Decision Making by Ensembles of Accumulators

Gregory E. Cox

University at Albany, State University of New York

Gordon D. Logan, Jeffrey D. Schall, and Thomas J. Palmeri

Vanderbilt University

Abstract

Evidence accumulation is a computational framework that accounts for behavior as well as the dynamics of individual neurons involved in decision making. Linking these two levels of description reveals a scaling paradox: How do choices and response times (RT) explained by models assuming single accumulators arise from a large ensemble of idiosyncratic accumulator neurons? We created a simulation model that makes decisions by aggregating across ensembles of accumulators, thereby instantiating the essential structure of neural ensembles that make decisions. Across different levels of simulated choice difficulty and speed-accuracy emphasis, choice proportions and RT distributions simulated by the ensembles are invariant to ensemble size and the accumulated evidence at RT is invariant across RT when the accumulators are at least moderately correlated in either baseline evidence or rates of accumulation and when RT is not governed by the most extreme accumulators. To explore the relationship between the low-level ensemble accumulators and high-level cognitive models, we fit simulated ensemble behavior with a standard LBA model. The standard LBA model generally recovered the core accumulator parameters (particularly drift rates and residual time) of individual ensemble accumulators with high accuracy, with variability parameters of the standard LBA modulating as a function of various ensemble parameters. Ensembles of accumulators also provide an alternative conception of speed-accuracy tradeoff without relying on varying thresholds of individual accumulators, instead by adjusting how ensembles of accumulators are aggregated or by how accumulators are correlated within ensembles. These results clarify relationships between neural and computational accounts of decision making.

Keywords: Response time; choice; speed-accuracy tradeoff; decision making; accumulator models; sequential sampling models.

Recent years have seen considerable progress in understanding the mechanisms involved in decision making. Progress in understanding these mechanisms has come from two directions. From one direction, efforts at understanding decision making behavior, namely

choice probabilities and response times (RTs), have led to sophisticated computational models of decision making in terms of the accumulation of evidence over time (Bogacz, Brown, Moehlis, Holmes, & Cohen, 2006; S. Brown & Heathcote, 2008; Busemeyer & Townsend, 1993; Ratcliff, 1978; Usher & McClelland, 2001). From the other direction, analyses of neurophysiology during decision making tasks have identified neurons in certain brain regions that appear to accumulate evidence over time to initiate an action, like shifting gaze (Gold & Shadlen, 2007; Hanes & Schall, 1996; O’Connell, Shadlen, Wong-Lin, & Kelly, 2018; Smith & Ratcliff, 2004). These two directions have converged on evidence accumulation as a common computational framework. This role for computational models as a bridge between behavior and the brain (cf. Marr, 1982) is at the core of the burgeoning field of model-based cognitive neuroscience (e.g., Forstmann, Wagenmakers, Eichele, Brown, & Serences, 2011; Forstmann & Wagenmakers, 2015; Palmeri, Schall, & Logan, 2015; Turner, Forstmann, Love, Palmeri, & Van Maanen, 2017).

The relationship between behavioral states and neural states takes the form of a *linking proposition* (Schall, 2004; Teller, 1984). In the case of decision making, a linking proposition identifies particular neurons with the computational process of evidence accumulation. While acceptance of this linking proposition is widespread (Gold & Shadlen, 2007; Hanes & Schall, 1996; O’Connell et al., 2018; Smith & Ratcliff, 2004), it entails a scaling paradox: On the one hand, no single neuron in a mammalian brain is responsible for decision behavior, so every choice is the outcome of an ensemble of neurons involved in decision making (e.g., Bonifazi, Ruaro, & Torre, 2005; J. W. Brown, Hanes, Schall, & Stuphorn, 2008; Tolhurst, Movshon, & Dean, 1983). On the other hand, computational models of decision making that assume single accumulators provide excellent accounts of decision making behavior as well as the activity of single neurons. How does a decision, the outcome and timing of which can be explained by a single model accumulator, arise from ensembles of redundant but idiosyncratic accumulator neurons, each of which individually appears to explain the variability of choice outcomes and response times?

We explored this paradox by developing a simulation model of choice decision making based on ensembles of up to a thousand accumulators (Zandbelt, Purcell, Palmeri, Logan, & Schall, 2014). Each choice alternative is associated with its own ensemble. Each of the individual accumulator elements produces its own unique dynamics governed by parameters that can be more or less similar to those of other accumulators within the same ensemble. The system makes decisions by aggregating across accumulators within each ensemble, selecting the alternative associated with the first ensemble to satisfy a particular stopping

We wish to thank Bram Zandbelt for providing the foundation for this research as well as valuable guidance.

This work was supported by National Institutes of Health grants NEI R01-EY021833, R01-EY08890, and P30-EY008126; the Advanced Computing Center for Research and Education (ACCRE) at Vanderbilt University, supported by NIH S10-OD023680; and by Robin and Richard Patton through the E. Bronson Ingram Chair in Neuroscience.

Simulation code and results are freely available via the Open Science Framework at <https://osf.io/vp65h/>.

Correspondence concerning this article should be addressed to Gregory E. Cox, Department of Psychology, University at Albany, Social Science 393, 1400 Washington Avenue, Albany, NY 12222. E-mail: gregcox7@gmail.com.

rule. Out of many evidence-accumulating neurons, one behavior emerges. In that spirit, we refer to the ensemble model as EPU-C for *E Pluribus Unum*¹-Choice (Zandbelt et al., 2014).

Because EPU-C embodies the core ensemble structure of the neural machinery thought to underlie decision making, we can use it to expose constraints to which real systems must adhere. We use these constraints to help resolve another paradox: Although computational models of decision making implement speed-accuracy tradeoff by modulating the threshold level of evidence, the neurons instantiating evidence accumulation do not change their level of activity at response in this manner (Hanks, Kiani, & Shadlen, 2014; Heitz & Schall, 2012, 2013; Servant, Tillman, Schall, Logan, & Palmeri, 2019; Thura & Cisek, 2016). If it cannot be the case that individual neural accumulators yield a speed-accuracy tradeoff by modulating their thresholds, speed-accuracy tradeoff must be implemented at the ensemble level by some other means. In this way, EPU-C allows us to identify ensemble properties that are capable of producing speed-accuracy tradeoff, helping to resolve a major discrepancy between computational models and the neural mechanisms of decision making.

Accumulator Models of Decision Making

Accumulator models have had a long and successful history explaining the details of decision making behavior (for recent reviews, see Bogacz et al., 2006; Forstmann, Ratcliff, & Wagenmakers, 2016; Ratcliff & Smith, 2004). Variants and extensions of these models account naturally for choice probabilities, response times (RTs), and tradeoffs between speed and accuracy. These models have been applied widely to understand perception, attention, memory, categorization, cognitive control, decision making, and more (S. Brown & Heathcote, 2008; Busemeyer & Townsend, 1993; Cox & Shiffrin, 2017; Holmes, Trueblood, & Heathcote, 2016; Link & Heath, 1975; Logan, 2002; Logan, Van Zandt, Verbruggen, & Wagenmakers, 2014; Nosofsky & Palmeri, 1997; Palmeri, 1997; Pike, 1973; Ratcliff, 1978; Ratcliff & Rouder, 1998; Polyn, Norman, & Kahana, 2009; Smith & Van Zandt, 2000; Usher & McClelland, 2001; Vickers, 1970). Although there is considerable diversity among accumulator models, they share the basic assumption that decisions are made by accumulating evidence for or against different alternatives over time. When the evidence supporting one alternative is sufficiently strong, that alternative is selected as the outcome of the decision. Accumulator models differ in how they approach the nature of the evidence being accumulated and how it may vary over time. Models also differ in the details of how evidence for different alternatives is integrated over time and across alternatives in order to make a final choice.

Although some accumulator models are incorporated as components of a larger theory that explicitly models the sources of evidence that are accumulated over time (e.g., Annis, Gauthier, & Palmeri, under review; Cox & Shiffrin, 2017; Logan, 2002; Nosofsky & Palmeri, 1997; Polyn et al., 2009), most accumulator models are agnostic with respect to the particular source of evidence, hence their wide applicability across domains. Instead of explicitly modeling the evidence generating process, these “generic” accumulator models treat the evidence in a statistical way, characterizing it using a standard parametric form.

¹Out of many, one.

Figure 1. Illustration of two basic features of decision making, difficulty and speed-accuracy emphasis. The left panel illustrates how difficulty and speed-accuracy emphasis are typically instantiated in evidence accumulation models. An *easy* choice is one in which one alternative has a much higher drift rate than the other, compared to a *hard* choice in which the two alternatives have similar drift rates. An emphasis on accuracy is typically instantiated as a higher response threshold in both accumulators, compared to an emphasis on speed in which thresholds are lower. The right panel illustrates the joint effects of difficulty and speed-accuracy emphasis on response time (RT) and choice probabilities. Accuracy emphasis is associated with responding that is slower but with the favored alternative (alternative 1 in this example) being selected more often, compared to speed emphasis. Within each level of speed emphasis, more difficult choices (leftmost points) are associated with slower responses that are less often in the direction of the favored alternative than easy choices (rightmost points).

Accumulator models all share three core parameters: *drift rate*, the rate at which evidence (for a particular alternative) accumulates over time; *threshold*, the amount of accumulated evidence required to select an alternative; and *residual time* or *non-decision time* that describes the time required to encode the alternatives and execute the actions associated with the selected alternative, neither of which are directly modeled as part of the accumulation process.

We illustrate these three parameters with of a particular accumulator model, the *Linear Ballistic Accumulator* (LBA; S. Brown & Heathcote, 2008, see Figure 1). We chose LBA for illustrative purposes both because of its relative simplicity and because it will serve as the basis for the ensemble model of choice that we develop later in this article. LBA assumes that each alternative is associated with its own accumulator. In the case of two alternatives, this might be deciding that an animal is a dog or a cat, a letter string is a word or a non-word, a memory probe is old or new, or whether a display contains a visual search target or not. After some initial time for detecting and encoding relevant stimuli (e.g., the object, string, probe), each accumulator begins to accrue evidence for its associated alternative at a particular drift rate. Eventually, the evidence in one of the accumulators will reach a threshold, causing the decision-maker to select the associated alternative and make the associated response. As a result, the chosen alternative can be viewed as the

winner of a race between accumulators. The predicted response time is the time needed for the winning accumulator to reach its threshold plus the residual time for encoding the alternatives and executing the response.

An important factor associated with decision making is variability, and different accumulator models embody different kinds of variability by how they parameterize the evidence accumulation process. This variability can operate on different timescales. Some accumulator models, like the diffusion model (Ratcliff, 1978) and the leaky competing accumulator model (Usher & McClelland, 2001), assume that the accumulated evidence itself is subject to variability at short timescales within a trial. LBA, as might be inferred from its name, is ballistic and assumes that accumulation within a trial is deterministic. Instead, LBA² assumes that the drift rates for each alternative vary *between* trials. For example, the same letter string in a lexical decision task might be associated with a high drift rate in the “word” accumulator on one trial and a low rate on a different trial due to this stochastic variability. It is also possible to have variability in the starting point of accumulation between trials; when starting point is higher, less evidence is needed to reach threshold. Starting point variability is a core aspect of the LBA (and can be present in the diffusion model; Ratcliff & Rouder, 1998). The joint effects of starting point and drift rate variability in the LBA are what leads it to predict variability in choices and response times. Finally, LBA (and other accumulator models) can allow trial-to-trial variability in residual time as well (though we do not explore that in this article).

Because accumulator models account for both choices and RTs, they are useful for understanding two important experimental factors that influence performance in decision making tasks: difficulty and speed-accuracy emphasis. The specifics of what might make a task easy or difficult depend on the nature of that task, but are typically modeled in terms of differences in drift rates, which represent the quality of the information driving a decision (e.g., S. Brown & Heathcote, 2008; Ratcliff & McKoon, 2008). In a race model like LBA, a more difficult choice is associated with a smaller difference in drift rates between accumulators, reflecting the idea that the evidence being accumulated over time does not distinguish between the alternatives very well. Speed-accuracy emphasis is often induced by instructions or reward incentives that encourage the decision maker either to take their time and emphasize accuracy or to emphasize speed and compromise accuracy (Heitz, 2014; Wickelgren, 1977). Behaviorally, this results in what is called a *speed-accuracy tradeoff*. Accumulator models generally emphasize speed or accuracy by adjusting thresholds, setting thresholds higher to favor accuracy over speed by requiring more accumulated evidence before responding, setting thresholds lower to favor speed over accuracy by allowing a response to be made on the basis of less evidence (but see Heitz & Schall, 2012; Rae, Heathcote, Donkin, Averell, & Brown, 2014; Servant, van Wouwe, Wylie, & Logan, 2018). Examples of predicted RTs and accuracy from LBA simulations manipulating difficulty via differences in drift rates and speed-accuracy emphasis via differences in threshold (details below) are illustrated in Figure 1. Later in this article, we will explore other avenues by which ensemble models can simulate a speed-accuracy tradeoff besides threshold adjustments.

²As well as the diffusion model introduced by Ratcliff (1978), and unlike earlier random walk models (Laming, 1968; Stone, 1960).

Relations between Accumulator Models and Neurons

Decision making has been one of the most productive applications of model-based cognitive neuroscience (e.g. Forstmann et al., 2016; Palmeri et al., 2015; Smith & Ratcliff, 2004). One of the earliest insights propelling the synergy of neuroscience and cognitive modeling in this domain was the observation that the pattern of spiking activity of certain individual neurons involved in decisions to move the eyes resembles an accumulation of evidence to threshold. Neurons of this kind have been identified in various brain regions, including the lateral intraparietal area (LIP; Colby, Duhamel, & Goldberg, 1996; Roitman & Shadlen, 2002; Thomas & Paré, 2007), frontal eye fields (FEF; Ding & Gold, 2012; Hanes, Patterson, & Schall, 1998; Lowe & Schall, 2018; Schall & Hanes, 1993; Woodman, Kang, Thompson, & Schall, 2008), and superior colliculus (SC; Reppert, Servant, Heitz, & Schall, 2018).

The activity of these neurons within a trial resembles that of an evidence accumulator model: Each neuron is associated with a particular decision, such as whether to move the eyes to a spot in the upper left or bottom right part of the visual field. At the beginning of a trial, these neurons have similar baseline firing rates but over the course of the trial, neurons associated with alternatives that are favored by the evidence provided on that trial (e.g., a region of the visual field containing features of a visual search target) ramp up their activity. When neurons associated with a particular alternative reach a critical firing rate, the decision-maker initiates the response associated with the selected alternative (e.g., directing the eyes toward a target stimulus).

The connection between neurons and accumulator models goes beyond mere resemblance³: Trials with longer RTs are associated with neurons reaching their critical firing rate more slowly than trials with shorter RTs (Hanes & Schall, 1996; Woodman et al., 2008), suggesting that these neurons play a causal role in decision making. When model accumulators are “fed” the same neural activity that presumably drives real accumulator neurons, these model accumulators yield behavioral and neural dynamics that closely match those observed in real accumulator neurons (Purcell et al., 2010; Purcell, Schall, Logan, & Palmeri, 2012). Finally, it is possible to jointly model the computational evidence accumulation process and properties of individual neural dynamics using similar parameters (P. J. Cassey, Gaut, Steyvers, & Brown, 2016; Ratcliff, Hasegawa, Hasegawa, Smith, & Segraves, 2007; Philiastides, Ratcliff, & Sajda, 2006). Taken together, these results support the linking proposition that certain neurons implement the evidence accumulation process posited by computational models of decision making (Gold & Shadlen, 2007; Palmeri et al., 2015; Schall, 2001, 2004, 2019; Smith & Ratcliff, 2004). Furthermore, neural dynamics at larger scales, measured via electroencephalography (EEG; O’Connell, Dockree, & Kelly, 2012; Schurger, Sitt, & Dehaene, 2012; van Ravenzwaaij, Provost, & Brown, 2017), magnetoencephalography (MEG; Smyrnis et al., 2012), and functional MRI (Forstmann et al., 2008; Turner, Rodriguez, Norcia, McClure, & Steyvers, 2016), can resemble the dynamics of evidence accumulator models. It is important to understand how individual neurons implement evidence accumulation in such a way that their ensemble dynamics also represent

³Although we note that there is some current debate regarding the nature of the dynamics of these neurons (Latimer, Yates, Meister, Huk, & Pillow, 2015; Shadlen et al., 2016; Zoltowski, Latimer, Yates, Huk, & Pillow, 2019).

an accumulating evidence signal.

Choice Decision Making by Ensembles of Accumulators

While it seems that the basic proposition linking neurons to evidence accumulation is sound, the details of how this proposition is realized in neural systems remain unclear: How can individual neurons be thought of as evidence accumulators if no single neuron is wholly responsible for behavior? And how can evidence accumulation models successfully explain decision behavior when they do not directly model the large ensembles of individual neurons that is in fact necessary to generate that behavior? In the context of explaining simple reaction time with only one alternative, Zandbelt et al. (2014) implemented an ensemble model called EPU (*e pluribus unum*, “out of many, one”). In EPU, a speeded response is the product of aggregating across many individual accumulator elements. Because each of these elements accumulates evidence to a threshold, EPU embodies the linking proposition between accumulators and individual neurons. But because each of these individual accumulators is embedded in a larger ensemble, no single accumulator is generally responsible for determining RT. Instead, RT is determined by *polling* the ensemble elements; only when a critical proportion of them reach threshold does EPU generate a response.

The model we explore in this paper extends EPU to choice (EPU-C), assuming multiple ensembles, each of which is associated with a different response alternative. EPU-C predicts both choice probabilities and RTs. The expanded range of behavior simulated by EPU-C, relative to the original EPU model, expands the range of potential internal dynamics that can arise from how ensembles compete against one another to produce a response.

EPU-C is a model in which *Ensemble Processing Units lead to Choice*. Each ensemble of accumulators associated with a particular choice alternative can be viewed collectively as a *unit* that operates by accumulating the results of its constituent accumulator elements. In this sense, each EPU-C unit is a *meta-accumulator*, operating at a higher level of processing than its individual accumulator elements. This view of EPU-C, as illustrated in Figure 2, represents a general processing structure that may embody how neural systems make decisions: While individual neurons have their own processing dynamics characterized as accumulation of evidence, behavioral responses result from aggregating across the ensemble of accumulators. This overall structure is the essence of EPU-C, which will receive a specific computational instantiation later.

Figure 2 provides a roadmap for our approach. We instantiate EPU-C using the LBA (S. Brown & Heathcote, 2008) to represent the individual accumulator elements out of which each ensemble is built. We refer to the ensemble elements as eLBA to emphasize that they are *elements* of a larger ensemble (Figure 2A). We assume that there are N eLBAs in each ensemble associated with each choice alternative. EPU-C makes a choice by selecting the first alternative for which a proportion p of its associated ensemble elements have reached threshold (Figure 2B). In addition to exploring the effects of different values of N and p on simulated choice decision making, this framework allows us to explore different ways in which the accumulators within an ensemble might be correlated with one another, an issue that naturally arises given that the evidence being accumulated by each element of the ensemble is drawn from the same source (e.g., the same visual display). We examine two types of correlation: correlation in drift rates and correlation in starting points. To the

D) Standard accumulator model (LBA) fit to ensemble behavior

C) Ensemble behavior (choices and response times)

B) Ensembles of accumulators (EPU-C)

A) Accumulator elements (eLBA)

Figure 2. Outline of the EPU-C modeling approach used in this article. Decision behavior (choice probabilities and RT distributions) arises from ensembles of accumulators, each of which is instantiated as a Linear Ballistic Accumulator element (eLBA). The behavior simulated by the ensemble can be fit by an accumulator model like LBA. See main text for details.

extent that drift rates are correlated within an ensemble, each of its component accumulators will accrue evidence at similar rates within a trial. To the extent that starting points are correlated within an ensemble, each of its component accumulators will tend to start with the same level of initial activity at the onset of accumulation.

We explore simulated ensemble behavior (Figure 2C) across different levels of ensemble size, aggregation rule, and drift rate and starting point correlations, examining the consequences of these factors for how an ensemble of accumulators responds to different levels of difficulty and speed-accuracy emphasis (e.g., Figure 1). These simulations explore the conditions that satisfy the linking proposition between neurons and evidence accumulation (Zandbelt et al., 2014):

Invariance of simulated behavior with ensemble size. Under some parameter regimes, the behavior of EPU-C is approximately invariant across ensemble size (above a certain ensemble size). This illustrates the robustness of the ensemble, an important consideration given that we do not know exactly how many elements are present in real neural ensembles. It indicates the regimes under which individual elements (or small numbers of elements) are “representative” of the whole ensemble, such that measurements of a small number of elements are informative of the behavior of the whole ensemble. This is important when considering real neurons given that it is not possible to simultaneously record activity from more than a modest number of neurons at once.

Invariance of activation at response time. An important feature of neurons involved in decision making is that the level of activity they reach at the time of a response is roughly constant irrespective of observed RT (Hanes & Schall, 1996). EPU-C allows us to discover what types of ensemble are able to maintain this invariance of activation at the level of individual elements, an important constraint on neural plausibility.

Invariance of ensemble element and accumulator model parameters. A viable ensemble model of decision making must produce behavior that can be reasonably accounted for by standard accumulator models (like the LBA), since it is the power of accumulator models to describe decision making behavior that underlies the link between accumulators and neural dynamics. By fitting the behavior simulated by EPU-C using a standard LBA model (Figure 2D), we discover not only the properties of ensembles that yield behavior consistent with an accumulator model of decision making, we discover the relationships between parameters of the LBA fitted to simulated ensemble behavior with those of the eLBA ensemble elements that produced that behavior, an important consideration for model-based cognitive neuroscience.

Our article closes by using the constraints we identify to address the other paradox we noted—namely that while computational models of decision making commonly implement speed-accuracy tradeoffs via threshold adjustments, neurons involved in decision making maintain the same threshold firing rate under speed and accuracy emphasis conditions. Given that the neural implementation of speed-accuracy tradeoffs seems not to rely on threshold adjustments at the level of individual neural accumulators, are there ensemble-level parameters that can realize a speed-accuracy tradeoff? These new insights give alternative ways of conceptualizing why certain choices are more difficult than others and how tradeoffs between speed and accuracy might be accomplished by neural ensembles in the brain.

EPU-C: An Ensemble Model of Choice Decision Making

Simulation Methods

Accumulator Elements of the Ensemble Model. Each element of EPU-C is modeled after the Linear Ballistic Accumulator (LBA; S. Brown & Heathcote, 2008)⁴. To distinguish these accumulator elements from the standard LBA that is typically fitted to behavioral data (which we will use in a later section), we refer to the individual ensemble elements as *eLBA*. Each ensemble of eLBAs is associated with a particular choice alternative; in this paper, we consider two alternatives, which we denote as alternative 1 and alternative 2. We assume that an equal number of eLBAs is associated with each alternative, with the number of accumulators in each ensemble denoted N . In the case of an ensemble of size $N = 1$, there is a single eLBA associated with each alternative and the ensemble model reduces mathematically to the standard LBA.

Single accumulator elements (eLBA). We first describe the operation of a single eLBA associated with a one alternative (within one ensemble) on a particular trial, as illustrated in Figure 2A. The initial level of activity x_0 (starting point) of a single eLBA on a particular trial is a random sample from a uniform distribution on the interval $[0, A]$, with parameter A representing the degree of variability associated with the starting point. The rate at which an eLBA accumulates evidence on a particular trial is denoted d and is also a random sample, in this case from a log-normal distribution with mean v and standard deviation s (v and s are on the log-scale). Note that, although many applications of LBA models assume a normal distribution of drift rates, we adopt the log-normal in order to avoid situations where accumulators have negative drift rates and never reach threshold; aside from this technical advantage, a log-normal distribution of drift rates produces similar behavioral predictions to a normal distribution (Terry et al., 2015)⁵. The level of activity of an eLBA begins at its starting point and accrues linearly over time, such that its level of activity at time t is simply given by $x(t) = x_0 + d \times t$. Accumulation stops when $x(t)$ reaches a threshold B (which is constrained to be greater than A , the maximum possible starting point). Across trials, the variability associated with the starting point x_0 and drift rate d leads to variability in the time needed for a single eLBA to reach its threshold; on some trials, the sampled drift rate may be high, leading to relatively rapid accumulation toward threshold; on other trials, the starting point might be high, meaning even a relatively low drift rate could lead to an eLBA reaching threshold quickly. We acknowledge that the strictly linear accumulator dynamics of an eLBA do not resemble the kinds of stochastically variable activity levels observed in actual neurons. This is in part because neurons express their measured activity level not directly but via discrete spikes generated probabilistically.

⁴We additionally explored alternative accumulator architectures as elements of the ensemble, such as racing diffusion models (Logan et al., 2014), but as in Zandbelt et al. (2014), we did not observe qualitative differences between these alternatives. Therefore, we chose to focus on an ensemble composed of LBAs, in part because of the simplicity of the LBA to describe, simulate, and fit; ensembles of stochastic accumulators require orders of magnitude more compute time to simulate.

⁵We note that our use of the log-normal distribution for drift rates differs from that in log-normal race models (Heathcote & Love, 2012; Ulrich & Miller, 1993). In particular, a race model that assumes that both starting points and drift rates are log-normally distributed makes it impossible to distinguish between these two types of variability. In our model, only drift rates are log-normally distributed, so their variability is distinct from that in starting points.

We take up this issue in more detail in the Discussion, but we note here that eLBA may be sufficient for approximating certain key aspects of neural dynamics because the kinds of between-trial variability incorporated in our eLBA can often explain apparent within-trial variability (Donkin, Brown, Heathcote, & Wagenmakers, 2011; Tillman, Van Zandt, & Logan, in press) and the apparent non-linear variability in neurons can be the product of much simpler latent, even linear, rate dynamics (e.g., P. J. Cassey et al., 2016; Zoltowski et al., 2019).

Ensembles of accumulators (EPU-C). Now, let us consider an ensemble of N eLBAs (Figure 2B) that is associated with a particular choice alternative. All of the accumulators in an ensemble have the same parameters. Even so, it is still the case that there is trial-to-trial variability in the drift rate and starting point of each of the N accumulators in the ensemble, so they will not necessarily have the same starting points and drift rates even within the same trial. The degree to which the trial-by-trial starting points and drift rates are correlated between accumulators within an ensemble is governed by r_A and r_d , respectively. If $r_d = 0$, then on any given trial, each of the N accumulators will have a drift rate sampled independently from a log-normal distribution with mean v and standard deviation s on the log scale; in other words, drift rate varies between accumulators *and* between trials. If $r_d = 1$, then all of the N accumulators will have the same drift rate on any given trial, but that common rate will be a random sample from a log-normal distribution with mean v and standard deviation s on the log scale; in other words, drift rate varies between trials but not between accumulators. These same principles hold for starting point variability.

Choices and response times. We explore with our simulations the case where there are two alternatives, each of which is associated with an ensemble of N accumulators functioning as described above. In general, it would be possible for each ensemble to have entirely different parameters defining their eLBA elements and their correlations in starting point and drift rate. To keep the simulations manageable, we assume that each ensemble has the same parameter values except for the mean drift rates (v_1 for alternative 1 and v_2 for alternative 2), reflecting differing amounts of incoming evidence in favor of each alternative. To select one of these alternatives and make a decision, we employ an *aggregation* rule⁶: The first ensemble for which a proportion p of its accumulators have reached threshold determines the final choice. The RT is the amount of time it takes for those accumulators to reach threshold, plus a residual non-decision time t_0 . For example, if $p = 0.1$ and $N = 100$, and time t is the first time at which 10 accumulators (i.e., 0.1×100 accumulators) associated with alternative 1 have reached threshold while fewer than 10 accumulators associated with alternative 2 have reached threshold, then the final choice is alternative 1 and the RT is equal to $t + t_0$.

Simulated experiment. The trial-by-trial variability in the starting points and drift rates of each accumulator element aggregates within the ensembles to yield trial-by-trial variability in behavior (choices and RTs). Unfortunately, the complexity at the ensemble level precludes analytic calculation of ensemble choice behavior, so we simulate a large number of trials (10,000) for a factorial combination of different parameter settings to understand the behavior of the ensemble model in detail.

⁶A mechanism akin to quorum sensing (Seeley, Visscher, & Passino, 2006).

Difficulty and speed-accuracy emphasis. To explore the range of ensemble behavior, we simulated an experiment with a factorial combination of difficulty and speed-accuracy emphasis that is typical of experiments in decision making. Consider, for example, an experiment where stimuli are random checkerboard patterns with varying numbers of white and black squares and the task involves deciding if the checkerboards are more white or more black (e.g., Middlebrooks, Zandbelt, Logan, Palmeri, & Schall, 2019; Ratcliff, Thapar, & McKoon, 2003). The two alternatives are “white” and “black”, and more difficult trials are those for which the proportion of white and black squares are closer to one another. Explicit instructions or reward incentives to respond accurately at the expense of speed or to respond quickly at the expense of accuracy induce different degrees of speed-accuracy emphasis.

Variants of this difficulty by speed-accuracy emphasis experimental structure are widely used to test accumulator models of decision making both empirically (e.g., S. Brown & Heathcote, 2008; Ratcliff & Rouder, 1998) and theoretically (e.g., Boehm et al., 2018; Dutilh et al., 2018). The canonical approach to applying accumulator models to experiments with this structure is to assume that variations in difficulty map onto variations in drift rate and that variations in speed-accuracy emphasis map onto variations in threshold. This is the approach we take in our simulations, though we emphasize that we are not claiming that real experimental manipulations of difficulty or speed-accuracy emphasis must selectively influence these particular model parameters (for evidence that they do not, see P. Cassey, Heathcote, & Brown, 2014; Heitz & Schall, 2012; Rae et al., 2014); rather, we wanted our simulations to be able to represent a variety of potential behaviors and varied these parameters to induce this variation.

Accumulator element parameters. For a given set of ensemble model parameters, a simulated experiment contains four levels of difficulty (four values of mean drift rates within each ensemble) crossed with two levels of speed-accuracy emphasis (two values of threshold). We chose parameter values for our simulations based on ranges of parameter values used in other simulations of the LBA (e.g., Donkin et al., 2011; Terry et al., 2015) and based on producing performance roughly in a range measured empirically in monkeys and humans in speeded decision making tasks (e.g., Pouget, Emeric, Stuphorn, Reis, & Schall, 2005).

The values we chose for all parameters in our simulations are given in Table 1. For simulation purposes, we defined drift rates in terms of the mean drift rate for accumulators in the ensemble associated with alternative 1 (v_1) and constrained the mean drift rate for accumulators in the ensemble associated with alternative 2 (v_2) such that $v_1 + v_2 = 1$. As a result, as the evidence favoring alternative 1 increases, the evidence favoring alternative 2 decreases by an equal amount. We did not simulate response bias or any other asymmetries between ensembles; although we simulate different values of threshold B , that threshold is always the same across both ensembles. Similarly, ensemble size (N), aggregation rule (p), drift rate correlation (r_d), and starting point correlation (r_A) were always equal across both ensembles, as were residual time (t_0) and the overall levels of starting point and drift rate variability (A and s , respectively).

Ensemble parameters. As described above, we conducted a wide-ranging set of simulations of EPU-C to examine its behavior in an experiment that factorially manipulated choice difficulty and speed-accuracy emphasis. We varied ensemble size across four

Table 1

Parameter values for EPU-C simulations. First five parameters refer to the individual eLBA's used as ensemble elements, final four parameters refer to the whole ensemble.

Parameter	Description	Values
A	Range of between-trial variability in starting point	2.8
B	Response threshold	3, 4.5
v_1	Mean drift rate (on a log scale) associated with alternative 1; mean drift rate for alternative 2 is given by $v_2 = 1 - v_1$	0, 0.2, 0.4, 0.8
s	Standard deviation of between-trial variability in drift rates	0.4
t_0	Residual (non-decision) time, in seconds	0.2
N	Number of accumulator elements in each ensemble	1, 10, 100, 1000
p	Proportion of accumulators in an ensemble that need to have reached threshold to select a response	0 to 1 in increments of 0.125.
r_d	Correlation in drift rates across accumulators within an ensemble	0 to 1 in increments of 0.125.
r_A	Correlation in starting point across accumulators within an ensemble	0 to 1 in increments of 0.25.

logarithmically-spaced levels ($N = 1, 10, 100$, and 1000), drift rate correlation across nine levels ($r_d = 0$ to 1 in 0.125 increments), starting point correlation across five levels ($r_A = 0$ to 1 in 0.25 increments), and aggregation rule across nine levels ($p = 0$ to 1 in 0.125 increments). An aggregation rule with $p = 0$ means that the very first accumulator to reach threshold within one of the ensembles determines the choice, whereas an aggregation rule with $p = 1$ means that the first ensemble for which all of its accumulators reach threshold determines the choice. 10,000 trials were simulated for each combination of parameter values at each level of threshold and drift rate, allowing for detailed characterization of distributions of response time and choice probability across all combinations of parameters.

Invariance of Ensemble Behavior With Ensemble Size

We first identify regions of the EPU-C parameter space that result in simulated behavior that is roughly invariant with ensemble size, either across all ensemble sizes or for ensembles that are sufficiently large. Invariance to ensemble size can be seen as a characterization of the robustness of ensemble models; ensembles that satisfy the linking proposition between accumulating neurons and computational models of decision making should be those in which changes in ensemble size do not change ensemble behavior. We will see that there are ranges of ensemble parameters whereby the predicted behavior is the same whether there are 10 accumulators or 1000 accumulators contributing to behavior; what this means practically is that ensemble models can be proposed without having to specify a precise size of the ensemble.

A recurring device in this article is the construction of visualizations of our large-scale simulation results (a factorial combination of many ensemble parameters, each combination of which simulates an experiment with a factorial combination of choice difficult and speed-

accuracy emphasis). These visualizations are necessary to communicate the essential features of the simulation results, but can be challenging both to construct and to interpret. As such, in each section, we have adopted certain conventions designed to highlight the aspects of our results that are germane to that section. In this section, we use these visualizations to highlight the ensemble parameters for which EPU-C behavior in our simulated experiment is roughly invariant to ensemble size.

Visualizing Predicted Performance

Figure 3 provides a general guide to how the simulation results in this section will be visualized. For a particular ensemble size ($N = 100$), drift rate correlation ($r_d = 0.25$), starting point correlation ($r_A = 1$), and aggregation rule ($p = 0.5$), EPU-C is used to simulate 10,000 trials at each level of threshold ($B = 3$ and 4.5) and each level of drift rate ($v_1 = 0, 0.2, 0.4$, or 0.8). The histograms of simulated RTs for choosing alternative 1 or alternative 2 are shown in Figure 3A. These histograms are displayed rotated by 90 degrees from their conventional orientation because another way of summarizing choice and RT is via a *quantile-probability plot*, as shown in Figure 3B; these plot choice probability on the horizontal axis and the 0.1, 0.3, 0.5, 0.7, and 0.9 quantiles of the RT distributions associated with that choice on the vertical axis. Each line in Figure 3B depicts how a particular RT quantile changes with choice probability as a function of drift rate (higher drift rates correspond to a higher proportion of choosing alternative 1), response threshold (lighter lines correspond to $B = 3$, darker to $B = 4.5$), and choice (solid lines for alternative 1, dashed for alternative 2).

These types of plots illustrate several key features of simulated EPU-C behavior: (1) Lower thresholds lead to faster responding that is closer to chance (lighter lines are lower and more tightly clustered around the midpoint of the horizontal axis) and higher drift rates lead to a greater probability of choosing alternative 1, and it takes less time to do so. (2) With high thresholds, higher drift rates are associated with slower RTs to alternative 2, the non-preferred alternative; these trials represent those in which the drift rate for alternative 1 happened to be weak thus allowing the slower-drifting alternative 2 to be chosen. (3) When threshold is low, higher drift rates are associated with faster RT's to alternative 2, reflecting those trials in which responding was primarily determined by variability in starting point rather than variability in drift (Ratcliff & Rouder, 1998). By plotting RT quantiles separately, we can see that predicted RT distributions are not merely shifted by a constant amount as drift rate and threshold change, but change their shape as well. It is clear that quantile-probability plots can convey a great deal of information about EPU-C simulated behavior. However, these plots are complex and are not be able to efficiently convey information about multiple parameter settings of ensemble models in a single plot (e.g., for different values of N).

As a result, we elected to simplify the quantile-probability plot into a form shown in Figure 3C, which plots choice probabilities against only the median RT, not the full RT distribution, and fills in the center to form a *speed-accuracy tradeoff polygon* (SATP). The shape of this polygon conveys both the inverted-U relationship between choice probability and RT as well as how this relation fans out (longer RTs and more extreme choice probabilities) with increasing threshold, thereby presenting information about predicted performance in a salient way with a minimum of visual clutter. Minimizing clutter is crucial as we embed

Figure 3. An illustration of how EPU-C behavior will be visualized across ensemble parameters. For each combination of ensemble parameters, we simulated 10,000 trials for each combination of drift rate (v_1) and threshold (B ; panel A presents histograms of choice frequency in different response time bins for $N = 100$, $r_d = 0.25$, $r_A = 1$, and $p = 0.5$). Although these predictions can be visualized in a quantile probability plot (B), we simplify this representation using just median response times and choice proportions to create a *speed-accuracy tradeoff polygon* (C). We then overlay polygons from different ensemble sizes and tile them according to different levels of p and r_d , shading each tile according to the ensemble size at which predicted behavior is approximately invariant to predicted behavior for $N = 1000$ (D). Panels E, F, and G illustrate particular tiles from the ensemble plot in panel D. See main text for additional details.

this polygon in a composite plot in Figure 3D, which summarizes our simulation results for the conditions in which starting point correlation is set to $r_A = 1$. Each sub-plot in this composite plot depicts four SATPs overlaid on one another, one for each level of ensemble size (differentiated by color). Thus, the degree to which performance is invariant across ensemble size within a sub-plot is visually conveyed by the extent to which these SATPs overlap. The complete composite plot contains 81 sub-plots of this type in an 9×9 grid, where moving from left to right within the grid corresponds to increasing levels of p (the aggregation rule), and moving from bottom to top within the grid corresponds to increasing levels of drift rate correlation r_d .

Quantifying and visualizing invariance with ensemble size. Before describing the individual sub-plots of the composite plot shown in Figure 3D, we need to explain how we quantified and further visualized the degree of invariance of EPU-C simulations as a function of ensemble size within each sub-plot. To quantify the degree of approximate invariance, we make use of the Komolgorov-Smirnoff (KS) statistic. The KS statistic is the largest absolute difference between two cumulative distribution functions⁷. This statistic is usually used to test the hypothesis that two distributions are equivalent (e.g., to determine whether two samples come from the same underlying population distribution) but for our purposes we use it purely as a method of quantifying the extent to which EPU-C simulations differ from one another as a function of ensemble size. Because the KS statistic is an absolute difference between distribution functions, it is bounded between 0 and 1 and for visualization purposes we consider two distributions to be *approximately invariant* if their KS statistic is less than 0.05⁸. We emphasize that this threshold is not meant to act as a hypothesis test, but only to help visualize regions of the EPU-C parameter space in which simulated behavior is more or less invariant as a function of ensemble size. Details about how we compute the KS statistic for joint distributions of choice probability and response time are given in Appendix A.

For each combination of parameter values in our simulations with ensemble size N less than 1000, we compute the KS statistic between its simulated joint choice/RT distribution and that simulated by EPU-C with $N = 1000$ (with all other parameters the same). For example, one of the KS statistics we compute is between predicted choice/RT distributions with $B = 3$, $v_1 = 0.4$, $r_A = 0.5$, $r_d = 0.125$, $p = 0.875$, between $N = 10$ and $N = 1000$; another is for those exact same ensemble parameters but between $N = 100$ and $N = 1000$. The resulting KS statistics indicates how different are simulations for the maximum ensemble size of $N = 1000$ from the simulated behavior with a smaller ensemble size. For any given sub-plot in the larger plot, which includes simulated performance for several values of threshold B and drift rate v_1 , we average the resulting KS statistics across threshold and drift rate for each ensemble size. As noted above, we consider predictions to be approximately invariant between ensemble sizes if their associated KS statistic, averaged

⁷One could also compute the integrated or average difference in distributions functions as a measure of discrepancy, the KS statistic makes it easier to compare our operational definition of *approximate invariance* with actual statistical tests that are performed to detect differences between distributions. Further, although the critical values of the KS statistic will depend on sample size in an empirical setting, the statistic itself will not.

⁸This is roughly the critical value for rejecting the null hypothesis that two distributions are equal at an alpha level of 0.05 if both distributions are estimated using samples of size 1475. This is a large number of trials for many paradigms, on the high end of the number of trials observed in a single session in nonhuman primate perceptual decision making tasks, so it is quite a conservative criterion for invariance.

over different conditions (defined by drift rate and threshold) is less than 0.05. For each combination of parameter values, we can then determine at what ensemble size N (< 1000), if any, predicted choice/RT behavior becomes approximately equivalent to that predicted from the maximum ensemble size of $N = 1000$. The result of this test is presented visually in terms of the background color of each sub-plot in each composite plot, which corresponds to the smallest ensemble size N at which simulated performance becomes approximately invariant with that simulated by $N = 1000$, such that the darker the background of each sub-plot, the larger the ensemble size needs to be to achieve this approximate invariance.

Examples when $r_A = 1$. We now turn our focus on a few sub-plots within the composite plot in Figure 3D before expanding our focus to the totality of our results. As illustrated in the top row of Figure 3D (of which Figure 3C forms a part), we can see that with perfect correlation of drift rates ($r_d = 1.0$) and starting points ($r_A = 1.0$) across units within an ensemble, EPU-C simulations completely overlap regardless of ensemble size N or aggregation rule p . This is because, under conditions of perfect correlation, EPU-C reduces to the (non-ensemble) standard LBA model, because each component accumulator acts in an identical fashion.

Looking at the second row of Figure 3D, we begin to see ensemble-like simulations where the drift rate correlation is no longer perfect (but starting point correlation remains fixed at $r_A = 1$). When $r_d = 0.875$ (the second row of Figure 3D), the two extreme aggregation rules ($p = 0$ or $p = 1$) clearly demonstrate a strong dependence on ensemble size, a dependence that persists whenever $r_d < 1$. This makes sense since these stopping rules are most sensitive to the degree of variability between accumulators, with simulated performance at $p = 0$ dominated by the fastest accumulator and simulated performance at $p = 1$ dominated by the slowest; the more accumulators in the ensemble, the faster will be the fastest one and the slower will be the slowest one. Conversely, the more variability there is between accumulators, the bigger will be the difference between the fastest/slowest accumulator and the others in the ensemble. This is evident in Figures 3E and 3G, which expand the sub-plots of the ensemble plot in Figure 3D that correspond to $p = 0$ (Figure 3E) and $p = 1$ (Figure 3G) when there is no drift rate correlation ($r_d = 0$). When $p = 0$ (Figure 3E), increasing ensemble size leads to faster responses overall while when $p = 1$ (Figure 3G), increasing ensemble size leads to overall slower responding.

Still looking at just the second row of Figure 3D, there is a band in which intermediate aggregation rules are approximately invariant across all ensemble sizes (roughly from $p = 0.375$ to $p = 0.625$), with the other (non-extreme) aggregation rules only yielding approximate invariance with more than one accumulator per ensemble. As drift rate correlation diminishes (i.e., as we look at lower and lower rows of Figure 3D), the “central band” in which performance is approximately invariant across all ensemble sizes gets narrower, its last vestige at $p = 0.5$ and $r_d = 0.625$. As drift rate correlation decreases, this central band of approximate invariance is gradually overtaken from the extreme aggregation rules such that, at moderate-to-high levels of drift rate correlation and aggregation rules farther from $p = 0.5$, approximate invariance is only achieved with more than one accumulator and at low levels of drift rate correlation ($r_d < 0.125$), approximate invariance can only be achieved with $N \geq 100$. One example of approximate invariance only for $N \geq 10$ is shown in Figure 3F, which shows how predictions for $N = 1$ and $N = 10$ are visibly different (non-overlapping) from predictions for $N = 100$ or $N = 1000$ which generally overlap one

another, though the differences between $N = 10$ and $N = 1000$ are sufficiently small to be considered approximately invariant according to our criterion.

The overall message from the results summarized in Figure 3D is that when start point correlation is perfect ($r_A = 1$), approximate invariance across all ensemble sizes can be achieved with moderate-to-high levels of drift rate correlation ($r_d > 0.5$) and intermediate aggregation rules (p close to 0.5), though approximate invariance can be found with $N = 10$ or 100 across a far broader range of ensemble parameter regimes. Only under extreme aggregation rules ($p = 0$ or $p = 1$) combined with non-perfect drift rate correlation ($r_d < 1$) was it impossible to find any invariance in simulated performance with ensemble size.

Invariance Results

Figure 4 depicts the full range of simulated EPU-C performance as start point correlation is systematically decreased from $r_A = 1$ to $r_A = 0$ in increments of 0.25 in each separate composite plot. The main feature of these simulations is that the central band of approximate invariance identified in Figure 3D for $r_A = 1$ shrinks as starting point correlation decreases, requiring even higher drift rate correlation and stricter adherence to intermediate aggregation rules until, when $r_A \leq 0.25$, approximate invariance is only observed with large ensemble sizes (i.e., $N = 100$) and sometimes not even then (as when $r_A = 0$ and $r_d = 0$). Reducing correlations between accumulators, either in terms of drift rate or start point, means that there is more variability across the ensemble, and our EPU-C simulations illustrate that this variability yields strikingly different predictions for large ensembles versus small ensembles, such that with no correlations at all (i.e., $r_A = r_d = 0$), it is not even possible for an ensemble with $N = 100$ to approximate the performance of one with $N = 1000$.

Interim Summary

In this section, we considered an important criterion for an ensemble to be able to satisfy the linking proposition between individual neurons (ensemble elements) and computational models of decision making: approximate invariance to ensemble size. If there is a point at which adding more elements to an ensemble fails to change its simulated behavior in an appreciable way, this means that the behavior of the ensemble is not dominated by single elements, and individual elements are representative of the behavior of the full ensemble. EPU-C allows us to characterize, in terms of its ensemble-level parameters, which types of ensemble are capable of achieving this approximate invariance. As we have seen, there are considerable regions of the EPU-C parameter space in which smaller ensembles of various sizes can reasonably approximate the behavior of large ensembles ($N = 1000$), and even some where single accumulators simulate performance that is similar to what would be expected from a large ensemble. Generally speaking, situations in which a small ensemble ($N = 1$ or 10) can approximate behavior from a large ensemble (with $N = 1000$) are those in which both starting point and drift rate are correlated across accumulators at moderate-to-high rates (i.e., r_A and r_d both at least 0.5); within these regimes, approximate invariance is easier to find with intermediate aggregation rules (p close to 0.5) but so long as the termination of aggregation is not completely dominated by either the first

Figure 4a. EPU-C simulated performance across ensemble size (N), aggregation rule (p), and drift rate correlation (r_d) while starting point correlation $r_A = 1$. Conventions are as defined in Figure 3; see main text for additional details (Figure continues on subsequent pages which correspond to decreasing values of r_A).

Figure 4b. (continued)

Figure 4c. (continued)

Figure 4d. (continued)

Figure 4e. (continued)

($p = 0$) or last ($p = 1$) to finish, small ensembles (with at least $N = 10$) are still capable of approximating the behavior of large ensembles.

Relating Ensemble Dynamics to Neural Dynamics

A hallmark signature of accumulator models is that accumulated evidence, which varies across trials, reaches a fixed threshold to trigger a choice response associated with the winning accumulator. It was the demonstration that this signature also holds for neurons in FEF (and related brain areas) that established the linking proposition between accumulator models and such neurons (Hanes & Schall, 1996). Specifically, the measured activity that movement-related neurons in FEF reach at the time of saccade is the same level of activity irrespective of response time (RT). Like accumulator models, activity at RT (A_{RT}) in FEF motor neurons is invariant with behavioral RT. As a result, this is another important constraint on the types of ensembles that can realize the linking proposition between neural systems and computational models of decision making.

In an ensemble model like EPU-C, the *measured activity* (accumulated evidence) at RT of any randomly selected individual accumulator element in an ensemble is not guaranteed to reach the same level because of the stochastic nature of the accumulation process and the aggregation rule. On some simulated trials, a randomly selected accumulator will be one of the elements to have reached threshold and contribute to the initiation of a response and its measured RT, while on other trials, that accumulator will not yet have reached threshold. While Zandbelt et al. (2014) demonstrated that an approximate invariance between A_{RT} and RT held for certain parameter regimes in an ensemble model of simple RT, it remains to be seen whether there are parameter regimes in an ensemble model of choice that will also demonstrate an invariance of activity with RT.

Varieties of A_{RT} and RT relationships

In much the same way that a neurophysiologist randomly selects a neuron in an area of the brain when they place an electrode, we randomly selected a single accumulator from EPU-C and “measured” its level of accumulated evidence at RT (the RT simulated by the ensemble). If the element has already reached threshold by this time, there are several possibilities for how its activation might have continued to evolve (see also Purcell & Palmeri, 2017): First, the accumulator might continue to accumulate even after crossing threshold; second, it might maintain activity at the threshold level; third, its level of activity might decay or be reset after hitting threshold. This final possibility does not seem computationally feasible, since it would require a separate ensemble to register when each accumulator element hit threshold since that information would no longer be encoded in the accumulators themselves; there would be no way to tell from its level of activation alone whether a below-threshold accumulator had already reached threshold earlier and since decayed or if it was still increasing and had yet to reach threshold. Further, although it may well be that accumulators do not increase activity at the same rate after crossing threshold as they do before, it seems improbable that they would be strictly bounded at threshold; these thresholds are parameters set to initiate a response, not the physiologically-determined maximum firing rate. Thus, we focus the majority of our attention on models where accumulators continue to accrue evidence even after reaching threshold. We note that, to ensure that

Figure 5. Examples of the range of potential relationships between accumulator activity at RT (A_{RT}) and response time (RT) across different EPU-C parameter settings. Plots collapse over different values of drift rate (v_1) and only include activity in an ensemble accumulator when the “correct” choice alternative associated with that ensemble was chosen.

our model analyses are consistent with neurophysiological analyses, we focus on activity for accumulators associated with “correct” response.

Figure 5 illustrates several examples of these analyses for a small subset of selected EPU-C parameter settings chosen to illustrate the range of possible results. The top row illustrates accumulator activity at RT (A_{RT}) under the assumption that activity is allowed to increase past threshold while the bottom row illustrates A_{RT} under the assumption that activity is bounded at the threshold. Each set of two columns corresponds to different aggregation rules (p), with the left two columns showing the simulated RT- A_{RT} relationship for $p = 0.25$, the middle two columns for $p = 0.5$, and the right two columns for $p = 0.75$. For now, we note several features of these relationships that will be important as we explore them across the full range of EPU-C parameters: First, it is clear that value A_{RT} is sensitive to the threshold B for obvious reasons (it is the activity at RT, which will be a function of accumulator threshold). Second, for unbounded accumulators (allowed to increase past threshold), there is a tendency to observe a negative correlation between RT and A_{RT} for lenient aggregation rules (e.g., $p = 0.25$), a positive correlation for stringent aggregation rules (e.g., $p = 0.75$), and near-zero correlation for intermediate aggregation rules (e.g., $p = 0.5$). Third, an important consequence of bounding accumulator activity at threshold is that it eliminates positive correlations between RT and A_{RT} by limiting the extent to which accumulator activity can increase over time. We will discuss these features, and the reasons for them, in more detail below.

A_{RT} Results

Similar to the composite plots we constructed to examine invariances in simulated behavior across ensemble size, we visualize the relationship between RT and A_{RT} in a series of composite plots as well. Each sub-plot summarizes the A_{RT} by RT relationship using the

line of best fit (illustrated in Figure 5) between RT and A_{RT} for each level of ensemble size (N). Although the composite plots for choice probability and response time incorporated both threshold settings in each sub-plot, because A_{RT} by definition varies with threshold, we produced two sets of composite plots, one for simulations with low threshold ($B = 3$) and one for simulations with high threshold ($B = 4.5$); because these are qualitatively similar, we present only the low threshold conditions in the main text. In addition, we only present in the main text composite plots that assume that accumulators are free to continue increasing in activity past their thresholds until a response is made; as we just illustrated, bounding accumulator activity at threshold merely suppresses correlations when $p > 0.5$ but otherwise does not exhibit any other qualitative differences, as shown in Appendix B.

To situate ourselves in these composite plots, we will look first at Figure 6a, which has the same composite structure as the plots used to visualize behavioral simulations of EPU-C, in that moving from top to bottom corresponds to lower levels of drift rate correlation (r_d) and moving from left to right corresponds to a greater proportion of accumulators needed to select an alternative (p). Figure 6a depicts results when starting point correlation is perfect ($r_A = 1$). The top rows of this Figure illustrate that when drift rate correlation is also perfect ($r_d = 1$) there is no variability at all in A_{RT} because all accumulators are identical; when one has reached threshold, they all have reached threshold. As drift rate correlation decreases (moving from top to bottom in the Figure), variability in A_{RT} increases, because now there is variability between accumulators. If the aggregation rule is exactly $p = 0.5$, such that half of the accumulators are above threshold and half are below at the time of response selection, there is no systematic correlation between A_{RT} and RT. This makes sense, since on any given trial, a randomly chosen accumulator could have activity distributed evenly around threshold. When aggregation rules differ from $p = 0.5$ and drift rate correlations are not perfect ($r_d < 1$), it is possible to detect correlations between A_{RT} and RT, with the magnitude of these correlations increasing the farther p is from 0.5. If $p < 0.5$, these correlations tend to be negative, such that the longer the RT, the lower is A_{RT} . These negative correlations arise from the fact that, for RT to be longer, the majority of accumulators would have to have had lower drift rates and would thus have not accumulated as much activation by the time enough other accumulators had reached (or exceeded) threshold⁹. If $p > 0.5$, correlations between RT and A_{RT} are positive; longer RT's are associated with having to wait for slower accumulators to finish which gives more time for other accumulators to keep increasing in activation.

As starting point correlation decreases (Figure 6b–e), the negative correlations between RT and A_{RT} present when $p < 0.5$ become less prominent and, similarly, the positive correlations when $p > 0.5$ also diminish until they are only present to any considerable degree when $p = 0$ or $p = 1$ exactly. This shows how increasing the between-accumulator variability in starting point effectively “washes out” the correlations between RT and any single accumulator.

⁹Although our focus is on the low threshold conditions, we note that this is also the reason why a negative correlation between RT and A_{RT} is less strong with higher threshold ($B = 4.5$): a higher threshold gives more opportunity for low drift accumulators to “recover” activity at longer RTs.

Figure 6a. Summary of the relationships between response time (RT) and the activity at RT (A_{RT}) of a randomly chosen eLBA accumulator across EPU-C parameters, under the assumption that accumulator activity is free to increase past threshold. Plot structure is analogous to that defined in Figure 3 and used in Figure 4. Relationships are summarized using the line of best fit, as illustrated for specific subsets of EPU-C parameters in Figure 5. Sub-plots are shaded according to the squared Pearson correlation between RT and A_{RT} in each sub-plot, with darker sub-plots corresponding to greater squared correlation (this shading is purely as a visual aid). See main text for additional details. Figure continues on subsequent pages, corresponding to decreasing values of starting point correlation (r_A).

Figure 6b. (continued)

Figure 6c. (continued)

Figure 6d. (continued)

Figure 6e. (continued)

Interim Summary

Our simulation results suggest that there are many parameter regimes of EPU-C under which an ensemble exhibits a negligible correlation between A_{RT} and RT, consistent with the neurophysiology of cells thought to be involved in decision making. Generally, the farther from $p = 0.5$ EPU-C parameters get, the stronger the predicted negative (if $p < 0.5$) or positive (if $p > 0.5$) correlation between RT and A_{RT} , though these correlations are mitigated to the extent that there is also between-accumulator variability in starting point ($r_A < 1$). This is a case where starting point correlation (r_A) has an opposite effect of drift rate correlation (r_d): While low starting point correlation decreases correlations between A_{RT} and RT, low drift rate correlation tends to increase correlations though only by a small amount. In addition, the fact that low starting point correlation increases the range of parameter settings consistent with a negligible correlation runs in the opposite direction from the tendency of low starting point correlation to make it harder for small ensembles to approximate large ones (as noted in the previous section).

Are there parameter regimes where EPU-C makes correct qualitative predictions about the correlation between RT and A_{RT} and also yields predictions that are approximately invariant to ensemble size? As noted above, invariance is easiest to find with high correlations in terms of both starting point ($r_A \geq 0.5$) and drift rate ($r_d \geq 0.5$) in conjunction with moderate aggregation rules (close to $p = 0.5$). Achieving approximately zero correlation between RT and A_{RT} as well as approximate behavioral invariance would seem to require keeping p close to 0.5, since otherwise the high inter-accumulator correlations in drift rates and starting points needed to achieve approximate invariance to ensemble size would also lead to strong correlations between RT and A_{RT} .

Relating Accumulator Model Parameters to Ensemble Model Parameters

In this section, we consider a final important constraint on the types of ensembles that can realize the linking proposition between neurons and computational models of decision making: whether ensemble behavior can be reasonably approximated by a typical computational model. This constraint is related to the approximate invariance across ensemble sizes, which effectively probes whether a small ensemble of elements produces behavior that is similar to a large ensemble whose elements have the same parameter values. Here, we are essentially asking if an “ensemble” of size $N = 1$ is capable of predicting behavior similar to that simulated by an ensemble with $N > 1$, but potentially with different parameter values from those that describe the individual accumulator elements of the ensemble. In studying this constraint, we also gain insight into how typical non-ensemble models of decision making (like LBA) might interpret changes in ensemble-level parameters like p , r_A , and r_d . Maintaining the connection between individual elements and accumulator models, we fit a typical LBA model (S. Brown & Heathcote, 2008; Terry et al., 2015) to the *behavior* (responses and RTs) simulated by EPU-C, exactly as if EPU-C were a participant (human or monkey) in a (simulated) experiment.

Model Fitting Methods

For each combination of EPU-C ensemble parameters (N , r_d , r_A , and p), we treated the thousands of simulated trials from 2 levels of threshold by 4 levels of drift rate as “data”

from a factorial experiment. We fit an LBA with one accumulator per choice alternative (the standard LBA) to each of these experiments separately by using gradient ascent to find a single set of LBA parameters that maximized the log-likelihood across all simulated trials¹⁰. The parameters of the standard LBA that we allowed to vary were a single value of starting point variability (A_{LBA}) for both accumulators, a single value of the summed drift rate between accumulators ($\bar{v}_{LBA}$), a single value of drift rate variability (s_{LBA}), a single value of residual time ($t_{0,LBA}$), two values of threshold corresponding to the two levels of speed-accuracy emphasis, and four values of drift rate for alternative 1 ($v_{1,LBA}$) corresponding to the four levels of difficulty, with the constraint that the drift rate for alternative 2 is $v_{2,LBA} = \bar{v}_{LBA} - v_{1,LBA}$. The LBA we fit to EPU-C simulated behavior thus had 10 free parameters in total, and these were the same set of parameters that we used to govern the individual accumulator elements in EPU-C. Note, however, that we distinguish the estimated LBA parameters with the subscript “LBA”, so as to differentiate these parameters from the eLBA parameters used to simulate EPU-C behavior.

Accumulator-to-Ensemble Fitting Results

Before presenting the results of our fits to simulated EPU-C behavior, we want to again emphasize that the estimated parameter values we report are with respect to a standard LBA whereas the generating parameter values are with respect to the eLBA accumulator elements within an ensemble of size N . As a result, these fits are not a parameter recovery exercise except in the case of $N = 1$, even though the parameters being fit are conceptually similar to the parameters used to govern the accumulator elements within EPU-C.

Quality of fit. To begin, Figure 7 illustrates the value of the mean KS statistic across conditions between the predicted choice and RT behavior of the best-fitting standard LBA and the simulations produced by EPU-C. This figure has a very similar composite structure to those we have used throughout this article, only now each sub-plot presents the mean KS statistic on the vertical axis and EPU-C starting point correlation (r_A) along the horizontal axis. The KS statistic is the same one we computed earlier to quantify the difference in predicted choice probabilities and RT distributions between different ensemble sizes within EPU-C; now we are using it to quantify differences in predicted choice and RT distributions between the best-fitting standard LBA and simulated EPU-C. What is striking about Figure 7 is the wide range of EPU-C parameters for which the mean KS statistic is quite small, generally less than the 0.05 threshold used earlier to delineate approximate invariance. This suggests that in many cases, the standard LBA, which is effectively an ensemble of size $N = 1$, can reasonably approximate behavior that was in fact simulated by a large ensemble of accumulators. This approximation fails when across-accumulator drift rate correlations and across-accumulator starting point correlations are both quite low and gets worse as ensemble size increases. But so long as at least one source of variability is reasonably strongly correlated between accumulators (i.e., either $r_d \geq 0.5$ or $r_A \geq 0.5$,

¹⁰Although we considered using Bayesian methods for model fitting, the large number of trials (80,000 per simulated experiment) being fit meant that the resulting posterior distribution was tightly clustered around the maximum likelihood estimate anyway. Further, because we assume all data come from a single (simulated) “participant” and different EPU-C variants are independent of one another, there is no hierarchical structure for which Bayesian estimation would be better suited.

Figure 7. Komolgorov-Smirnoff (KS) statistic as a measure of misfit between a standard LBA model and EPU-C simulations.

roughly speaking) then the standard LBA can reasonably approximate ensemble behavior, even across large differences in ensemble size and aggregation rule.

Comparison of LBA and eLBA parameters. Even in situations where the standard LBA yields good quantitative fits to behavior produced by an ensemble, there is no guarantee that it does so using parameter values that are in any way close to those that describe the individual eLBA ensemble elements. Indeed, the simulation results we have already presented show that it is generally not possible for the same individual element parameters to yield approximately invariant ensemble predictions between ensembles with $N = 1$ and those with $N > 1$ except within a fairly narrow range of parameter settings. So how does the typical LBA (with, by definition, $N = 1$) do it?

First, we note that the best-fitting LBA parameters match almost exactly those of the eLBA for “ensembles” of size $N = 1$ (Figure 8). The only parameter for which there appears to be any divergence between the best-fitting LBA and the eLBA is the drift rate in the chance condition, i.e., $v_{1,1}$, the true value of which is zero. This divergence, which is confined to within 0.01 of the true zero drift rate, reflects the slight deviations from exactly 50/50 responding that are due to having a finite number of simulated EPU-C trials.

Figures 9, 10, and 11, show differences between estimated LBA parameters and the parameters governing the eLBA accumulator elements. These Figures only plot parameter differences for LBA fits that provide a good approximation to EPU-C behavior (for which the KS statistic is less than 0.05), so as to focus on *how* the LBA achieves a good approximation to EPU-C. Thus, in keeping with the fit results shown in Figure 7, there are fewer points plotted for larger ensemble sizes reflecting the fact that the LBA has a harder time approximating the behavior of larger ensembles under certain ensemble parameter regimes. We also note that the qualitative relationships between LBA and EPU-C parameters appear consistent across ensemble sizes that are greater than one, conditioned on the LBA providing a good fit to EPU-C behavior. As a result, we will discuss the relationships between LBA and eLBA parameters generally without reference to N , since Figures 9, 10, and 11 are qualitatively similar.

We now discuss each parameter in the order in which they are plotted, from top to bottom in Figures 9, 10, and 11:

Residual time (t_0). The best-fitting standard LBA residual time parameter generally matches the EPU-C residual time parameter (except when EPU-C has a more stringent aggregation rule, $p > 0.5$). In these cases, the estimated LBA t_0 is on average *less* than the corresponding EPU-C parameter, despite the fact that higher p generally leads to slower responses. This indicates that the LBA accommodates higher p not by a simple shift of response times, but in terms of variability associated with the accumulation process itself.

Drift rates ($v_{1,1}, v_{1,2}, v_{1,3}, v_{1,4}$). The standard LBA also effectively “recovers” the relative balance of evidence across conditions as reflected in the v_1 parameters that give the difference in drift rates between alternative 1 and alternative 2. Again, when $p > 0.5$, the LBA drift rates tend to be lower than those of the eLBA, accounting in part for the additional time involved with a more stringent EPU-C aggregation rule. Nonetheless, this divergence is equivalent across all three above-chance levels of drift ($v_{1,2}$, $v_{1,3}$, and $v_{1,4}$), such that the LBA still respects the relative levels of evidence implied by these drift rates.

Starting point variability (A). While the residual times and drift rates of the LBA are more sensitive to the EPU-C aggregation rule (p) than to either drift rate or

Figure 8. Differences between the eLBA parameters used to simulate EPU-C behavior and the best-fitting corresponding standard LBA parameters. Differences are plotted only for LBA fits for which the KS statistic (Figure 7) is less than 0.05.

Figure 9. (continued from Figure 8)

starting point correlation or ensemble size, the LBA starting point variability is sensitive to EPU-C starting point correlation (r_A). Across all conditions, lower r_A leads to lower estimates of A_{LBA} . One might have expected the opposite—less starting point correlation between ensemble elements produces more variability in starting point values across the ensemble. Instead, LBA starting point variability reflects the average effect of starting point variability across eLBA elements; the variance of an average of random variables (like the random starting points across an ensemble) has lower variance than those variables to the extent that they are uncorrelated.

Response thresholds (B_1 , B_2). EPU-C starting point correlation (r_A) has the same effect on LBA response thresholds as it does on LBA starting point variability. As r_A increases, so do *both* starting point variability (A_{LBA}) and threshold (B_{LBA}). As we subsequently explore in more detail, this hints at a potential role for r_A in governing speed-

Figure 10. (continued from Figure 9)

accuracy tradeoffs. Similarly, increasing p is associated with higher LBA threshold estimates, which can be most clearly seen by comparing situations when $r_A = 1$. Finally, there is a hint of a nonmonotonic relationship between LBA response thresholds and drift rate correlation, such that increasing r_d leads to increasing threshold up to about $r_d = 0.25$ but increasing $r_d > 0.25$ leads to *lower* estimates of B_{LBA} and A_{LBA} .

Drift rate variability (s). Thus far, none of the LBA parameters we have explored have been strongly monotonically related to EPU-C drift rate correlation (r_d), but LBA drift rate variability is. As r_d increases, so does s_{LBA} , reflecting greater between-trial variability in the estimated LBA drift rates. As with the relationship between r_A and A_{LBA} , this may seem counterintuitive at first, but reflects the fact that the LBA captures the *average* across ensemble elements, and that the average of a set of correlated quantities has higher variance than an average of a set of uncorrelated quantities. More heterogeneous ensembles,

Figure 11. (continued from Figure 10)

regardless of aggregation rule, act to amplify the signal-to-noise ratio of the evidence being accumulated by the individual eLBA elements, and this is reflected in the LBA in terms of both starting point (A_{LBA}) and drift rate (s_{LBA}) variability.

Sum of drift rates between alternatives ($\bar{v}$). Recall that the v_1 parameters reflect the *difference* between the drift rate for alternative 1 versus alternative 2, while $\bar{v}$ reflects the sum of those drift rates. In our EPU-C simulations, the sum of the drift rates was constrained to be 1, but we allowed it to be a free parameter when fitting the LBA to simulated EPU-C behavior. The sum of drift rates $\bar{v}$ can be thought of as the total amount of “energy” or “evidence” driving the system. This parameter appears to be very sensitive to EPU-C aggregation rule (p), such that more lenient aggregation rules are associated with higher $\bar{v}$, in part accounting for the faster responding associated with lower values of p because *both* alternatives accumulate evidence faster.

Interim Summary

The standard LBA is capable of predicting behavior similar to that simulated by EPU-C, so long as the ensemble’s elements display sufficiently strong correlation in terms of either drift rates or starting points. Indeed, the situations in which the LBA is capable of fitting EPU-C behavior overlap considerably with those in which it is possible to achieve approximate invariance to ensemble size with a sufficiently large ensemble: There should be at least some degree of correlation ($r_d > 0$ and/or $r_A > 0$) and aggregation rules should not be too extreme ($p > 0$ and $p < 1$). In terms of the parameters of the standard LBA, it closely recovers the balance of evidence feeding into the ensemble—the differences in drift rates between alternatives—even in situations where other parameters differ from those of the generating eLBA components. Meanwhile, the standard LBA accounts for ensemble-level EPU-C parameters via a number of routes. Larger ensemble sizes (N) lead to lower estimates of variability parameters (A_{LBA} and s_{LBA}) to the extent that ensemble elements are uncorrelated along the relevant dimensions (r_A and r_d , respectively). The effects of other ensemble parameters on estimated LBA parameters are summarized in Figure 12, which focuses on just the fits for ensembles of size $N = 1000$, including a depiction of the ratio of A_{LBA}/B_{LBA} since these two LBA parameters were found to be correlated above. The effect of increasing the proportion of accumulators required for responding (p) emerges in several standard LBA parameters, including increasing thresholds (B_{LBA}) relative to the range of starting point variability (A_{LBA}/B_{LBA}), but especially leads to a decrease in the overall drift rate in both accumulators ($\bar{v}_{LBA}$). Increasing drift rate correlation (r_d) leads to higher estimates of drift variability (s_{LBA}), while increasing starting point correlation (r_A) is associated not just with higher starting point variability relative to threshold (A_{LBA}/B_{LBA}), but with higher response thresholds overall.

Speed-Accuracy Tradeoffs by Ensembles

As noted earlier, evidence suggests that the neural implementation of speed-accuracy tradeoff (SAT) cannot rely on individual ensemble elements adjusting their thresholds (Heitz & Schall, 2013; Servant et al., 2019), so SAT could arise from other parametric changes within an ensemble of accumulators. Our EPU-C simulations varied four ensemble properties: size (N), aggregation rule (p), starting point correlation (r_A), and drift rate correlation (r_d). To illustrate whether any of these parameters yield SAT, Figure 13 presents EPU-C simulations in a specific condition that most clearly demonstrates the differences in behavior between different values of ensemble parameters (other conditions are qualitatively identical). Because ensemble parameters can interact with one another, we plot the same behavior in three different ways, each of which makes it easier to see how a particular ensemble parameter is jointly related to choice probabilities and RTs. A speed-accuracy tradeoff is evident whenever the probability of choosing the “correct” alternative decreases (lower accuracy) while RT also decreases (higher speed). Finally, Figure 13 is constructed so as to emphasize those combinations of EPU-C parameters that are jointly consistent with the three constraints, since any ensemble-based explanation of SAT must lie within these parameter regions if it is to plausibly realize the linking proposition between neurons and evidence accumulation.

Figure 12. Relationships between EPU-C ensemble parameters (aggregation rule p , drift rate correlation r_d , and starting point correlation r_A) and both behavior (choice probability and median response time [RT]) and best-fitting standard LBA parameters. Plots are restricted to just conditions with ensemble size $N = 1000$, low threshold ($B = 3$), and EPU-C drift rate $v_1 = 0.2$. Dashed horizontal lines in parameter plots show the value of the corresponding parameter in the eLBA components of EPU-C.

Figure 13a. Behavioral effects of varying EPU-C ensemble parameters: aggregation rule p , drift rate correlation r_d , and starting point correlation r_A . For illustration, plots are limited to conditions with ensemble size $N = 1000$, low threshold ($B = 3$), and drift rate $v_1 = 0.2$. Plotting colors and symbols emphasize the combinations of EPU-C parameters that are consistent with different constraints: approximate invariance between small and large ensembles; absolute value of the correlation between response time (RT) and activity at response time (A_{RT}), denoted $|\text{cor}(RT, A_{RT})|$; and whether the LBA successfully fits EPU-C behavior. Each subplot depicts the same data, just reorganized to emphasize effects of different ensemble parameters on predicted behavior. This plot places aggregation rule (p) along the horizontal axis.

SAT via aggregation rule

As we noted, EPU-C can be viewed in as a “meta-accumulator” model, in that its decisions are made by accumulating the number of accumulators associated with each alternative to reach threshold. The aggregation rule in EPU-C thus plays the role of a threshold, but at a higher level than that of the individual accumulator elements out of which EPU-C is built. Consistent with this interpretation, when EPU-C behavior is fit by an LBA, higher aggregation rules map in part onto higher estimates of the LBA threshold parameter (see Figure 12). It is also the case that lower values of p lead to faster responses that are, on average, closer to chance levels, i.e., a speed-accuracy tradeoff (Figure 13a). Except when both drift rate and starting point correlation are exactly one, increasing p leads to an increase in the whole RT distribution; this is only accompanied by a substantial increase in accuracy, however, if starting point correlation is neither perfect ($r_A = 1$) nor absent ($r_A = 0$). Increasing p does not help to increase accuracy if there is too much variability in starting point relative to threshold ($r_A = 0$) such that individual accumulator activity

Figure 13b. (continued). Drift rate correlation (r_d) along the horizontal axis. Figure 13c. (continued). Starting point correlation (r_A) along the horizontal axis.

is dominated by starting point variability rather than accumulated evidence. Conversely, increasing p does not help to “wash out” any extraneous starting point variability if there is no such variability within the ensemble ($r_A = 1$).

Because extreme aggregation rules ($p = 0$ or $p = 1$) are inconsistent with any of the constraints we investigated (except when both $r_A = 1$ and $r_d = 1$, such that the ensemble is equivalent to a single accumulator), p cannot be adjusted across its full range in order to implement SAT in a strategic way. Similarly, the range over which p might be varied while remaining consistent with either or both of these linking propositions depends on the degree of correlation within the ensembles, with essentially the full range (excluding the extremes) available if both $r_A \geq 0.5$ and $r_d \geq 0.5$. As a result, in situations where correlations between accumulator elements are expected to be moderate or high, modulating aggregation rule in terms of p offers a potent strategy for modulating SAT without adjusting the thresholds of the individual accumulators (below, we discuss how these correlations might arise and themselves be modulated). Modulating p could be accomplished by adjusting the threshold of a higher-level *ensemble processing unit* that keeps track of how many input elements have reached threshold.

SAT via drift rate correlation

Our LBA fits to EPU-C simulations suggested only a weak relationship between fitted threshold (B_{LBA}) and EPU-C drift rate correlation (r_d), with changes in r_d being accounted for instead by the LBA’s drift rate variability parameter (s_{LBA}). That said, changes in drift rate correlation do lead to changes in choice proportion and response time that sometimes represent a speed-accuracy tradeoff (Figure 13b). Increasing drift rate correlation always leads to decreasing accuracy, an effect that we noted above is due to the fact that increasing correlation reduces the amount of information contributed by each individual ensemble element. However, for more lenient aggregation rules ($p < 0.5$), increasing drift rate correlation also leads to *increasing* response time; this is not a case of speed-accuracy tradeoff. Only when $p \geq 0.5$ does median RT decrease along with accuracy, while the tail of the RT distribution (represented by the 0.9 quantile) shows a somewhat more complex pattern. We saw hints of this in the non-monotonic relationship between r_d and estimated response thresholds in the LBA, above.

Why is it that drift rate correlation (and, as we shall see, starting point correlation as well) leads to faster and less accurate responses with more stringent aggregation rules ($p \geq 0.5$) but to slower and less accurate responses with more lenient aggregation rules ($p < 0.5$)? In short, it is because these different aggregation rules are sensitive to different kinds of variability between accumulators: The speed of a more stringent aggregation rule ($p \geq 0.5$) is sensitive to the slowest accumulators, and increasing correlation means the slowest accumulators are not much slower than any of the others, such that the final response time is faster. The speed of a more lenient aggregation rule ($p < 0.5$) is sensitive to the fastest accumulators, and increasing correlation means that the fastest accumulators are not much faster than any of the others, such that the final response time is slower. By reducing variability between accumulators, correlations cause the fastest accumulators to be (relatively) slower and the slowest accumulators to be (relatively) faster; lenient aggregation rules ($p < 0.5$) are sensitive to the low end of the distribution of accumulator speeds, while stringent aggregation rules ($p \geq 0.5$) are sensitive to the high end. Correlation pulls together

both the upper and lower ends of the distribution of finishing times across accumulators within an ensemble, speeding responses when $p \geq 0.5$ but slowing them when $p < 0.5$.

The coupling between drift rate correlation and aggregation rule means that modulating r_d is only a viable way to induce SAT if $p \geq 0.5$. Within this regime, there are still a wide range of parameters that are consistent with the constraints we identified, though again the range of possible variation is broader when both $r_A \geq 0.5$, and especially if $r_d \geq 0.5$ as well. Later in the General Discussion, as we relate our model to other models in the literature, we note that there are various ways that drift rate correlations might arise, though not all of them would seem to be capable of being modulated strategically. One important source of drift rate correlations is the source of stimulation; if neurons are receiving inputs generated from the same source, it is natural that this shared input would lead to shared drift rates, though of course a decision maker in an experiment would rarely have direct control over the stimuli provided by the experimenter. A decision maker could, however, modulate the input to ensemble elements by adapting receptive fields to allow for more or less overlap (this is a core mechanism in the Neural Theory of Visual Attention, described below; Bundesen, Habekost, & Kyllingsbæk, 2005) and could also induce correlations via interactions between or within ensembles (Salinas & Sejnowski, 2001). In addition, even if a decision maker could not *increase* drift rate correlation, they might be able to *decrease* it (leading to higher accuracy and slower responding), consistent with accounts that associate attention with a reduction in neural correlations (M. R. Cohen & Maunsell, 2009).

SAT via starting point correlation

Starting point correlation corresponds to a different theoretical construct that operates on a different time scale from drift rate correlation: Drift rate correlation implies that ensemble elements receive similar evidence signals or interact with one another during the course of the accumulation process, such that their activity *within* a trial is similar. Starting point correlation implies that ensemble elements start at similar levels of initial activation, suggesting similar variability *between* trials. In fitting the standard LBA model to the simulated behavior produced by EPU-C, we found that starting point correlations lead to lower estimates of threshold parameters, suggesting that, like p , changes in r_A could manifest as a speed-accuracy tradeoff. Despite its conceptual difference from drift rate correlation, starting point correlation has similar effects on simulated behavior (Figure 13c). Increasing r_A always leads to decreases in accuracy, but only if $p \geq 0.5$ does it lead to increases in speed as well (at least as measured by the median RT; the tail of the RT distribution is only sped up if $p \geq 0.875$). So much like r_d , modulating r_A would only lead to SAT in a specific range of aggregation rules where $p \geq 0.5$. And, again, the full range of r_A is only available for modulation if drift rate correlation is moderate or high ($r_d \geq 0.5$).

The reason why higher r_A leads to SAT when $p \geq 0.5$ is conceptually similar to why increasing baseline or lowering thresholds induces SAT in single accumulators: Increasing starting point correlation within an ensemble increases the influence of the initial state of the system on its eventual behavior, relative to the influence of the incoming evidence (represented as drift rates). If starting points are uncorrelated, they are effectively noise that eventually gets overwhelmed by incoming evidence signals (drift rates) such that the final state of the system is more dependent on the quality of evidence than its initial state. In an ensemble with uncorrelated starting points, a single accumulator starting at a high

level by chance would, under most aggregation rules (with sufficiently high p), have little influence on the final choice outcome because it would almost certainly get balanced out by other accumulators with lower (random) starting points. But in an ensemble with correlated starting points, all accumulators would tend to start at the same level such that, regardless of aggregation rule, a high starting point would lead to a fast response that left little room for evidence to be accumulated.

But how might ensembles modulate their degree of starting point correlation? One way to induce a correlation in starting points within an ensemble would be to inject each ensemble element with a shared input signal; because this input is shared across elements but is not affected by evidence, it would induce correlations in starting levels but not necessarily in drift rates. Moreover, if this shared input were excitatory, this would lead to both higher baseline neural activity and higher starting point correlation. Simulations of a population coding ensemble choice model (Standage, Wang, & Blohm, 2014) similarly found that a shared excitatory input could modulate SAT, but our simulation results suggest that it is not the higher activity *per se* but the correlation it induces that is most important. Further, empirical recordings of movement-related neurons responsible for choosing saccade targets have found that inputs from other regions to those neurons are indeed increased under speed stress, consistent with a shared excitatory input being used to induce starting point correlations (Heitz & Schall, 2013). And while Heitz and Schall (2013) found reduced spike-field coherence under speed stress, which might be considered evidence against an correlation-based mechanism for SAT, this was only for higher frequency bands (around 40 Hz); they found increased coherence under speed stress in lower frequency bands (below 10 Hz) which would correspond to the slower rate at which start points might be expected to vary relative to the fast rate at which sensory evidence evolves over time.

Summary

To an extent, all three ensemble parameters we investigated are capable of modulating SAT without adjusting the thresholds of individual ensemble elements, thus presenting possible solutions to the paradox of how neurons implement SAT relative to how SAT is conceptualized in models of decision making like LBA. But the constraints required to satisfy the linking proposition regarding the mapping between individual neurons and computational models of decision making, as well as the fact that the two correlation parameters only result in SAT under certain aggregation rules, mean that there are limits to how well each of these parameters can explain SAT. Although aggregation rule p is conceptually easy to understand as a kind of “meta-threshold”, being able to modulate aggregation rule beyond roughly $p = 0.5$ requires that at least one form of correlation be present to a sufficient degree ($r_A \geq 0.5$ or $r_d \geq 0.5$). Drift rate and starting point correlation can also modulate SAT, but only yield a true tradeoff (rather than just a decrease in accuracy) when $p \geq 0.5$ and being able to modulate them over their full range requires that the other form of correlation be sufficiently high as well.

As we discuss later, many theories of simple decisions would seem consistent with the presence of drift rate correlations when extended to an ensemble, largely because each element of an ensemble accumulates evidence from the same source such that any variability in the underlying stimulus representation will tend to be mirrored in the variability of the neural evidence driving a decision. While this means that the constraint requiring at

least one form of correlation is likely to be satisfied in most situations (because it is likely that $r_d \geq 0.5$), this also means that it is unclear whether r_d could be modulated enough to achieve a substantial SAT. On the other hand, the presence of drift rate correlations enables both p and r_A to be modulated over a broader range, making it easier for these ensemble parameters to implement a broader range of SAT. Modulating p could be accomplished by modulating the thresholds not of individual elements, but of the higher-order ensemble processing units that keep track of how many of these elements have reached threshold. Because starting point correlation reflects the state of the system prior to the onset of evidence accumulation, it could be modulated by entraining elements to a common oscillatory rhythm (a mechanism closely linked to attention and expectation; de Oliveira, Thiele, & Hoffmann, 1997; Fiebelkorn & Kastner, 2019) or, as noted above, by exciting all elements of an ensemble with a common shared input (Standage et al., 2014), which has the added benefit of being consistent with an increase in baseline activity for decision-related neurons when under speed stress (Heitz & Schall, 2012; Reppert et al., 2018).

Of course, there is no *a priori* reason why only one of these parameters need be modulated by a decision maker in order to achieve a SAT by an ensemble. For example, it seems likely that the mechanisms proposed above for modulating correlations would likely impact *both* starting point and drift rate correlation in the same way. Nonetheless, each of these parameters represents a potential solution to the paradox of how neural systems implement a speed-accuracy tradeoff without adjusting the thresholds of individual accumulator elements and deserve further empirical investigation and can serve as the basis for new theories of this fundamental cognitive phenomenon.

General Discussion

In this article, we confronted a paradox that arises out of the linking proposition that neural systems implement the evidence accumulation processes that lead to decision behavior: How is it that individual neurons act like evidence accumulators when no single neuron is responsible for behavior? And how can computational models of evidence accumulation account so well for decision making behavior when they do not directly represent an ensemble? We confronted this paradox by using an ensemble model of decision making, EPU-C, to explore the types of ensembles that would allow this paradox to be resolved. We addressed three particular constraints on any solution to this paradox: an ensemble must be capable of displaying approximate invariance to ensemble size past a certain point, lest its elements no longer be representative of the whole ensemble; an ensemble must yield approximately zero relationship between the activity of its elements at RT and behavioral RT to be consistent with observed neurophysiology; finally, an ensemble must generate simulated behavior that must be well accounted for by current computational models of decision making. We observed that ensembles consistent with all three of these constraints are those for which aggregation rules are not too extreme, such that decisions are not dominated by either the fastest or slowest accumulator elements, and for which there is at least one source of correlation between elements in an ensemble, either in terms of drift rates or starting points. The range of plausible ensembles thus closely resembles those identified by Zandbelt et al. (2014) in the context of simple reaction time. By extending those prior results to impose constraints on ensemble models of decision making during choice, we could then turn to another paradox: While computational models of decision making account for

speed-accuracy tradeoff via threshold adjustments, neurons do not appear to adjust their critical firing rates in response to speed-accuracy emphasis (Heitz & Schall, 2012, 2013; Reppert et al., 2018). Within the realm of ensembles that remain consistent with the constraints above, we showed that it is possible to achieve a speed-accuracy tradeoff not by adjusting the thresholds of the individual accumulator elements, but by modulating various parameters that govern how those elements are integrated within an ensemble.

Relations To Other Decision Making Models

An advantage of EPU-C is that it makes it possible to transparently instantiate and manipulate ensemble-level parameters like aggregation rule, starting point correlation, and drift rate correlation. A disadvantage of our approach is that, by exploring these ensemble properties in an abstract way, it may not be obvious how they could arise in any specific instance. In this section, we relate the structures of EPU-C to other models of decision making, with a focus on how EPU-C’s ensemble parameters can be related to theoretical constructs in other models.

Neural population decisions. One clear analogue to the present work comes from Xiao-Jing Wang and colleagues (e.g., Wang, 2002; Wong & Wang, 2006) using spiking neural network models. They have thoroughly explored the dynamics of decision making in a model that assumes each alternative is associated with a pool of neurons that are self-excitatory but mutually inhibitory, such that these pools compete with one another until the aggregate activity of one pool exceeds a threshold. This model shares with EPU-C the assumption that different alternatives are associated with an ensemble of elements that individually act as accumulators, but differs in important respects. Our implementation of EPU-C is built around eLBA elements that do not exhibit variability within trials, just across trials, whereas the model neurons in the models of Wang and colleagues exhibit variability within trials, not across trials. EPU-C does not directly implement inhibitory or excitatory connections between elements.

Beyond these technical differences, EPU-C shares important connections with this and other population decision models (such as the Ising model; Verdonck & Tuerlinckx, 2014). Although EPU-C does not directly model connections between accumulator elements, *correlations* between elements in EPU-C could well reflect (at a more fine-grained level) mutually excitatory connections within ensembles or shared inhibitory signals across ensembles. Furthermore, although EPU-C’s aggregation rule depends on a proportion of elements having reached threshold, rather than the overall activity across elements, a rule that depends on the overall activity crossing a threshold should be quite similar to an EPU-C aggregation rule with $p = 0.5$. Although in principle a small number of highly active elements might skew the overall activity upward, across trials this would happen about as often as a small number of strongly inactive elements skewing the activity downward; the close resemblance between such a activity pooling model and one with $p = 0.5$ was previously noted in our prior work modeling simple reaction times (Zandbelt et al., 2014). As we showed in our simulations, aggregation rules close to $p = 0.5$ more often show approximate invariance across ensemble size as well as the correct lack of correlation between activity at RT and response time. This is particularly so when ensemble elements are strongly correlated (as they would be if they were mutually excitatory within a population). Thus, we believe that despite the many surface differences between EPU-C and spiking neural

population models, our results still speak to the circumstances under which they are viable ways of realizing the linking proposition between neural activity and choice behavior.

Continuous response models. Our exploration of ensemble models can be viewed on a continuum of ensemble sizes, with the typical type of accumulator model applied to behavioral data sitting on the lowest end of the continuum when $N = 1$. Although we explored very large values of N further along this continuum, recently several models of decision making with continuous responses (Ratcliff, 2018; Smith, 2016; Smith & Corbett, 2018) have explored what might be considered the limit under which $N \rightarrow \infty$ (Kvam, 2019). These models consider situations in which responses are on a continuous scale, such as a point on a number line, a point on the rim of a circle, or a spot on a screen. As a result, the number of alternatives is effectively unlimited and these models can be viewed (and are often implemented computationally) as assuming that each of the effectively infinite response locations is associated with its own accumulator. While at first glance, this is quite different from our approach where there are a small number of alternatives each of which is associated with many accumulators, it is possible to divide up the response regions into different sectors such that selecting any point within that sector results in the same response, thus assigning many accumulators to a small number of alternatives.

Like EPU-C, these models, particularly that of Ratcliff (2018), allow for correlations between accumulators, where the degree of correlation is a function of the spatial distance between alternatives; that is, accumulators associated with nearby alternatives will be more strongly correlated than accumulators associated with more distant alternatives. The crucial difference between EPU-C and continuous response models is the aggregation rule: In the models of both Ratcliff (2018) and Smith (2016), the first accumulator to reach threshold determines the response; in EPU-C parlance, this is $p = 0$. While the structure of our simulations does not admit easy comparisons between EPU-C and these continuous response models, their aggregation rule suggests they would not yield approximate invariance to ensemble size except with strong correlations and would certainly have difficulty predicting neural activity at RT that is properly uncorrelated with response time (as we showed, $p = 0$ means that longer response times are generally associated with lower measured activity of a randomly selected element). That said, these potential difficulties would need to be investigated and this class of models is sufficiently new that there are many possible variants left to be explored that might overcome any of these potential problems.

Neural theory of visual attention. The Theory of Visual Attention (TVA; Bundesen, 1990; Logan, 2002) posits that rapid perceptual and, to an extent, memory-based decisions are the outcome of a race between an ensemble of accumulators corresponding to different possible categorizations of perceptual objects. When this theory was extended from its original mathematical form to account for neural correlates of attention (nTVA; Bundesen et al., 2005), a core component of the neural theory was that the individual neurons implementing racing evidence could be dynamically reconfigured such that their receptive fields overlapped to different degrees depending on the pertinence of different categorizations. Within the TVA framework, “categorizations” roughly correspond to different responses that can be made on the basis of the features of a perceptual object, such that nTVA posits that the receptive fields of decision-related neurons are sensitive to task demands. If true, this suggests a way by which drift rates between accumulators in an ensemble might become correlated if they are “assigned” to the same categorization—

essentially, they would have overlapping receptive fields and thus share inputs, which would lead to correlated drift rates.

Counter models. Taking a very high-level view, EPU-C can be viewed as a type of *counter model*, in that it “counts” the number of accumulators in each ensemble having reached threshold. Counter models, as a broad class, have a long history in models of decision making (Luce, 1986; Townsend & Ashby, 1983; Vickers, 1979), but like most models of decision making, they rarely model directly the processes that generate the counts that are being accumulated. Instead, these processes are often modeled statistically, typically in terms of Poisson processes that emit counts at different rates for different alternatives, with these rates potentially changing over time (Smith & Van Zandt, 2000). EPU-C instead treats each “count” as itself the outcome of an accumulation process, such that the rate at which counts accumulate over time is an emergent property of these ensemble component processes. Thus, EPU-C offers a way to go beyond the abstractions involved in most counter models, assigning meaning to what each “count” reflects (a neural accumulator reaching threshold) and to the dynamics with which they accumulate (as an emergent property of how accumulators mutually vary among one another).

Perspectives and Extensions

As we noted in our introduction, although we made specific technical assumptions in order to explore the behavior of EPU-C, the general structure of EPU-C exists apart from the specific version used in this article (an ensemble of eLBAs). This general structure is one in which choices are made by assigning *ensemble processing units* to each alternative, each of which is comprised of many individual accumulators, the results of which are themselves accumulated (counted) at a higher level. As we described earlier, this general processing structure is quite similar to the neural population models of Wang and colleagues (e.g., Wang, 2002; Wong & Wang, 2006) such that, even though the specific technical assumptions we made in this article are quite different from such spiking neural network models, the similar computational architecture means that the insights derived in this paper can be applied even to these otherwise dissimilar models. In addition, this shared computational architecture points to a variety of extensions of the EPU-C approach that we believe will provide additional insights into how behaviors arise from ensembles of accumulators.

Alternative accumulator elements. Perhaps the most obvious extension beyond our current implementation of EPU-C would be to swap out the eLBA for a different type of accumulator. Based on our previous work (Zandbelt et al., 2014) and on the wide variety of approximate equivalences between LBA-type models and diffusion-type models (Donkin et al., 2011; Jones & Dzhafarov, 2014; Tillman et al., in press), we do not expect this will have a substantive effect on the qualitative results we report, though the quantitative details (e.g., the precise degree of correlation or aggregation rule required for approximate invariance to ensemble size) will almost certainly differ with different assumptions regarding accumulator elements. One reason to prefer accumulator elements besides eLBA would be to increase the resemblance between accumulator elements and neural dynamics which are clearly not strictly linear. That said, it is entirely possible that the variable and nonlinear firing rate differences observed in neurons during decision making are actually the product of latent linear dynamics (e.g., P. J. Cassey et al., 2016; Zoltowski et al., 2019).

However, there is another important reason to explore different variants for ensemble elements, namely, different forms of correlation. The work we presented here highlights the important role for correlations between elements not just for maintaining an approximate equivalence across ensemble sizes but potentially for modulating speed-accuracy tradeoff as well. However, the forms of correlation we explored in this work were limited to correlations that occur between trials, albeit from different sources: Starting point correlation reflects the degree to which ensemble elements share their initial states while drift rate correlation reflects the degree to which the inputs to ensemble elements are shared across the ensemble. Although we noted that drift rate correlations might arise from interactions between ensemble elements during a trial, a full exploration of correlated *within*-trial variability would require a different model for the individual ensemble elements, like a diffusion. Of course, such an exploration would also be orders of magnitude more difficult computationally, particularly if other ensemble level parameters were varied as in this paper, but we believe efforts along those lines would enrich the picture given by the present EPU-C modeling.

Alternative decision rules. The decision rules we explored in this article assumed that choices and response times were determined by which ensemble first achieved a certain proportion of elements reaching threshold. One extension to this approach would be to allow different ensembles to have different proportions required to initiate a response, which would be a way to implement a response bias at the ensemble level: Generally speaking, an ensemble that required fewer threshold-crossings would tend to “win” more often than an ensemble that required more elements to reach threshold. Another extension to this approach might involve stochastic decision rules—instead of a fixed proportion of elements required on each trial, perhaps the proportion required varies from trial to trial (analogous to threshold or starting point variability at the individual accumulator level) or perhaps each threshold crossing merely increases the likelihood of selecting the associated alternative but does not guarantee that a choice is made at a certain proportion. Different varieties of stochastic decision rules deserve exploration since it is unlikely that real neural systems “keep track” of the individual neurons involved in processing specific stimuli/alternatives on each trial. That said, a stochastic decision rule would amount to a mixture between the fixed proportion rules we explored in this paper, such that the predictions of stochastic models could be derived from the results in this paper as a weighted mixture of the predictions of fixed proportion models.

Interactions between ensembles. The forms of correlation we explored in this article were limited to those that might arise from interactions between elements within an ensemble, though we noted earlier how shared inhibitory signals from other ensembles might also induce drift rate correlations within an ensemble. Of course, there are still a variety of ways that ensembles might interact with one another. Although inhibitory interactions between accumulators are common in some computational models of decision making (e.g., Boucher, Palmeri, Logan, & Schall, 2007; Mordkoff & Yantis, 1991; Purcell et al., 2012; Usher & McClelland, 2001) and have been explored in ensemble models (Wong & Wang, 2006), facilitatory interactions are also eminently plausible, particularly in situations where alternatives are similar to one another or are based on similar sources of evidence (Bussemeyer & Townsend, 1993; Cox & Criss, 2017, 2019, in press; Townsend & Wenger, 2004). As with our explorations of correlated within-trial variability within an ensemble, exploring the all the ways in which ensembles might interact with one another would be

a complex undertaking, but the results presented here can help guide research along these lines by helping to exclude certain model variants that are *a priori* implausible with respect to describing the neural underpinnings of decision making.

Implications For Model-Based Cognitive Neuroscience

We couched the present work in the context of *linking propositions*, theoretical claims about how the cognitive processes that lead to behavior map onto observable properties of the neural systems that implement those processes. Linking propositions are central to model-based cognitive neuroscience, in that the success of this burgeoning field depends on the extent to which the processes and representations instantiated in a cognitive model are linked in a clear and testable way with measurable properties of the brain. For example, an important project in model-based cognitive neuroscience is to use neural measures to inform estimates of parameters of cognitive models (e.g., Turner, Forstmann, & Steyvers, 2019). But this approach will only lead to insights with a proper understanding of the details of the linking propositions between model parameters and neural measures. Our work suggests that although there are ensembles for which the parameters describing individual neurons would be similar to those of a model of behavior (i.e., those for which approximate invariance to ensemble size holds for any $N \geq 1$), there are many plausible ensembles for which this simple mapping is unlikely to obtain. Instead, some aspects of decision making, like speed-accuracy tradeoff, may need to be explained in terms of ensemble-level properties of the system, like inter-element correlations or aggregation rules, that have no direct analog in most standard computational models of decision making. That said, our simulations also illustrated that it was possible to estimate the relative drift rates of eLBA elements when fitting EPU-C behavior using a standard LBA, suggesting that there are still important parameters related to decision making that can be, in principle, inferred directly from the activity of individual neurons.

Given that a complete understanding of the linking proposition between neural systems and computational models of decision making seems to entail both individual- and ensemble-level parameters, combining measurements of individual neural activity with more spatially diffuse measures like EEG could allow for simultaneous inference about both of these levels. That this coupling is at least plausible is supported by the fact that single-unit recordings and EEG often convey similar information (Purcell, Schall, & Woodman, 2013; Servant et al., 2019; Woodman, Kang, Rossi, & Schall, 2007). In addition, multiunit recordings may reveal more details about exactly how neurons within an ensemble may be correlated with one another (cf. E. N. Brown, Kass, & Mitra, 2004; J. Y. Cohen et al., 2010; Maier, Schall, & Woodman, 2018). These empirical advances will put us in a better position to understand the properties of ensembles involved in decision making, which will help us to extend the constraints on the propositions that link choice decision making behavior, the neural activity that gives rise to it, and the computational cognitive models that bridge the gap between brain and behavior.

References

- Annis, J., Gauthier, I., & Palmeri, T. J. (under review). Combining convolutional neural networks and cognitive models to predict novel object recognition in humans. *Manuscript under review*.

- Boehm, U., Annis, J., Frank, M. J., Hawkins, G. E., Heathcote, A., Kellen, D., ... Wagenmakers, E.-J. (2018). Estimating across-trial variability parameters of the diffusion decision model: Expert advice and recommendations. *Journal of Mathematical Psychology*, 87, 46–75. doi: 10.1016/j.jmp.2018.09.004
- Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (2006). The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review*, 113(4), 700–765.
- Bonifazi, P., Ruaro, M. E., & Torre, V. (2005). Statistical properties of information processing in neuronal networks. *European Journal of Neuroscience*, 22, 2953–2964.
- Boucher, L., Palmeri, T. J., Logan, G. D., & Schall, J. D. (2007). Inhibitory control in mind and brain: An interactive race model of countermanding saccades. *Psychological Review*, 114(2), 376–397. doi: 10.1037/0033-295X.114.2.376
- Brown, E. N., Kass, R. E., & Mitra, P. P. (2004). Multiple neural spike train data analysis: state-of-the-art and future challenges. *Nature Neuroscience*, 7(5), 456–461.
- Brown, J. W., Hanes, D. P., Schall, J. D., & Stuphorn, V. (2008). Relation of frontal eye field activity to saccade initiation during a countermanding task. *Experimental Brain Research*, 190, 135–151.
- Brown, S., & Heathcote, A. (2008). The simplest complete model of choice response time: Linear ballistic accumulation. *Cognitive Psychology*, 57, 153–178.
- Bundesden, C. (1990). A theory of visual attention. *Psychological review*, 97(4), 523–47. doi: 10.1037/0033-295X.97.4.523
- Bundesden, C., Habekost, T., & Kyllingsbæk, S. (2005). A neural theory of visual attention: Bridging cognition and neurophysiology. *Psychological Review*, 112(2), 291–328. doi: 10.1037/0033-295X.112.2.291
- Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, 100(3), 432–459. doi: 10.1037/0033-295X.100.3.432
- Cassey, P., Heathcote, A., & Brown, S. D. (2014). Brain and behavior in decision-making. *PLoS Computational Biology*, 10(7), e1003700. doi: 10.1371/journal.pcbi.1003700
- Cassey, P. J., Gaut, G., Steyvers, M., & Brown, S. D. (2016). A generative joint model for spike trains and saccades during perceptual decision-making. *Psychonomic Bulletin & Review*, 23, 1757–1778. doi: 10.3758/s13423-016-1056-z
- Cohen, J. Y., Crowder, E. A., Heitz, R. P., Subraveti, C. R., Thompson, K. G., Woodman, G. F., & Schall, J. D. (2010). Cooperation and competition among frontal eye field neurons during visual target selection. *The Journal of Neuroscience*, 30(9), 3227–3238.
- Cohen, M. R., & Maunsell, J. H. R. (2009). Attention improves performance primarily by reducing interneuronal correlations. *Nature Neuroscience*, 12(12), 1594–1600.
- Colby, C. L., Duhamel, J.-R., & Goldberg, M. E. (1996). Visual, presaccadic, and cognitive activation of single neurons in monkey lateral intraparietal area. *Journal of Neurophysiology*, 76(5), 2841–2852.
- Cox, G. E., & Criss, A. H. (2017). Parallel interactive retrieval of item and associative information from event memory. *Cognitive psychology*, 97, 31–61. doi: 10.1016/j.cogpsych.2017.05.004
- Cox, G. E., & Criss, A. H. (2019). Parametric supplements to systems factorial analysis: Identifying interactive parallel processing using systems of accumulators. *Journal of Mathematical Psychology*, 92(102247), 1–19.
- Cox, G. E., & Criss, A. H. (in press). Similarity leads to correlated processing: A dynamic model of encoding and recognition of episodic associations. *Psychological Review*.
- Cox, G. E., & Shiffrin, R. M. (2017). A dynamic approach to recognition memory. *Psychological Review*, 124(6), 795–860. doi: 10.1037/rev0000076
- de Oliveira, S. C., Thiele, A., & Hoffmann, K. (1997). Synchronization of neuronal activity during stimulus expectation in a direction discrimination task. *The Journal of Neuroscience*, 17(23),

- 9248–9260.
- Ding, L., & Gold, J. I. (2012). Neural correlates of perceptual decision making before, during, and after decision commitment in monkey frontal eye field. *Cerebral Cortex*, *22*, 1052–1067.
- Donkin, C., Brown, S., Heathcote, A., & Wagenmakers, E.-J. (2011). Diffusion versus linear ballistic accumulation: different models but the same conclusions about psychological processes? *Psychonomic Bulletin & Review*, *18*(1), 61–69. doi: 10.3758/s13423-010-0022-4
- Dutilh, G., Annis, J., Brown, S. D., Cassey, P., Evans, N. J., Grasman, R. P. P. P., ... Donkin, C. (2018). The quality of response time data inference: A blinded, collaborative assessment of the validity of cognitive models. *Psychonomic Bulletin & Review*. doi: 10.3758/s13423-017-1417-2
- Fiebelkorn, I. C., & Kastner, S. (2019). A rhythmic theory of attention. *Trends in Cognitive Sciences*, *23*(2), 87–101.
- Forstmann, B. U., Dutilh, G., Brown, S. D., Neumann, J., von Cramon, D. Y., Ridderinkhof, K. R., & Wagenmakers, E.-J. (2008). Striatum and pre-sma facilitate decision-making under time pressure. *Proceedings of the National Academy of Sciences*, *105*(45), 17538–17542. doi: 10.1073/pnas.0805903105
- Forstmann, B. U., Ratcliff, R., & Wagenmakers, E.-J. (2016). Sequential sampling models in cognitive neuroscience: Advantages, applications, and extensions. *Annual Review of Psychology*, *67*(1), 641–666. doi: 10.1146/annurev-psych-122414-033645
- Forstmann, B. U., & Wagenmakers, E.-J. (2015). *An introduction to model-based cognitive neuroscience*. New York: Springer.
- Forstmann, B. U., Wagenmakers, E.-J., Eichele, T., Brown, S., & Serences, J. T. (2011). Reciprocal relations between cognitive neuroscience and formal cognitive models: opposites attract? *Trends in cognitive sciences*, *15*(6), 272–9. doi: 10.1016/j.tics.2011.04.002
- Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, *30*(1), 535–574. doi: 10.1146/annurev.neuro.29.051605.113038
- Hanes, D. P., Patterson, W. F., & Schall, J. D. (1998). Role of frontal eye fields in countermanding saccades: Visual, movement, and fixation activity. *Journal of Neurophysiology*, *79*(2), 817–834. doi: 10.1152/jn.1998.79.2.817
- Hanes, D. P., & Schall, J. D. (1996). Neural control of voluntary movement initiation. *Science*, *274*(5286), 427–430. doi: 10.1126/science.274.5286.427
- Hanks, T., Kiani, R., & Shadlen, M. N. (2014). A neural mechanism of speed-accuracy tradeoff in macaque area LIP. *eLife*, *3*(e02260), 1–17.
- Heathcote, A., & Love, J. (2012). Linear deterministic accumulator models of simple choice. *Frontiers in Psychology*, *3*(292), 1–19.
- Heitz, R. P. (2014). The speed-accuracy tradeoff: history, physiology, methodology, and behavior. *Frontiers in Neuroscience*, *8*. doi: 10.3389/fnins.2014.00150
- Heitz, R. P., & Schall, J. D. (2012). Neural mechanisms of speed-accuracy tradeoff. *Neuron*, *76*(3), 616–628. doi: 10.1016/j.neuron.2012.08.030
- Heitz, R. P., & Schall, J. D. (2013). Neural chronometry and coherency across speed-accuracy demands reveal lack of homomorphism between computational and neural mechanisms of evidence accumulation. *Philosophical Transactions of the Royal Society B*, *368*(20130071), 1–11.
- Holmes, W. R., Trueblood, J. S., & Heathcote, A. (2016). A new framework for modeling decisions about changing information: The Piecewise Linear Ballistic Accumulator model. *Cognitive Psychology*, *85*, 1–29.
- Jones, M., & Dzharov, E. N. (2014). Unfalsifiability and mutual translatability of major modeling schemes for choice reaction time. *Psychological Review*, *121*(1), 1–32.
- Kvam, P. D. (2019). A geometric framework for modeling dynamic decisions among arbitrarily many alternatives. *Journal of Mathematical Psychology*, *91*, 14–37.
- Laming, D. R. J. (1968). *Information theory of choice-reaction times*. London: Academic Press.

- Latimer, K. W., Yates, J. L., Meister, M. L. R., Huk, A. C., & Pillow, J. W. (2015). Single-trial spike trains in parietal cortex reveal discrete steps during decision-making. *Science*, *349*(6244), 184–187.
- Link, S. W., & Heath, R. A. (1975). A sequential theory of psychological discrimination. *Psychometrika*, *40*, 77–105.
- Logan, G. D. (2002). An instance theory of attention and memory. *Psychological Review*, *109*(2), 376–400. doi: 10.1037/0033-295X.109.2.376
- Logan, G. D., Van Zandt, T., Verbruggen, F., & Wagenmakers, E.-J. (2014). On the ability to inhibit thought and action: General and special theories of an act of control. *Psychological Review*, *121*(1), 66–95. doi: 10.1037/a0035230
- Lowe, K. A., & Schall, J. D. (2018). Functional categories of visuomotor neurons in macaque frontal eye field. *eNeuro*, *5*(5), 1–21. doi: 10.1523/ENEURO.0131-18.2018
- Luce, R. D. (1986). *Response times: Their role in inferring elementary mental organization*. New York: Oxford University Press.
- Maier, A., Schall, J. D., & Woodman, G. F. (2018). Neural recordings at multiple scales. In E.-J. Wagenmakers & J. T. Wixted (Eds.), *Stevens' handbook of experimental psychology and cognitive neuroscience: V. methodology*. Hoboken, New Jersey: John Wiley & Sons.
- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. San Francisco: W.H. Freeman.
- Middlebrooks, P. G., Zandbelt, B. B., Logan, G. D., Palmeri, T. J., & Schall, J. D. (2019). Unification of countermanding and perceptual decision-making. doi: 10.1101/158170
- Mordkoff, J. T., & Yantis, S. (1991). An interactive race model of divided attention. *Journal of Experimental Psychology: Human Perception and Performance*, *17*(2), 520–538.
- Nosofsky, R. M., & Palmeri, T. J. (1997). An exemplar-based random walk model of speeded classification. *Psychological Review*, *104*(2), 266–300. doi: 10.1037/0033-295X.104.2.266
- O'Connell, R. G., Dockree, P. M., & Kelly, S. P. (2012). A supramodal accumulation-to-bound signal that determines perceptual decisions in humans. *Nature Neuroscience*, *15*(12), 1729–1735. doi: 10.1038/nn.3248
- O'Connell, R. G., Shadlen, M. N., Wong-Lin, K., & Kelly, S. P. (2018). Bridging neural and computational viewpoints on perceptual decision-making. *Trends in Neurosciences*, *41*(11), 838–852.
- Palmeri, T. J. (1997). Exemplar similarity and the development of automaticity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *23*(2), 324–354. doi: 10.1037/0278-7393.23.2.324
- Palmeri, T. J., Schall, J. D., & Logan, G. D. (2015). *Neurocognitive modeling of perceptual decision making* (Vol. 1; J. R. Busemeyer, Z. Wang, J. T. Townsend, & A. Eidels, Eds.). Oxford University Press. doi: 10.1093/oxfordhb/9780199957996.013.15
- Philiastides, M. G., Ratcliff, R., & Sajda, P. (2006). Neural representation of task difficulty and decision making during perceptual categorization: A timing diagram. *Journal of Neuroscience*, *26*(35), 8965–8975. doi: 10.1523/JNEUROSCI.1655-06.2006
- Pike, R. (1973). Response latency models for signal detection. *Psychological Review*, *80*(1), 53–68.
- Polyn, S. M., Norman, K. A., & Kahana, M. J. (2009). A context maintenance and retrieval model of organizational processes in free recall. *Psychological Review*, *116*(1), 129–156.
- Pouget, P., Emeric, E. E., Stuphorn, V., Reis, K., & Schall, J. D. (2005). Chronometry of visual responses in frontal eye field, supplementary eye field, and anterior cingulate cortex. *Journal of Neurophysiology*, *94*, 2086–2092.
- Purcell, B. A., Heitz, R. P., Cohen, J. Y., Schall, J. D., Logan, G. D., & Palmeri, T. J. (2010). Neurally constrained modeling of perceptual decision making. *Psychological Review*, *117*(4), 1113–1143. doi: 10.1037/a0020311

- Purcell, B. A., & Palmeri, T. J. (2017). Relating accumulator model parameters and neural dynamics. *Journal of Mathematical Psychology*, 76, 156–171. doi: 10.1016/j.jmp.2016.07.001
- Purcell, B. A., Schall, J. D., Logan, G. D., & Palmeri, T. J. (2012). From salience to saccades: Multiple-alternative gated stochastic accumulator model of visual search. *Journal of Neuroscience*, 32(10), 3433–3446. doi: 10.1523/JNEUROSCI.4622-11.2012
- Purcell, B. A., Schall, J. D., & Woodman, G. F. (2013). On the origin of event-related potentials indexing covert attentional selection during visual search: timing of selection by macaque frontal eye field and event-related potentials during pop-out search. *Journal of Neurophysiology*, 109(2), 557–569. doi: 10.1152/jn.00549.2012
- Rae, B., Heathcote, A., Donkin, C., Averell, L., & Brown, S. (2014). The hare and the tortoise: Emphasizing speed can change the evidence used to make decisions. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 40(5), 1226–1243.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85(2), 59–108. doi: 10.1037/0033-295X.85.2.59
- Ratcliff, R. (2018). Decision making on spatially continuous scales. *Psychological Review*, 125(6), 888–935.
- Ratcliff, R., Hasegawa, Y. T., Hasegawa, R. P., Smith, P. L., & Segraves, M. A. (2007). Dual diffusion model for single-cell recording data from the superior colliculus in a brightness-discrimination task. *Journal of Neurophysiology*, 97(2), 1756–1774. doi: 10.1152/jn.00393.2006
- Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation*, 20, 873–922.
- Ratcliff, R., & Rouder, J. N. (1998). Modeling response times for two-choice decisions. *Psychological Science*, 9(5), 347–356. doi: 10.1111/1467-9280.00067
- Ratcliff, R., & Smith, P. L. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychological Review*, 111(2), 333–367. doi: 10.1037/0033-295X.111.2.333
- Ratcliff, R., Thapar, A., & McKoon, G. (2003). A diffusion model analysis of the effects of aging on brightness discrimination. *Perception & Psychophysics*, 65(4), 523–535.
- Reppert, T. R., Servant, M., Heitz, R. P., & Schall, J. D. (2018). Neural mechanisms of speed-accuracy tradeoff of visual search: saccade vigor, the origin of targeting errors, and comparison of the superior colliculus and frontal eye field. *Journal of Neurophysiology*, 120, 372–384.
- Roitman, J. D., & Shadlen, M. N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *The Journal of Neuroscience*, 22(21), 9475–9489.
- Salinas, E., & Sejnowski, T. J. (2001). Correlated neuronal activity and the flow of neural information. *Nature Reviews Neuroscience*, 2, 539–550.
- Schall, J. D. (2001). Neural basis of deciding, choosing and acting. *Nature Reviews Neuroscience*, 2(1), 33–42. doi: 10.1038/35049054
- Schall, J. D. (2004). On building a bridge between brain and behavior. *Annual Review of Psychology*, 55(1), 23–50. doi: 10.1146/annurev.psych.55.090902.141907
- Schall, J. D. (2019). Accumulators, neurons, and response time. *Trends in Neurosciences*, 42(12), 848–860.
- Schall, J. D., & Hanes, D. P. (1993). Neural basis of saccade target selection in frontal eye field during visual search. *Nature*, 366(6454), 467–469. doi: 10.1038/366467a0
- Schurger, A., Sitt, J. D., & Dehaene, S. (2012). An accumulator model for spontaneous neural activity prior to self-initiated movement. *Proceedings of the National Academy of Sciences*, 109(42), E2904–E2913. doi: 10.1073/pnas.1210467109
- Seeley, T. D., Visscher, P. K., & Passino, K. M. (2006). Group decision making in honey bee swarms: When 10,000 bees go house hunting, how do they cooperatively choose their new nesting site? *American Scientist*, 94(3), 220–229.
- Servant, M., Tillman, G., Schall, J. D., Logan, G. D., & Palmeri, T. J. (2019). Neurally-constrained

- modeling of speed-accuracy tradeoff during visual search: Gated accumulation of modulated evidence. *Journal of Neurophysiology*. doi: 10.1152/jn.00507.2018
- Servant, M., van Wouwe, N., Wylie, S. A., & Logan, G. D. (2018). A model-based quantification of action control deficits in parkinson's disease. *Neuropsychologia*, *111*, 26–35. doi: 10.1016/j.neuropsychologia.2018.01.014
- Shadlen, M. N., Kiani, R., Newsome, W. T., Gold, J. I., Wolpert, D. M., Zylberberg, A., ... Roitman, J. (2016). Comment on “single-trial spike trains in parietal cortex reveal discrete steps during decision-making”. *Science*, *351*(6280), 1406-b.
- Smith, P. L. (2016). Diffusion theory of decision making in continuous report. *Psychological Review*, *123*(4), 425–451.
- Smith, P. L., & Corbett, E. A. (2018). Speeded multielement decision-making as diffusion in a hypersphere: Theory and application to double-target detection. *Psychonomic Bulletin & Review*.
- Smith, P. L., & Ratcliff, R. (2004). Psychology and neurobiology of simple decisions. *Trends in Neurosciences*, *27*(3), 161–168. doi: 10.1016/j.tins.2004.01.006
- Smith, P. L., & Van Zandt, T. (2000). Time-dependent Poisson counter models of response latency in simple judgment. *British Journal of Mathematical and Statistical Psychology*, *53*, 293–315.
- Smyrnis, N., Mylonas, D. S., Rezaie, R., Siettos, C. I., Ventouras, E., Ktonas, P. Y., ... Papanicolaou, A. C. (2012). Single-trial magnetoencephalography signals encoded as an unfolding decision process. *NeuroImage*, *59*(4), 3604–3610. doi: 10.1016/j.neuroimage.2011.10.091
- Standage, D., Wang, D.-H., & Blohm, G. (2014). Neural dynamics implement a flexible decision bound with a fixed firing rate for choice: a model-based hypothesis. *Frontiers in Neuroscience*, *8*(318), 1–9.
- Stone, M. (1960). Models for choice-reaction time. *Psychometrika*, *25*, 251–260.
- Teller, D. Y. (1984). Linking propositions. *Vision Research*, *24*(10), 1233–1246. doi: 10.1016/0042-6989(84)90178-0
- Terry, A., Marley, A. A. J., Barnwal, A., Wagenmakers, E.-J., Heathcote, A., & Brown, S. D. (2015). Generalising the drift rate distribution for linear ballistic accumulators. *Journal of Mathematical Psychology*, *68–69*, 49–58.
- Thomas, N. W. D., & Paré, M. (2007). Temporal processing of saccade targets in parietal cortex area LIP during visual search. *Journal of Neurophysiology*, *97*, 942–947.
- Thura, D., & Cisek, P. (2016). Modulation of premotor and primary motor cortical activity during volitional adjustments of speed-accuracy trade-offs. *The Journal of Neuroscience*, *36*(3), 938–956.
- Tillman, G., Van Zandt, T., & Logan, G. D. (in press). Sequential sampling models without random between-trial variability: The racing diffusion model of speeded decision making. *Psychonomic Bulletin & Review*.
- Tolhurst, D. J., Movshon, J. A., & Dean, A. F. (1983). The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Research*, *23*(8), 775–785.
- Townsend, J. T., & Ashby, F. G. (1983). *Stochastic modeling of elementary psychological processes*. Cambridge: Cambridge University Press.
- Townsend, J. T., & Wenger, M. J. (2004). A theory of interactive parallel processing: New capacity measures and predictions for a response time inequality series. *Psychological Review*, *111*(4), 1003–1035. doi: 10.1037/0033-295X.111.4.1003
- 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. doi: 10.1016/j.jmp.2016.01.001
- Turner, B. M., Forstmann, B. U., & Steyvers, M. (2019). *Joint models of neural and behavioral data*. Cham, Switzerland: Springer.
- Turner, B. M., Rodriguez, C. A., Norcia, T. M., McClure, S. M., & Steyvers, M. (2016). Why

- more is better: Simultaneous modeling of EEG, fMRI, and behavioral data. *NeuroImage*, 128, 96–115.
- Ulrich, R., & Miller, J. (1993). Information processing models generating lognormally distributed response times. *Journal of Mathematical Psychology*, 37, 513–525.
- Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review*, 108(3), 550–592. doi: 10.1037/0033-295X.108.3.550
- van Ravenzwaaij, D., Provost, A., & Brown, S. D. (2017). A confirmatory approach for integrating neural and behavioral data into a single model. *Journal of Mathematical Psychology*, 76, 131–141.
- Verdonck, S., & Tuerlinckx, F. (2014). The Ising decision maker: A binary stochastic network for choice response time. *Psychological Review*, 121(3), 422–462.
- Vickers, D. (1970). Evidence for an accumulator model of psychophysical discrimination. *Ergonomics*, 13(1), 37–58. doi: 10.1080/00140137008931117
- Vickers, D. (1979). *Decision processes in visual perception*. New York: Academic Press.
- Wang, X.-J. (2002). Probabilistic decision making by slow reverberation in cortical circuits. *Neuron*, 36(5), 955–968. doi: 10.1016/S0896-6273(02)01092-9
- Wickelgren, W. A. (1977). Speed-accuracy tradeoff and information processing dynamics. *Acta Psychologica*, 41, 67–85.
- Wong, K.-F., & Wang, X.-J. (2006). A recurrent network mechanism of time integration in perceptual decisions. *Journal of Neuroscience*, 26(4), 1314–1328. doi: 10.1523/JNEUROSCI.3733-05.2006
- Woodman, G. F., Kang, M.-S., Rossi, A. F., & Schall, J. D. (2007). Nonhuman primate event-related potentials indexing covert shifts of attention. *Proceedings of the National Academy of Sciences*, 104(38), 15111–15116. doi: 10.1073/pnas.0703477104
- Woodman, G. F., Kang, M.-S., Thompson, K., & Schall, J. D. (2008). The effect of visual search efficiency on response preparation. *Psychological Science*, 19(2), 128–136. doi: 10.1111/j.1467-9280.2008.02058.x
- Zandbelt, B., Purcell, B. A., Palmeri, T. J., Logan, G. D., & Schall, J. D. (2014). Response times from ensembles of accumulators. *Proceedings of the National Academy of Sciences*, 111(7), 2848–2853. doi: 10.1073/pnas.1310577111
- Zoltowski, D. M., Latimer, K. W., Yates, J. L., Huk, A. C., & Pillow, J. W. (2019). Discrete stepping and nonlinear ramping dynamics underlie spiking responses of LIP neurons during decision-making. *Neuron*, 102, 1249–1258. doi: 10.1016/j.neuron.2019.04.031

Appendix A

Komolgorov-Smirnoff Statistic for Joint RT-Choice Distributions

The Komolgorov-Smirnoff (KS) statistic can only be computed between pairs of univariate distributions, but EPU-C predicts two variables (choice and response time), so we need to project EPU-C predictions onto a single dimension. To do so, we make use of the facts that, in our current application, there are only two choice alternatives and that RT’s can only be positive. These properties allow us to project EPU-C predictions onto the full real line by treating response times to alternative 1 as positive and responses to alternative 2 as “negative” response times, as illustrated in Figure A1. The resulting KS statistic thus jointly tests for differences in both choice probabilities and RT distributions, making use of the complete set of predictions (not just the summary values we plot).

Appendix B

Figure A1. Visual illustration of how the Komolgorov-Smirnoff (KS) statistic is computed, which is used to quantify the difference in predictions between levels of ensemble size (N). We transform predicted behavior (top panels) into single dimension by treating response times to alternative 2 as negative and those to alternative 1 as positive (middle panels). The KS statistic is then the largest absolute difference in the empirical cumulative distribution functions (CDF) based on these transformed predictions (bottom panel). See main text for additional detail.

Figure B1a. Summary of the relationships between response time (RT) and the activity of a randomly chosen eLBA accumulator at RT (A_{RT}) across EPU-C parameters, under the assumption that accumulator activity is bounded at past threshold. Plot structure is analogous to that already used in Figure 6. Figure continues on subsequent pages, corresponding to decreasing values of startpoint correlation (r_A).

Figure B1b. (continued)

Figure B1c. (continued)

Figure B1d. (continued)

Figure B1e. (continued)

Bounded A_{RT} -RT Correlations

It is clear from Figure B1 that many of the same features hold for stopping rules with $p < 0.5$, namely, a negative correlation between RT and A_{RT} that increases with between-accumulator startpoint correlation. On the other hand, there is little evidence for positive correlations between RT and A_{RT} when $p > 0.5$, which makes sense given that these are situations in which the majority of accumulators have reached threshold by the time a response is selected. If activity is bounded at threshold, these “finished” accumulators have no variability in activation that could correlate with RT; put another way, for any individual accumulator, it is not possible to predict the RT solely from the fact that its activation is at threshold.