers can record neural responses with the high level of temporal resolution that is crucial for capturing perceptual decisions that occur within short time windows. Decisions on perceptual tasks such as the random dot motion task, where subjects are asked to determine which direction a cloud of dots is moving, are typically made within 1000 ms of stimulus onset. With typical event related potential (ERP) analyses, researchers can observe the time course of voltage activity at multiple scalp electrode locations in response to the appearance of a stimulus, and they can dissociate the peaks of this voltage activity at the millisecond level.

Philiastides and Sajda (2006a) used this ERP approach in a line of studies that identified two EEG components that were relevant to distinguishing between face and car stimuli, with one occurring at 170 ms post-stimulus and the other at 300 ms (see also Philiastides, Ratcliff, & Sajda, 2006; Philiastides & Sajda, 2006b). In the original study, stimuli were combined images of faces and cars with varying levels of coherence, such that some images were easily recognizable as a car, some were easily recognizable as a face, and others were more ambiguous. The amplitudes of both the early and late ERPs varied with the coherence of the stimuli at the trial level, with higher voltage amplitudes corresponding to higher coherence for each stimulus-type. The authors concluded that these ERP components tracked stimulus quality at the sensory level.

Ratcliff, Philiastides, and Sajda (2009) reanalyzed the data from Philiastides and Sajda (2006a) using a DDM framework (Ratcliff, 1978; Ratcliff & Rouder, 1998) in an effort to link the ERPs to an evidence accumulation process. Binning the data by high and low peak amplitude for each component revealed differences in estimates of drift rate for the late component, such that drift rates were estimated to be more positive (reflecting faster response times for correct responses) for the subset of high-amplitude peaks occurring 300 ms post-stimulus relative to the subset of low-amplitude peaks. This model-fitting endeavor was an important stride for neuroscience in the decision-making literature in that it revealed an ERP link between stimulus quality and subsequent accuracy and response times on a perceptual task, helping to mitigate the black box problem. Further, the fact that only the late ERP component produced different drift rate estimates, depending on peak amplitude, indicates a dissociation between perceptual encoding and post-sensory processing,

which are two stages of perceptual decision-making that have been difficult to disentangle on a quantitative level in EEG, functional magnetic resonance imaging (fMRI), and single-unit recording studies (see Kelly & O'Connell, 2013, for review).

The authors of Ratcliff et al. (2009) were transparent about the fact that their model fits would not have been possible without the inclusion of a trial-to-trial (i.e., between-trial) variability parameter. They upheld trial-to-trial variability as a crucial aspect of systems neuroscience and suggested that decoding the noise could reveal aspects of decision models that are not presently well-understood. The obvious question that arises from this work, however, concerns the mechanisms underlying trial-to-trial variability. There is, at present, no satisfying theoretical explanation for what would be happening in the brain to result in this variability.

Urgency

Unsatisfied by the lack of a mechanistically plausible account of trial-to-trial variability in the decision process, Ditterich (2006a) proposed a time-variant model that assumed a gain function on accuracy as time increases within a trial, represented as a logistic curve. Specifically, the gain in accuracy is accompanied by an implicit loss in taking longer to complete a trial. In this regime, accumulating evidence has an initially higher gain to avoid fast errors, but as time increases, the gain asymptotes such that accumulating more evidence does not necessarily produce increases in accuracy.

To this end, Ditterich (2006a) fit a variant of the DDM that assumed (1) no between-trial variability, (2) between-trial variability in the drift rate, and (3) the time-variant model described above to both behavioral and neural data from a RDM task (Roitman & Shadlen, 2002). Ditterich (2006a) found that only the time-variant model could account for both the correct and error response time distributions and

accurately predict neural firing rates. However, as a model including between-trial variability in nondecision time and/or starting point was never fit to the data, it is difficult to determine if the best-fitting time-variant model was truly a better representation of the decision-making process than the time-invariant diffusion model.

Following Ditterich (2006a), spectral analyses accompanying human EEG studies have supported the notion of time-variant mechanisms in decision making. Using data from a standard RDM task in human subjects, van Vugt, Simen, Nystrom, Holmes, and Cohen (2014) identified a correlation between 4-9 Hz theta band activity and drift rates fit within a DDM framework. The authors further found that oscillatory power decreased as the decision interval increased, which is contrary to the expectation when assuming that evidence accumulates until it reaches a particular bound. Among other possibilities, the authors suggest that these theta oscillations reflect a signal related to the time within a trial rather than tracking the actual accumulation of evidence (Cisek, Puskas, & El-Murr, 2009), providing additional evidence for the role of time-variant mechanisms in the decision process.

Wyart, V., and Summerfield (2012) further investigated the assumption of linear accumulation of evidence across time in a study that analyzed spectral power oscillations during a Gabor pattern orientation detection task. The main result was that evidence accumulation, as tracked by EEG oscillations in the delta band (1-3 Hz), does not occur linearly over time, but rather fluctuates in a pattern that may point to a weighting and updating mechanism. This suggests an additional cognitive process which evaluates sensory input prior to treating it as evidence in the direction of a particular decision, much like the time-dependent gating described in Ditterich (2006a)

and similar to the gating mechanisms used in Purcell et al. (2010) and Purcell, Schall, Logan, and Palmeri (2012).

Cisek et al. (2009) offered additional insight into the possibility of time-dependent gating by analyzing data from a novel task created to dissociate the predictions made by a standard DDM and an urgency gating model. The authors pointed out that in typical decision-making paradigms, stimulus coherence varies from trial to trial, but the overall signal-to-noise ratio within a stimulus on any individual trial always remains constant. To dissociate predictions between a model that incorporates time-dependent integration from one that does not, Cisek et al. (2009), developed a task that dynamically altered the sensory evidence within a trial. At first, three circles were displayed to the subjects. Next, some number of tokens began in the center circle and then moved one at a time to one of the outer circles during the trial. The participants' task was to choose which of the outer circles would contain the most tokens at the end of the trial. The authors fit several variations of the DDM to the data, including versions with leak and a low-pass filter, but all of them were out-performed by an urgency-gating model. While the claim that predictions of the two models can only be discriminated using time-varying stimuli was ultimately disproven by Hawkins, Wagenmakers, Ratcliff, and Brown (2015), Cisek et al.'s study provided initial support to the notion of time invariance because in their task, strength of evidence alone is not what determined how decisions were made, but rather the strength of evidence over time combined with the urgency to make a choice determined the choice

Thura, Beauregard-Racine, Fradet, and Cisek (2012) extended the task presented in Cisek et al. (2009) to the RDM, where motion coherence changed within each

trial, rather than across trials. The authors speculated that the use of dynamically-changing stimuli adds a degree of environmental realism that is not present with traditional decision-making tasks. In this task, subjects initially responded quickly, ignoring the additional information that occurred later in the trial. However, after a few trials, subjects learned that waiting for the additional sensory evidence was advantageous in terms of overall accuracy. Mirroring the results of Cisek et al. (2009), response time and accuracy distributions were successfully modeled in an urgency gating framework, but not by the standard DDM. Given that urgency gating models succeed under environmentally plausible conditions where the DDM fails, the urgency gating mechanism may serve as a more plausible explanation for how the evidence accumulation process evolves over time.

Collapsing Bounds

Recently, a similar mechanism for time variance has been proposed, where the bounds of the DDM collapse (i.e., move toward the starting point) with increases in time (Bowman, Kording, & Gottfried, 2015; Shadlen & Kiani, 2013). These models are mathematically similar to the urgency models above and provide similarly effective accounts of behavioral data. Similar to the urgency models, as time increases, less information is needed to reach a decision as the bounds collapse, causing the evidence accumulation process to terminate quicker than in a fixed bounds framework. The effective difference is that for more difficult decisions, responses are faster when the bounds collapse than when they do not. This small change produces response time distributions that are less skewed (Hawkins, Forstmann, Wagenmakers, Ratcliff, & Brown, 2015). Much like the success of the urgency gating model, models with

collapsing decision boundaries have proven effective in fitting specific neurophysiological and behavioral data better than time-invariant competitors (e.g., Bowman et al., 2015; Cisek et al., 2009; Ditterich, 2006a).

The plausibility of the collapsing boundary assumption garnered support by measuring choice confidence through post decision wagering tasks (Hampton, 2001; Kiani & Shadlen, 2009; Shields, Smith, & Washburn, 1997). In these tasks, after certain trials, subjects are given the choice to "opt out" of a prior decision to receive a guaranteed small reward or continue with their original decision to receive a larger reward, if correct. Under this paradigm, choice confidence can be assessed by comparing the accuracy of trials where the subject chose to elicit a response against trials in which the subject opted out.

Kiani and Shadlen (2009) used the post decision wagering task to explore the neural mechanisms underlying choice certainty in primates within a random dot motion task that varied both motion strength and viewing duration across trials. To measure neural activity, the authors recorded the neural firing rates from motion sensitive neurons in the lateral intraparietal cortex (LIP). As mentioned in the introduction, previous research has found that the LIP response to sensory stimuli is unique in that the buildup in firing not only scales with the quality and strength of evidence, but it also predicts motor responses (see Roitman & Shadlen, 2002), indicating that the LIP tracks an implicit decision variable related to the accumulation of evidence across time (see Kelly & O'Connell, 2013, for review).

Interestingly, when comparing choice accuracy on trials where the primate had the option to opt out but elicited a response to trials when the primates were not given the chance to opt out, accuracy was better on trials where they waived the

right to opt out. This pattern was consistent across all motion strengths and viewing durations, including equal coherency conditions, leading the authors to assume that both the properties of the stimulus and an internal sense of uncertainty determined how reliable the evidence was to the decision maker, and this presumed reliability regulated confidence across trials (Kiani & Shadlen, 2009).

This assumption was further supported in the neural data. By measuring the difference in spike rates between pools of neurons in the LIP sensitive to the two directions of motion in the RDM task, which was thought to reflect the amount of evidence accumulated across time, the authors found that the primate was more likely to opt for the sure bet when the difference between the motion sensitive neurons was minimal, such as in low coherency conditions (i.e., unreliable stimuli). This led the authors to assume that the LIP response reflects the state of certainty in decisions, and that the mechanisms underlying this representation are linked to the same evidence accumulation mechanisms underlying choice and response time (Kiani & Shadlen, 2009).

Based on this relationship, Kiani and Shadlen (2009) used an extension of a simple diffusion model framework to map the LIP response (i.e., accumulated evidence) to the probability of a correct response, and found that this framework could jointly explain choice, response time, and confidence across time. More importantly, however, is this mapping revealed insight into the relationship between the reliability of evidence and evidence accumulation in the sequential sampling framework. Based on this relationship, the authors posited that if the reliability of the stimulus is unknown, such as when motion strengths vary across trials in an RDM task, it is optimal to adopt a collapsing decision boundary across time (Shadlen & Kiani, 2013). The idea

is that, in the presence of unreliable stimuli, where the amount of evidence accumulation remains close to the neutral point (i.e., equal activation between motion-sensitive neurons), an observer should collapse their decision boundaries to cross a threshold quicker and initiate a response.

Recently, van Maanen, Fontanesi, Hawkins, and Forstmann (2016) provided both behavioral and neural evidence supporting the idea of collapsing decision thresholds and urgency in an expanded judgement task. In this task, they presented subjects with two stacks of bricks that accumulated height in discrete time steps. By manipulating the rate of evidence accumulation for each stack, the authors could "expand" the average decision time by decreasing the rate of accumulation. Using this task, van Maanen et al. (2016) found that as the average rate of accumulation decreased and the average decision time increased, participants were more willing to form a response on less overall stimulus information according to an ideal observer model. The pattern of evidence needed to form a response conditional on the time of the response formed an approximately linear function, suggesting time-variance similar to the other models discussed in this section (e.g., Bowman et al., 2015; Cisek et al., 2009; Drugowitsch, Moreno-Bote, Churchland, Shadlen, & Pouget, 2012). Additionally, they found that the slope of this linear decrease in evidence with time negatively correlated with striatal activity across subjects, suggesting that the striatum may play a role in the degree of urgency across time (van Maanen et al., 2016).

1.1.5 Push Back

Although the evidence is building for the presence of time-variant decision policies, there is also evidence continuing to support time invariance. In the most extensive

comparison of time-variant and time-invariant models to date, Hawkins, Forstmann, et al. (2015) fit variants of a diffusion model with fixed boundaries, variants of a diffusion model with collapsing boundaries, and a variant of the urgency gating model to a variety of human and primate data sets to determine which of the three models best characterized the decision-making process of each subject across all the tasks. The results were mixed, at best, with data from the human subjects being best accounted for by the time-invariant models, and data from the non-human primates being best accounted for by the time-variant models. However, the discrepancies among the model fits by primate type were explained by the use of strategy, practice effects, reward contingencies, or differences in physiology.

In a similar type of analytic strategy, Voskuilen, Ratcliff, and Smith (2016) investigated time-variant and time-invariant models by fitting them to several sets of data from numerosity judgment and RDM tasks. Similar to Hawkins, Forstmann, et al. (2015), Voskuilen et al. (2016) argued that the properties of the stimuli themselves in tasks like the expanded judgement tasks (e.g., Cisek et al., 2009; van Maanen et al., 2016) were a likely contributor to the evidence supporting time-variant models. After finding that each model class fit the data well with no clear "winner" emerging from a comparison using traditional performance measures (e.g., Bayesian Information Criterion; Schwarz, 1978), they used a more novel model comparison technique (i.e., the parametric bootstrap cross-fitting method; Wagenmakers, Ratcliff, Gomez, & Iverson, 2004). Voskuilen et al. (2016), which is thought to be more sensitive to the functional form of each model, to account for the model mimicry observed in their results. Using this technique, they found that the time-invariant model provided a

better fit to the majority of data, further reinforcing the idea that time-invariant assumptions are more appropriate for speeded judgments, such as those in numerosity or RDM tasks, whereas time-variant mechanisms are better suited for stimuli with dynamic time-varying features, such as in the expanded judgment task.

1.1.6 Summary and Outline

While the debate between the specific mechanisms involved in the decision-making process has intensified in recent years, there is no denying that the time-invariant diffusion model has been extremely successful in fitting a variety of choice and response time data from a wide range of tasks, and it has even gained theoretical support from studies in human EEG and monkey neurophysiology. There still remains, however, the black box problem described by Ditterich (2006a), such that there is little proof of a quantitative link between diffusion-based parameter estimates and the actual parameter values that can be obtained from neural firing rate and EEG spectral band oscillations as evidence is accumulated toward a decision bound. Here, we have compiled evidence from behavioral and neural research that a new class of time-variant models have the potential to bridge the gap and provide a more cohesive picture of how the decision process occurs in the brain. However, much more research is necessary before we can disambiguate the evidence accumulation process and make any substantial claims as to which class of models best characterizes how decisions are made.

The goal of the present article is to examine the impact of task characteristics on the decision making process in a two-alternative forced choice framework. Specifically, the article will explore the strategies used in the decision making process when the

interrogation paradigm is mixed with a free response paradigm in an RDM task, where evidence is continuously flowing at a fixed rate throughout the trial (see Voskuilen et al., 2016). Theoretically, we will evaluate the relative merits of using a fixed or collapsing boundary within the DDM. At the outset, our hypothesis is that by mixing the trial types, we establish an expectation that there is a time cost associated with continuing to gather information, because at any point, the stimulus might disappear. This manipulation, combined factorially with a coherency manipulation, should mandate a collapsing boundary policy in order to maximize accuracy but minimize cost (Malhotra, Leslie, H., & Bogacz, 2017).

The outline of the paper is as follows. First, we present the technical details of the diffusion decision model, adapted for the interrogation and free response paradigms. Here, we discuss some candidates for collapsing boundaries, and then compare the relative predictions of the fixed and collapsing-bound models in the context of our task. As we discuss below, the task is designed to differentiate the models, and so it is less susceptible to concerns about specific model fit statistics (e.g., Voskuilen et al., 2016). We then discuss the model variants examined in this article, whose construction depends on a factorial manipulation of various mechanisms meant to control the time-accuracy curves in the interrogation paradigm. Second, we discuss the methods used to collect the data and fit the models to these data. Third, we present the raw data from our task, followed by the results from our model-fitting exercise. We then explain why a model provided the best account of a given subject by examining the specific representations inferred by the two models. Finally, we close with a brief discussion of why collapsing boundaries are the best description of the behavior observed in our task.

Chapter 2: Models

2.1 The Diffusion Decision Framework

In this section, we will describe the details of the diffusion decision model. First, we discuss how sensory evidence accumulates within the model in the absence of a decision boundary. Second, we discuss how the model can be extended to account for the free response paradigm, where choices are made once an observer has accumulated enough evidence to make a decision. Third, we discuss how decision boundaries can be collapsed over time. Fourth, we briefly discuss the theoretical implications of including a collapsing bound in accounting for data within our mixed RDM task. Finally, we identify the many combinations of parameter settings that will be investigated in this article.

2.1.1 Interrogation Paradigm

The diffusion decision model (DDM; Ratcliff, 1978; Ratcliff & Rouder, 1998) assumes that, on the presentation of a stimulus S (e.g., a random dot display), an observer accumulates sensory evidence as noisy samples and integrates these samples over the course of a trial. The integration process can be described in terms of the

differential equation

$$de\left(t\right) = \mu_s dt + \sigma_w dW, \tag{2.1}$$

where $e\left(t\right)$ is the value of the integrated evidence at time t, μ_s is the mean of the accumulated samples (i.e., drift rate), and σ_w represents the moment-to-moment variability in the drift rate (Ratcliff, 1978; Ratcliff & Rouder, 1998; Turner, Gao, et al., 2017). The parameter μ_s dictates the strength of evidence, so in the dot motion task, μ_s would indicate the strength of coherence toward one direction or another. However, due to some perceptual biases, the motion coherence may not map directly onto the diffusion process in Equation 2.1. Denoting the coherence variable as C, and v as the coherence scaling factor,

$$\mu_s = Cv.$$

The values of C can be positive, arbitrarily indicating strength of motion to the right, or negative, indicating motion to the left. If we assume that no decision bound interferes with the integration process (e.g., in the interrogation paradigm), Equation 2.1 dictates how observers integrate evidence from the initial presentation of the stimulus until the time it is removed from the screen. If W in Equation 2.1 denotes a Wiener process, as it does in the DDM (see Smith, 2000; Stone, 1960, for a detailed overview), the sensory evidence $e(t)$ is assumed to be sampled continuously over time from a normal distribution such that $e(t) \sim N\left(\mu\left(t\right), \sigma\left(t\right)\right)$ with mean $\mu(t)$ standard deviation $\sigma(t)$, such that

$$\begin{aligned} \mu(t) &= \mu_s t \\ \sigma(t) &= \sigma_w \sqrt{t}. \end{aligned} \tag{2.2}$$

When modeling data from an interrogation paradigm, the typical approach is to compare the state of evidence at the interrogation time t to some criterion $c(t)$. Given our choices about how positive evidence relates to rightward motion, a rightward response would be generated if $e(t) > c(t)$. Otherwise, a leftward response would be made.

As mentioned in the introduction, the DDM also includes several sources of variability beyond simple sensory integration that complicate Equation 2.1. These extra sources of variability have proven effective in accounting for a number of empirical benchmarks in decision-making tasks (Ratcliff & Rouder, 1998). The first source of variability is the within-trial variability term σ_w already discussed in Equation 2.1. The second source of variability is trial-to-trial (i.e., between-trial) variability in the drift rate, which results from variability in the mean stimulus value of the stimuli used in the task. In line with Equation 2.1, this variability in implemented by perturbing the trial-specific value of μ, typically from a normal distribution such that $\mu \sim \mathcal{N}(\mu_s, \sigma_b)$. Here, the standard deviation parameter σ_b represents the between-trial variability in drift rate. The third source of variability is variability in the starting point of the diffusion process. This source is typically thought of as being independent of the sensory properties of the stimulus itself, and instead is internal to the observer. Starting point variability is implemented in a similar manner as the between-trial variability in drift: on each trial, the initial evidence value is perturbed by a normal distribution centered at z_0 with a standard deviation σ_0, such that $e(t = 0) \sim \mathcal{N}(z_0, \sigma_0)$. Considering these three sources of variability, the mean

$\mu(t)$ and standard deviation $\sigma(t)$ from Equation 2.2 becomes

$$\begin{aligned} \mu(t) &= \mu_s t + z_0 \\ \sigma(t) &= \sqrt{\sigma_0^2 + \sigma_w^2 t + \sigma_b^2 t^2}. \end{aligned} \tag{2.3}$$

2.1.2 Free Response Paradigm

While the above equations accurately describe the evolution of sensory evidence over time in an interrogation paradigm, another prevalent decision-making paradigm is the free response paradigm. In this paradigm, subjects are asked to elicit a response once enough evidence has been acquired. The typical approach to extending the DDM to account for data from this paradigm has been to assume the presence of a decision boundary. Within the DDM, the decision boundary is assumed to be symmetric about zero, such that if the amount of evidence required to make a rightward response is a, the amount of evidence required to make a leftward response is $-a$. Although the DDM is primarily designed to capture data from a two-alternative forced choice task, extensions to multiple alternatives are possible (see Diederich & Oswald, 2014; Leite & Ratcliff, 2010).

The central focus of this article is on how best to impose a decision boundary. The two classes of decision boundaries are fixed bounds (i.e. time-invariant), and collapsing bounds (i.e., time-variant). We now discuss each of these two classes in turn.

Fixed Boundaries

In assuming the presence of a fixed boundary, suppose the threshold amount of evidence required to make a decision is a. We can express the amount of bias in the

decision by relating the mean of the starting point z_0 to the threshold by setting

$$\omega = \frac{z_0}{a} \tag{2.4}$$

(Navarro & Fuss, 2009; Turner et al., 2015; Vandekerckhove, Tuerlinckx, & Lee, 2008). In this parameterization, we can define the probability density function describing the distribution of response times associated with a rightward choice as

$$\begin{aligned} f(t|a,\omega,\mu_s) &= \frac{\pi}{a^2}\exp\left(-\mu_s a\omega - \frac{\mu_s^2(t)}{2}\right) \\ &\quad \times \sum_{m=1}^{\infty} m\exp\left(-\frac{m^2\pi^2(t)}{2a^2}\right)\sin(m\pi\omega) \end{aligned} \tag{2.5}$$

(Feller, 1968; Navarro & Fuss, 2009; Tuerlinckx, 2004). To describe the distribution of response times associated with a leftward choice, we simply flip the drift rate parameter μ_s in Equation 2.5 to make it negative (i.e., $-\mu_s$). As an additional layer of complexity, we can assume the presence of a nondecision time parameter τ_{er}, related to nondecision processes such as perceptual encoding and motor execution. The inclusion of this parameter results in a linear shift in the leading edge of the response time distribution shown in Equation 2.5, which can be implemented by replacing t with $t - \tau_{er}$.

Collapsing Boundaries

An alternative way to impose a decision boundary is to assume that the amount of evidence needed to make a decision depends on the time that has elapsed. These "collapsing" boundary models are viewed as extensions of the simple fixed bound version of the DDM discussed above, but feature decisions boundaries whose collapsing functions are dictated by another set of model parameters. With these additional parameters, collapsing bounds models can explain a variety of data that traditional fixed

bounds models cannot, namely the neurophysiological and decision making behavior of primates (e.g., Ditterich, 2006a; Kiani & Shadlen, 2009; Shadlen & Kiani, 2013). Additionally, researchers have found that adopting a collapsing decision boundary is optimal in situations where the reliability of the source of evidence is unknown (Shadlen & Kiani, 2013), when there is an effort cost of deliberation time (Busemeyer & Rapoport, 1988; Drugowitsch et al., 2012; Rapoport & Burkheimer, 1971), or when one is attempting to maximize reward after withholding a response for a prespecified amount of time (e.g., Malhotra et al., 2017; Thura et al., 2012)

One conventionally assumed functional form specifies that the upper boundary u decreases from its initial value a across time t, such that

$$u(t) = a - \left(1 - \exp\left(-\left(\frac{t}{\lambda}\right)^{k}\right)\right)\left(\frac{1}{2}a - a'\right) \tag{2.6}$$

(Hawkins, Forstmann, et al., 2015), where a is the initial starting point of the boundary, a' is the asymptotic boundary setting, λ is a scaling parameter similar to that of the Weibull distribution, and k is the shape parameter, also similar to that of the Weibull distribution

The specific shape of the collapse is determined by the values of the parameters a', λ, and k. The asymptotic boundary setting a' controls the extent to which the boundaries collapse. Larger values of a' lead to a larger separation between the collapsing upper and lower and boundaries (i.e., straighter boundaries), while smaller values of a' induce a larger collapse of the decision boundaries. When $a' = 0$, the boundaries "completely collapse," where the boundaries decrease to a critical value c that is half of its initial staring point $c = a/2$ (Hawkins, Forstmann, et al., 2015). The left panel of Figure 2.1 illustrates how the parameter a' affects the shape of the decision boundary. Here, the scaling parameter is fixed at $\lambda = 1$, and a' varies from

0 to 1.5. When $a' = 0$, a complete collapse ensues such that the boundaries collapse to half of the starting boundary value: $c = a/2 = 3/2 = 1.5$. As the value of a' increases, the collapse becomes less severe, resembling a fixed bound.

The scaling parameter λ determines the stage (i.e., time) of the boundary collapse. As λ decreases, the bounds collapse earlier with respect to time. The effect of λ on the decision boundary is illustrated in the right panel of Figure 2.1. Here, a' is held constant at $a' = 1$ and λ varies from 0.5 to 1.5. When $\lambda = 0.5$, the bounds collapse earlier in the decision process, reaching the asymptotic collapse much earlier than when λ is larger.

Finally, the shape parameter k determines the shape of the boundary collapse. Depending on the value of k, the boundary could collapse very early in the decision process (early collapse), gradually throughout the decision process (gradual collapse), or later in the decision process (late collapse). For the purposes of this article, the shape parameter was fixed at $k = 3$, which produces a "late" collapse (see Hawkins, Forstmann, et al., 2015, for a demonstration).

2.1.3 Model Predictions

The shape of the boundary has a major impact on how the diffusion model characterizes the decision-making process. In the time-invariant model, the decision boundaries are fixed, which implies that the amount of accumulated evidence needed to reach a decision remains fixed over time. By contrast, in the time-variant model, the decision boundaries collapse toward zero, meaning that less evidence is required as time increases.

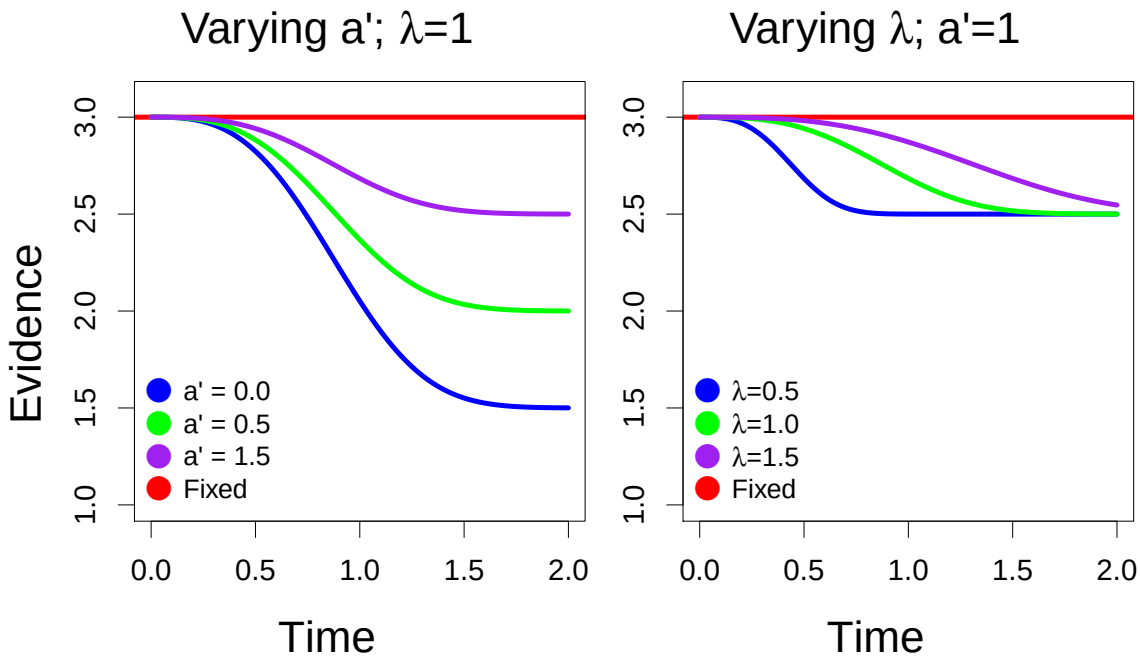

Figure 2.1: **Effects of Parameters on the Shape of the Collapse.** A demonstration of how varying either a' (left panel) or the scaling parameter λ (right panel) affects the shape of the boundary collapse.

Figure 2.2 illustrates how the models lead to different predictions for the response times of a given choice. Here, the accumulation path of one simulation is shown (i.e., blue line) along with two hypothetical boundaries (i.e., black lines). The fixed boundary is invariant with respect to time and stays constant at $a = 3$. The collapsing boundary starts at $a = 3$ and completely collapses to $a/2 = 1.5$. The red circles indicate the point at which a response would be made from the two model variants. The y value associated with each of these circles indicates the amount of evidence that would be needed, whereas the x value indicates the response time. In this simulation, and in general, the amount of time to make a response is smaller in the collapsing bound model than the fixed bound model.

Although Figure 2.2 illustrates how the two models can predict different response times, it is somewhat abstract, and so one may wonder whether the models make different predictions for our mixed RDM task. To investigate this, we simulated the two models 10,000 times and recorded important response probabilities. For the time-invariant model, we set the decision threshold to $a = 3$. For the time-variant model, we fixed the starting point of the threshold $a = 3$, the asymptotic boundary parameter $a' = 0$, the scaling parameter $\lambda = 0.5$, and the shape parameter $k = 3$. All remaining parameters were equivalent across models: the starting point $z_0 = 0$, within-trial variability in drift $s = 0.1$, nondecision time $t_{er} = 10$ (in milliseconds), between-trial variability in nondecision time to $s_\tau = 0$, between-trial variability in starting point to $s_0 = 5$, and between-trial variability in drift to $\eta = 0.2$. In line with our mixed RDM task, we examined five different interrogation times: 0.1 s, 0.3 s, 0.5 s, 0.7 s, and 0.9 s. We also examined three different coherencies: 0%, 0.25%, and

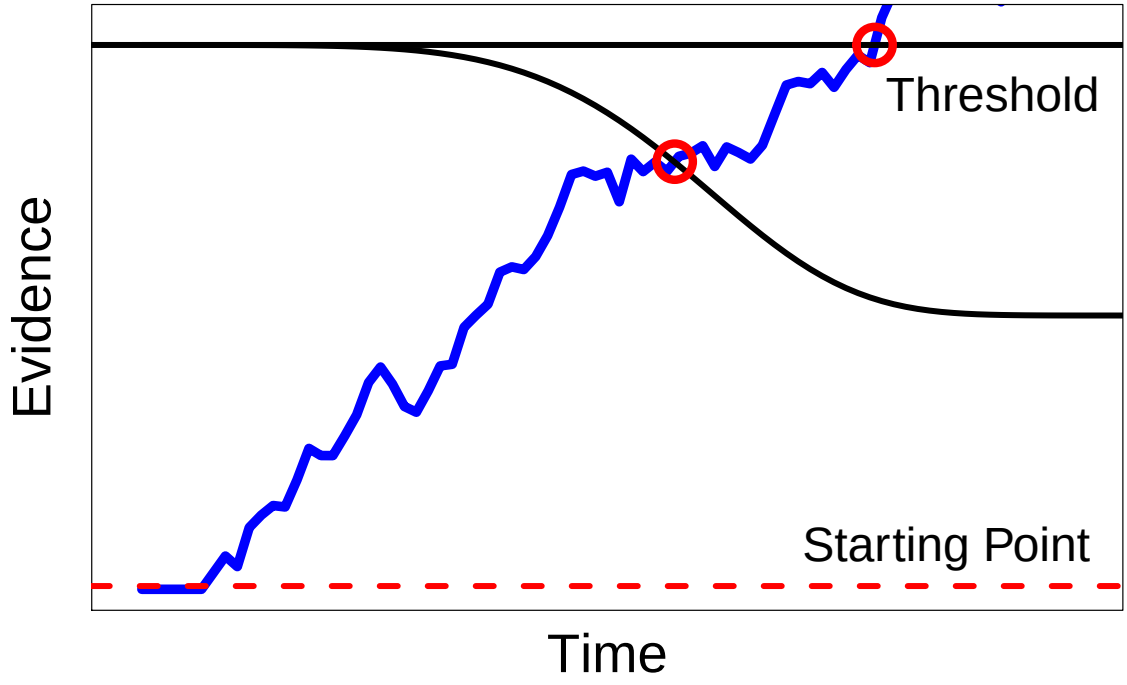

Figure 2.2: **Differences in Accumulation to Bound.** A demonstration of the differences in model predictions for response time in the presence of either a fixed or collapsing bound.

0.50%. All parameters were chosen to highlight the potential differences in predictions between the two models.

Figure 2.3 shows the response probabilities for important statistics in our mixed RDM task. The time-variant model is illustrated as the lines with "X"s, whereas the time-invariant model is illustrated as the lines with open circles. The three coherencies are illustrated as black (0%), red (0.25%), and blue (0.50%) lines. The left panel shows each model's predictions for the probability of responding prior to the disappearance of the stimulus for each interrogation time. Both models make similar predictions for the first interrogation time, but as time increases, the predictions diverge such that the time-variant model predicts a larger probability of early responding (i.e., prior to the disappearance of the stimulus). As one might predict from Figure 2.2, the probability of an early response increases by virtue of the decreased need for more evidence with increases in time. Furthermore, the degree of separation is modulated by the strength of coherence. Namely, at low coherencies, the differences in the models' predictions are larger than when the coherencies are high.

The middle panel of Figure 2.3 shows the accuracy of each model's predictions conditional on an "early" response. In our task, an early response is defined as a response that is made prior to the go-cue. By this metric, the two models make very similar predictions, with the exception of the 0.25% coherency in the last three interrogation times. Here, the middle panel shows that the fixed bound model is more accurate relative to the collapsing bound model. As in the left panel, the boundary has a clear effect on the processing of the stimulus information: because the bound collapses by the last few interrogation times, less stimulus information goes into the sensory evidence variable by the time a decision is made. The consequence is that

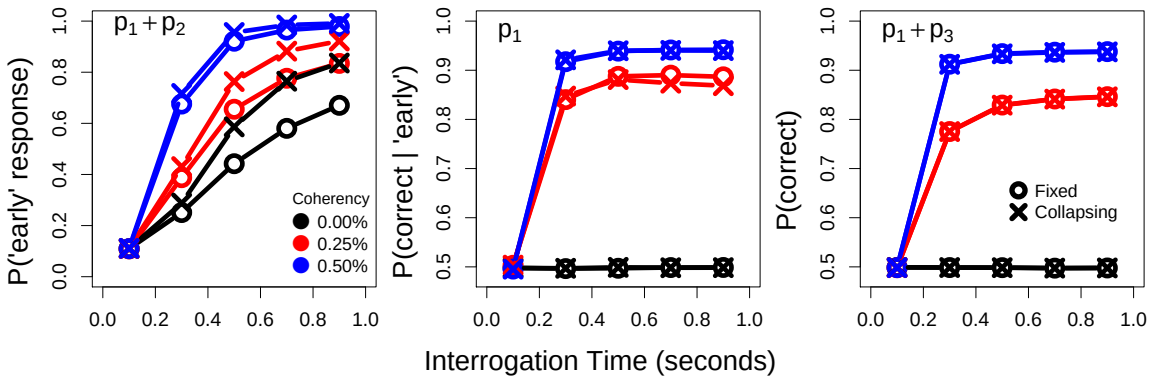

Figure 2.3: **Differences in Response Probabilities.** The time-invariant and time-variant models make different predictions for important response probabilities in the mixed RDM task. The left panel shows that the time-variant model predicts a larger probability of an "early" response (i.e., a response made prior to the disappearance of the stimulus). The middle panel shows the accuracy of the model's predictions for trials in which an early response was made, whereas the right panel shows the accuracy across both types of responding (i.e., early and not early). Predictions for the time-variant and time-invariant models are shown as lines with either open circles or "X" symbols, respectively. The color of each pair of lines illustrates differences in the strength of motion coherence: 0% (black), 0.25% (red), and 0.50% (blue).

the collapsing bound model produces faster response times at the cost of being less accurate.

The right panel of Figure 2.3 shows the total accuracy of the model's predictions, aggregating across whether or not the simulated accumulation path hit a boundary. By this metric, both models make identical predictions for accuracy, modulated only by the strength of motion coherence.

The purpose of the simulation above is to show that mixing interrogation and free response paradigms, along with a coherence manipulation, can create differences in the model predictions. it is important to note that, in this task, the three predictions in Figure 2.3 will not be useful in discriminating between the time-invariant and

time-variant models when analyzed independently, but rather, they should be used in conjunction to aid in discrimination. Hence, the experimental design will play a crucial role in the model comparison analyses we report below (I. J. Myung & Pitt, 2002, 1997; J. I. Myung & Pitt, 2009). That is not to say that either the collapsing bound or fixed bound models cannot adjust other parameters to compensate for predictions that are not supported by the data. For example, the fixed bound model can easily produce increases in the probability of early responses (i.e., the left panel of Figure 2.3) by increasing the between-trial variability term. As Equation 2.3 shows, increases in the between-trial variability term cause increases in the variance of the state of sensory evidence, and in the limit, the between-trial variability term has the largest impact on how the variance increases with time. When assuming the presence of a bound, the end result is an increase in (early) terminations that also lead to a decrease in accuracy, on average (e.g., see the middle panel of Figure 2.3). The question is whether the adjustments made by either model produce deficiencies in other aspects of the data. Our experiment was designed to capture these model adjustments, while providing enough constraint to detect deficiencies.

2.1.4 Model Variants

To test the relative merits of collapsing and fixed boundaries, we created 16 different variants by allowing different combinations of between-trial sources of variability and collapsing bound parameters. As a guide, Table 2.1 provides a list of all model parameters, and Table 2.2 provides each model variant along with a summary of which parameters were freely estimated. For simplicity, we can divide these 16 models into four classes. Class 1 (models 1-4) consists of four variants assuming fixed boundaries

that differ by which sources of between-trial variability are allowed to be estimated. Classes 2, 3, and 4 are time-variant models that simply allow different combinations of the collapsing bound parameters to be freely estimated. All models freely estimate the mean of the nondecision time parameter τ_{er}, the between trial variability in nondecision time S_τ, an initial threshold a, the mean of the starting point z_0, and coherence scaling parameter v. The response criterion $c(t) = \xi$, which was used to determine a rightward or leftward response should a response not be made prior to the stimulus being removed from the screen, was fixed at $\xi = 0$. The within-trial variability in drift was fixed at $s = 0.1$ for identifiability purposes. We now discuss the remaining specifications intrinsic to each model variant within each class.

Table 2.1: **Diffusion Model Parameters.** Notation of the parameters used throughout this article, along with a description of their functionality.

Parameter	Description
a	initial response threshold value
v	drift rate
z_0	starting point
s	within-trial variability in drift rate
τ_{er}	nondecision time
S_τ	between-trial variability in nondecision time
S_0	between-trial variability in starting point
η	between-trial variability in drift rate
a'	extent of the boundary collapse
λ	stage of the boundary collapse
k	shape of the boundary collapse

Class 1 is composed of four models that systematically fix and free the between-trial variability in starting point parameter S_0 and the between-trial variability in drift rate η. Specifically, model 1 assumes that both of these terms are fixed to be

zero (i.e., $S_0 = \eta = 0$), model 2 allows η to be freely estimated (i.e., $S_0 = 0$), model 3 allows S_0 to be freely estimated (i.e., $\eta = 0$), and model 4 allows both model parameters to be freely estimated.

Classes 2 (models 5-8), 3 (models 9-12), and 4 (models 13-16) are derivatives of the four models comprising Class 1, but systematically fix and free the asymptotic boundary setting a' and the scaling parameter λ. For all three of these classes, we set the shape parameter $k = 3$ to reduce both the complexity of the time-variant models and the number of potential model variants. Class 2 freely estimates a' but fixes $\lambda = 1$. Class 3 freely estimates λ, but fixes $a' = 0$. Finally, Class 4 freely estimates both a' and λ.

The arrangement of parameters featured in each of the 16 variants is displayed in Table 2.2. Model variants were determined by the type of boundary (time-invariant versus time-variant) and the combination of freely estimated and fixed parameters. If a parameter is freely estimated in the model, it is denoted as "FE" in the table. If the parameter is fixed, it is denoted as "F." If a parameter is absent from the model (e.g, the parameters associated with the collapsing boundary that are not included in the time-variant models), the space is left blank. This model order (1-16) will remain consistent in the following sections, and models will be referred to by their model number denoted in Table 2.2.

Table 2.2: **Model Variants.** Each row corresponds to a specific model we fit to data, where parameters (i.e., columns) were either freely estimated (FE) or fixed (F). If a parameter is not applicable, such as in the time-invariant models, the space is left blank.

Model	Decision Boundaries	a	v	z_0	s	τ_{er}	S_τ	S_0	η	a'	λ	k	Free Parameters
1	Fixed	FE	FE	FE	F	FE	FE	F	F				5
2	Fixed	FE	FE	FE	F	FE	FE	F	FE				6
3	Fixed	FE	FE	FE	F	FE	FE	FE	F				6
4	Fixed	FE	FE	FE	F	FE	FE	FE	FE				7
5	Collapsing	FE	FE	FE	F	FE	FE	F	F	FE	F	F	6
6	Collapsing	FE	FE	FE	F	FE	FE	F	FE	FE	F	F	7
7	Collapsing	FE	FE	FE	F	FE	FE	FE	F	FE	F	F	7
8	Collapsing	FE	FE	FE	F	FE	FE	FE	FE	FE	F	F	8
9	Collapsing	FE	FE	FE	F	FE	FE	F	F	F	FE	F	6
10	Collapsing	FE	FE	FE	F	FE	FE	F	FE	F	FE	F	7
11	Collapsing	FE	FE	FE	F	FE	FE	FE	F	F	FE	F	7
12	Collapsing	FE	FE	FE	F	FE	FE	FE	FE	F	FE	F	8
13	Collapsing	FE	FE	FE	F	FE	FE	F	F	FE	FE	F	7
14	Collapsing	FE	FE	FE	F	FE	FE	F	FE	FE	FE	F	8
15	Collapsing	FE	FE	FE	F	FE	FE	FE	F	FE	FE	F	8
16	Collapsing	FE	FE	FE	F	FE	FE	FE	FE	FE	FE	F	9

Chapter 3: Methods

3.1 Methods

In this section, we describe our perceptual decision making task. As discussed in the introduction, our strategy was to use a mixture of interrogation and free response paradigms so that the amount of stimulus processing time could be treated as an independent variable in our experiment. In addition, the levels of processing time were crossed with the coherency of the dot motion, creating a fully factorial design of strength and duration of evidence.

3.1.1 Subjects

14 healthy subjects were recruited from The Ohio State University. All subjects provided written consent in accordance with the university's institutional review board and were paid $10 per hour of their time. All subjects had normal or corrected-to-normal vision. With the exception of two subjects, all subjects completed the full design consisting of 650 trials. Of the two who failed to complete the study, one only completed 543 of the 650 trials due to scheduling conflicts, and the other abandoned the task at trial 220 of 650. However, both subjects' data were included in all analyses.

3.1.2 Stimuli and Equipment

The random dot motion task was created using a custom program named the State Machine Interface Library for Experiments (SMILE; `https://github.com/compmem/smile`) , which is an open source, python-based experiment building library. Subjects completed the task on a desktop computer with a 24" LCD Display, running at 120 Hz, in a darkened and sound attenuated room.

3.1.3 Design

The mixed random dot motion task used a 5×5 (motion strength $\times$ interrogation time) within-subject factorial design. The levels of motion coherency were 0%, 5%, 15%, 25%, or 35% in either the rightward or leftward direction. The condition of 0% coherency was used as a control condition to enhance our ability to infer perceptual biases that might occur for each subject. The interrogation time could occur at any of the following times: 100 ms, 300 ms, 500 ms, 700 ms, or 900 ms. The interrogation times serve as the control for the stimulus duration. That is, the stimuli were presented for a certain number of milliseconds, after which time the dots were removed from the display. Subjects were trained to treat the removal of the dots as a "go-cue" such that a response should be initiated. However, subjects were also instructed that they could initiate a response while the dots were still on the screen. The full factorial design of our task is displayed in Table 3.1.

On each trial, subjects were presented with clouds of dots constructed using one of the motion strength and interrogation time pairs listed in Table 3.1, and they were required to make a direction of motion (i.e., either left or right) decision based on

which direction most of the dots were moving by pressing either the "D" key with their left hand or the "K" key with their right hand.

Subjects completed one practice block before completing 13 additional blocks, consisting of 50 trials per block. This resulted in 650 total trials. On each block, each motion-strength-interrogation-pair in Table 3.1 was randomly presented once for rightward moving dots and once for leftward moving dots.

3.1.4 Procedure

After obtaining informed consent, subjects sat in front of the computer and were provided with brief instructions by the experimenter. These instructions informed the subject that they would be completing several blocks of a random dot motion task, and that the dot stimuli would only be on the screen for a short period of time before being removed. After explaining the task, the experimenter stressed that while the dots would be removed after a short period of time, the subject did not have to wait for the dots to disappear and could respond at any time. Subjects were informed that they could respond as soon as they were aware of which direction the dots were moving.

After the instruction period, the experimenter started the program and left the room. Each block began with an instruction screen that provided a description of the task, the key mapping, and an illustration of leftward and rightward moving dots. Once subjects felt they were ready to begin the task, they pressed the ENTER key and the task began. Each trial began with the presentation of a fixation cross that remained on screen for 100 ms. Then, a cloud of randomly moving dots were presented to the subject for a prespecified amount of time, and the subjects were asked to make

Table 3.1: **Experimental design.** Columns represent the levels of stimulus durations (i.e., interrogation times), whereas rows represent the levels of coherency. In the Direction blocks, each stimulus duration and motion strength pair was randomly selected and presented twice within a block.

	Stimulus Duration (seconds)				
Coherency	0.1	0.3	0.5	0.7	0.9
0	(0, 0.1)	(0, 0.3)	(0, 0.5)	(0, 0.7)	(0, 0.9)
0.05	(0.05, 0.1)	(0.05, 0.3)	(0.05, 0.5)	(0.05, 0.7)	(0.05, 0.9)
0.15	(0.15, 0.1)	(0.15, 0.3)	(0.15, 0.5)	(0.15, 0.7)	(0.15, 0.9)
0.25	(0.25, 0.1)	(0.25, 0.3)	(0.25, 0.5)	(0.25, 0.7)	(0.25, 0.9)
0.35	(0.35, 0.1)	(0.35, 0.3)	(0.35, 0.5)	(0.35, 0.7)	(0.35, 0.9)

a direction decision. The motion strength and interrogation times were randomly chosen and counterbalanced across trials. Once a response was made, feedback was presented for 100 ms in the form of a green checkmark for correct answers and a red X for errors. Finally, the fixation cross reappeared, denoting a new trial. Subjects were also given feedback in the form of "too fast" or "too slow" if they responded prior to 100 ms after stimulus onset or after 2500 ms post interrogation time, respectively.

3.1.5 Fitting the Models

We fit each of the 16 models at the subject level using likelihood-free Bayesian methods. In the next section, we will provide a brief overview of fitting models in a Bayesian framework, specifically focusing on how the likelihood function of each model was approximated.

3.1.6 Likelihood Free Inference

In Bayesian estimation, we strive to update our belief about some parameters θ after observing some data D by means of the posterior distribution $p(\theta|D)$. To obtain this posterior distribution, we must specify two things: 1) the likelihood function $L(\theta|D)$, which can be thought of as the probability of observing data D given the parameters θ, and the prior distribution $p(\theta)$, which reflects our uncertainty regarding the true value of θ prior to observing any data.

To specify the likelihood function $L(\theta|D)$, we need a statement describing the relationship between θ and the observed data $D = \{D_1, D_2, ...D_N\}$, which we assume are independent and identically distributed (i.i.d). To begin, we assume that this i.i.d. data arises from some model with parameters θ such that $D \sim Model(\theta)$. Using this notation, we can write the probability density function describing the relationship

between any observation D_i and the parameters θ as

$$\text{Model}\left(D_i|\theta\right). \tag{3.1}$$

With this probability density function, we simply take the product of all the densities evaluated at each observation D_i

$$L\left(\theta|D\right) = \prod_{i=1}^{N} \text{Model}\left(D_i|\theta\right). \tag{3.2}$$

To specify a prior over θ, we simply chose a probability distribution that represents our knowledge of the distribution of θ prior to observing any data. Once a distribution is chosen, we can evaluate the posterior distribution using Bayes' rule

$$p\left(\theta|D\right) = \frac{L\left(\theta|D\right)p\left(\theta\right)}{p\left(D\right)}, \tag{3.3}$$

where $p\left(D\right)$ is the marginal probability of the observed data D. In practice, as $p\left(D\right)$ does not depend on parameters θ, it can be thought of as a normalizing constant and dropped from the equation. As such, Bayes' rule can be rewritten as

$$p\left(\theta|D\right) \propto p\left(D|\theta\right) \times p\left(\theta\right). \tag{3.4}$$

In other words, the posterior distribution of the parameters θ given the data D is proportional to the likelihood times the prior.

For simple models, where the likelihood function is known, computing Bayes' rule is trivial. However, with more complex models, where the closed-form of solution of the likelihood function is not explicitly known or tractable (such as some of the models in this paper), we cannot approximate Bayes' rule directly. As such, we must result to simulating these models using likelihood-free methods, such as approximate Bayesian computation (Turner & Sederberg, 2012; Turner & Van Zandt, 2012, 2014),

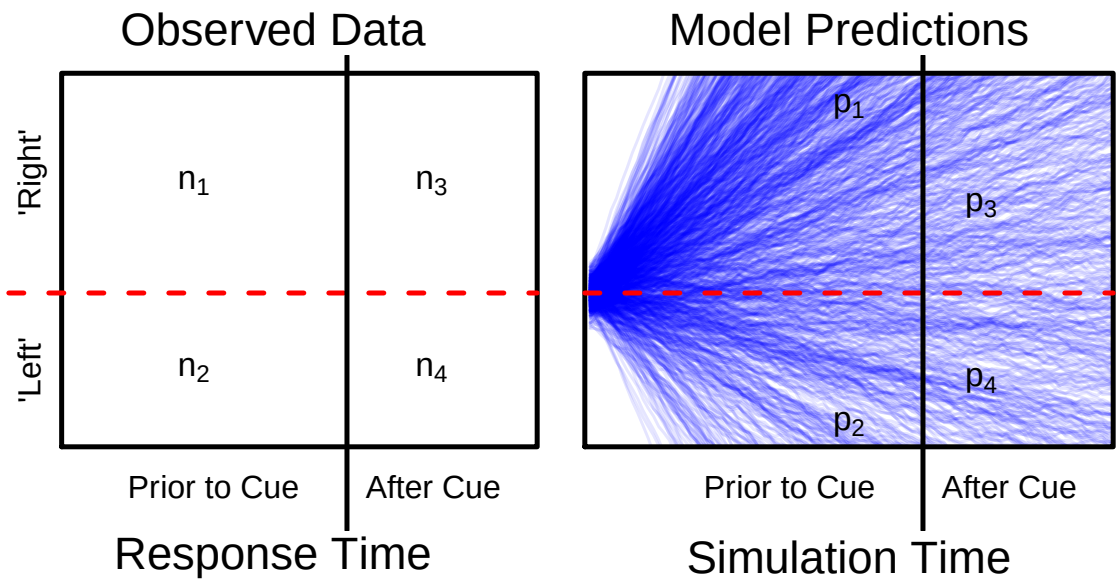

Figure 3.1: **Model Fitting Strategy.** For a given stimulus coherence and duration (i.e., black vertical line), the figure shows how the data (left panel) and model simulations (right panel) were organized to evaluate the suitability of a given model parameter. The left panel shows how the data were discretized into four contingencies: rightward (n_1) or leftward (n_2) response prior to the cue disappearing, and a rightward (n_3) or leftward (n_4) response after the interrogation time. The right side illustrates an analogous discretization of the model simulations. The red horizontal line represents the right/left criterion. Note that in the model simulations, the criterion is only realized if a bound is not hit prior to the disappearance of the cue.

which allow us evaluate these models in a fully Bayesian context. In the next section, we will describe in greater detail how we used likelihood free methods to fit our models to data.

3.1.7 Model Fitting

Figure 3.1 shows the basic strategy we used for fitting each of the 16 models to data by demonstrating how the data were organized for a given stimulus coherency

i and duration time j (see Table 3.1). For each cell in the factorial design, we organized the data $D_{i,j} = \{n_1, n_2, n_3, n_4\}$ based on the number of times four unique events occurred: a rightward (n_1) or leftward (n_2) response prior to the dots disappearing, and a rightward (n_3) or leftward (n_4) response after the dots disappeared (i.e., after the interrogation time), where $\sum_k n_k = N_{i,j}$. From a modeling perspective, this distinction is important because it separates out two different types of accumulation: in one type, enough evidence was acquired prior to an interrogation time, whereas in another, introspection is necessary to determine the state of evidence acquired conditional on the stimulus duration. By factorially manipulating the stimulus coherency with stimulus duration, we can decipher how coherency interacts with duration, and provide greater constraints on the models.

To fit our models to data, we require a method of assessing the relative accuracy of a model's predictions relative to the data contingencies illustrated in the left panel of Figure 3.1. If the likelihood function were known, we could easily calculate the probability of observing each response contingency shown in the left panel of Figure 3.1 by plugging in a set of model parameters into an equation. As discussed above, for the time-invariant model, the joint distribution of choice response time has been derived, and this joint probability density function could be combined with the additional contingency of not arriving at a bound prior to the interrogation time by calculating the expected state of evidence after some time t (see Equation 3.8).

However, some of the models under investigation in this article do not yet have analytic likelihood functions, because introducing a collapsing bound makes Equation 2.5 inappropriate for our analyses. As a workaround, we can instead simulate the model many times to obtain an approximation of how likely the data in the left

panel of Figure 3.1 are under a given set of model parameters (i.e., the likelihood function). Various methods for using model simulations to approximate the likelihood function have been developed and shown to be accurate in several different modeling applications (e.g., Hawkins, Forstmann, et al., 2015; Heathcote, Brown, & Mewhort, 2002; Palestro et al., 2017; Turner, Dennis, & Van Zandt, 2013; Turner, Schley, Muller, & Tsetsos, 2017; Turner & Sederberg, 2012, 2014; Turner & Van Zandt, 2012, 2014). The right panel of Figure 3.1 shows simulated trajectories from one of the models for a given set of model parameters. Here, we simulated the model many times for a fixed amount of time (i.e., the longest interrogation time in our experiment), and then analyzed the accumulation paths to separate the event contingencies in an analogous way to the data arrangement in the left panel.

To separate the event contingencies, we simply track the state of each accumulation path. Letting $x(t)$ denote a given accumulation path, the first contingency is whether a bound was hit prior to the stimulus disappearance. If we let $a(t)$ denote the state of the boundary at a given time t, and t^* denote the interrogation time, we can define the first point t_0 such that a boundary was crossed as

$$t_0 = \min\left[t \in (0, \infty) : |x(t)| \geq a(t)\right].$$

Because we cannot realistically simulate the model for an infinite amount of time, it is possible that the state $x(t)$ never crosses the boundary $a(t)$. In this case, we set $t_0 = t^*$ to indicate that a boundary was not hit prior to the stimulus disappearing.[1] Using the boundary crossing time t_0 and the boundary symmetry assumed within the DDM (i.e., the boundaries are symmetrical about 0), we can define the probability

[1]Of course, any arbitrary value of t_0 such that $t_0 \geq t^*$ will satisfy the law of total probability for the event contingencies $\sum_k p_k = 1$.

of making a right (p_1) or left (p_2) response prior to the stimulus disappearing as

$$\begin{aligned} p_1 &= p\left(t_0 < t^* \cap x(t = t_0) > 0\right), \text{ and} \\ p_2 &= p\left(t_0 < t^* \cap x(t = t_0) < 0\right), \end{aligned}$$

respectively. The second contingency is whether a decision was made after the stimulus disappeared. For simplicity, we assume that when the stimulus disappears, no more information can be processed, although there are many other proposals for post-stimulus processing (Moran, 2015; Pleskac & Busemeyer, 2010). Because we assume that no more information can be processed, to form estimates of the final contingencies, we examine the states of evidence at the interrogation time t^*. To make a left/right decision, we assume the observer evaluates the state of evidence at the interrogation time relative to a criterion parameter ξ. If the state of evidence is greater than the criterion, a rightward decision is made, whereas if the state of evidence is less than the criterion, a leftward decision is made. Formally, the probabilities of leftward (p_3) and rightward (p_4) responses following an interrogation time are

$$\begin{aligned} p_3 &= p\left(t_0 \geq t^* \cap x(t = t^*) > \xi\right), \text{ and} \\ p_4 &= p\left(t_0 \geq t^* \cap x(t = t^*) < \xi\right), \end{aligned}$$

respectively. Together, the above decision rules enforce the constraint $\sum_k p_k = 1$. Each event probability p_k is shown in the right panel of Figure 3.1, and their relationship to important probabilities in the mixed RDM task is shown in Figure 2.3. The red horizontal line represents the criterion parameter ξ, but it should be noted that the criterion only plays a role when decisions are made after the stimulus has disappeared. Similarly, the response threshold $a(t)$, illustrated as the boundary box

surrounding the trajectories, only plays a role when decisions are made prior to the disappearance of the stimuli.

Once each of the event probabilities have been obtained, we need to evaluate how closely the distribution of event probabilities match the observed data. One way to do this is through a simplified version of the probability density approximation (PDA; Turner & Sederberg, 2014) method for data of discrete type. Here, we conceive of the data observations as a multinomial draw with probabilities dictated by the model simulations. For a given stimulus coherency i and viewing duration j, we can evaluate the probability of observing the data $D_{i,j}$ given a set of model parameter θ as

$$\pi(D_{i,j}|\theta) = \frac{N!}{n_1!, n_2!, n_3!, n_4!} p_1^{n_1} p_2^{n_2} p_3^{n_3} p_4^{n_4},$$

where

$$(p_1, p_2, p_3, p_4) \sim \text{Model}(\theta).$$

The notation Model(θ) denotes the model simulation process and path analysis described above, producing the predicted model probabilities p_k. Note that to simulate the model, design information such as stimulus coherency and viewing duration are used explicitly in the model to generate the paths and produce each p_k. For computational convenience, we simulated the model 1000 times for a given parameter proposal θ using the longest viewing duration in our experiment (i.e., 900 ms). We then calculated each of the p_ks by analyzing the trajectories for each viewing duration by setting t^* accordingly in the process described above. To form an approximation of the likelihood function $\mathcal{L}(D|\theta)$ of the data D given a parameter proposal θ, we need only calculate the distribution of p relative to the distribution of n across all

stimulus coherencies and viewing durations using the following equation:

$$\mathcal{L}(\theta|D) = \prod_i \prod_j \pi\left(D_{i,j}|\theta\right).$$

To generate parameter proposals θ, we used a genetic sampling algorithm known as differential evolution with Markov chain Monte Carlo (DE-MCMC ter Braak, 2006; Turner & Sederberg, 2012; Turner, Sederberg, Brown, & Steyvers, 2013). Specifically, we used the "burn in" mode of the ABCDE algorithm (Turner & Sederberg, 2012), as we were only interested in maximizing the approximate posterior density (i.e., the approximate maximum a posteriori estimate), rather than sampling full posterior distributions. In brief, DE-MCMC is a population Monte Carlo algorithm, which is a special type of MCMC algorithm that makes use of particle filtering to guide the next generated proposal based on the current state of the estimated posterior. The communication between the chains allows the algorithm to automatically tune to the shape of the posterior, making it especially useful for models with highly correlated parameters (such as the DDMs investigated here; Turner & Sederberg, 2012). For each model, we implemented this sampler with 32 chains for 300 iterations. The in-group migration probability was set to 0.15, and the jump scaling factor was set to $b = 0.001$.

3.1.8 Prior Specification

With a suitable approximation of the likelihood function in hand, the final step in estimating the joint posterior distribution is to specify prior distributions for each

of the model parameters. We specified the following priors:

$$\begin{aligned} \log(v), log(z_0), \log(t_{er}), \log(S_t), \log(\eta), \log(a'), \log(\lambda) &\sim \mathcal{N}(0,1) \\ \log(a) &\sim \mathcal{N}(1, 0.5) \\ \log(S_0) &\sim \mathcal{N}(0.5, 1). \end{aligned}$$

While these priors may seem informative, as most of the parameters are on the log scale, the constraint is justified as once these variables are transformed, the range of the prior is only mildly informative. For example, the 95% credible range of a log normal prior with mean 0 and standard deviation of 1 is $(0.14, 7.13)$.

3.1.9 Model Comparison

To compare the relative fit of the 16 models, we computed three different metrics: the Bayesian information criterion (BIC; Schwarz, 1978), the approximated posterior model probabilities (Wasserman, 2000), and the Bayes factor. The BIC is computed for each model using the equation

$$\text{BIC} = -2\log\left(L\left(\hat{\theta}|D\right)\right) + \log(N)\, p, \tag{3.5}$$

where $L\left(\hat{\theta}|D\right)$ represents the log of the posterior density at the parameter value that maximized it (i.e., the maximum a posteriori (MAP) estimate), p represents the number of parameters for a given model, and N is the number of data points for a given subject. The BIC penalizes models based on complexity, with models of higher complexity (i.e., models with more parameters) receiving a stronger penalty than models of lower complexity. As such, the inclusion of additional parameters must substantially improve the fit to the data to overcome the penalty incurred for adding them.

We can use the BIC values calculated for each candidate model to approximate each model's posterior model probability relative to the other 15 models. This approximation, as given by Wasserman (2000) is as follows:

$$P_{BIC}\left(M_i|Data\right) = \frac{\exp\left(-\frac{1}{2}\mathrm{BIC}\left(M_i\right)\right)}{\sum_{j=1}^{m}\exp\left(-\frac{1}{2}\mathrm{BIC}\left(M_j\right)\right)}. \tag{3.6}$$

We can calculate the Bayes factor $BF_{q,r}$ between two model candidates M_q and M_r using

$$BF_{q,r} = \frac{p(D|M_q)}{p(D|M_r)}, \tag{3.7}$$

where $p\left(D|M_i\right)$ is the marginal likelihood of the data D under a model M_i. While simple in theory, to calculate the Bayes factor, it is preferred that the likelihood function for each candidate model be analytically tractable (see Gelman, Carlin, Stern, & Rubin, 2004; Liu & Aitkin, 2008, for greater detail). Unfortunately, while the likelihood function for the time-invariant diffusion decision model has been analytically derived (Feller, 1968; Navarro & Fuss, 2009; Ratcliff, 1978; Tuerlinckx, 2004; Voss, Voss, & Lerche, 2015), the collapsing decision boundary complicates these equations enough to render the likelihood intractable. As the likelihoods must be approximated, so too must the marginal likelihoods in Equation 3.7.

To approximate the Bayes factor, we can use a method suggested in Kass and Raftery (1995), which estimates the Bayes factor by comparing the BIC value calculated for each candidate model. Kass and Raftery (1995) demonstrated that the difference between BIC values for candidate Models q and r asymptotically approximates $-2\log\left(BF_{qr}\right)$ as N approaches ∞. Hence, we can approximate the Bayes factor

in Equation 3.7 by using the BIC values from Equation 3.5:

$$BF_{q,r} \approx \exp\left[-\frac{1}{2}\left(\mathrm{BIC}_q - \mathrm{BIC}_r\right)\right], \tag{3.8}$$

where BIC_i denotes the BIC score for Model i.

Chapter 4: Results

4.1 Results

We present the results in three stages. First, we show the raw behavioral data to ensure that our task manipulations were effective. Second, we provide details of the model comparison by reporting the Bayesian Information Criterion (BIC; Schwarz, 1978), an approximation of each model's posterior model probability (Wasserman, 2000), and an approximation to the Bayes factor (Kass & Raftery, 1995). Third, we provide some insight into the model comparison results by examining the representations inferred from the model fits.

4.1.1 Raw Behavioral Data

To explore the effectiveness of our task manipulation, we first examined the choice response time distributions as a function of the levels of our independent variables. Recall that we explicitly manipulated the length of time that a given stimulus was presented and the strength of motion coherence. Our hypothesis was that each of these variables should affect both the accuracy and the response time associated with each decision, as they both uniquely contribute to the amount of evidence for a choice. Namely, the strength of motion coherence should contribute positively to the quality

of the decision, where increases in motion coherence should increase accuracy and decrease response times. Interrogation times should also contribute positively to the quality of the decision, where increases in the length of the viewing times should increase accuracy. The interaction between interrogation times and response times is complicated in our task, as subjects had the ability to respond before the stimulus was taken off the screen. However, a "premature" (i.e., relative to the interrogation time) response does give us information about the perceived strength of evidence as well as the amount of evidence required by an observer to make a choice. As we will discuss later, these features of the data were helpful in evaluating performance among the models.

Figure 4.1 shows the choice response time distributions for each level of the two independent variables in our task. The panels of Figure 4.1 are organized by the time at which subjects were interrogated (columns) and strength of motion coherency (rows). Within each panel, response times associated with correct responses are shown on the positive axis, whereas response times associated with incorrect responses are shown on the negative axis. As a reference, a response time of zero seconds is illustrated as a red vertical line, and the interrogation times (also indicated by the column) are shown as the blue vertical lines. Response times were trimmed to be greater than 100 ms and less than 4000 ms, resulting in a removal of approximately 0.7% of trials. Again, as the response times depend to some extent on the interrogation times, this particular feature of our data is difficult to visualize in Figure 4.1. Within each panel, if response times appear between the red vertical line and the blue vertical line (on either axis), then the choice was made prior to the stimulus disappearing. By

contrast, if a response time is larger than the interrogation times (i.e., blue vertical lines), then a response was made after the stimulus disappeared.

Figure 4.1 reveals a few interesting interactions between the independent and dependent variables. To get a sense of accuracy, we must compare the relative heights of the two response time distributions within a panel, where a larger density of the response time distribution on the positive axis indicates greater accuracy. Comparing across the rows, for a given interrogation time, accuracy increases with increasing coherency, and the response times tend to decrease. Together, these results suggest that coherency had a large impact on the strength of evidence for the correct choice. Comparing across columns, for a given coherency, accuracy increases with increasing interrogation times, suggesting that with longer viewing durations, a stronger overall strength of evidence could be appreciated by the subjects. The interrogation times do not seem to have a direct impact on the response times, although as interrogation times increased, more responses were made prior to the stimulus disappearing. Also, some distributions do appear to be bimodal, especially at longer interrogation times. One potential explanation for this bimodality is that longer interrogation times induce different response modalities across subjects: some subjects tend to respond prior to the go-cue (creating one mode), whereas others prefer to wait until the go-cue before initiating a response, producing to the bimodal shape.

The left panel of Figure 4.2 displays the accuracy data in another way, where here the accuracy for each experimental cell was calculated by collapsing over response time. The accuracy (y-axis) is shown as a bar plot for each interrogation time (x-axis) and for each level of interrogation time (columns). In the 0% coherency

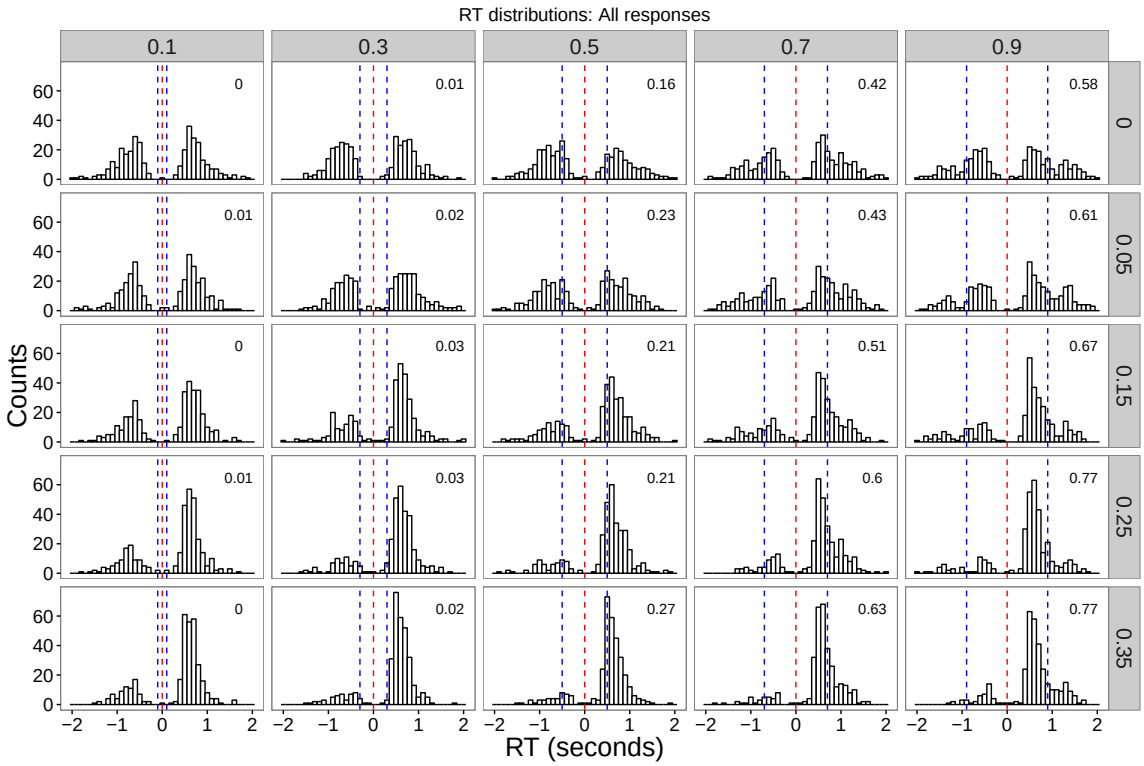

Figure 4.1: **Choice Response Time Distributions.** Each panel shows the choice response time distributions from our experiment, collapsed across subjects, where response times associated with correct responses are shown on the positive axis, and incorrect responses on the negative axis. The panels are organized by the time at which subjects were interrogated (columns) and strength of motion coherency (rows). As a reference, a response time of zero seconds is illustrated as the red vertical line, and the level of interrogation time is illustrated as the blue vertical line. The statistic in the upper right corner of each panel provides the probability that a response was made prior to the stimulus disappearing.

condition, accuracy centers around chance responding (i.e., 0.50). As coherency increases, Figure 4.2 shows that the proportion of correct responses increases for each level of interrogation time. For a given coherency, the accuracy of the responses tends to increase with increases in the interrogation time. The right panel of Figure 4.2 shows the accuracy data in the same way, but separated by whether the response was made prior to (i.e., top panel) or after (i.e., bottom panel) the stimulus disappeared. Here, we see that when a response was executed prior to the cue disappearing, it tended to be modulated by the interrogation time, where longer interrogation times tend to increase the accuracy. However, this trend is not necessarily observed when responses are made after the cue disappeared (i.e., bottom panel).

Taken together, Figures 4.1 and 4.2 suggest that the data from the mixed RDM task match our expectations about how each independent variable should interact with both accuracy and response time.

4.1.2 Model Comparison

The most important aspect of our results is the comparison in fit statistics across models. We used the approach detailed in Section 3.1.5 to fit each of the 16 model variants to our data. We fit each subject independently and obtained a single maximum a posteriori (MAP) estimate that maximized the posterior density. The MAP estimate was then used to evaluate the log likelihood obtained, so that the Bayesian predictive information (BIC; Schwarz, 1978) could be calculated.

Figure 4.3 shows the resulting BIC values obtained by fitting each model (i.e., rows) to each subject's data (i.e., columns 1-14). For visual purposes and because the BIC values cannot be interpreted in an absolute sense, we z-scored the BICs (i.e.,

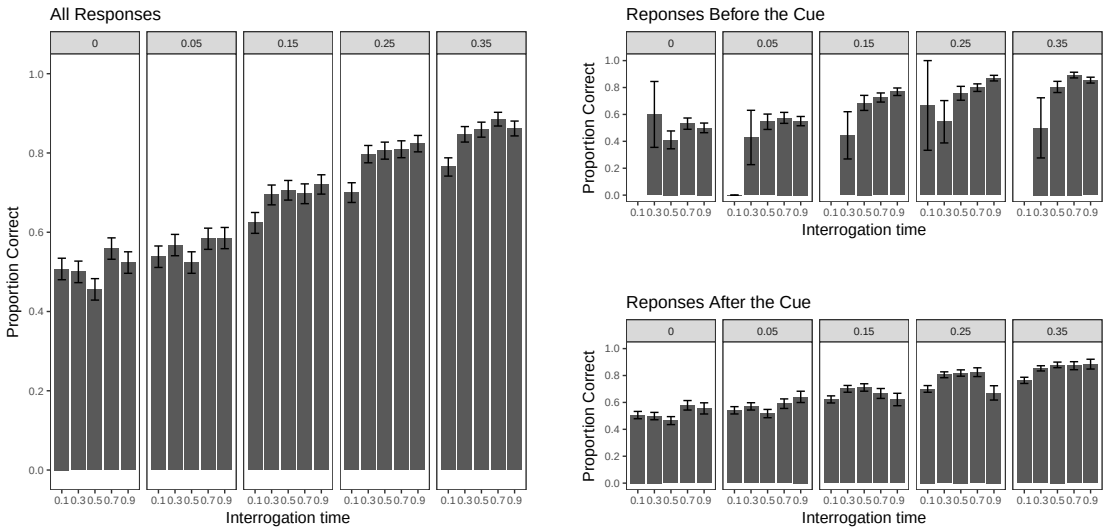

Figure 4.2: **Accuracy Across Independent Variables.** The left panel shows the accuracy (y-axis) as a bar plot for each interrogation time (x-axis), and for each coherency (panels), collapsed across response times. The right panel shows the same information as in the left panel, but separated based on whether the response was made prior to (top panel) or after (bottom panel) the stimulus disappeared. Error bars represent the standard error in the proportion correct.

the zBIC) across models for a given subject so that model performance could be assessed more easily. To do this, we subtracted each BIC value in a given column from the mean BIC value in that column and divided by the standard deviation of the BIC values within that column. The final column shows the performance of each model by averaging the BICs across subjects. Although these group values do not technically reflect how well each model fit to the entire set of data, they provide some information about the average performance across subjects. Each model performance statistic is color coded according to the legend on the right hand side; for the zBIC (and the BIC), lower values reflect better model performance, which are illustrated with cooler (i.e., bluish) colors. The left panel of Figure 4.3 provides a schematic of the various models investigated here. Parameters are represented as nodes in each column, where parameters that were freely estimated are empty, parameters that were fixed are solid, and parameters that are not applicable are shown as "x"s. From left to right, the columns are the model numbers, between-trial variability in starting point S_Z, between-trial variability in drift η, the asymptotic boundary setting a', and the scaling parameter λ. For ease of discussion, the models were grouped into four classes mentioned in Section 2.1.4: Models 1-4 (Class 1; red) are the time-invariant models, Models 5-8 (Class 2; blue) are the time-variant models with a' free and λ fixed, Models 9-12 (Class 3; green) are the time-variant models with a' fixed and λ free, and Models 13-16 (Class 4; purple) are the time-variant models with both a' and λ free.

Figure 4.3 shows that, across subjects (i.e., going down the group column in Figure 4.3), lower zBIC values are obtained with time-variant models, specifically the Class 3 models that allow λ to be free. The Class 4 models also perform well, presumably

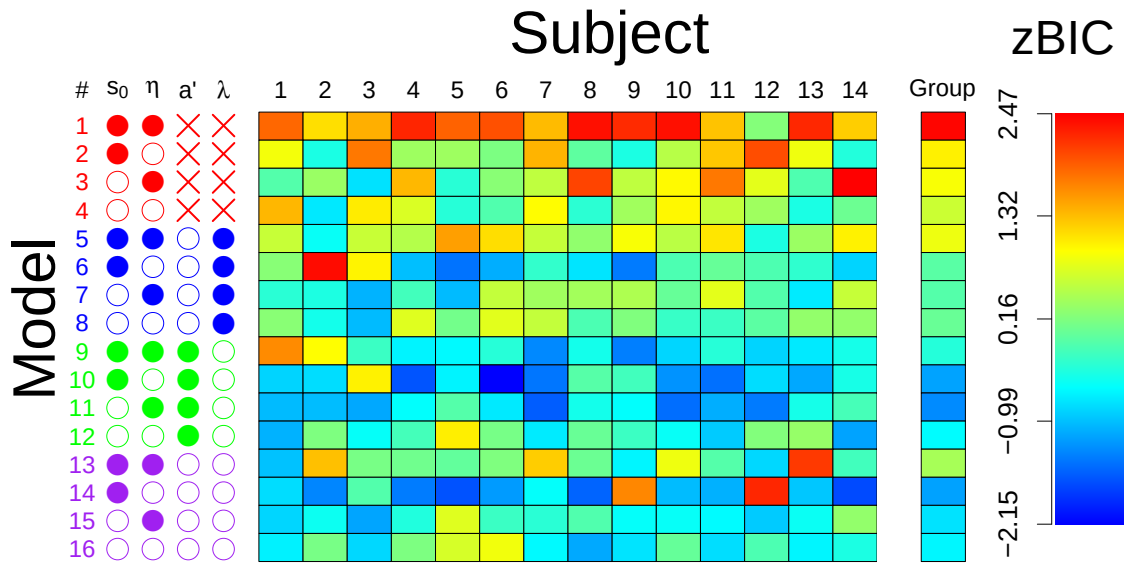

Figure 4.3: **Model Comparison.** Each box illustrates the z-scored BIC value obtained for each subject (columns) and model (rows) combination, color coded according to the legend on the right-hand side. Lower BIC values (i.e., bluish colors) correspond to better model performance. The panel on the left summarizes the models investigated here, where the columns correspond to the model number and the specific model parameters that were either fixed (filled circle), free to vary (empty circle) or not applicable (an "x" symbol). For convenience, the models were also grouped into classes: time-invariant (Models 1-4; red), time-variant models with a' free (Models 5-8; blue), time-variant models with λ free (Models 9-12; green), and time-variant models with both a' and λ free (Models 13-16; purple).

because λ is free, and the penalty incurred by freeing a' is not enough to hinder the performance of these models. Of course, for some subjects this is not the case, such as with Subject 12 and 13, who are poorly fit by Class 4 models, but are still fit well by Class 3 models. In this case, the penalty term for freeing a' was enough to rule out Class 4 models. The clear worst performing class is Class 1, with the worst performer within this class being the one that is most constrained (i.e., having only nondecision time variability).

Within a class, it is more difficult to say which parameters should be free to vary, as some inconsistencies emerge. For example, in Class 2, having between-trial variability in starting point performs about as well as having between-trial variability in drift or both free to vary. However, in Class 3, having either one of these parameters free to vary performs better than having neither one free or both free. For Class 4, having between-trial variability in drift is better than every other combination.

Tables 4.1 and 4.2 provide the model fit statistics on their original scale for each of the 16 models and each subject. The best fitting model for each subject is in boldface type. For all subjects, the best fitting model is a variant of the time-variant diffusion model with collapsing boundaries, with Class 3 (λ free) accounting for the most subjects overall. The second best class of models was Class 4, which freely estimated both a' and λ, best accounting for 5 out of 14 subjects. The remaining two subjects were best accounted for by either Class 1 (i.e., the time-invariant class) or Class 2, which freely estimated a'.

Table 4.1: **Model fit statistics for Subjects 1-7.** boldface type represents the lowest BIC value for each subject.

Model	1	2	3	4	5	6	7
1	1049.27	1042.37	1093.19	1008.76	803.54	919.73	982.01
2	953.62	928.95	1134.93	863.97	784.63	807.44	986.10
3	885.27	932.26	781.04	824.48	801.99	839.36	910.01
4	1001.95	908.22	1046.31	891.85	775.69	791.2	945.51
5	926.87	920.86	991.25	874.36	798.97	867.23	913.18
6	889.56	1134.45	1040.52	761.77	762.4	732.06	832.11
7	833.55	930.09	780.56	821.79	767.78	833.68	892.91
8	890.09	925.9	786.16	894.18	781.46	844.88	912.16
9	1028.06	1028.89	882.41	786.15	772.47	775.89	739.83
10	783.61	902.80	1041.80	**712.44**	772.06	**668.09**	729.12
11	769.1	889.06	772.27	791.82	779.11	754.18	**715.78**
12	**763.61**	972.59	841.95	822.66	793.02	805.73	793.89
13	773.29	1055.65	930.19	842.7	780.21	808.74	971.92
14	788.95	**864.94**	900.79	731.19	**759.90**	725.36	806.56
15	784.72	922.83	**769.63**	805.99	789.51	783.16	827.23
16	801.43	970.06	808.69	849.1	788.87	850.85	801.99

Table 4.2: **Model fit statistics for Subjects 8-14.** boldface type represents the lowest BIC value for each subject.

Model	8	9	10	11	12	13	14
1	775.71	991.92	724.68	992.44	501.22	915.63	769.59
2	634.18	883.91	687.46	989.68	529.26	814.55	680.70
3	767.44	959.03	706.02	947.96	499.26	730.39	758.59
4	616.46	916.72	696.79	916.11	503.74	720.85	705.31
5	652.17	938.47	687.78	970.18	491.42	777.46	757.84
6	591.90	**844.78**	674.59	857.79	496.09	730.25	655.59
7	658.51	919.95	677.78	937.88	496.51	700.61	734.95
8	627.44	908.64	672.07	829.66	497.42	774.14	718.03
9	607.86	845.82	660.56	817.11	484.81	700.03	676.48
10	631.55	893.08	652.55	**702.74**	485.65	**671.05**	677.13
11	607.87	877.35	**648.10**	741.94	**476.54**	718.87	692.81
12	637.92	891.72	666.32	759.94	501.03	775.38	639.13
13	638.97	875.00	694.13	846.25	485.23	907.59	691.22
14	**546.29**	968.94	657.48	743.69	532.9	684.89	**610.12**
15	629.47	878.16	666.24	791.53	484.00	714.01	718.07
16	570.64	870.55	677.58	772.04	495.96	705.65	678.64

4.1.3 Posterior Model Probabilities

In addition to comparing subjects using standardized BIC scores, we also compared subjects by calculating the approximate posterior model probabilities as given by Wasserman (2000). Figure 4.4 shows the resulting posterior model probability values obtained. Each model performance statistic is color coded according to the legend on the right hand side. For the approximate posterior model probability, values closer to one (warmer colors) indicate the better fit, relative to the 15 other model variants for each subject. The left panel of Figure 4.4 provides a schematic of the various models investigated here. Parameters are represented as nodes in each column, where parameters that were freely estimated are empty, parameters that were fixed are solid, and parameters that are not applicable are shown as "x"s. From left to right, the columns are the model numbers, between-trial variability in starting point S_Z, between-trial variability in drift η, the asymptotic boundary setting a', and the scaling parameter of the collapsing decision boundary λ. For ease of discussion, the models were grouped into four classes: Models 1-4 (Class 1; red) are the time-invariant models, Models 5-8 (Class 2; blue) are the time-variant models with a' free and λ fixed, Models 9-12 (Class 3; green) are the time-variant models with a' fixed and λ free, and Models 13-16 (Class 4; purple) are the time-variant models with both a' and λ free.

Figure 4.4 clearly shows that all subject's data are best explained by one or two model variants, and these model variants are some variant of the time-invariant diffusion model. Across subjects, most of the models with high posterior probabilities belong to Class 3 (λ is free), but Class 4 also performs well. These results suggest that freeing λ (Class 3) substantially improves model performance for the majority of

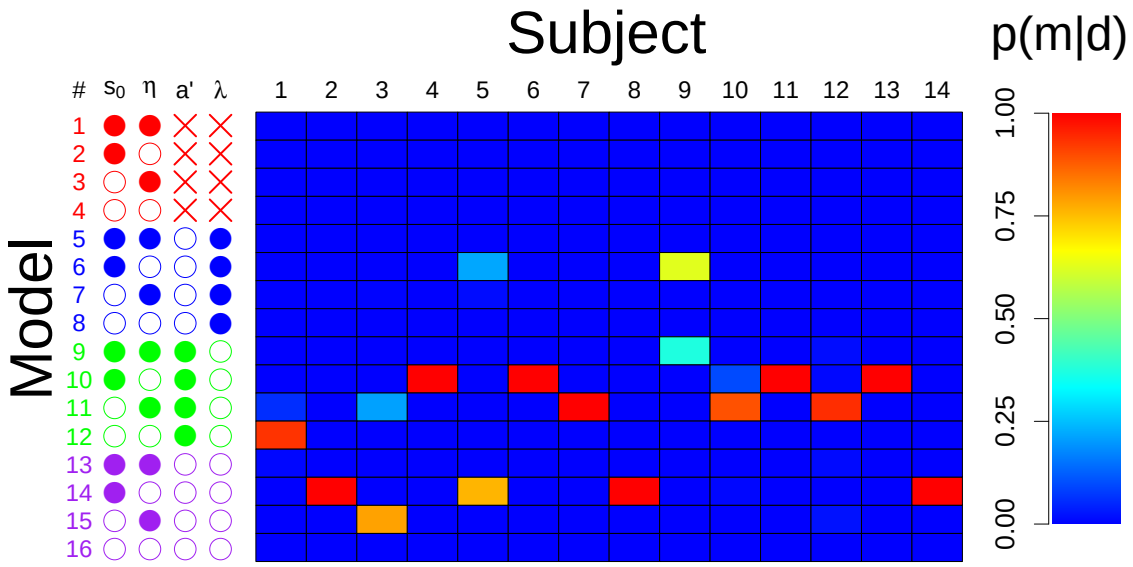

Figure 4.4: **Model Comparison.** Each box illustrates the posterior model probability obtained for each subject (columns) and model (rows) combination, color coded according to the legend on the right-hand side. Higher posterior probabilities (i.e., warmer colors) correspond to better model performance. The panel on the left summarizes the models investigated here, where the columns correspond to the model number and the specific model parameters that were either fixed (filled circle), free to vary (empty circle) or not applicable (an "x" symbol). For convenience, the models were also grouped into classes: time-invariant (Models 1-4; red), time-variant models with a' free (Models 5-8; blue), time-variant models with λ free (Models 9-12; green), and time-variant models with both a' and λ free (Models 13-16; purple).

subjects in our study. However, freeing a' in addition to λ (Class 4) improved model fit in such a way that that the penalty induced by the additional parameter was not enough to damage the performance of the model. Again, while the posterior model probability lacks the finer gradient of the standardized BIC model comparison, the clear worst performer is Class 1, with all four models showing 0% probability for all subjects.

Approximating the Bayes Factor

Once each BIC value had been obtained (see Figure 4.3), we used Equation 3.7 to approximate the Bayes factor (Kass & Raftery, 1995). Rather than comparing all 16 models, Figure 4.5 shows the Bayes factor for the best fitting collapsing bound model (model 11, which includes between-trial variability in starting point and the shape parameter λ) relative to the best fitting fixed bound model (model 4, which includes all three sources of variability). Figure 4.5 shows the Bayes factor for each subject, sorted according to the amount of evidence provided for the collapsing bound variant. The Bayes factor is shown on the log scale, so subjects near zero are best described by neither model preferentially. However, for 11 out of the 14 subjects, the Bayes factor for the collapsing bound model is larger than 100, indicating strong preference.

Table 4.3 displays the numerical Bayes factor scores for each subject. Again, 13 of 14 subjects showed a preference for the time-variant model, with the majority showing "strong" evidence for the time-variant model.

4.1.4 Inferred Task Representations

Thus far, our results suggest that all of our subjects are best accounted for by some variant of a collapsing bound model. However, model fit comparisons such as

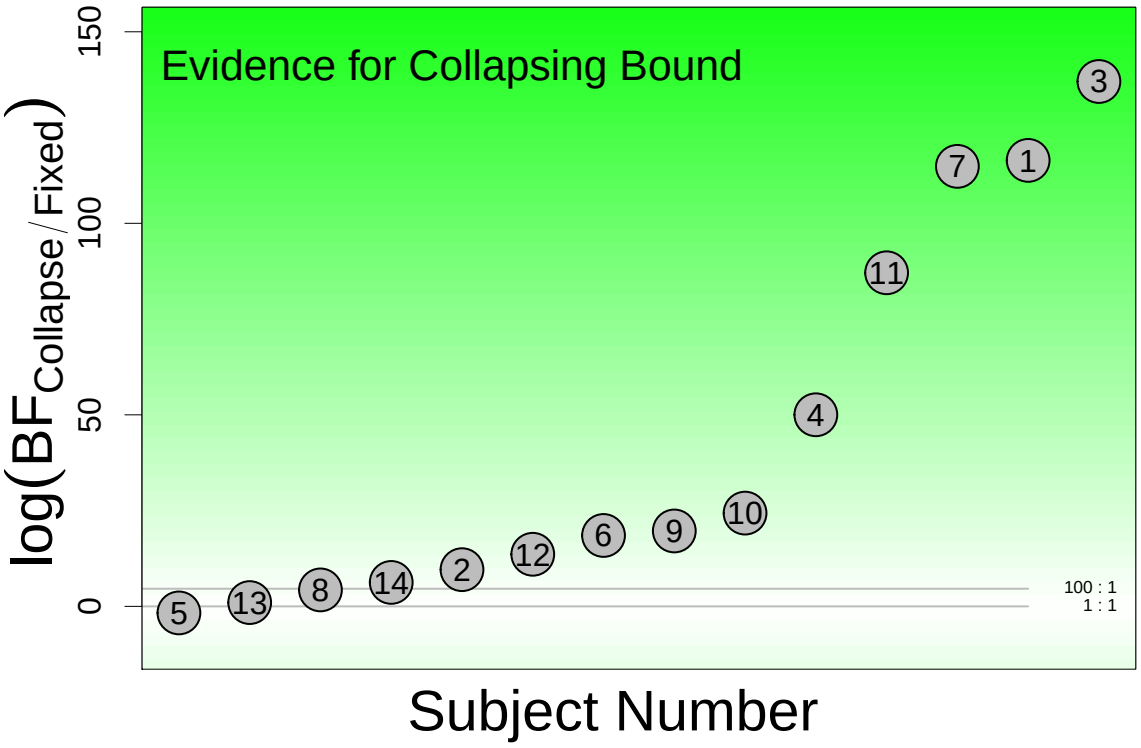

Figure 4.5: **Bayes factor Approximation by Subject.** Each point represents the Bayes factor comparing the best fitting time-invariant model to the best fitting time-variant model for each subject.

Table 4.3: **Bayes factor comparison**. Bayes factor approximations comparing the best fitting time-variant model (i.e., model 11, which includes between-trial variability in starting point and freely estimates λ) to the best fitting time-invariant model (i.e., model 4, where both sources of trial-to-trial variability are included) for each of the subjects in our data.

Subject	Log(Bayes Factor)	Strength of Evidence	Model
1	116.43	Strong	Time-variant
2	9.58	Strong	Time-variant
3	137.02	Strong	Time-variant
4	50.02	Strong	Time-variant
5	-1.71	Strong	Time-invariant
6	18.51	Strong	Time-variant
7	114.87	Strong	Time-variant
8	4.3	Strong	Time-variant
9	19.69	Strong	Time-variant
10	24.35	Strong	Time-variant
11	57.09	Strong	Time-variant
12	13.6	Strong	Time-variant
13	0.99	Moderate	Time-variant
14	6.25	Strong	Time-variant

the ones presented in Figure 4.3 and Figure 4.5 do not tell us much about why one model architecture provides a better fit over another. To investigate this, we can look to the task representations that might have been used by each subject during the task. Examining the inferred task representations has the advantage of projecting several parameter dimensions down into something that is easily visualized, so that we can gain clearer insight into how the various parameters interact with one another.

To illustrate the differences in model fit, Figure 4.6 shows two subjects who obtained two different results in the model fitting comparison. The left column shows Subject 6, whose data were best accounted for by the "modern" time-invariant model (model 4; includes both sources of between-trial variability), whereas the right column shows Subject 8, whose data were best accounted for by its time-variant derivative (model 16; includes both sources of between-trial variability and freely estimates a' and λ). The top row shows the empirical response time distributions, collapsed across choice and condition information. The middle row shows the inferred task representation used by a fixed bound model, whereas the bottom row shows the inferred task representations used by time-variant model. These subjects were chosen because (1) they yielded different model fitting results, and (2) their general pattern of choice response time distributions were similar.

To generate the task representations, we simply recreated plots such as the one presented in Figure 2.3 by using the MAP estimates obtained during the model fitting process. The relevant MAP estimates for Figure 4.6 are the threshold, drift rates, nondecision times, and collapsing bound parameters (i.e., for the bottom row). For the drift rates, we plotted the average trajectory that was inferred from fitting the model to data for each stimulus coherency condition. For Subject 6, Figure 4.6 shows

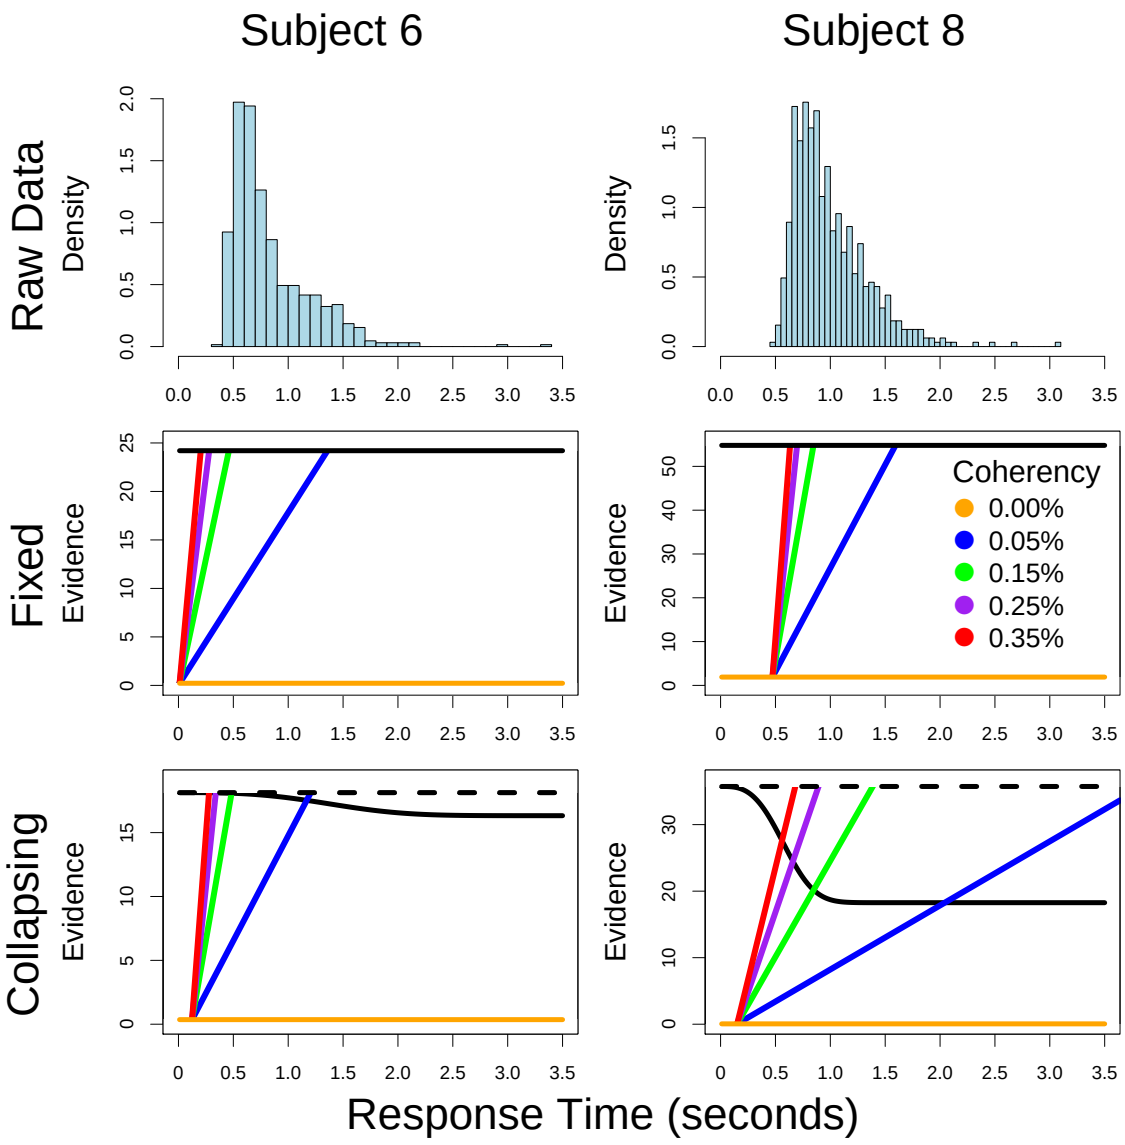

Figure 4.6: **Inferred Task Representations.** The columns correspond to two subjects whose data were either fit best by a fixed boundary model (left column), or a collapsing boundary model (right column). The top row shows the empirical response time distribution collapsed across choice and stimulus condition, the middle row shows the task representation inferred by a fixed boundary model, whereas the bottom row shows the task representation inferred by a collapsing boundary model. The drift rates for each coherency condition are shown, color coded according to the key in the middle right panel.

that for the time-variant model, the inferred boundary does not collapse until late in the decision-making process (i.e., after most of the drift rate trajectories have crossed the boundary), and the inferred collapse in the bottom panel is minimal, suggesting that the collapse played no role in the model's ability to account for the data of this subject. As such, the time-invariant model provided a better fit once balancing for the number of parameters.

The right panel of Figure 4.6 shows the inferred representations for Subject 8, who was best captured by the time-variant model. In contrast to the left panel, the inferred boundary collapses completely and early in the decision making process, causing the drift rates to terminate at the collapsing boundary rather than the fixed boundary. For this subject, the interaction of the drift rates and the stage of the collapse allow the collapsing bound model to account for the tail end of the response time distribution slightly better. Across all subjects, the pattern of drift rates interacting with the collapsing bound parameters was a strong predictor in determining which model would ultimately provide the best fit. Namely, when the drift rates of one or more coherency condition hit a boundary during a collapse, the collapsing boundary provided the better fit.

Inferred task representation plots for the remaining 12 subjects are displayed below. In these plots, the best fitting model is indicated by the color of the response time distribution in the top row of each figure. If the distribution is red, then the subject's data is better fit by the time-invariant model (model 4). If the distribution is blue, then the subject's data is better fit by the time-variant model (model 16).

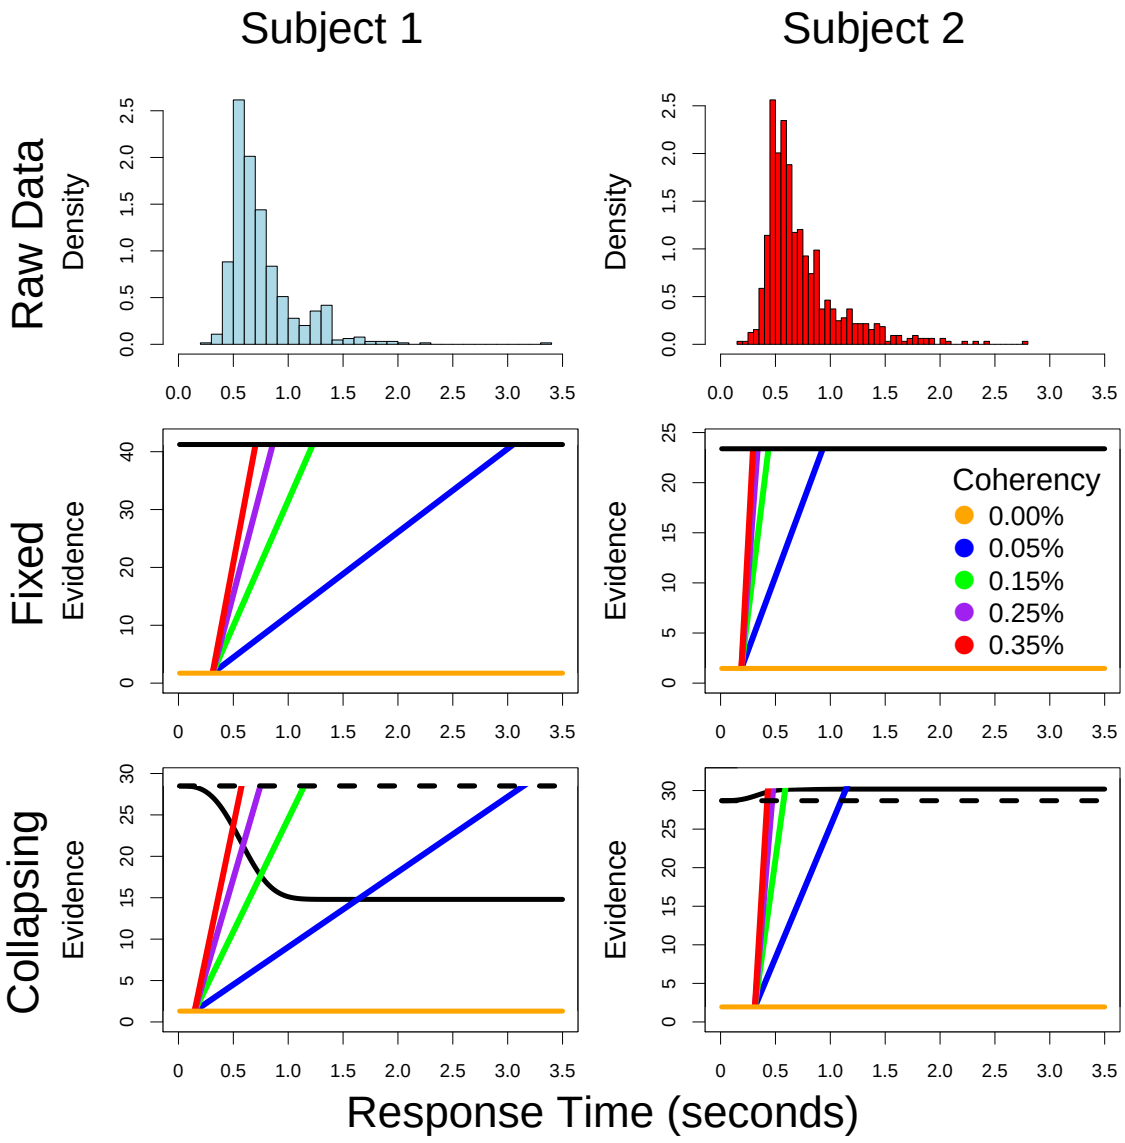

Figure 4.7: **Inferred Task Representations for subjects 1 and 2.** The columns correspond to two subjects' data. The top row shows the empirical response time distribution collapsed across choice and stimulus conditions. If the response time distribution is red, the subjects' data was best fit by the time-invariant model (model 4). If the response time distribution is light blue, the subjects data was best fit by the time-variant model (model 16). The middle row shows the task representation inferred by a fixed boundary model, whereas the bottom row shows the task representation inferred by a collapsing boundary model. The drift rates for each coherency condition are shown, color coded according to the key in the middle right panel.

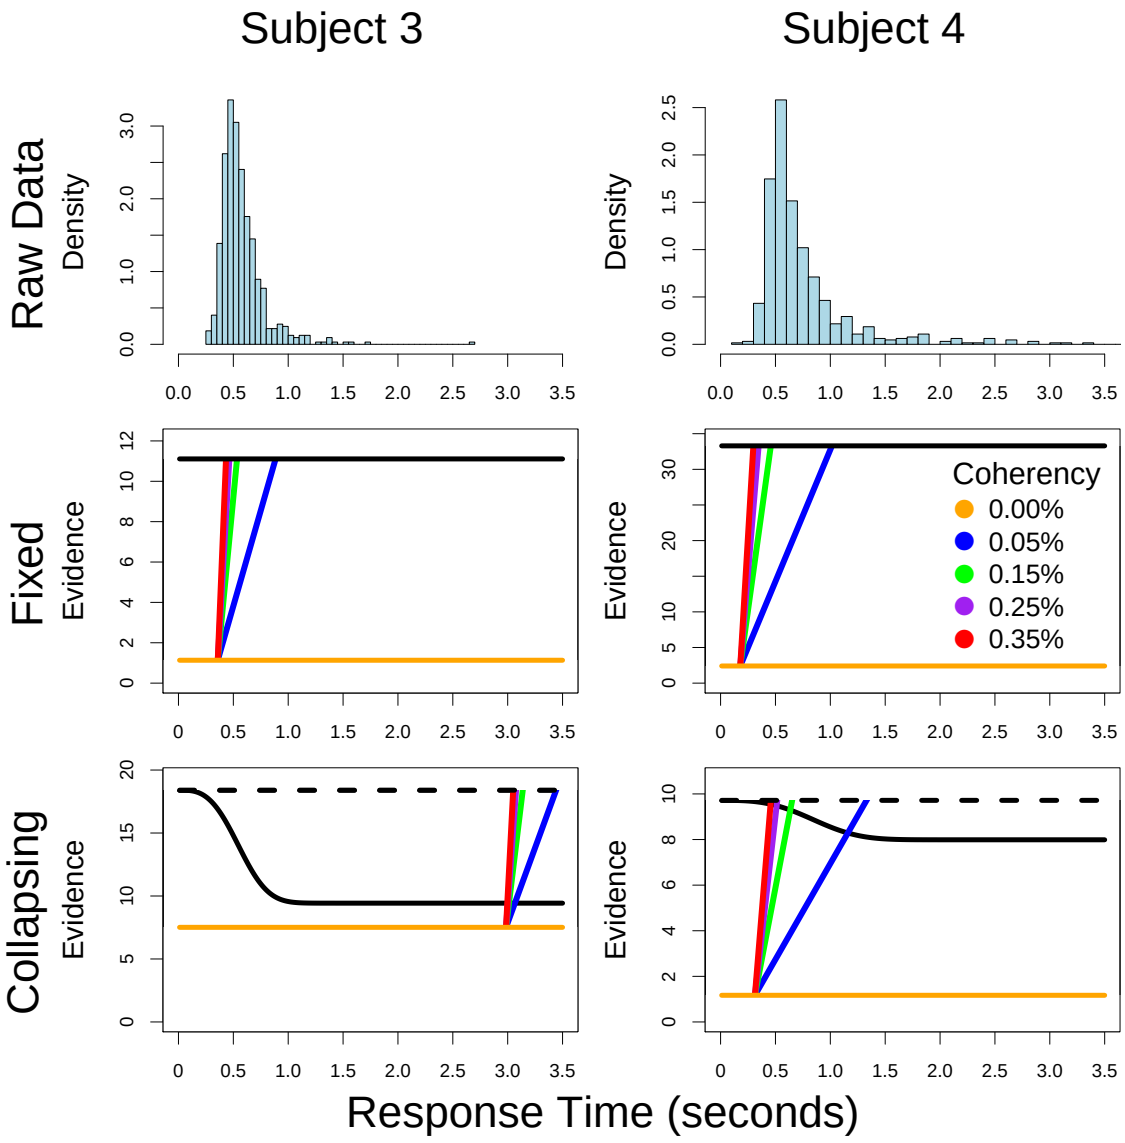

Figure 4.8: **Inferred Task Representations for Subjects 3 and 4.** The columns correspond to two subjects' data. The top row shows the empirical response time distribution collapsed across choice and stimulus conditions. If the response time distribution is red, the subjects' data was best fit by the time-invariant model (model 4). If the response time distribution is light blue, the subjects data was best fit by the time-variant model (model 16). The middle row shows the task representation inferred by a fixed boundary model, whereas the bottom row shows the task representation inferred by a collapsing boundary model. The drift rates for each coherency condition are shown, color coded according to the key in the middle right panel.

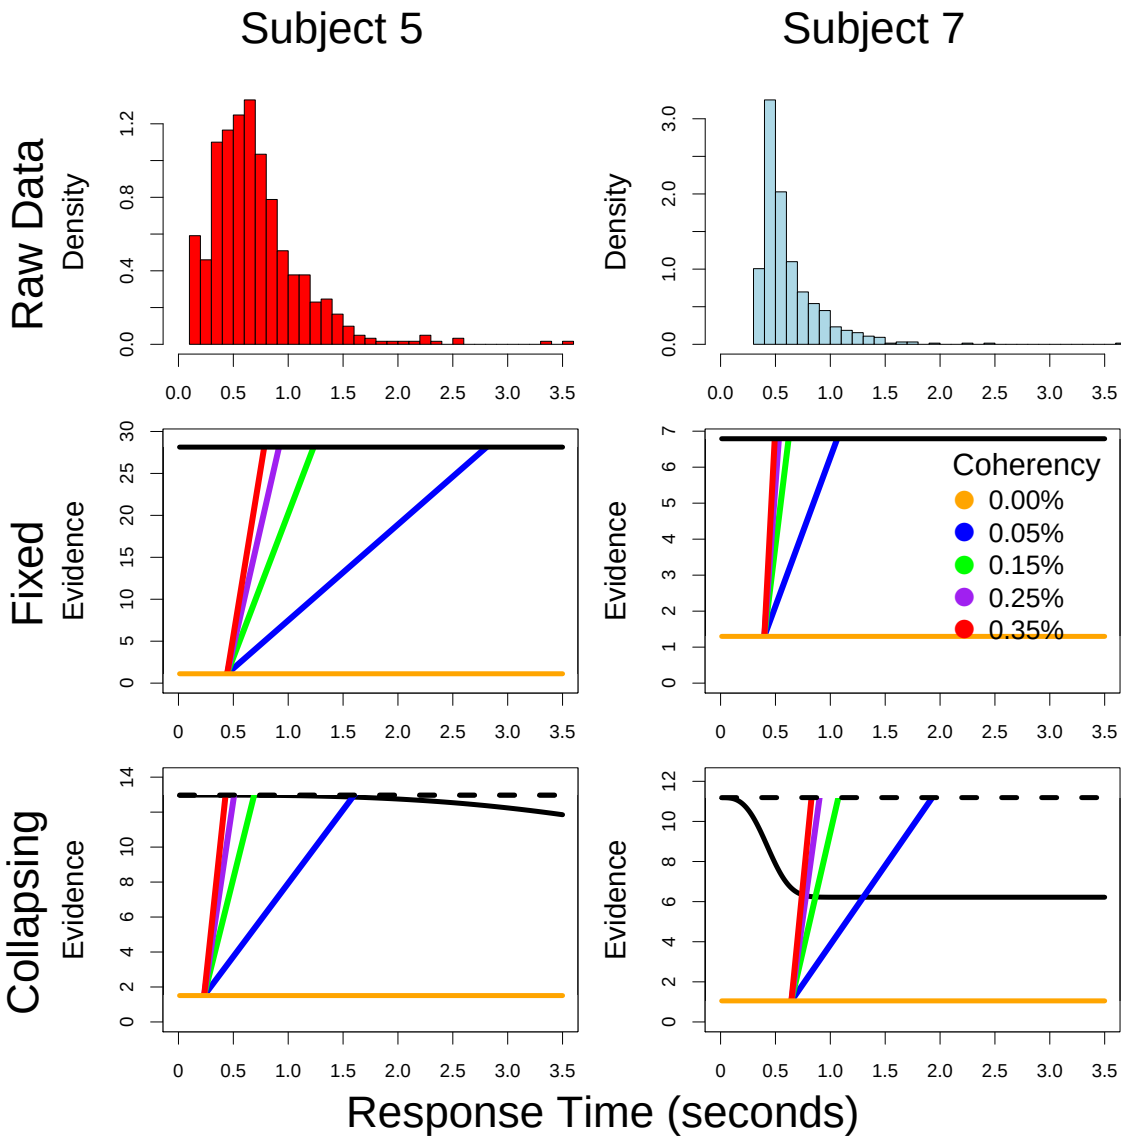

Figure 4.9: **Inferred Task Representations for Subjects 5 and 7.** The columns correspond to two subjects' data. The top row shows the empirical response time distribution collapsed across choice and stimulus conditions. If the response time distribution is red, the subjects' data was best fit by the time-invariant model (model 4). If the response time distribution is light blue, the subjects data was best fit by the time-variant model (model 16). The middle row shows the task representation inferred by a fixed boundary model, whereas the bottom row shows the task representation inferred by a collapsing boundary model. The drift rates for each coherency condition are shown, color coded according to the key in the middle right panel.

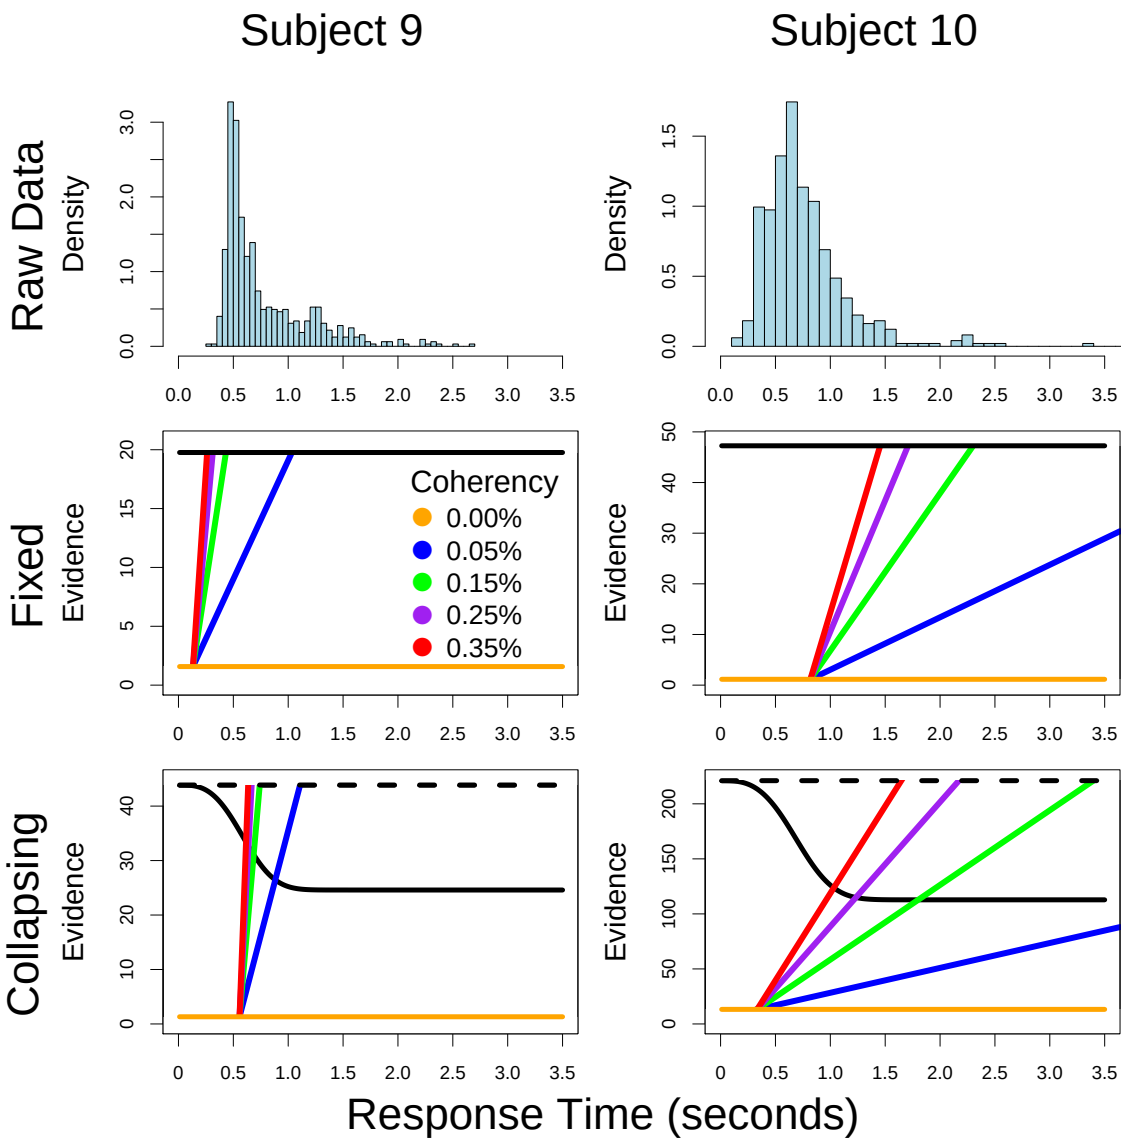

Figure 4.10: **Inferred Task Representations for Subjects 9 and 10.** The columns correspond to two subjects' data. The top row shows the empirical response time distribution collapsed across choice and stimulus conditions. If the response time distribution is red, the subjects' data was best fit by the time-invariant model (model 4). If the response time distribution is light blue, the subjects data was best fit by the time-variant model (model 16). The middle row shows the task representation inferred by a fixed boundary model, whereas the bottom row shows the task representation inferred by a collapsing boundary model. The drift rates for each coherency condition are shown, color coded according to the key in the middle right panel.

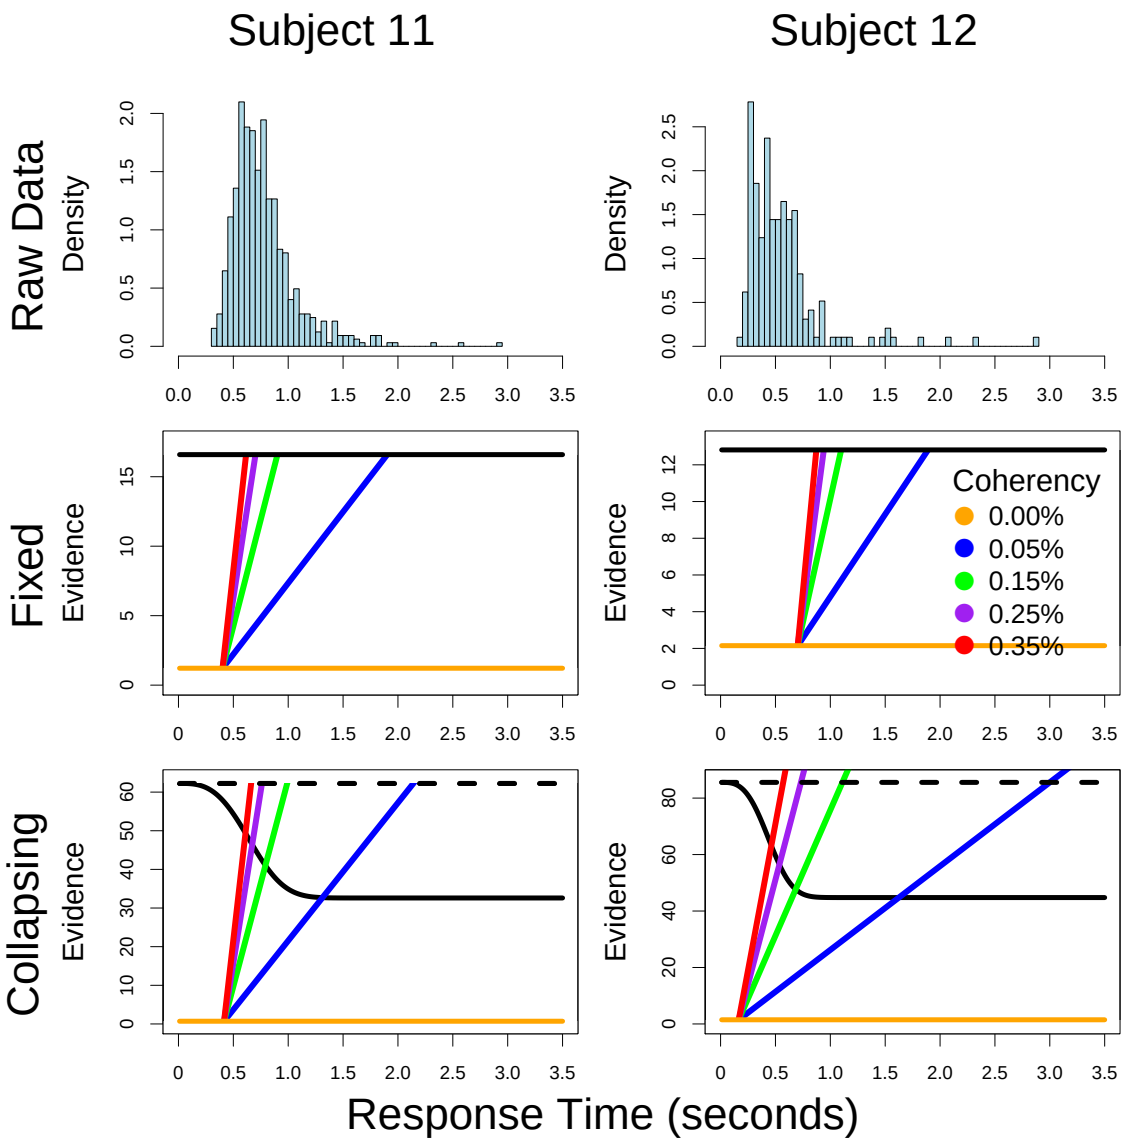

Figure 4.11: **Inferred Task Representations for Subjects 11 and 12.** The columns correspond to two subjects' data. The top row shows the empirical response time distribution collapsed across choice and stimulus conditions. If the response time distribution is red, the subjects' data was best fit by the time-invariant model (model 4). If the response time distribution is light blue, the subjects data was best fit by the time-variant model (model 16). The middle row shows the task representation inferred by a fixed boundary model, whereas the bottom row shows the task representation inferred by a collapsing boundary model. The drift rates for each coherency condition are shown, color coded according to the key in the middle right panel.

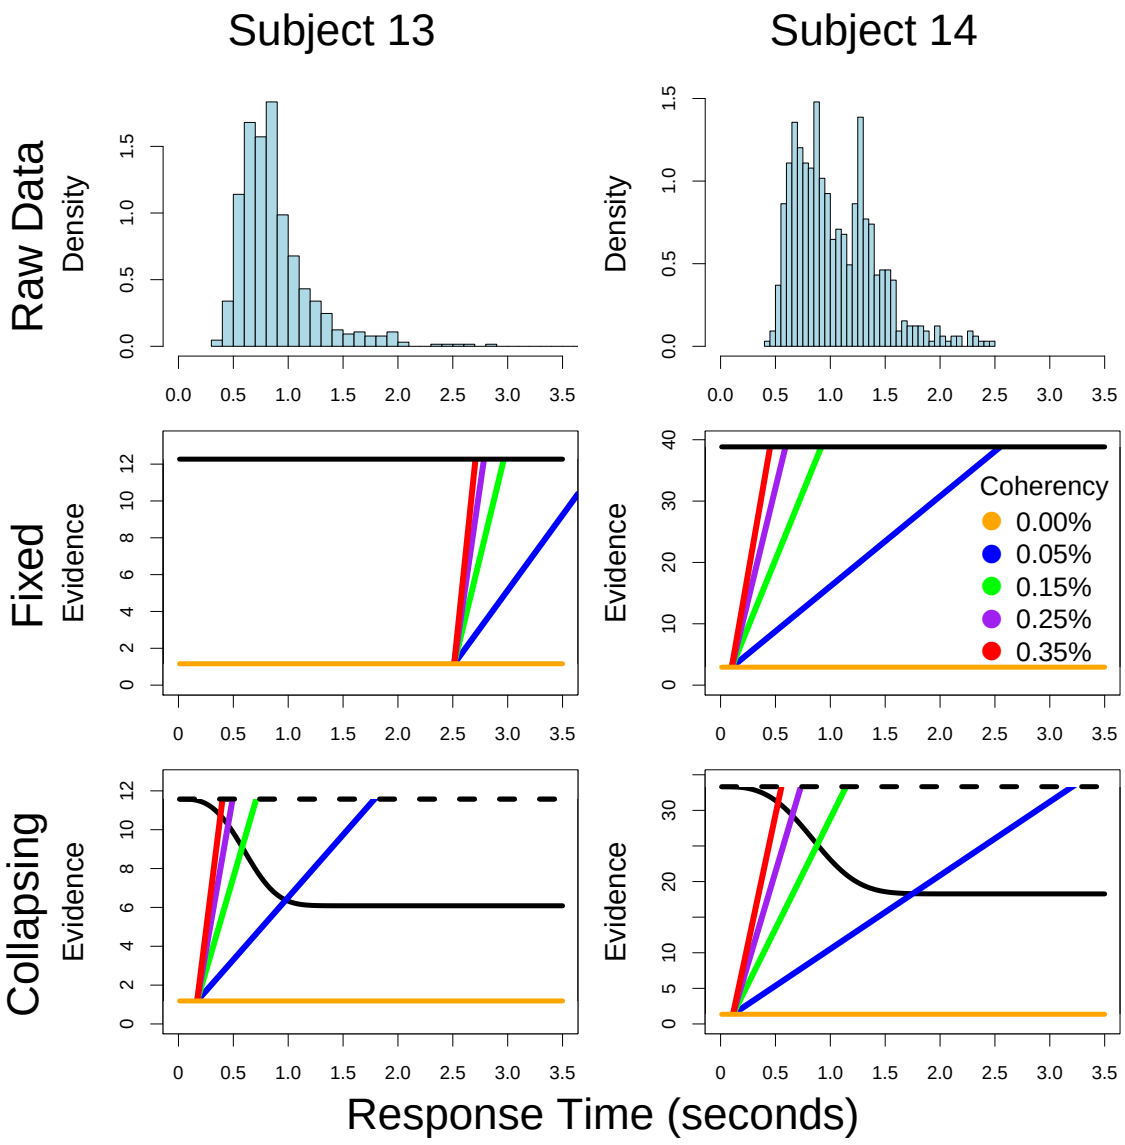

Figure 4.12: **Inferred Task Representations for Subjects 13 and 14.** The columns correspond to two subjects' data. The top row shows the empirical response time distribution collapsed across choice and stimulus conditions. If the response time distribution is red, the subjects' data was best fit by the time-invariant model (model 4). If the response time distribution is light blue, the subjects data was best fit by the time-variant model (model 16). The middle row shows the task representation inferred by a fixed boundary model, whereas the bottom row shows the task representation inferred by a collapsing boundary model. The drift rates for each coherency condition are shown, color coded according to the key in the middle right panel.

Chapter 5: Discussion

5.1 Discussion

The standard DDM with time-invariant bounds has proven exemplary in accounting for a variety of empirical data across a number of domains (see Forstmann, Ratcliff, & Wagenmakers, 2015, for a review). Recently, the standard DDM has come under scrutiny with the focal point being the manner in which the model deals with the timing of the choice. In its place, diffusion models with additional mechanisms, such as urgency or collapsing boundaries have been proposed, and these models have proven effective in accounting for both human and primate data from similar perceptual decision-making tasks (Ditterich, 2006a, 2006b; Kiani & Shadlen, 2009). While there is considerable controversy surrounding the nature of the task demands, stimulus used, and even the species of the subject under study, if these mechanisms are proven generally effective in capturing data, they provide an interesting and exciting alternative to a long standing tradition in models instantiating sequential sampling theory.

The purpose of this article was to further investigate how stimulus characteristics interact with task demands. While other reports have been criticized for using stimuli that may inherently bias subjects toward a collapsing bound policy, our task centers

on one of the most well-studied stimuli in perceptual decision making history: the random dot motion task. Starting from a well-agreed upon stimulus, we asked whether factorially manipulating the strength of coherent motion as well as the length of the stimulus exposure could induce a collapsing boundary policy.

To test this hypothesis, we compared the relative fit of 16 diffusion model variants with fixed and collapsing decision thresholds to data from our modified random dot motion task. We compared each model's ability to fit data from the mixed RDM task by calculating three different model fit statistics: the Bayesian information criterion (Schwarz, 1978), approximated posterior model probabilities (Wasserman, 2000), and an approximation to the Bayes factor (Kass & Raftery, 1995). The BIC measures suggested that a variant with collapsing decision boundaries was preferred across all 14 subjects. Collapsing across subjects, we found that time-variant models that freely estimate the scaling parameter λ in the collapsing decision boundary provide the best fits overall. The approximate posterior model probabilities showed similar results, with time-variant models that freely estimate the scaling parameter λ having higher posterior model probabilities across all subjects. The Bayes factor analyses further confirmed this point in that they provided definitive and consistent evidence for time-variant models. Taken together, these results suggest that the mechanisms in the time-variant model provide a more viable explanation for the strategies subjects use to complete our mixed RDM task.

5.1.1 Why Collapse?

Given that we find evidence supporting collapsing boundaries, it seems productive to speculate on the reason for adopting a collapse for our particular task. As other

researchers have compared models using time-variant and time-invariant mechanisms (e.g., Ditterich, 2006a; Hawkins, Forstmann, et al., 2015; Voskuilen et al., 2016), we can follow some of their rationales in interpreting our results. This section discusses three such rationales: species differences, stimuli differences, and task differences.

Species Differences

Much of the support for time-variant models comes from neurophysiological literature exploring the decision-making behavior of primates (e.g., Chuchland, Kiani, & Shadlen, 2008; Hanks, Mazurek, Kiana, Hopp, & Shadlen, 2011) and behavioral literature comparing the relative fits of time-variant and time-invariant models to primate data (e.g., Ditterich, 2006a). In the most broad comparison to date, Hawkins, Forstmann, et al. (2015) tested whether human and primate data was best captured by time-variant or time-invariant models and found that the majority of the primate data sets were best accounted for by time-variant models. However, when applying the same set of models to human data, Hawkins, Forstmann, et al. (2015) found that the majority of human data sets were best accounted for by time-invariant models.

One of the major reasons for the differences in results across species was the effect of practice. Hawkins, Forstmann, et al. (2015) explained that non-human primates are subject to many practice sessions, often on the order of months or years. Due to this extensive practice, it is of course reasonable that the distributions of response times change. The field of automaticity articulates the effect of practice on the shape of a response time distribution via the so-called "power law of practice" (Logan, 1988). Here, a common empirical result is that the mean and the standard deviation of the response time distribution decrease with the number of practice sessions according to a power function. If non-human primates are severely over practiced, then it should

follow that the shape of the response time distributions would appear more Gaussian, according to the power law. The practice effect explanation seems to hold for non-human primates and is likely the reason Hawkins, Forstmann, et al. (2015) found such compelling evidence for collapsing boundaries in this particular species, but not humans.

While the shapes of our subject's response time distributions do not appear Gaussian (see Figure 4.6 as an example), an analysis of the mean and standard deviation of the response times across each experimental block in our task reveals they decrease according to a power function, suggesting that practice did have an effect on performance. Figure 5.1 illustrates the differences in fit between a linear function and a power function for the mean (left panel) and the standard deviation (right panel) of response times across each of the 13 experimental blocks. Due to the large difference in response times between the first and second block, both the mean and standard deviations appear to be best fit by a power function ($BIC_{mean} = 132.70$; $BIC_{std} = 122.33$) versus a linear function ($BIC_{mean} = 142.66$; $BIC_{std} = 131.10$).

These analyses suggest that practice effects clearly exist in our data, which implies that subjects were not overly practiced on the task. As one of the arguments for the particular shapes of response time distributions was that the non-human primates were overly practiced, the fact that our subjects were not implies that the evidence we found for collapsing bound models is not explainable through only the result of practice. Because our human subjects were not as practiced as the non-human primates investigated in Hawkins, Forstmann, et al. (2015), this rationale is not applicable to the data presented here.

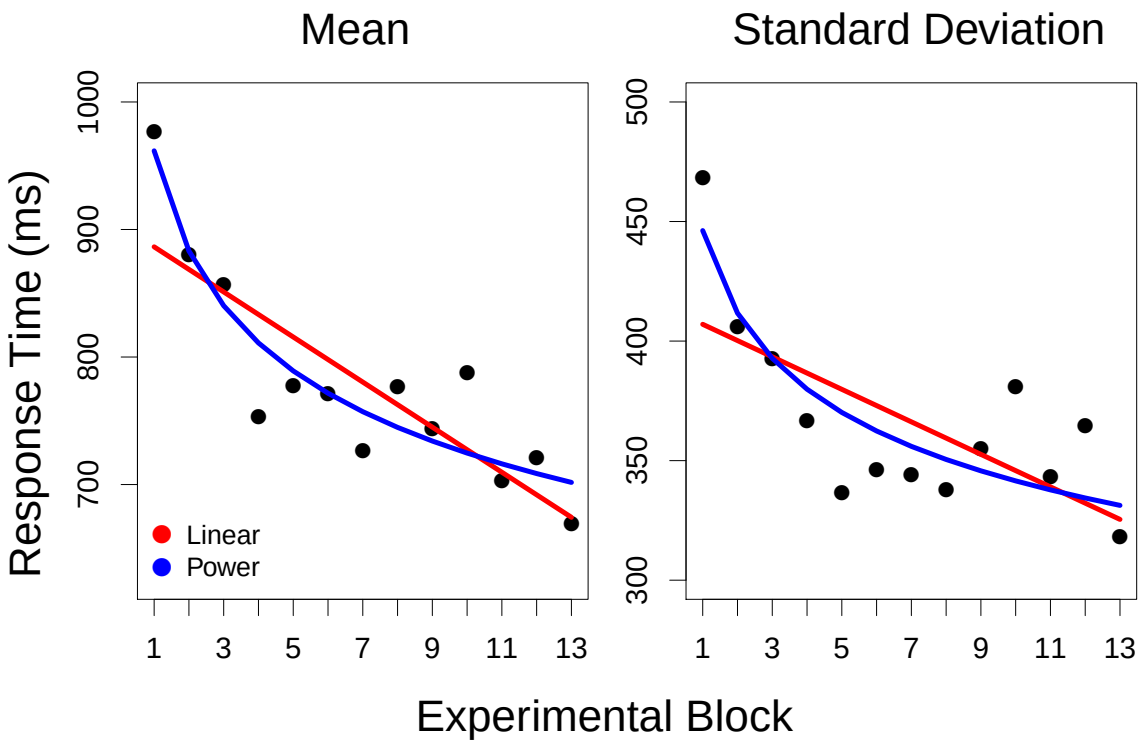

Figure 5.1: **Analysis of Practice Effects.** The left panel shows the best fitting linear (red) and power (blue) function to the mean response time across blocks. Similarly, the right panel shows the best fitting linear (red) and power (blue) function to the standard deviation of the response times across blocks.

Stimuli Differences

Another potential explanation for the adaption of a collapsing boundary focuses on the properties of the stimuli used in decision-making tasks. Voskuilen et al. (2016) point out that many of the studies that conclude time-variant models provide a better explanation of decision-making behavior have analyzed data from expanded judgement tasks (e.g., Busemeyer & Rapoport, 1988; Drugowitsch et al., 2012; van Maanen et al., 2016), where the rate of evidence accumulation across time is controlled by the properties of the stimuli. The idea is that as the rate of evidence accumulation decreases and the average decision time increases, subjects are more willing to base their decisions on less evidence at later times (i.e., they feel a sense of urgency to respond). Thus, one tenable argument against evidence for collapsing boundaries is that the properties of the stimuli used in these expanded judgement tasks induce a bias toward a time-variant policy, severely limiting this mechanism's generalizability.

By and large, previous literature does support the notion that the properties of a stimulus can affect a subject's decision threshold. Namely, time-invariant models appear to better explain data from speeded judgement tasks, which use simple stimuli to induce a fast perceptual judgement, while time-variant models appear to be more suitable for tasks using dynamic stimuli that manipulate the rate of evidence accumulation and the average decision time (e.g., expanded judgement tasks). While, at present, it is unclear whether this division in the literature is due to actual differences in strategy or can be explained by the novelty of time-variant models and their lack of application to speeded judgement data, there is currently little evidence to suggest that a boundary collapse is necessary for speeded judgments. However, given we find evidence that subjects collapse their decision boundaries in our task, which featured

stimuli more closely aligned with traditional speeded judgement tasks, it seems unlikely that the boundary policy adopted by our subjects can be explained solely by the properties of the stimuli we used.

Task Differences

Hawkins, Forstmann, et al. (2015) argued that the mechanisms underlying decision-making behavior may vary as a function of task constraint. In their study, they found that two main task characteristics - the administration of reward and the level of practice - served as strong predictors for which model (time-variant versus time-invariant) best described a subject's data. Regarding reward, in tasks where subjects are required to withhold their response for a certain amount of time, researchers find that it is optimal to adapt a collapsing decision threshold to maximize reward (e.g., Busemeyer & Rapoport, 1988; Ditterich, 2006a; Drugowitsch et al., 2012; Malhotra et al., 2017; Thura et al., 2012). Regarding practice, subjects who complete extensive training sessions prior to commencing a task are often better characterized by time-variant models, while subjects with fewer or no practice sessions are better characterized by time-invariant models, suggesting that collapsing thresholds may result from high levels of practice (Balci et al., 2011; Hawkins, Forstmann, et al., 2015; Starns & Ratcliff, 2010).

While these two task constraints served as a viable explanation for the results of Hawkins, Forstmann, et al. (2015), they are not applicable to the current study. As subjects received a flat reward at the end of the task that was not contingent on performance and only completed one short practice block, neither reward nor extensive practice can explain the evidence of the collapsing boundary in our results.

However, we do believe that the task demands have induced the need for a collapsing boundary policy in our task. Namely, the coupling of the interrogation and free response paradigms may have induced subjects to attend to the evidence at earlier time points more than at later ones, because on some trials, the stimulus was very brief. It seems plausible that the manipulation of stimulus duration induced a cost on integration time, as integrating for longer periods of time was associated with the risk of the stimulus being taken away. Furthermore, the particular sequence of coherencies we investigated here may have also contributed to the time-variant decision policy according to the optimality analyses presented in Malhotra et al. (2017). In short, subjects learn that some trials are difficult and some are easy. Hence, it seems reasonable to engage in different approaches to solving these different types of problems. On easy trials, because it is easier to engage in the integration process, subjects may use a time-invariant boundary. However, on hard trials, once the subject realizes that the motion signal is very weak, the policy might change toward collapsing their boundary so that they can move on to the next trial in order to reduce the overall amount of energy expended. Such a behavioral dynamic has already been observed in the context of multi-alternative choice (Hawkins, Brown, Steyvers, & Wagenmakers, 2012), and so we feel that this is the best explanation for why subjects adapt a collapsing decision boundary in our mixed RDM task.

5.2 Conclusions

This article contributes to the growing body of literature on the nature of decision processes with respect to time. The purpose of this article was to test whether manipulating the coherency of motion in a speeded judgment task and the duration

of stimulus exposure could induce a time-varying decision policy. We found definitive evidence supporting time-variant models for all 14 of our subjects, leading us to conclude that this particular experimental manipulation was successful in inducing this decision bias. As we have discussed, explanations such as the type of stimulus, degree of practice, and reward rate maximization do not seem to be applicable here. Instead, we suspect that the coupling of interrogation and free response paradigms is the most likely explanation of the evidence we have found for a time-varying decision policy.

References

Afraz, S. R., Kiani, R., & Esteky, H. (2006). Microstimulation of inferotemporal cortex influences face categorization. *Nature*, *442*, 692–695.

Balci, F., Simen, P., Niyogi, R., Saxe, A., Hughes, J. A., Holmes, P., & et al. (2011). Acquisition of decision making criteria: Reward rate ultimately beats accuracy. *Attention, Perception and Psychophysics*, *73*, 640–657.

Bowman, N. E., Kording, K. P., & Gottfried, J. A. (2015). Temporal integration of olfactory perceptual evidence in human orbitofrontal cortex. *Neuron*, *75*, 916–927.

Britten, W. T., K. H.and Newsome, Shadlen, M. N., Celebrini, S., & Movshon, J. A. (1996). relationship between behavioral choice and the visual responses of neurons in macaque MT. *Visual Neuroscience*, *13*, 87–100.

Busemeyer, J., & Rapoport, A. (1988). Psychological models of deferred decision making. *Journal of Mathematical Psychology*, *109*, 91–134.

Chuchland, A. K., Kiani, R., & Shadlen, M. N. (2008). Decision-making with multiple alternatives. *Nature Neuroscience*, *11*, 693–702.

Churchland, A., Kiani, R., Chaudhuri, R., Wang, X., Pouget, A., & Shadlen, M. (2011). Variance as a signature of neural computations during decision making. *Neuron*, *69*, 818–831.

Cisek, P., Puskas, G. A., & El-Murr, S. (2009). Decisions in changing conditions: the urgency-gating model. *J Neurosci*, *29*, 11560–11517.

Cook, E. P., & Maunsell, J. (2002). Dynamics of neuronal responses in macaque MT and VIP during motion detection. *Nature Neuroscience*, *24*, 985–994.

Diederich, A., & Oswald, P. (2014). Sequential Sampling Model for Multiattribute Choice Alternatives with Random Attention Time and Processing Order. *frontiers in Human Neuroscience*, *8*, 119–135.

Ditterich, J. (2006a). Evidence of time-variant decision making. *European Journal of Neuroscience*, *24*, 3628-3641.

Ditterich, J. (2006b). Stochastic models of decisions about motion direction: Behavior and physiology. *Neural Networks*, *19*, 981-1012.

Drugowitsch, J., Moreno-Bote, R., Churchland, A. K., Shadlen, M. N., & Pouget, A. (2012). The cost of accumulating evidence in perceptual decision making. *Journal of Neuroscience*, *32*, 3612–3628.

Feller, W. (1968). *An introduction to probability theory and its applications* (Vol. 1). John Wiley: New York.

Forstmann, B. U., Ratcliff, R., & Wagenmakers, E.-J. (2015). Sequential Sampling Models in Cognitive Neuroscience: Advantages, Applications, and Extensions. *Annual Review of Psychology*, *67*, 641–666.

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

Garrett, H. E. (1922). A study of the relation of accuracy to speed. *Archives of Psychology*, *56*, 1–105.

Gelman, A., Carlin, J. B., Stern, H. S., & Rubin, D. B. (2004). *Bayesian Data Analysis*. New York, NY: Chapman and Hall.

Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Reviews in Neuroscience*, *30*, 535–574.

Hampton, R. (2001). Rhesus monkeys know when they remember. *Proceedings of the National Academy of Sciences*, *98*, 5359–5362.

Hanks, T. D., Mazurek, M. E., Kiana, R., Hopp, E., & Shadlen, M. N. (2011). Elapsed decision time affects the weighting of prior probability in a perceptual decision. *Journal of Neuroscience*, 6339–6351.

Hawkins, G. E., Brown, S. D., Steyvers, M., & Wagenmakers, E. J. (2012). An optimal adjustment procedure to minimize experiment time in decisions with multiple alternatives. *Psychonomic Bulletin $ Review*, *19*, 339–348.

Hawkins, G. E., Forstmann, B. U., Wagenmakers, E. J., Ratcliff, R., & Brown, S. D. (2015). Revisiting the evidence for collapsing boundaries and urgency signals in perceptual decision-making. *The Journal of Neuroscience*, *35*, 2476-2484.

Hawkins, G. E., Wagenmakers, E. J., Ratcliff, R., & Brown, S. D. (2015). Discriminating evidence accumulation from urgency signals in speeded decision making. *Journal of Neurophysiology*, *114*, 40–47.

Heathcote, A., Brown, S. D., & Mewhort, D. J. (2002). Quantile maximum likelihood estimation of response time distributions. *Psychonomic Bulletin and Review*, *9*, 394–401.

Hernandex, A., Zainos, A., & Romo, R. (2002). Temporal evolution of a decision-making process in medial premotor cortex. *Neuron*, *33*, 959–972.

Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, *90*, 773–795.

Kelly, S. P., & O'Connell, R. G. (2013). Internal and external influences on the rate of sensory evidence accumulation in the human brain. *Journal of Neuroscience*, *33*, 219434–19441.

Kiani, R., & Shadlen, M. N. (2009). Representations of confidence associated with a decision by neurons in the Parietal Cortex. *Science*, *324*, 3017–3029.

Laming, D. R. (1968). *Information theory of choice reaction time*. New York, NY: Wiley Press.

Leite, F. P., & Ratcliff, R. (2010). Modeling reaction time and accuracy of multiple-alternative decisions. *Attention, Perception, & Psychophysics*, *17*, 246–273.

Liu, C. C., & Aitkin, M. (2008). Bayes factors: Prior sensitivity and model generalizability. *Journal of Mathematical Psychology*, *52*, 362–375.

Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review*, *95*, 492–527.

Luna, R., Hernndez, A., Brody, C. D., & Romo, R. (2005). Neural codes for perceptual discrimination in primary somatosensory cortex. *Nature neuroscience*, *8*, 1210–1219.

Malhotra, G., Leslie, D. S., H., L. C. J., & Bogacz, R. (2017). *Time-varying decision boundaries: insights from optimality analysis.* (https://doi.org/10.3758/s13423-017-1340-6)

Mazurek, M. E., Roitman, J. D., Ditterich, J., & Shadlen, M. N. (2003). A role for neural integrators in perceptual decision making. *Cerebral Cortex*, *13*, 1257–1269.

Moran, R. (2015). Optimal decision making in heterogeneous and biased environment . *Psychonomic Bulletin & Review*, *22*, 38–53.

Moran, R., Teogorescu, A., & Usher, M. (2015). Post choice information integration as a causal determinant of confidence: novel data and a computational account. *Cognitive Psychology*, *78*, 99–147.

Myung, I. J., & Pitt, M. (2002). Mathematical modeling. In H. Pashler & J. Wixted (Eds.), *Stevens' Handbook of Experimental Psychology* (3rd ed., pp. 429–460). John Wiley and Sons: New York.

Myung, I. J., & Pitt, M. A. (1997). Applying Occam's razor in modeling cognition: A Bayesian approach. *Psychonomic Bulletin and Review*, *4*, 79–95.

Myung, J. I., & Pitt, M. A. (2009). Optimal experimental design for model discrimination. *Psychological Review*, *116*, 499–518.

Navarro, D. J., & Fuss, I. G. (2009). Fast and accurate calculations for first-passage times in Wiener diffusion models. *Journal of Mathematical Psychology*, *53*, 222–230.

Newsome, W., & Pare, E. (1988). A selective impairment of motion perception following lesions of the middle temporal visual area (MT). *The Journal of Neuroscience*, *8*, 2201–2211.

Palestro, J. J., Bahg, G., Sederberg, P. B., Lu, Z. L., Steyvers, M., & Turner, B. M. (2017). *A tutorial on Joint Models of Neural and Behavioral Measures of Cognition.* (In Press)

Philiastides, M. G., Ratcliff, R., & Sajda, P. (2006). Neural representation of task difficult and decision making during perceptual categorization: A timing diagram. *Journal of Neuroscience*, *26*, 8965–8975.

Philiastides, M. G., & Sajda, P. (2006a). Temporal characterization of the neural correlates of perceptual decision making in the human brain. *Cerebral Cortex*, *16*, 509–518.

Philiastides, M. G., & Sajda, P. (2006b). Temporal characterization of the neural correlates of perceptual decision making in the human brain. *The Journal of Neuroscience*, *16*, 13082–13091.

Pleskac, T. J., & Busemeyer, J. R. (2010). Two Stage Dynamic Signal Detection Theory: A Dynamic and Stochastic Theory of Confidence, Choice, and Response

Time. *Psychological Review*, *117*, 864–901.

Purcell, B., Heitz, R., Cohen, J., Schall, J., Logan, G., & Palmeri, T. (2010). Neurally-constrained modeling of perceptual decision making. *Psychological Review*, *117*, 1113–1143.

Purcell, B., Schall, J., Logan, G., & Palmeri, T. (2012). Gated stochastic accumulator model of visual search decisions in FEF. *Journal of Neuroscience*, *32*, 3433–3446.

Rapoport, A., & Burkheimer, G. (1971). Models for deferred decision making. *Journal of Mathematical Psychology*, *8*, 508–538.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*, 59–108.

Ratcliff, R. (2006). Modeling response signal and response time data. *Cognitive Psychology*, *53*, 195–237.

Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. *Neural Computation*, *20*, 873–922.

Ratcliff, R., Philiastides, M. G., & Sajda, P. (2009). Quality of evidence for perceptual decision making is indexed by trial-to-trial variability of the EEG. *Proceedings of the National Academy of Sciences of the United States*, *106*, 6539–6544.

Ratcliff, R., & Rouder, J. N. (1998). Modeling response times for two-choice decisions. *Psychological Science*, *9*, 347–356.

Ratcliff, R., Thapar, A., & McKoon, G. (2006). Againg, practice, and perceptual tasks: a diffusion model analysi. *Psychological and Aging*, *21*, 353–371.

Ratcliff, R., & Tuerlinckx, F. (2002). Estimating parameters of the diffusion model: Approaches to dealing with contaminant reaction time and parameter variability. *Psychonomic Bulletin and Review*, *9*, 438–481.

Roitman, J. D., & Shadlen, M. N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J Neurosci*, *22*, 9475–9489.

Romo, R., Hernandez, A., & Zainos, A. (2004). Neuronal correlates of a perceptual decision in ventral premotor cortex. *Neuron*, *41*, 165–173.

Schall, J. D. (2004). On Building a Bridge Between Brain and Behavior. *Annual Review of Psychology*, *55*, 23–50.

Schall, J. D., & Thompson, K. G. (1999). Neural selection and control of visually guided eye movements. *Annual review of neuroscience*, *22*, 241–259.

Schwarz, G. (1978). Estimating the Dimension of a Model. *The Annals of Statistics*, *6*, 461–464.

Shadlen, M. N., & Kiani, R. (2013). Decision making as a window on cognition. *Neuron*, *80*.

Shadlen, M. N., & Newsome, W. T. (1996). Motion perception: Seeing and deciding. *Proceedings of the National Academy of Sciences of the United States*, *93*, 628–633.

Shields, W., Smith, J., & Washburn, D. (1997). Uncertain responses by humans and rhesus monkeys (Macaca mulatta) in a psychophysical same-different task. *Journal of Experimental Psychology*, *126*, 147–164.

Smith, P. L. (2000). Stochastic dynamic models of response time and accuracy: A foundational primer. *Journal of Mathematical Psychology*, *44*, 408–463.

Starns, J. J., & Ratcliff, R. (2010). The effects of aging on the speed-accuracy compromise: boundary optimality in the diffusion model. *Psychological Aging*, *25*, 377–390.

Stone, M. (1960). Models for choice reaction time. *Psychometrika*, *25*, 251–260.

ter Braak, C. J. F. (2006). A Markov Chain Monte Carlo version of the genetic algorithm Differential Evolution: easy Bayesian computing for real parameter spaces. *Statistics and Computing*, *16*, 239–249.

Thura, D., Beauregard-Racine, J., Fradet, C. W., & Cisek, P. (2012). A law of comparative judgement. *Journal of Neurophysiology*, *108*, 2912–2930.

Tuerlinckx, F. (2004). The efficient computation of the cumulative distribution and probability density functions in the diffusion model. *Behavior Research Methods, Instruments, & Computers*, *36*, 702–716.

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

Turner, B. M., Forstmann, B. U., Love, B. C., Palmeri, T. J., & Van Maanen, L. (2017). Approaches to analysis in model-based cognitive neuroscience. *Journal of Mathematical Psychology*, *76*, 65-79.

Turner, B. M., Gao, J., Koenig, S., Palfy, D., & McClelland, J. L. (2017). *The dynamics of multimodal integration: The averaging diffusion model.* (In Press) doi: 10.3758/s13423-017-1255-2

Turner, B. M., Schley, D. R., Muller, C., & Tsetsos, K. (2017). *Competing models of multi-attribute, multi-alternative preferential choice.* (In Press)

Turner, B. M., & Sederberg, P. B. (2012). Approximate Bayesian Computation with Differential Evolution. *Journal of Mathematical Psychology*, *56*, 375–385.

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

Turner, B. M., Sederberg, P. B., Brown, S., & Steyvers, M. (2013). A method for efficiently sampling from distributions with correlated dimensions. *Psychological Methods*, *18*, 368–384.

Turner, B. M., Van Maanen, L., & Forstmann, B. U. (2015). Combining Cognitive Abstractions with Neurophysiology: The Neural Drift Diffusion Model. *Psychological Review*, *122*, 312–336.

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

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

Usher, M., & McClelland, J. L. (2001). On the Time Course of Perceptual choice: The leaky competing accumulator model. *Psychological Review*, *108*, 550–592.

van Maanen, L., Fontanesi, L., Hawkins, G. E., & Forstmann, B. U. (2016). Striatal activation reflects urgency in perceptual decision making. *Neuroimage*, *16*, 294–303.

Vandekerckhove, J., Tuerlinckx, F., & Lee, M. D. (2008). A Bayesian approach to diffusion process models of decision-making. In V. M. Sloutsky, B. C. Love, & K. McRae (Eds.), *Proceedings of the 30rd Annual Conference of the Cognitive Science Society* (pp. 1429–1434). Austin, TX: Cognitive Science Society.

van Vugt, M. K., Simen, P., Nystrom, L. E., Holmes, P., & Cohen, J. D. (2014). EEG oscillations reveal neural correlates of evidence accumulation. *Dynamics of decision making: from evidence to preference and belief*, *108*.

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.

Voss, A., Voss, J., & Lerche, V. (2015). Assessing cognitive processes with diffusion model analyses: a tutorial based on fast-dm-30. *Frontiers in Psychology*, *6*.

Wagenmakers, E. J., Ratcliff, R., Gomez, P., & Iverson, G. J. (2004). Assessing model mimicry using the parametric bootstrap. *Journal of Mathematical Psychology*, *48*, 28–50.

Wasserman, L. (2000). Bayesian model selection and model averaging. *Journal of Mathematical Psychology*, *44*, 92–107.

White, C., Ratcliff, R., Vasey, M., & McKoon, G. (2009). Dysphoria and memory for emotional material: A diffusion model analysis. *Cognition and Emotion*, *23*, 181–205.

Wickelgren, W. A. (1977). Speed-accuracy tradeoff and information processing dynamics. *Acta Psychologica*, *41*, 67–85.

Wyart, D. G., V., V., J., Scholl, & Summerfield, C. (2012). Rhythmic fluctuations in evidence accumulation during decision making in the human brain. *Neuron*, *76*, 847-858.