The Autocorrelated Bayesian Sampler: A Rational Process for Probability Judgments,
Estimates, Confidence Intervals, Choices, Confidence Judgments, and Response Times

Jian-Qiao Zhu¹, Joakim Sundh², Jake Spicer¹, Nick Chater³, and Adam N Sanborn¹

¹Department of Psychology, University of Warwick

²Department of Psychology, Uppsala University

³Warwick Business School, University of Warwick

Author Note

Zhu, Sundh, Spicer, and Sanborn were supported by a European Research Council consolidator grant (817492-SAMPLING). Chater was supported by the ESRC Network for Integrated Behavioural Science (ES/P008976/1). Earlier versions of the paper are available on the PsyArXiv preprint server via the following link: <https://psyarxiv.com/3qxf7/>. This work was presented at seminars at University of Birmingham, TU Darmstadt, Princeton University, Central European University, and Cambridge University.

Abstract

Normative models of decision-making that optimally transform noisy (sensory) information into categorical decisions qualitatively mismatch human behavior. Indeed, leading computational models have only achieved high empirical corroboration by adding task-specific assumptions that deviate from normative principles. In response, we offer a Bayesian approach that implicitly produces a posterior distribution of possible answers (hypotheses) in response to sensory information. But we assume that the brain has no direct access to this posterior, but can only *sample* hypotheses according to their posterior probabilities. Accordingly, we argue that the primary problem of normative concern in decision-making is integrating stochastic *hypotheses*, rather than stochastic sensory information, to make categorical decisions. This implies that human response variability arises mainly from posterior sampling rather than sensory noise. Because human hypothesis generation is serially correlated, hypothesis samples will be autocorrelated. Guided by this new problem formulation, we develop a new process, the Autocorrelated Bayesian Sampler (ABS), which grounds autocorrelated hypothesis generation in a sophisticated sampling algorithm. The ABS provides a single mechanism qualitatively explains many empirical effects of probability judgments, estimates, confidence intervals, choice, confidence judgments, response times, and their relationships. Our analysis demonstrates the unifying power of a perspective shift in the exploration of normative models. It also exemplifies the proposal that the “Bayesian brain” operates using samples not probabilities, and that variability in human behavior may primarily reflect computational rather than sensory noise.

Keywords: Behavioral Science, Sampling, Bayesian Models of Cognition, Normative Model, Rational Analysis.

The Autocorrelated Bayesian Sampler: A Rational Process for Probability Judgments, Estimates, Confidence Intervals, Choices, Confidence Judgments, and Response Times

Human judgment and decision making has been studied using a wide variety of measures. Participants are asked to provide probability judgments (e.g., Sloman et al., 2004; Fox & Rottenstreich, 2003; Costello & Watts, 2014; Dasgupta, Schultz, & Gershman, 2017; Zhu, Sanborn, & Chater, 2020), estimates of physical quantities (e.g., Gilden et al., 1995; Jazayeri & Movshon, 2007) with associated confidence intervals (e.g., Juslin & Olsson, 1997; Juslin, Winman, & Hansson, 2007), and choices (e.g., Tversky & Kahneman, 1974; Usher & McClelland, 2004; Wyart & Koechlin, 2016; Fehr & Rangel, 2011) with their associated responses times (e.g., Ratcliff & Rouder, 1998; Krajovich & Rangel, 2011; Ratcliff & Starns, 2009; Blurton et al., 2020) and confidence judgments (e.g., Baranski & Petrusic, 1998; Juslin et al., 2007; Li & Ma, 2020; Pleskac & Busemeyer, 2010). Yet while each measurement has been subject to an enormous amount of empirical and modeling work in psychology, a unified theoretical framework that can provide an integrated account of human performance across all six measures (i.e., probability judgments, estimates, confidence intervals, choices, confidence judgments, and response times) is currently lacking.

Theorists have taken steps towards a unified model from two starting points, normative and descriptive. Existing normative models are elegant, parsimonious and are easily extendable to all six measures. But these models fail to provide a satisfactory account of many qualitative effects that have been observed in empirical data. By contrast, various descriptive models, which systematically deviate from normative assumptions, capture the empirical effects both qualitatively and quantitatively for up to three of these six measures, but no single model makes predictions across all measures.

Here, we develop a simple and consistent process that can account for a surprising variety of qualitative findings across all six measures. To achieve this goal, we build on a

strong normative foundation for all six measures, rooted in a sampling approximation to Bayesian inference. This approach also implies a radical shift in viewpoint concerning the nature of the decision-making process and the origin of variability in human behavior. In the perceptual decision-making literature, existing normative models generally operate on noisy sensory information and evaluate the relative probability that this noisy information is generated by the different hypotheses (corresponding to choice options) (e.g., Green & Swets, 1966; Ratcliff & Rouder, 1998).

Our starting point is that exact Bayesian computations are generally intractable; and hence a Bayesian brain can, at best, only approximate these computations. One of the most widely used approach to such approximation in statistics and machine learning assumes that the brain draws samples from posterior probabilities. Inspired by this approach, many theorists in the Bayesian tradition have argued that the cognitive processes thus operate over these samples, rather than representations of probabilities (Griffiths, Vul, & Sanborn, 2012; Vul et al., 2014; Sanborn & Chater, 2016; Lieder et al., 2018; Dasgupta et al., 2017; Zhu, Sanborn, & Chater, 2020). But this process of sampling will inevitably be noisy—different samples will be drawn on different occasions. Thus, in this type of model, the main source of variability does not arise from sensory noise, but from computational noise caused by the process of sampling. That is, instead of evaluating evidence from the sensory system or memory, we propose that the cognitive system operates on stochastically generated hypotheses.

Our aim in this paper is to outline a general process that can be applied to a wide variety of measures and tasks, when equipped with a task-specific representation. Our focus is to show that this process provides a unified framework which captures a wide range of qualitative phenomena across measures and tasks, rather than to produce a comprehensive quantitative model of a particular task. The following sections are structured as follows. First,

we review the traditional probabilistic view of normative decision making and note its limitations in explaining psychological data. Then we propose an alternative sampling-based approximation approach to alleviate the computational burden associated with the normative models, which in turn suggests a shift in the target problem of normative concern from accumulating sensory data to integrating stochastic hypotheses. We demonstrate the unifying power of the perspective shift by applying a rational process model across the six behavioral measures, emphasizing on qualitative model behaviors. Finally, we explore how to create these complete models after exploring the general judgement and decision-making process in detail.

Overview of Probabilistic Decision Making

The idea that human decision-making process is an optimal, perhaps Bayesian, process is attractive because given a problem description it prescribes the best behavior, and can therefore be justified as an ultimate state of evolution and/or learning (Green & Swets, 1966; Ratcliff & Rouder, 1998; Gold & Shadlen, 2002; Pleskac & Busemeyer, 2010; Bogacz et al., 2006; Moran, 2015; Drugowitsch et al., 2019; Hawkins & Heathcote, 2021; Tickle et al., 2021). There are also a wide variety of task-specific Bayesian models in psychology, but in the area of cognitive and perceptual decision making, they are often elaborations of the general decision process of Signal Detection Theory (SDT), which describes how sensory evidence can be transformed into optimal behavior (Green & Swets, 1966).

To illustrate our discussion, we shall consider the following trial in a perceptual task as a running example: a cloud of 24 dots briefly appears on-screen (i.e., stimulus, s)¹.

¹ We will only occasionally touch on the large and distinct literature on risky decision making based on verbally or numerically problems in the fields of judgment and decision making and

Participants might be asked to report a probability that the number of dots falls within a certain window (i. probability judgments). They may also be asked to estimate the exact number of dots (ii. estimates) or to provide a confidence interval for the estimate (iii. confidence interval). Alternatively, participants could be asked to decide whether the number of dots was greater or smaller than some predefined boundary (iv. choices) and the elapsed times for making such decision (vi. response times), for example, by asking participants whether or not there were greater than 25 dots on the screen. Finally, it is also possible to elicit their confidence rating for either of these responses (v. confidence in the decision).

Figure 1. Illustrations of the scope of behavioral measures for a single task. After the presentation of sensory stimulus, people can be asked a wide range of questions and their responses lead to corresponding behavioral measures.

Importantly, while this choice of example of numerosity judgment is intended to provide a simple and concrete illustration, and one that connects naturally to existing models such as SDT, the general approach applies quite generally. For a wide range of tasks, the six

behavioral economics (e.g., Tversky & Kahneman, 1974), which has largely developed as a separate tradition.

behaviors above can be collected and modelled. So, for example, participants might be asked memory-based questions about how much their last grocery bill was (e.g., “Your last grocery bill exceeded £150. What is the probability that this proposition is correct?”), or even asking participants about one-off future events such as how many years they expect to live. Thus, while we use the numerosity example in Figure 1 because it is simple and straightforward to relate to SDT, our approach applies to cognitive as well as perceptual tasks, as we will see below.

More formally, in the SDT, we wish to choose between option A and B , based on a total of T units of sensory input $(s_1, s_2, \dots, s_T)$ typically assumed to be accumulated over time. Assuming that both options are equally likely a priori (i.e., $P(A) = P(B)$), the key variable is the summed log-likelihood ratio over the evidence from each individual unit of sensory input:

$$L(T) = \sum_{t=1}^T \log \frac{P(s_t|A)}{P(s_t|B)} \quad (1)$$

where $P(s_t|A)$ is the likelihood of sensory evidence s_t when option A is the correct choice (similarly for $P(s_t|B)$). The probability of choosing A over B should be a function of the summed log likelihood ratios. Provided with imperfect evidence (e.g., detecting a ship on a noisy radar image), SDT is a principled way to filter out irrelevant sensory noise to pick out the useful signal (e.g., whether the image contains ship). The approach can be applied to a wide range of domains in psychology, including memory, perception, and reasoning (Kellen et al., 2021; Rotello, 2018; Trippas et al., 2018; McCarley & Benjamin, 2013).

SDT, however, makes no explicit commitment on the time course of how evidence is generated and/or collected, and so makes no predictions for response time. This issue can be addressed with a dynamic extension of SDT: the sequential probability ratio test (SPRT), which postulates that the stream of sensory evidence arrives steadily and sequentially over time (Edwards, 1965; Laming, 1968; Bogacz et al., 2006). To deal optimally with the incoming sensory evidence in, for example, binary choice, the evidence should be

continuously integrated into the log-likelihood ratio between the two options until a fixed threshold is reached, and response times are predicted to depend on the amount of evidence accumulated. More formally, the log-likelihood ratio for choosing option A over B is recursively updated after the arrival of each new piece of sensory evidence (s_T):

$$L(T) = L(T - 1) + \log \frac{P(s_T|A)}{P(s_T|B)} \quad (2)$$

Once the log-likelihood ratio reaches a threshold (assuming symmetric thresholds: $L(T) > \delta$ or $L(T) < -\delta$), the evidence accumulation process stops and the response depends on whether the positive or negative threshold is reached. Increasing the magnitude of the threshold (δ) produces a slower but more accurate response as more sensory evidence, on average, is accumulated before either threshold is reached. The SPRT is optimal in the sense that the expected amount of evidence (i.e., T) is minimized for any fixed probability of deciding incorrectly (Wald & Wolfowitz, 1948). In other words, following the SPRT allows for the fastest response time for a particular error rate. Because the sensory inputs are assumed to be independent of one another, the SPRT is a random walk model whose starting point is $L(0) = 0$ and with two absorbing states: $-\delta$ and δ (see Figure 2A).

Figure 2. Schematic illustrations of the computational mechanisms and potential behavioral outputs of the SPRT (A) and the ABS (B). A typical trial of a numerosity task is visualized in

which 24 dots were briefly presented on-screen as the stimulus. The SPRT draws sequential samples from the noisy sensory representation (e.g., corrupted images), while the ABS draws autocorrelated samples of hypotheses (e.g., numbers of dots).

Figure 3. Further illustrations of the autocorrelated sampling process (top) and the Bayesian Monte Carlo process (bottom), expanding the illustrative example of numerosity in Figure 2B. Here, the sampler was automatically terminated when 5 samples were generated, while the dashed lines denote potential future samples if continued. Samples were compared to a decision boundary of 25 (red dots: evidence for lower than 25 dots, blue dots: evidence for greater than or equal to 25). The 5 samples were then integrated with a prior on responses (here used an asymmetric prior, Beta(2,1)), reaching a posterior of Beta(5,3). The mean of this posterior on responses was then used to generate probability judgments or confidence judgments in decision-making.

While intuitive and simple, the framework of the SPRT also makes imperfect decisions that take time, as people do, which is an advantage over SDT in modeling empirical data. Indeed, the SPRT can produce human-like speed accuracy tradeoffs: requiring faster decisions reduces accuracy, while requiring more accurate decisions reduces speed. This is captured in the model by assuming that people control the magnitude of the thresholds to suit their objectives. In response to experimental emphasis on speed (accuracy), people are assumed to be able to decrease (increase) the decision threshold; the model's guarantee of optimal performance implies that these two measures will trade off against one another.

Unfortunately, the SPRT does not easily explain other psychological relationships between choice and response times. In binary choice, for example, the SPRT predicts identical response time distributions for choosing either of the two options (assuming an unbiased starting point, $L(0) = 0$, and symmetric thresholds), contradicting the empirical observation that mean response times differ for correct and incorrect decisions (Stone, 1960; Ratcliff & Rouder, 1998). This is far from the only issue: Table 1 summarizes several qualitative effects of choice, response time, and confidence, the majority of which cannot be accommodated by the SPRT.

These stylized facts have been used to motivate descriptive models, including the family of models known as Drift Diffusion Models (DDMs) that relaxes the normative SPRT framework to better describe human data, specifically regarding three key measures: choice, response time, and confidence. While such approaches have been highly successful, our focus here remains on approaches closely tied to normative depictions of behavior, though we return to DDMs and other common descriptive models below.

A Representation for Producing Estimates and Confidence Intervals

The categorical-hypotheses representations used in SPRT can produce choice, response time, and confidence measures, but are not fine-grained enough to produce probability judgments (e.g., judge the probability that the number of dots was greater than 25), estimates (e.g., how many dots are there on the screen), or confidence intervals (e.g., placing a 95% confidence interval around the estimate). What is needed is an extension of the hypothesis space beyond the categorical hypotheses used when making a choice. In principle, within a Bayesian framework, this is straightforward, although the resulting model looks very different. Instead of simply using two categorical hypotheses (e.g., whether or not there are more than 25 dots on the screen), the model can instead represent the fine-grained hypotheses relevant for estimates (e.g., the exact number of dots on the screen). With such a representation, estimates and confidence intervals can simply be a function (e.g., the mean and quantiles respectively) of this distribution. The probabilities of categorical hypotheses used to produce choices, confidence judgments, and response times can be calculated simply by summing up the posterior probability of the fine-grained hypotheses that are consistent with each choice (e.g., summing the probability of all the hypotheses in which the number of dots is more than 25).

This representational change, however, does not allow a probabilistic model to account for many empirical effects found with estimates and confidence intervals. For estimates, anchoring effects demonstrate a dependence on preceding choices even when the choice question transparently contains no information (Tversky & Kahneman, 1974). Moreover, estimated confidence intervals are empirically far too narrow and are strikingly different depending on whether participants produce or evaluate them (Juslin et al., 2007). In addition, a long line of empirical work shows that probability judgments are systematically biased and incoherent (e.g., subadditivity, conjunction fallacies, partition dependence), which

argues against *all* purely probabilistic models (e.g., Tversky & Kahneman, 1983; Tversky & Koehler, 1994; Tentori et al., 2013; Dasgupta et al., 2017; Zhu et al., 2020).

Exact probabilistic models also show fundamental mismatches with the results of recent investigations into the source of noise in human judgment and decision making. While probabilistic models assume a noise-free inference process using precise probabilities, there is growing empirical evidence suggesting that much, or even most, variability in decision making in fact arises from “computational noise” (i.e., variability in precision and approximation used to perform inference) rather than “sensory noise” (i.e., variability in relevant sensory features) or “decision noise” (i.e., variability associated with action selection) (Drugowitsch et al., 2016; Wyart & Koechlin, 2016; Findling & Wyart, 2021; Stengård & van den Berg, 2019). Clearly, then, there are problems with the descriptive adequacy of all probabilistic models, including SDT and the SPRT, which may stem from the psychologically implausible assumption of exact calculation of probabilities and the lack of mechanism to account for the stochasticity arising from the inference process. In the next section, we propose how to address these fundamental problems, before evaluating how far the proposed solution produces a better qualitative match to a wide range of regularities in human behavior.

A Sampling-based Approximation Perspective for Rational Decision Making

Assuming imprecise probabilities does not necessarily mean abandoning probabilistic models. While exact Bayesian computation is often out of reach for real-world computational mechanisms, including the human brain (Anderson, 1991; Aragonés et al., 2005; Kwisthout et al., 2011), computer scientists and statisticians have devised a number of sophisticated, general-purpose approximations for producing useful answers with a more reasonable amount of computational time and effort. It is therefore interesting to explore whether the brain has

hit on similar solutions. One major family of general-purpose approximations in computer science and statistics are sampling techniques².

Following the Bayesian approach, we propose that people solve cognitive tasks by building an internal model and posterior distribution over fine-grained hypotheses, which can support the responses for all of the aforementioned six behavioral measures. But, because the exact representation of the posterior probabilities of hypotheses is typically computationally intractable, we further hypothesize that the posterior probability distribution is not computed exactly, but is approximated by drawing representative samples from that distribution. Sampling-based approximations to the posterior are appealing as a psychological mechanism because (i) some sampling algorithms (e.g., MCMC: Brooks et al., 2011) need only local knowledge of the target posterior distribution and can represent only one or a few hypotheses at a time, lending these algorithms psychological plausibility (Anderson, 1991; Sanborn, Griffiths, & Navarro, 2010), (ii) sampling algorithms show much of the same behavioral variability and deviations from ideal probabilistic inference as observed in people across a range of domains (Griffiths, Vul, & Sanborn, 2012; Sanborn & Chater, 2016; Dasgupta et al., 2017; Lieder et al., 2018; Vul et al., 2014; Zhu et al., 2022b), and (iii) the variability of sampling algorithms has been found to match neural variability in the cortex (Hoyer & Hyvarinen, 2003; Haefner et al., 2016; Fiser et al., 2010). These suggest that the sampling-based explanations can connect with all three of Marr's (1982) celebrated explanatory levels: computational (through implementing Bayesian inference), algorithmic (via a tractable

² We do not further consider other general-purpose approximation algorithms, such as variational inference, in which a simpler distribution is used to approximate a more complex one, and the statistical distances between the two distributions are minimized by optimisation algorithms. Such algorithms provide an alternative source of explanations for human behaviour (Gershman & Beck, 2017; Sanborn, 2017).

computational mechanism), and implementational (through potentially mapping on to neural activity).

Taking a sampling-based approximation perspective to model choices suggests decision-making should be conceptualized as the problem of integrating a sequence of stochastic *hypotheses* into categorical decisions. The key distinction with other probabilistic models such as SDT and the SPRT is that we specifically define the ‘evidence’ as samples of hypotheses, abstracting away from noisy sensory percepts or memory traces (see Figure 2B) and implying that it is computational noise in the inference process that is the primary source of variability in behavior.

The Autocorrelated Bayesian Sampler

Here we outline a rational process for producing probability judgments, estimates, confidence intervals, choices, confidence judgments, and response times based on a sampling approximation of the posterior probability of fine-grained hypotheses, which we call the Autocorrelated Bayesian Sampler (ABS). Our key theoretical contributions are links between the sampling process and each of the six behavioral measures. This is possible because samples of the fine-grained hypotheses contain all of the relevant information to produce these (and indeed many other) aspects of behaviors.

Continuing our numerosity example (see Figure 1 and 2), the ABS produces behavior based on the posterior probability of the hypotheses, $P(h|s)$, which is calculated using Bayes rule:

$$P(h|s) = \frac{P(s|h)P(h)}{P(s)} \quad (3)$$

where h is a hypothesis, s is a stimulus, $P(s|h)$ is the likelihood of a stimulus given a hypothesis, $P(h)$ is the prior probability of a hypothesis³, and $P(s)$ is the overall probability of observing the stimuli across all possible hypotheses included in the internal model. In the numerosity task, for example, the hypothesis space reflects all possible numbers of dots that may have appeared on-screen, while the posterior distribution could be represented as a Gaussian distribution with mean equal to 24—the correct number of dots. The variance of the distribution may stem from all kinds of uncertainties including, for example, perceptual noise and/or uncertainties associated with the processing of information.

This general framework applies far beyond our numerosity task illustration. For example, it can be applied to intuitive physics when h is a complete object trajectory and s is the initial movement of an object (e.g., Hamrick et al., 2015; Battaglia, Hamrick, & Tenenbaum, 2013; Sanborn, Mansinghka, & Griffiths, 2013), language production when h is the next word in a sentence and s are the preceding words (e.g., Chater & Manning, 2006; Levy, Reali, & Griffiths, 2008), and common-sense reasoning when h is a social goal of other agents and s is a sequence of actions performed by those agents (e.g., Baker, Saxe, & Tenenbaum, 2009; Baker, Goodman, & Tenenbaum, 2008). Similarly, Bayesian models have also been successfully implemented in explaining effects in other areas of psychology such as vision (e.g., Yuille & Kersten, 2006), motor control (Kording & Wolpert, 2004), causal reasoning (e.g., Abbott & Griffiths, 2011; Bramley et al., 2017), and learning (e.g., Courville & Daw, 2007; Gershman, Blei, & Niv, 2010).

³ Note that this prior reflects the prior knowledge of the hypothesis space and therefore should be amenable to feedback and experimental instructions about the hypothesis space. For the running numerosity example, the experimenter could explicitly inform participants, or they could learn through experiences, that the numbers of dots appearing on-screen across trials are uniformly distributed in the range of [21, 30]. Indeed, previous work has shown that participants can quickly acquire an accurate prior from feedback in a numerosity task (Sanborn & Beierholm, 2016). In this case, for simplicity, we should use a uniform distribution as the prior for hypotheses.

Next, a set of hypotheses are sampled from the fine-grained posterior distribution, and these samples directly and straightforwardly support all six of our measures (Figure 2B). Probability judgments are based on the relative proportion of the samples (e.g., the number of samples with numerosities greater than 25). Estimates are based on a summary statistic of the samples (e.g., the mean sampled numerosity or the value of the latest sample). Confidence intervals are based on the quantiles of the samples (e.g., ordering five samples and using the numerosities of the 2nd and 4th sample as the bounds of a 60% confidence interval). Choices are based on the preponderance of the samples (e.g., depending on whether more than half the samples have numerosities greater than 25). Confidence judgments are (like probability judgments) based on the relatively proportion of the samples that agree with the choice. Response times are a function of the number of samples drawn (e.g., on average drawing four samples takes longer than three).

To generate concrete predictions from the model, and assess the match with human behavior, we need to outline three further aspects of the model: the choices of *sampling algorithm*, *prior on responses*, and *stopping rule*, to which we now turn.

The Sampling Algorithm

We assume that the mind conducts sampling-based approximations by drawing samples of hypotheses in proportion to the posterior probabilities associated with each hypothesis (e.g., Griffiths, Vul, & Sanborn, 2012; Vul et al., 2014; Dasgupta et al., 2017; Zhu et al., 2020; Chater et al., 2020). Rather than reviewing the extensive literature on sampling algorithms in statistics and computer science (see Andrieu et al., 2003 for an overview), we focus on algorithms that have been previously shown to match human behavior in some domains of psychology.

The simplest sampling algorithm is direct sampling in which independent and identically distributed (i.i.d.) samples are drawn (e.g., Vul et al., 2014). However, a lot must

be known about the target distribution to draw i.i.d. samples: people would have to (at least implicitly) know the posterior probability of every hypothesis, which fails to alleviate the intractability problem that motivates the need for sampling approximations. Another difficulty with direct sampling is descriptive. Human hypothesis generation is not a process of drawing independent samples, as the direct sampling requires. Instead, what comes to mind now depends on what came to mind in the past (Gilden et al., 1995; Dasgupta et al., 2017; Zhu et al., 2022c).

In light of this, researchers have recently started to explore a family of more sophisticated sampling algorithms called Markov Chain Monte Carlo (MCMC; Robert & Casella, 2004). MCMC algorithms explore the hypothesis space using only local knowledge about the probability distribution, greatly reducing the knowledge required to generate samples. The key idea of MCMC is that, in its simplest form, it represents only a single hypothesis at a time, and probabilistically transitions between hypotheses in proportion to their posterior probabilities. The local transitions induce a serial dependence between samples, akin to the local transitions in human hypothesis generation (Bramley et al., 2017; Franken, Theodoropoulos, & Bramley, 2022).

In our own work, we have found that an extension of MCMC, named MC³, provides a close match to the dynamics of repeated human judgement, capturing the observed long-range autocorrelations between estimates, as well as the heavy-tailed distribution of changes between responses (Zhu, Sanborn, & Chater, 2019; Zhu et al., 2022c; Zhu et al., 2021; Spicer et al., 2022). We therefore use MC³ as the sampling algorithm in the present model, though the specific mechanics of this algorithm beyond dependencies between samples are not necessary for almost all of the behaviors targeted here (see Appendix A for algorithmic details). That is, with the exception of explaining the cross-trial autocorrelation results which requires quantitative characterizations of the dependence in samples, the key condition for a

sampler to reproduce the qualitative model behaviors (e.g., comparing average model behaviors between experimental conditions) is simply that sampling is local and autocorrelated. Thus, most model predictions can be replicated using many other MCMC sampling algorithms, including the widely-used Random Walk Metropolis algorithm, and many others, so long as the generated samples are positively correlated across time.

Using dependent samples influenced our choice for how the ABS produces estimates. In past work, estimates have been based on the most recently generated sample or by averaging over samples (Vul et al., 2014; Lieder et al., 2018). While the mean of a set of independent samples is clearly a better estimate of the underlying mean than a single sample, with dependent samples, earlier samples are more likely to be biased by the starting point than later samples. For this reason, we chose to use the last sample as our estimate. However, these two approaches do not predict qualitatively different behavior on aggregate (see Appendix E for details).

Producing confidence intervals, however, requires more than a single sample, and instead can reflect statistics of the entire set of samples: for example, the 2.5% and 97.5% quantiles of the samples can represent a 95% confidence interval of the target distribution. This approach can only be applied directly for large samples. With small samples, we produce more fine-grained intervals by following Juslin et al. (2007) and use linear interpolation to fill in the gap between the two quantiles of the samples.

We also assume that sampling takes time. For simplicity, we model the time necessary to produce N samples as a Poisson process: while time taken to produce a new sample is random, the samples are generated at a constant rate (λ samples per sec). In a Poisson process, the waiting time between samples is exponentially distributed, and the time necessary to generate N samples follows an Erlang distribution:

$$f(t) \sim \text{Erlang}(N, \lambda) \tag{4}$$

The mean and variance of RT for a sample size of N are then defined as $\mathbb{E}[t] = \frac{N}{\lambda}$ and $\mathbb{V}[t] = \frac{N}{\lambda^2}$ respectively. Using a Poisson process allowed us to more closely link our approach to existing models such as the Poisson random walk model (Blurton et al., 2020; discussed below), though the results in this paper would be qualitatively the same under a wide variety of assumptions of how long it takes to generate each additional sample. This is because many empirical results only require assuming the samples were generated sequentially and the time to generate a sample is non-zero. Exponential waiting times are assumed here to explain the shape of RT distributions, particularly the observation that the response times for probability judgments (which we assume to have been produced using a fixed number of samples) have heavy tails (see Appendix F)⁴.

The Bayesian Monte Carlo Prior on Responses

Samples of the fine-grained hypotheses generated from our sampling algorithm can be readily used to make a choice. In our numerosity example, if asked to decide whether the number of dots that appeared on-screen is greater than 25, the hypothesis space should be partitioned into two subspaces with 25 on the boundary line. Samples that indicate greater than 25 dots or not should be categorized as supporting evidence for the corresponding hypotheses. That is, evidence is directly translated from the samples, here taking one of two values. And also, unlike the evidence used in SPRT, there is no inherent uncertainty about which alternative each sample supports. For the numerosity example, the generated sample can denote any number of dots in the hypothesis space, but it can only support one alternative in decisions: if

⁴ Non-decision time, as is commonly used in models of response time, would be necessary to include to quantitatively reproduce empirical RT distributions, but it is not needed to produce the qualitative effects we evaluate below.

the sample was 23, it only supports the hypothesis that there were less than 25 dots on screen. Similarly for M -alternative choices ($M > 2$), the hypothesis space should be partitioned into M subspaces with hypothesis samples from each subspace supporting the corresponding alternative.

These samples implicitly carry information about the probability that each choice alternative is correct. For example, when asked about the probability that the number of dots is greater than 25, the relative frequency of evidence in favor of the event should inform the probability estimate. But, as explored in Zhu et al. (2020), people should not directly use the relative frequency of the hypotheses as a probability estimate. This is especially true when sample sizes are small because the relative frequency tends to be extreme. Indeed, a single sample would lead to a probability estimate of either 0 or 1. This problem can be solved by incorporating a prior on responses to temper the relative frequency in the estimates of the probability that each choice alternative is correct, an approach that in statistics is called Bayesian Monte Carlo (Gelman et al., 2013; Rasmussen & Ghahramani, 2002)⁵.

The Bayesian Sampler model of Zhu et al. (2020) used a fixed prior on responses, and for mathematical simplicity, this was chosen to be a Beta distribution, because this is the conjugate prior for probability estimates. The Beta distribution is bounded by 0 and 1, and has two parameters, α_0 and α_1 , which determine its shape: when both parameters exceed 1,

⁵ The Bayesian Monte Carlo prior on responses is different from the prior on hypothesis in Equation 3. Specifically, the Bayesian Monte Carlo prior should capture full or partial information about the frequencies of the relevant behavioral outcomes in past trials. Consider again the binary choice in the numerosity example where people were asked to judge whether the number of dots is greater than 25. In this case, the prior on responses should reflect, to some extent, prior belief in different probabilities that each response is correct; and this prior knowledge could be acquired through feedback of correctly choosing greater-than-25 and of correctly choosing the alternative less-or-equal-than-25. An additional difference is that when knowledge of the probabilities is precise, even if there is uncertainty in the hypotheses, the effect of the Bayesian Monte Carlo prior reduces. Thus, in Equation 5, as the sample size, N , approaches infinity then $\hat{P}(A)$ tends to $P(A)$.

the Beta distribution is unimodal with a peak in the middle of the range (i.e., at $\frac{\alpha_1-1}{\alpha_0+\alpha_1-2}$); when both parameters equal 1, it is uniform; and when both parameters are less than 1, it is bimodal with peaks at both 0 and 1. Most critically, using the Beta distribution as the prior enables evidence to act as pseudocounts in the parameters. For $S(A)$ pieces of evidence of event A , $S(\neg A)$ of event not- A , and a $\text{Beta}(\alpha_0, \alpha_1)$ prior, people will have a posterior distribution for probability estimates that is distributed according to $\text{Beta}(\alpha_0 + S(A), \alpha_1 + S(\neg A))$. The Bayesian Sampler model used the expected value of this posterior distribution as its probability estimate, which is also simple to calculate:

$$\hat{P}(A) = \frac{\alpha_0 + S(A)}{\alpha_0 + S(A) + \alpha_1 + S(\neg A)} = \frac{\alpha_0 + S(A)}{N + \alpha_0 + \alpha_1} \quad (5)$$

where $N = S(A) + S(\neg A)$ denotes the total amount of samples that were generated and translated into evidence. Both the prior parameters (which affect α_0, α_1) and the sampling process (which affects $S(A)$ and N) affect the expected value. As the prior parameters are defined to be non-negative (i.e., $\alpha_0, \alpha_1 \geq 0$), the Bayesian Sampler's estimated probabilities tend to avoid extreme values and regress to the mean of the prior (i.e., $\alpha_0/(\alpha_0 + \alpha_1)$).

Here, we generalize the prior on responses used in the original Bayesian Sampler in two ways. The first is to make it multivariate: in many situations, people can be asked to judge a multivariate event where the hypothesis space should be partitioned into many subspaces. For example, when asked “what is the probability that the hottest day of the week will be Sunday?”, there are seven comparable events for the seven days in a week (“Sunday hottest”, “Monday hottest”, and so on). In this case, the Dirichlet distribution, a multivariate generalization of the Beta distribution, is the natural conjugate prior. For an M -variate Dirichlet prior, $\text{Dir}(\boldsymbol{\alpha})$, with $\boldsymbol{\alpha} = (\alpha_0, \alpha_1, \dots, \alpha_{M-1})$, people report the mean posterior distribution as their probability estimates:

$$\hat{P}(A) = \frac{\alpha_0 + S(A)}{N + \sum_{m=0}^{M-1} \alpha_m} \quad (6)$$

This view of probability estimates implies an indifference point (when the underlying probability and the estimated probability matches) that depends on the number of alternatives (see Figure 4A).

Figure 4. Relationship between the underlying probabilities of event A and the average probability judgments of A predicted by the ABS. **(A)** From left to right, the number of alternatives, M , varies from 2 to 7. Within each panel, simulated sample sizes, N , range from 1 to 9 in increments of 2. While uniform Dirichlet priors were used in the simulation, the indifferent points are always located at $(1/M, 1/M)$ for symmetric Dirichlet priors. **(B)** As produced by the ABS, the empirical indifference points between mean probability judgments and objective probabilities are related to the inverse of the number of alternatives (M). Indifference points were directly reported by the data analyses in Fox & Rottenstreich (2004), Bardolet, Fox, & Lovallo (2004), and Varey, Mellers, & Birnbaum (1990), and were inferred from the regression in Attneave (1953).

The second way in which we generalize the prior on responses of the Bayesian Sampler is to allow it to adapt to experience (e.g., the trial history in an experiment). In Bayesian data analysis, when no prior information is available, a default prior is typically recommended (Gelman et al., 2013). However, for many real-world applications and especially for everyday cognitive tasks, historical data (e.g., past experiences of the same

task, data from similar previous tasks or from observing others' performing the task) are available which can help people can construct an appropriate prior. For example, if repeatedly choosing between the same two alternatives, historical choice data should provide useful information such as the base rate, which in turn can help construct a prior on responses to guide future decisions. How to construct an adaptive prior based on historical data is a topic of debate in statistics and computer science because it is difficult to determine how much to generalize previous experience to new situations (Ibrahim et al., 2015; Chen et al., 2000; Diaconis & Ylvisaker, 1979)⁶. For simplicity, we assume that people only use information from the immediately previous trial to develop their adaptive prior for the present trial: in binary choice, a non-informative, uniform prior ($\text{Beta}(1,1)$) is adjusted to favor the option the feedback indicated was correct, becoming either $\text{Beta}(2,1)$ or $\text{Beta}(1,2)$.

The adaptive prior on responses, in conjunction with the generated samples, then determines the model's estimated probability of a categorical alternative being correct. This estimated probability is used both as the model's probability estimate *and* its confidence judgment in whether a choice is correct⁷. The equivalence between the two is not unique to our model— it has been previously posited as the Bayesian Confidence Hypothesis (Mamassian, 2016; Kepecs & Mainen, 2012; Pouget et al., 2016), and has attracted both support (Calder-Travis et al., 2020) and criticism (Li & Ma, 2020).

⁶ Incorporating historical information into new situations is known as power prior in the statistics literature, which is also closely related to the ideas of meta-learning (or learning-to-learn) and hierarchical Bayesian modeling.

⁷ Confidence judgments are often made on various ordinal rather than probability scales, though analyses of confidence judgments often just assume that they are monotonically related (Li & Ma, 2020; Shekhar & Rahnev, 2021a), as we do here. For comparability across different ordinal scales we present all of the model predictions on the probability scale rather than specifying those relationships.

The Stopping Rule

Any model of judgment or decision that depends on the sequential accumulation of evidence needs a rule determining at what point to stop collecting evidence and make a decision. When to stop drawing samples should depend on both the costs (e.g., metabolic, opportunity, etc.) of sampling as well as the task-specific benefits of additional samples for providing a good response. For probability judgments, estimates, and confidence intervals, in the absence of a reason to do otherwise, we make the simplest possible assumption: that a fixed number of samples are drawn to answer each question.

A fixed number of samples will allow the model to produce indifferent probability judgments (e.g., judging a binary event to have a probability of 0.5) as is often observed in the human data⁸. However, for making decisions a fixed sample is likely to be too simple. If the samples so far leave the evidence finely balanced regarding which decision to make, then it is likely that more data will be collected. While it is possible in principle to derive an optimal stopping rule for the sampling process in this model, unlike with the SPRT, the optimal rule is not analytically tractable and can instead only be computed using dynamic programming (see Appendix C). So, again for simplicity, we use a well-known heuristic stopping rule instead: the max-minus-next rule, which counts the difference in evidence between the top two hypotheses, and terminates the sampling process whenever the accumulated difference exceeds a threshold. This simple heuristic stopping rule has also been shown to approach the performance of an optimal SPRT even in multi-alternative settings (Dragalin et al., 1999; 2000). For binary choices, this reduces to just the difference in the number of samples in favor of each alternative, which has been proposed in past work

⁸ A reanalysis of the four probability judgment experiments presented in Sundh et al. (2021) shows that participants produced indifferent probability judgments (i.e., exactly 50 on a 0-100 scale) on about 6% of trials.

(Hamrick et al, 2015; Vul et al, 2014). The decision-making panels of Figure 2B

demonstrates the max-minus-next stopping rule with a threshold value of 2.

While the choice of stopping rule does not change how samples are used to produce the different measures, it does influence the content of the samples and the variability of the sample size and hence responses times. So, for example, in our model, while the response times for a *probability judgment* which assumes fixed sample sizes will follow an Erlang distribution (see Appendix F for further justification), response times for a *choice* (which assumes optional stopping) will follow a mixture of Erlang distributions (see Appendix B for details).

Table 1

Key Empirical Effects of the Six Behavioral Measures

	Empirical Effects	Description	Example References	ABS Explanations
Probability Judgments	Conservatism and probabilistic identities	Probability judgments show a linear bias toward 0.5 for binary events.	Erev et al. (1994); Costello & Watts (2014)	Incorporating a symmetric prior on responses pulls probability judgments towards 0.5 for binary events.
	Mean-variance relationship	The variance of probability judgments is highest for moderate probabilities and goes to zero before reaching the limits of the scale.	Sundh et al. (2021)	Same as above.
	Explicit subadditivity	When assessed separately, the sum of the probability judgments of the unpacked descriptors exceeds that of the packed descriptor.	Tversky & Koehler (1994); Fischhoff et al. (1978); Olson (1976); Mehle et al.	Same as above.

		(1981); Russo & Kolzow (1992); Koehler & Tversky (1993)	
Conjunction fallacy	The judged probability of the conjunctive event is higher than that of its constituents.	Tversky & Kahneman (1983); Costello & Watts (2014)	Because of their added complexity, fewer conjunctive event samples are generated, resulting in a greater pull toward 0.5 for binary events.
Implicit subadditivity in typical unpacking	When assessed jointly, the probability judgment of the unpacked descriptor exceeds that of the packed descriptor when the unpacked descriptor includes high probability events.	Sloman et al. (2004); Dasgupta et al. (2017)	The hypothesis that the autocorrelated sampler initializes on tends to be probable (i.e., typical examples).
Implicit superadditivity in atypical unpacking	When assessed jointly, the probability judgment of the unpacked descriptor is smaller than that of the packed descriptor when the unpacked descriptor includes low probability events.	Sloman et al. (2004); Dasgupta et al. (2017)	The hypothesis that the autocorrelated sampler initializes on tends to be improbable (i.e., atypical examples).
Partition dependence	Probability judgments are biased toward an even allocation across all possible events.	Fox & Rottenstreich (2004); Attneave (1953); Bardole, Fox, & Lovallo	Incorporating a symmetric Dirichlet prior on responses (i.e., a multivariate generalization of the Beta prior) pulls probability judgments towards $1/M$ for M -alternative events.

			(2004); Varey, Mellers, & Birnbbaum (1990)	
Decisions Affecting Later Responses	Anchoring	Estimates are pulled toward the boundary used in the preceding decision when the boundary is far from the correct answer.	Tversky & Kahneman (1974); Epley & Gilovich (2006); Spicer et al. (2022)	The initial hypothesis used by the autocorrelated sampler will often be the comparison value used in the decision task (i.e., the anchored hypothesis). Moreover, the samples used to reach the decision are assumed to be reused in the estimation task.
	Repulsion	Estimates are pushed away from the boundary used in the preceding decision when the boundary is close to the correct answer.	Jazayeri & Movshon (2007); Luu & Stocker (2018); Spicer et al. (2022)	The optional stopping rule only terminates the sampler whenever sufficiently more evidence supporting one alternative is obtained. As above, samples used to reach the decision got reused in the estimation task.
Accuracy and Response Times	Speed- accuracy trade-off	More accurate decisions are associated with longer decision times such that people can trade accuracy for speed.	Garrett (1922); Johnson (1939)	The sequential arrival of samples predicts that, on average, a greater sample size requires more time to generate but results in more accurate choices. This tradeoff is controlled by the decision threshold, in which people can balance the cost of waiting and the benefits of a refined approximation by adjusting the threshold.
	Slow errors	For difficult conditions and when accuracy is emphasized, mean decision times for incorrect choices are slower than those for correct choices.	Ratcliff & Rouder (1998); Vickers (1979); Townsend & Ashby (1983)	Optionally terminating an autocorrelated sampling process will increase the proportions of slower responses for error responses (vice versa for correct responses).

	Fast errors	For easy conditions and when speed is emphasized, mean decision times for incorrect choices are faster than those for correct choices.	Ratcliff & Rouder (1998); Townsend & Ashby (1983)	Adaptively updating the prior on responses to reflect previous choice outcomes will increase the proportions for fast responses that are errors.
	Near-linear relationship of RT quantiles with a fan shape	The quantiles of the RTs from two difficulty levels are close to linear.	Ratcliff et al. (2003); Ratcliff et al. (2015)	Optionally terminating an autocorrelated sampling process and the waiting time between two consecutive samples are exponentially distributed make the tails of RT distributions more spread out for harder tasks.
Confidence in Decisions	Positive relationship between confidence and stimuli discriminability	Confidence increases as stimuli increases.	Baranski & Petrusic (1998); Garrett (1922); Vickers (1979); Johnson (1939); Vickers & Packer (1982)	Higher stimuli discriminability partitions the hypothesis space to more unevenly favoring the correct response. Thus, the hypothesis samples that were converted into evidence become more homogenous with increasing discriminability.
	Resolution of confidence	There is a positive relationship between choice accuracy and confidence judgments.	Baranski & Petrusic (1998); Vickers (1979); Garrett (1922); Ariely et al. (2000); Johnson (1939); Vickers & Packer (1982)	Optionally terminate an autocorrelated sampling process will make the correct responses faster and thereby greater confidences.
	Metacognitive inefficiency	Metacognitive efficiency (as measured by <i>meta_d'/d'</i>)	Shekhar & Rahnev (2021a; 2021b)	Confidence based on the proportion of response-consistent binomial samples (see Appendix D for detail).

		decreases for higher confidence ratings.		The sample autocorrelation amplifies the metacognitive inefficiency.
	Negative (cross-trial) relationship between confidence and RT	Within the same condition, confidence decreases as RT increases.	Henmon (1911); Baranski & Petrusic (1998); Vickers & Packer (1982); Johnson (1939)	Confidence based on the proportion of response-consistent samples with an optional stopping rule.
	Positive (cross-condition) relationship between confidence and RT	People are more confident in conditions in which they take more time to make a choice.	Vickers & Packer (1982)	The decision threshold is greater in the accuracy condition than in the speed condition, so choices will take longer while the greater difference in sample counts at threshold leads to higher confidence.
Confidence Intervals	Strong overconfidence in self-produced confidence intervals	Self-generated confidence intervals are far too narrow	Juslin et al. (2007); Juslin et al. (2003); Juslin & Persson (2002)	Sample quantiles are too narrow to represent distributional intervals.
	Little or no overconfidence in evaluation of confidence intervals	Evaluations of confidence intervals are well calibrated.	Juslin et al. (2007); Juslin et al. (2003); Juslin & Persson (2002)	Samples are relatively representative of the target distribution on aggregate.
Cross-trial autocorrelation in RT and Estimates	Long-range autocorrelations in estimation time series	Sequences of estimates show autocorrelations that are proportional to frequency ($1/f$ noise).	Gilden et al. (1995); Gilden (2001)	The autocorrelations are generated by the local nature of sampling.

Lesser long-range autocorrelations in RT time series	Sequences of RTs also show a certain degree of autocorrelation but are closer to independent than are sequences of estimates.	Wagenmakers et al. (2004)	Basing RT on the number of autocorrelated samples until threshold (i.e., counting autocorrelated samples) reduces the degree of autocorrelation in samples.
--	---	---------------------------	---

Explaining Key Empirical Targets in Probability Judgments, Estimates, Confidence Intervals, Choices, Confidence Judgments, and Response Times

We now demonstrate the explanatory power of the ABS. We focus on behavioral results that deviate from the Bayesian ideal embodied in models like the SPRT (see Table 1), simulating these using a consistent set of parameters (detailed in Appendix A). To facilitate understanding of the active ingredients of the model, we also show results from three restricted variants of the full ABS model. The *no-prior variant* removes the adaptive prior (i.e., equivalent to fixing the prior to $\text{Beta}(0,0)$) while keeping the remaining components. The *direct-sampling variant* uses independent samples instead of the autocorrelated samples while keeping the remaining components. The *fixed-sample-size variant* always uses a fixed number of samples ($N=5$) to form behaviors while keeping the remaining components.

Biases in Probability Judgments

Biases in probability judgments are perhaps the most direct evidence against purely normative probabilistic models. We find that the prior on responses and local sampling algorithm of the ABS, which help to reduce the computational burden of the model by reusing old and useful computations and by using only local knowledge of the posterior distribution respectively, suffice to produce many of these biases. In the ABS, there are no biases in the underlying posterior probabilities; biases arise solely from the algorithmic process by which the posterior is sampled and judgments and decisions are generated.

Using prior knowledge to temper the probability estimates was the basis of the Bayesian Sampler model (Zhu et al., 2020). The ABS works in the same way, except that it uses autocorrelated, rather than independent, samples⁹. As shown in Figure 4A, the prior in the Bayesian Sampler produces a linear bias toward conservative judgments (Zhu et al., 2020) where people avoid the extremes in their probability judgments (Fiedler, 1991; Peterson & Beach, 1967; Erev, Wallsten, & Budescu, 1994). This type of conservatism captures the results of a series of probabilistic identities investigated by Costello, Watts, and colleagues (Costello & Watts, 2014; Costello & Watts, 2018; Costello, Watts, & Fisher, 2018), which were constructed by adding and subtracting various mean judgments across combinations of events. While all these identities would equal zero if participants reported coherent probabilities (even if just on average), mean judgments were zero for some identities and substantially different from zero for others. The results from the entire set of identities, including conditional probability judgments of dependent events, were well fit by the Bayesian Sampler's linear conservatism bias (Zhu et al., 2020). As the average behavior of the ABS is approximated by the Bayesian Sampler especially when the effects of local sampling are not strong (e.g., random initializations of the local sampler), the ABS will produce these results as well.

The sample size and prior on responses of the ABS can be dissociated by examining the mean-variance relationship in probability judgments. When probability judgments are

⁹ The autocorrelation of the samples does not qualitatively alter the overall model predictions for probability judgments (with the exception of implicit subadditivity and implicit superadditivity, discussed below) because autocorrelated sample sizes can be corrected to produce the effective sample size of independent samples using the following equation: $N_{\text{eff}} = \frac{N}{1 + 2 \sum_{t=1}^{\infty} \rho_t}$ where ρ_t is the degree of autocorrelation at lag t . For the parameters we used in the simulations, the effective sample size is on average 16.80% of the autocorrelated sample size (95% CI, [16.40%, 17.21%]). In addition, in the studies we refer to, there is very rarely any feedback. Without feedback, we assume the ABS prior does not change from trial to trial, making it identical to the fixed prior of the Bayesian Sampler for binary events.

based on sampled outcomes, the relationship between the mean probability estimates and the variance of the estimates will constitute an inverse U-shaped (“rainbow-shaped”) curve (see Figure 5A). The prior on responses then constrains the range of possible probability estimates that an agent can produce, thereby lowering the relative variance and pulling the curve both inward and downward (see Figure 5B). For example, for a binary event with a uniform prior, if a single sample is drawn, probability judgments will be either 0.33 or 0.67, and total variance will be relatively lower than for the pure proportions of sampled outcomes (taking now account of the prior on responses). Overall, the Bayesian Sampler predicts a shrinkage of the mean-variance curve for probability judgments, and this was empirically validated in four experiments (Sundh et al., 2021). For the same reasons as the Bayesian Sampler, the ABS model predicts this shrinkage of the mean-variance curve as well (see Figure 5B).

Figure 5. (A) Empirical results based on the 4 experiments of Zhu et al. (2020) and Sundh et al. (2021), showing inverted U-shaped relationships between means and variances of probability estimates, such that extreme probability estimates (very near 0 or 1) are ruled out. Solid curves are the mixed-effect regression models fitted on individual-level probability estimates. *(B)* Analytic approximations of the mean-variance relationship predicted by the

ABS. The model predicts an inverted-U shape with stronger priors on responses moving the curve inward and more samples moving the curves downward.

Moreover, the phenomenon of explicit subadditivity in probability judgments also occurs as a direct consequence of using the prior on responses. Explicit subadditivity is when the estimated probability of an event (A_0) is lower than the sum of estimated probabilities for events ($A_1, A_2, \dots, A_{M'}$) where A_0 is the disjunction of those M' mutually exclusive events. That is:

$$\hat{P}(A_0) < \hat{P}(A_1) + \hat{P}(A_2) + \dots + \hat{P}(A_{M'}) \quad (7)$$

where probability theory requires that these should be equal. An explicit subadditivity bias was observed generally in between-participant designs in which participants were asked explicitly to judge the probability of each of the M' events and their disjunction, A_0 , so that a total of $M' + 1$ probability estimates were recorded (e.g., Tversky & Koehler, 1994; Tversky & Fox, 1994; Fox et al., 1994). According to the sampling account, for each query, because participants do not know the full range of questions that are asked, they will treat the event to be judged as a binary event; that is, participants will sample instances and non-instances of that event (that is, A_m vs. not- A_m), thus requiring a Beta prior on responses. The resulting estimate of each $P(A_m)$ will therefore be inflated by regression to the mean. The regression-to-mean effect then applies multiple times on aggregate to the right-hand side of Equation 7 and only once to the left-hand-side, predicting a subadditivity bias for low probability events.

As a corollary, more probability judgments queried on the right-hand-side of Equation 7 should associate with a greater degree of subadditivity bias. For M' component hypotheses, the predicted difference between the sum of the M' probability estimates and the probability estimates of the disjunction can be derived as follows:

$$\begin{aligned}
& \hat{P}(A_1) + \hat{P}(A_2) + \dots + \hat{P}(A_{M'}) - \hat{P}(A_0) \\
&= \sum_{m=1}^{M'} \left[\frac{N}{N + 2\alpha_0} P(A_m) + \frac{\alpha_0}{N + 2\alpha_0} \right] - \left[\frac{N}{N + 2\alpha_0} P(A_0) + \frac{\alpha_0}{N + 2\alpha_0} \right] \\
&= (M' - 1) \frac{\alpha_0}{N + 2\alpha_0}
\end{aligned}$$

where the assumptions were fixed sample size (N) and symmetric prior on responses, $\text{Beta}(\alpha_0, \alpha_0)$. Indeed, the empirical findings suggest a positive relationship between M' and the degree of explicit subadditivity bias, and the Bayesian Sampler correctly captures the relationship (see Figure 6). Moreover, when the disjunction of M' mutually exclusive events was exactly 1, participants were sometimes asked to only judge the probabilities of M' component hypotheses but not their disjunction. In this case, model predictions can be analytically approximated as $(M' - 2) \frac{\alpha_0}{N + 2\alpha_0}$ under the same assumption as before. This prediction also matches the empirical pattern known as the binary complementarity: on average, no subadditivity bias was observed for mutually exhaustive events when $M' = 2$ (Tversky & Koehler, 1994). The ABS model inherits these predictions from the Bayesian Sampler.

Figure 6. Bias in explicit subadditivity, computed as the sum of the probability estimates of each component hypothesis minus the probability estimates of their disjunction, increases as the number of component hypotheses increases. This empirical effect is correctly captured by the Bayesian Sampler (solid line). The sample sizes were set at $N=5$, and the prior on responses was $\text{Beta}(1,1)$.

Similarly, this regression-to-mean effect predicts the conjunction fallacy (Costello & Watts, 2016; 2017; Zhu et al., 2020). The conjunction fallacy arises where the estimated probability for a conjunctive event is greater than that for its constituent events $\hat{P}(A_0 \cap A_1) \geq \hat{P}(A_0)$, whereas the probability theory requires the probability of conjunctive events to be less or equal with their constituents, A_0 and A_1 (Tversky & Kahneman, 1983). The conjunction fallacy occurs in the Bayesian Sampler when the regression-to-mean applies more to the conjunctive event than to the constituent events. Specifically, it is assumed that fewer samples of the more-complex conjunctive events can be generated or tallied in a fixed amount of time; and the prior produces a greater regression-to-mean effect for smaller sizes

(see Equation 6). This allows the Bayesian Sampler to predict above-chance levels of conjunction fallacies when the conjunction and constituent event both have low probability (Zhu et al., 2020), as is often the case in empirical work (Costello & Watts, 2016). The ABS model also inherits this prediction from the Bayesian Sampler.

In contrast with the explicit judgments of $M'+I$ probabilities above, both subadditivity and its opposite effect, superadditivity, have been observed in so-called implicit experimental designs. In implicit designs, only two probability judgments are made: one for the unpacked descriptor (e.g., “baby bottles and other bottles made of glass”) and one for the simple disjunctive descriptor (e.g., “bottles made of glass”) (Dasgupta et al., 2017; Sloman et al., 2004). Unpacking to typical examples (e.g., a baby bottle in the category of bottles made of glass) leads to subadditivity: $\hat{P}(A_0) \leq \hat{P}(A_1 \cap A_2 \cap \dots \cap A_M)$, whereas unpacking to atypical examples (e.g., a shampoo bottle in the category of bottles made of glass) leads to superadditivity: $\hat{P}(A_0) \geq \hat{P}(A_1 \cap A_2 \cap \dots \cap A_M)$ (Dasgupta et al., 2017; Sloman et al., 2004). Again, since A_0 was unpacked into M mutually exclusive events ($A_0 = A_1 \cap A_2 \cap \dots \cap A_M$), probability theory requires the two probability estimates to be equal. Previous work with autocorrelated sampling models (Dasgupta et al., 2017; Sanborn & Chater, 2016) accounted for this effect by assuming that the descriptor influenced the local sampler’s starting point: typical unpacking initializes the sampler in a high probability region of the hypothesis space, while atypical unpacking initializes it in a low probability region. As a result, the proportion of hypotheses supporting the event’s occurrence will be highest for typical unpacking, intermediate for the simple disjunctive descriptor (assuming it results in a random starting point), and lowest for atypical unpacking. We believe that the ABS will inherit this prediction because it produces autocorrelated samples, though we do not reproduce it here because auxiliary assumptions about the locations and probabilities of hypotheses are needed to do so. This explanation of implicit subadditivity and superadditivity depends on local sampling.

They cannot be predicted by the Bayesian Sampler model (Zhu et al., 2020; see a similar argument against a “regressive model” in Tversky and Koehler, 1994), which assumes independent sampling.

Interestingly, people’s probability estimates are also found to exhibit so-called “partition dependence.” That is, they regress to $\frac{1}{M}$ where M is the number of alternatives that people are encouraged to consider (see Figure 4B for a summary; Attneave, 1953; Fox & Rottenstreich, 2003; 2004; Bardolet et al., 2004; Varey et al., 1990). For example, asking “what is the probability that Sunday will be hotter than any other day next week?” encouraged participants to treat the event as binary, and their estimates were observed to be biased toward $\frac{1}{2}$, while asking, “what is the probability that the hottest day of the week will be Sunday?”, encouraged participants to consider seven possible outcomes, and estimates were observed to be biased toward $\frac{1}{7}$ (Fox & Rottenstreich, 2003). In ABS, framing the probability query as judging an M -variant event invokes a Dirichlet prior with M parameters, $\text{Dir}(\alpha_0, \alpha_1, \dots, \alpha_{M-1})$, which for a binary event reduces to a Beta prior, $\text{Beta}(\alpha_0, \alpha_1)$. Partition dependence effects can be explained by assuming that people have no *a priori* reason to believe one event occurs more often than another event: $\alpha_0 = \alpha_1 = \dots = \alpha_{M-1}$ and so probability estimates are predicted to be biased toward $\frac{\alpha_0}{\sum_{m=0}^{M-1} \alpha_k} = \frac{1}{M}$ (see Figure 4A). In the ABS, the impact of this non-informative prior should be more pronounced in situations where people are less knowledgeable about the probability estimation task or less confident in a learning context (reflecting fewer samples), matching the empirical results (Fox & Rottenstreich, 2006).

Choice Accuracy and Response Times

The Bayesian Monte Carlo process for choice and RT correctly predicts four key relationships between choice and RT. First, and in common with many other evidence accumulation models, the ABS predicts a trade-off between accuracy and speed where increasing decision thresholds lead to, on average, more evidence being accumulated (and thus higher accuracy) as well as longer response times. This trade-off between accuracy and speed has been widely documented in the literature (Garrett, 1922; Johnson, 1939; Pachella, 1974; Wickelgren, 1977; Ratcliff & Rouder, 1998).

Second, unlike many models, the ABS predicts that correct and incorrect responses have unequal average response times. The empirical result is that, when accuracy is emphasized (or in difficult tasks), errors are usually slower than correct responses. By contrast, when speed is emphasized (or in easy tasks), errors are usually faster (Luce, 1986; Swensson, 1972; Ratcliff & Rouder, 1998; Ratcliff, Thapar, & McKoon, 2003). This empirical pattern is surprisingly difficult to match for models that accumulate relative evidence to symmetric bounds: these models predict that the response time distributions for correct responses and errors will always be the same, regardless of choice accuracy (Link & Heath, 1975; Vickers, 1979). To produce slow errors, the usual route is to add variability to the strength of the “signal”, or the drift rate in DDMs (Ratcliff & Rouder, 1998). While both strong signals and weak signals will produce equal mean response times, weak signals are both more error-prone and slower. So, with an equal mixture of strong and weak signals, there will be more slow errors and more fast correct responses.

The ABS produces slow errors in a different way. Instead of adding cross-trial variation to the signal strength, or independent within-trial variation to the signal strength (Diederich & Oswald, 2016), slow errors result from the local sampling algorithm producing autocorrelated samples. For example, if the sampling algorithm begins far above the decision

boundary (e.g., the red subspace in the posterior of hypotheses illustrated in Figure 2B), then the initial samples will almost all favor the correct response, while if the sampling algorithm begins far below the decision boundary (e.g., the blue subspace in the posterior of hypotheses) then the proportion of correct samples will almost all favor the incorrect response. Slow errors also require optional stopping, because with a fixed stopping rule the response distribution is itself fixed. This can be seen in the simulation in Figure 7B: both autocorrelation and optional stopping (i.e., the no prior variant) are needed to produce errors that are on-average slower than correct responses.

Fast errors, often found in easy tasks, are produced in a different way. The usual route to producing fast errors is to assume variability in the starting point of the evidence accumulation process (Laming, 1968; Ratcliff & Rouder, 1998; Ratcliff et al., 2003). In the ABS, the adaptive prior on responses is assumed to change in response to the outcomes of the preceding trial. This encourages repeating past successes, but also introduces cross-trial variability in the starting point of the accumulator. This is because the accumulator will be biased toward whichever response was correct on the last trial, and assuming that (as is usual in experiments) the correct response randomly varies between trials, it will sometimes be closer to the correct threshold and sometimes closer to the error threshold. For those latter trials, the amount of evidence required to reach the error threshold is reduced, leading to a shortened mean response times for errors. As with slow errors, optional stopping is also necessary: only with both the adaptive prior on responses and optional stopping (i.e., the direct sampling variant) do fast errors appear in the stimulation in Figure 7D.

The differences between the simulations of the ‘difficult-accuracy’ condition (Figure 7B) and the ‘easy-speed’ condition (Figure 7D) track the conditions in which slow errors and fast errors are found. We assume that: (i) greater emphasis on accuracy causes the threshold to be higher, and consequentially more pieces of evidence are needed to terminate the

sampling algorithms, and (ii) easier stimuli makes the evidence more homogenous (e.g., samples are more likely to point to the same response)¹⁰. As a result, the ‘easy-speed’ condition involves integration over homogenous but smaller amounts of evidence than in the ‘difficult-accuracy’ condition. In other words, the starting point of the accumulator has more influences, while the degree of autocorrelation has less influence, on determining the predicted behavior in the ‘easy-speed’ condition than in the ‘difficult-accuracy’. Across Figure 7B and 7D, only the full ABS model matches the empirical observations that slow errors are more common in the ‘difficult-accuracy’ condition, while fast errors are more common in the ‘easy-speed’ condition.

Figure 7. (A) Empirical choice outcomes and response-time distributions in the difficult-accuracy condition. **(B)** Simulated choice outcomes and RT distributions in the difficult-accuracy condition. RTs were fitted with Gamma distributions with the best-fitting

¹⁰ A reduction in trial difficulty can lead to either less variable posterior of hypotheses (e.g., people are more certain about the number of dots), the decision boundary that partitions the hypothesis space moving to one extreme (e.g., people are asked to judge whether the number of dots are less than 1000 while only 24 dots appeared onscreen), or both. Overall, these model implications of the difficulty reduction all contribute to an increase in proportions of samples that support the correct response.

distribution shown as solid lines for correct (in blue) and error (in red) responses. Overlaid dots and their horizontal error bars denote mean RTs and 95% confidence intervals respectively. Similarly, **(C)** and **(D)** are respectively empirical data and model simulations for the easy-speed condition. Across the different variants, only the full ABS model reproduces both slow and fast errors in the correct experimental conditions. All predicted RT distributions were unimodal and positively skewed. The full model also correctly reproduces an RT distribution that becomes more positively skewed and spread out with an increase in decision threshold. Empirical data were adapted from Ratcliff, Thapar, and McKoon (2003).

As in other models (e.g., Blurton et al., 2020), the assumptions of exponential waiting time between consecutive samples and the optional stopping rule correctly reproduce many distributional properties of RTs including (i) that there tends to be one mode in the distribution and (ii) distributions with higher means more positively skewed. Further regularities in the shapes of RT distributions were stressed by Ratcliff et al. (2015) using quantile-quantile (Q-Q) plots (see Figure 8A). Plotting the quantiles of RT from one difficulty condition against the quantiles from another difficulty condition, the empirical Q-Q plots reveals near-linear relationships and a fan shape: increasing task difficulty has its greatest impact on the tails of the distribution with the near linearity suggesting similar RT distribution shapes across difficulty conditions. As shown in Figure 8B, the ABS captures the fan shape and near-linear regularity. The direct-sampling variant shows results that are closer to linear, as would be expected if the autocorrelation in samples causes the upper tails in RT distributions to spread out even more in harder tasks. Also of interest is the fixed-sample-size variant, which because it always collects the same number of samples for all difficulty levels,

produces identical quantiles between RTs from one level of difficulty and those from another, and thus doesn't match the empirical data.

Figure 8. Sample quantile-quantile plots of response-times distributions. **(A)** An example Q-Q plot of empirical RT distributions adapted from Ratcliff, Thapar, and McKoon (2003). One difficulty level was selected to compute its quantiles and then quantiles of the other four difficulty levels were plotted against the first condition. The rank of a condition depends on its mean RTs. **(B)** Q-Q plots of RT distributions produced by the ABS model and its variants.

Confidence in Decisions

From a Bayesian perspective, it is natural to map decision confidence onto the posterior probability that the decision is correct, a mapping which has been called the Bayesian Confidence Hypothesis (Mamassian, 2016; Kepecs & Mainen, 2012; Pouget et al., 2016). For the SPRT, posterior probability is updated as sensory samples are observed, and its posterior probability at the time of choice is simple: it is the posterior probability when the threshold is reached, because as evidence collection stops once this occurs (see Figure 2A confidence). The SPRT thus predicts that decision confidence only relates to the threshold values, because the threshold captures the amount of evidence favoring one hypothesis or the other. Given that the threshold value is fixed prior to, and independent of, the characteristics of a particular

trial, this means that confidence will be the same for all trials on which the same hypothesis is chosen¹¹.

Unlike the SPRT, the ABS does not have direct access to its posterior probability that a response is correct (i.e., its confidence). Instead it needs to estimate this probability given a set of samples (see Figure 2B confidence). Fortunately, the form of the adaptive prior on responses (a Beta distribution in the case of binary choice) makes this estimate easy to update as samples are sequentially generated. At the start of the trial, the adaptive prior on responses reflects the prior belief in different probabilities that each response is correct. Using the binary choice example, assume a prior for choice A of $\text{Beta}(i, j)$ (and a prior for choice B of $\text{Beta}(j, i)$). When coming to a decision, samples in favor of each response, $S(A)$ and $S(B)$, are sequentially collected until the decision process is terminated by the stopping rule. The confidence after N samples (i.e., $N = S(A) + S(B)$) is then

$$\begin{aligned} \text{Conf}_A &= \frac{S(A) + i}{N + i + j}, \text{ in } A \\ \text{Conf}_B &= \frac{S(B) + j}{N + i + j}, \text{ in } B. \end{aligned} \quad (8)$$

The max-minus-next heuristic stopping rule terminates the sampling algorithm when the quantity of evidence favoring one choice exceeds a threshold, $\Delta = |i + S(A) - (j + S(B))| > 0$. The final decision confidence can then be rewritten as follows:

$$\text{Conf}_A = 1 - \text{Conf}_B = \frac{i + j + N + \Delta}{2(i + j + N)}, \text{ if } A \text{ was chosen} \quad (9)$$

¹¹ While this characterizes an SPRT that stops at the threshold (e.g., as discussed in Pleskac and Busemeyer, 2010, and Vickers, 1979), there are variants that would make different predictions. For example, the boundaries do not necessarily have to be symmetric, or they could collapse. Also, it is possible that the SPRT's confidence would reflect not just the threshold but the posterior probability of the all the sensory samples that were observed before stopping, with confidence then being greater than or equal to the threshold confidence. All these variants would produce variable posterior probabilities and hence variable confidence judgments. However, none of these variants would produce estimates, confidence intervals, or other judgments that are generated by the ABS, but which are beyond the scope of the SPRT.

where the confidence judgments predicted by the ABS are decided by both the threshold values (Δ) and the amount of evidence accumulated (N ; this is same as the number of samples generated because evidence is directly mapped from hypothesis samples; for example, a sample of 27 dots will be converted into a piece of evidence for the proposition that the number of dots are greater than 25): the greater the number of samples generated before a decision is reached, the lower the confidence in that decision. This is because the ABS embodies a prior over the strength of signal in the Bayesian Monte Carlo process, and the longer the sampling process continues, the more likely the signal is weak, and so the posterior probability that the decision is correct correspondingly decreases. This is in contrast to the SPRT: in the SPRT, confidence does not change with additional sampling because its confidence is determined by a fixed decision threshold.

The decreasing decision confidence of the ABS with an increasing number of samples allows it to capture four key empirical phenomena which are not accommodated by the SPRT described above: the positive relationship between confidence and the discriminability of the stimuli (Baranski & Petrusic, 1998; Vickers, 1979; Vickers & Packer, 1982, Figure 9A), the “resolution of confidence” effect (Ariely et al., 2000; Baranski & Petrusic, 1998; Vickers, 1979; Garrett, 1922; Vickers, 2014; Vickers & Packer, 1982, Figure 9B), so-called “metacognitive inefficiency” (Shekhar & Rahnev, 2021a; 2021b, Figure 9C), and the complex relationship between RT and confidence (Baranski & Petrusic, 1998; Vickers & Packer, 1982, Figure 9D).

Figure 9. ABS simulations showing effects of confidence in decisions. **(A)** Positive relationship between stimulus discriminability and average confidence. Empirical data adapted from Vickers & Packer (1982). **(B)** Resolution of confidence in which average confidence is higher for correct responses than for incorrect responses. Empirical data replotted from Vickers & Packer (1982). **(C)** Degree of metacognitive efficiency showing extreme confidence ratings are less informative about the accuracy of a choice. Empirical plot adapted from Shekhar & Rahnev (2021a). **(D)** Negative (cross-trials) relationship between confidence and RT and positive (cross-conditions) relationship between confidence and RT.

Each dot denotes a level of difficulty. Empirical data adapted from Vickers & Packer (1982). Error bars denote 95% confidence intervals of the model simulations.

The first of these effects, the positive relationship between confidence and stimulus discriminability, follows from the strength of the signal in the ABS. More discriminable stimuli will result in more homogenous evidence supporting one alternative (i.e., samples will more consistently support one response alternative over the other), and because decision confidence is a transformation of the proportion of samples that support the chosen response, more discriminable stimuli will on average produce higher confidence judgments (see Equation 8). Conversely, on more difficult trials, the evidence will be more heterogeneous and so the ABS predicts lower average decision confidence. Figure 9A shows this qualitative effect arising in ABS model simulations, in which confidence is expressed on a probability scale which is ordinally related to the scale with which the empirical data were collected, and it is produced by all model variants (see Table 2).

Second, average confidence ratings tend to be higher for correct responses than for incorrect responses (e.g., Ariely et al., 2000; Baranski & Petrusic, 1998; Vickers, 1979; Vickers, 2014; Vickers & Packer, 1982). This so-called “resolution-of-confidence” effect also holds true even if stimulus difficulty is held constant (Baranski & Petrusic, 1998) and even if choice and confidence are simultaneously elicited from participants (Ratcliff & Starns, 2009; Van Zandt, 2000; Kiani et al., 2014). Once again, the SPRT cannot properly explain this effect given that its thresholds are fixed prior to, and independently from, the characteristics of particular trials (e.g., it is constant across all trials or randomly drawn from a fixed distribution). However, if we make the assumption that people have the correct generative model of the task (i.e., the probability of generating a sample that supports the correct alternative is the largest among all other alternatives), the ABS predicts that correct

responses will on average be made with higher confidence. This is tied to the explanation for slower errors above: autocorrelations cause errors to be slower on average, and slower responses produce lower confidence judgments (see Equation 9). Therefore, the ABS predicts a resolution-of-confidence effect in experimental conditions that produce slow errors (see Figure 9B). As this effect requires both optional stopping and autocorrelated samples, as also are required to produce slow errors, only the full model and the no prior variant produce it (see Table 2).

Third, studies have shown that the meta-cognitive judgments in confidence ratings generally carry less information about the accuracy of a decision than would be predicted by a purely normative account like the SPRT. Thus, there seems to be a systematic deficit in “metacognitive efficiency” (Shekhar & Rahnev, 2021a; 2021b). To give an intuition, imagine a participant is asked to make a decision whether to respond *A* or *B* to a stimulus. The participant’s ability to discriminate between the alternatives (i.e., d') can be calculated, based on SDT, by using the percentage of *A* stimuli that are correctly identified (i.e., hits) and the percentage of *B* stimuli that are incorrectly identified as *A* stimuli (i.e., false alarms). This standard d' measure can also be extended to metacognition by choosing a confidence criterion and recalculating the hit and false alarm rates from confidence judgments that exceed this criterion to produce a *meta_* d' ¹². SDT predicts that d' equals *meta_* d' for any confidence criterion and so predicts that metacognitive judgments are always efficient (while

¹² Informativeness of choices and confidence ratings are measured as stimulus sensitivity d' and *meta_* d' respectively (Maniscalco & Lau, 2012; Fleming & Lau, 2014). More specifically, $d' = \phi^{-1}(\text{hit rate}) - \phi^{-1}(\text{false alarm rate})$ where ϕ^{-1} is the inverse of the cumulative Gaussian distribution. *meta_* d' is calculated in the same manner but with the hit rate and the false alarm rate tallied according to a criterion value that partitions confidence ratings. More specifically, the hit rate is the proportion of trials in which participants reported high confidence given a correct response, whereas the false alarm rate is the proportion of trials in which participants reported high confidence given an incorrect response; and the confidence criterion value determines whether a confidence judgment is considered high or low.

the SPRT predicts constant confidence judgments and so cannot be evaluated using this measure). By contrast, a value of $meta_d'/d' < 1$ would indicate that information available for the decision is lost in part or in whole when making confidence judgments. Empirically, metacognition has been found to be inefficient, and moreover $meta_d'$ decreases relative to d' as the confidence criterion increases, meaning that higher confidence ratings are less informative than lower confidence ratings (Shekhar & Rahnev, 2021a; 2021b). Metacognitive inefficiency has been explained by adding additional noise to confidence judgments (Shekhar & Rahnev, 2021a).

While it would be straightforward to add noise to the ABS confidence judgments, surprisingly this additional noise is not necessary to produce such metacognitive inefficiency; in fact, there are multiple routes for the ABS to produce this effect already offered in the current specification. A first route derives from more informed decisions being overall less confident decisions. For example, imagine using a stopping rule with $\Delta = 2$ and a symmetric Beta(1,1) prior on responses. If a decision is made based on only a total of 2 samples, then both will have to be in favor of the chosen response and confidence will be 75% (i.e., plugging these values in Equation 8: $\frac{2+1}{2+2} = 75\%$). However, if a decision is made based on a total of 100 samples then only 51 can have supported the chosen alternative (because the stopping rule requires $\Delta = 51 - 49 = 2$) and confidence will be about 51% (i.e., plugging these values in Equation 8: $\frac{51+1}{100+2} \approx 51\%$). Thus, with optional stopping, lower confidence decisions will be based on more samples (and so have longer RTs) and hence will be more informative (see Equation 9). A second route derives from basing confidence judgments on discrete samples of hypothesis counts rather than the Gaussian distributed sensory evidence assumed by SDT; this applies even if samples are independent, a fixed number of samples are generated, and no prior is used (see Appendix D). Therefore, the ABS predicts decreasing

metacognitive efficiency for more extreme confidence judgments not only for the full model (see Figure 9C) but also for all its variants (see Table 2).

Finally, confidence is empirically observed to systematically vary with RTs, with positive (cross-condition) and negative (cross-trial) relationships between confidence and RTs (see Figure 9D). When people are forced to respond more quickly in a particular experimental condition, their confidence reduces, which is in line with the standard speed-accuracy trade-off, assuming the confidence positively co-varies with accuracy (Vickers & Packer, 1982; Irwin et al., 1956). Both the SPRT and the ABS can capture the positive (cross-condition) relationship simply by varying the threshold according to experimental conditions: emphasizing accuracy moves the threshold further away from the starting point of the accumulator (and the opposite is true for the speed condition). Higher threshold values in the SPRT lead to more extreme final log odds and therefore higher confidence readouts. Higher threshold values in the ABS (i.e., larger Δ) naturally lead to higher confidence as shown in Equation 9.

However, within a condition, people are *more* confident in decisions they reach quickly – intuitively, the “easy” trials are decided quickly and with high confidence (Baranski & Petrusic, 1998; Vickers & Packer, 1982). Crucially, the SPRT cannot explain this because the strength of evidence at which a decision is made depends only on the threshold, which is determined prior to, and hence independently from, the characteristics of any particular trial. The ABS can explain this negative (cross-trial) relationship because earlier termination (for a fixed threshold Δ) implies that there will be a higher proportion of evidence supporting the chosen alternative. As a result, the ABS predicts that within a condition, faster decisions will be given with higher confidence.

Confidence Intervals

So far, we have considered confidence in decisions. But confidence reports can also be elicited for estimates by asking for confidence intervals. Commonly a participant is given a probability first and then asked to produce an interval (by giving upper and lower bounds) that correspond to the probability (e.g., ‘give the smallest interval which you are 60% certain to include the number of dots which appeared onscreen: between ____ and ____ dots’). However, this procedure can also be reversed: participants can be shown an interval of some quantity of interest and then asked to evaluate the probability of that interval (e.g., ‘what is the probability that the number of dots which appeared onscreen falls in the range of 23 to 25?’; Juslin & Persson, 2002).

In the ABS, confidence interval production and evaluation both are driven by very similar mechanisms to those underlying the naïve intuitive statistician model of Juslin, Winman, & Hansson (2007). Taking a set of samples, a confidence interval can be produced by using the lower and upper bounds of the sample coverage (i.e., empirical quantiles of the samples) as the lower and upper bounds of the confidence interval. When the values of the quantiles are not explicitly represented in the sample (e.g., deriving a 93% CI based on 5 samples), linear interpolation was assumed to fill in the gap between the samples (Juslin et al., 2007). This mechanism correctly predicts the considerable overconfidence in interval production found empirically (see Figure 10A dots; Juslin et al., 2003; Juslin et al., 2007). This is because for small sample sizes, the empirical quantile of the sample will have a shorter range than the confidence interval from the posterior because distributional tails tend to be underrepresented within a few samples. Therefore, the proportion generated from the sample will be too small, producing an overconfident interval in our simulations (see Figure 10B dots). One might then question why interval production overconfidence is not corrected in the same manner described for probability judgments above where useful prior knowledge

is incorporated – this lack of correction for confidence interval production is what was “naïve” about the naïve intuitive statistician model. Corrections for intervals, however, depend on the functional form of the distribution, so that a general correction process is difficult to establish in the ABS. While the standard computation of a confidence interval assumes a Gaussian distribution, for unknown distributions confidence intervals are usually produced by bootstrapping. For the purposes of producing the confidence interval for a sample, as opposed to producing the confidence interval for a mean, bootstrapping is essentially what the ABS does.

In contrast to confidence interval production, confidence interval evaluation shows very different empirical results: here there is little to no overconfidence with only a small degree of conservatism at the extremes of the subjective probability (see Figure 10A squares; Juslin et al., 2003). This arises in the ABS (see Figure 10B squares), using the simplest possible assumption (and following Juslin et al., 2007) that people answer this question by generating samples and calculating the proportion that fall within the provided interval. As noted by Juslin et al., (2007), this proportion is an unbiased estimator, and hence shows good calibration¹³.

¹³ There is only slight overconfidence predicted by the naïve intuitive statistician model when the internal generative model does not perfectly describe the data generating process (Juslin et al., 2007), and the ABS would show the same effect with an imperfect model of the data generating process.

Figure 10. (A) Empirical data for interval evaluation (i.e., probability judgment) and interval production, adapted from Juslin et al. (2003). Interval evaluations were relatively well calibrated while substantial overconfidence was observed in interval production. The dashed line illustrates perfect calibration. *(B)* ABS predictions of confidence interval production and evaluation: strong overconfidence in interval production (blue dots) and no overconfidence in interval evaluation (red squares). The horizontal axis indicates either the requested interval coverage (production) or the judged probability of the interval (evaluation), while the vertical axis indicates the empirical proportion of events covered by the interval.

Decisions Affecting Later Estimates

Besides eliciting confidence judgments after choices, experimenters have often asked participants to provide separate secondary responses to the same stimulus. One example is the decision-estimation task where people were asked to first choose, say, whether the number of dots which appeared on-screen was greater or smaller than 25, and then are asked, immediately following the choice, to estimate the number of dots. In this setting, an estimate can be influenced by the preceding choice (e.g., Tversky & Kahneman, 1974; Jazayeri & Movshon, 2007). In cognitive judgments, estimate have often been observed to be pulled

towards a preceding arbitrary value—a phenomenon better known as the *anchoring* bias (Tversky & Kahneman, 1974; Epley & Gilovich, 2006). For example, in the famous study of Tversky and Kahneman (1974), participants were first asked to choose whether the percentage of African countries in the United Nations was higher or lower than a value (h^*), and then give an estimate of that percentage. The comparison value used in the choice, h^* , was generated randomly and so should have been irrelevant to the distribution of hypotheses (and thus irrelevant to the estimate too), but estimates were biased toward h^* .

However, in an almost identical paradigm of decision-estimation tasks, perceptual judgments showed the opposite effect: estimates of aspects such as dot orientation or direction of motion have been observed to be pushed away from h^* (Jazayeri & Movshon, 2007; Zamboni et al., 2016; Luu & Stocker, 2018). The phenomenon is better known as the *repulsion* effect found in perceptual tasks.

Existing models of anchoring cannot predict the repulsion effect and vice versa. This is because they only predict one direction of bias (e.g., Tversky & Kahneman, 1974; Strack & Mussweiler, 1997; Jazayeri & Movshon, 2007; Luu & Stocker, 2018), and thus fail to capture the co-occurrence of anchoring and repulsion. While this would be tenable if anchoring and repulsion were each specific to their respective (cognitive or perceptual) domains, a recent empirical investigation suggests otherwise (Spicer et al., 2022). In this work, we noted that the location of the comparison value, h^* , relative to the distribution of hypotheses has not typically been the same across cognitive and perceptual paradigms. Indeed, it was found empirically that the relative location of h^* determines whether the subsequent estimates will be pulled toward or pushed away in both cognitive and perceptual tasks. Specifically, estimates are drawn towards distant values of h^* (replicating the anchoring effect) but pushed away from nearby values of h^* (replicating the repulsion effect; Spicer et al., 2022). This

finding is consistent with a common general-purpose algorithm underlying decision-making in both cognition and perception.

The anchoring effect, the repulsion effect, and their dependence on the relative position of h^* are captured by the ABS assuming that the set of samples generated to make the choice is then reused to produce the estimate. To explain anchoring, the ABS follows the approach of Lieder et al. (2018) and assumes that local sampling algorithm uses h^* as an initial hypothesis. For a small number of iterations, the local sampler will then be biased toward the initial hypothesis. Anchoring is then produced in our simulations for the full model and all variants except the direct sampling variant (see Figure 11 and Table 2).

To explain repulsion, we first note that in the ABS h^* effectively partitions the hypothesis space into two binary response regions. The sampling algorithm is adaptively terminated when a sufficient number of samples support one alternative over the other, with the amount determined by the threshold parameter (i.e., Δ). This adaptive stopping rule produces a repulsion bias if the estimate is also based on the same set of samples (Zhu et al., 2019), because the sampling process will terminate when the weight of evidence favors one hypothesis rather than when the evidence is finely balanced: In effect, optional stopping for choice biases the subsequent estimate away from indifference (i.e., the decision boundary). Thus, repulsion is produced by the full model and all of the variants except for the fixed sample size variant (see Figure 11 and Table 2).

Figure 11. (A) Experimental data on the decision-estimation task. For the region of correct hypotheses in the range of [21, 30], the estimates were pushed away from a nearby comparison value (25.5) used in the preceding decision task (blue bars), while pulled toward a far-off comparison value (75.5; red bars). The data was replotted from Spicer et al. (2022). **(B)** Simulating the ABS (*Right*), its direct sampling variant (*Left*), and fixed sample size variant (*Middle*) on decision-estimation tasks. The full ABS model predicts both anchoring (for far-away target stimuli; red bars) and repulsion effects (for close-by target stimuli; blue bars), whereas the direct-sampling variant only predicts the repulsion effect and the fixed-sample-size variant only predicts the anchoring bias. Target Gaussian distributions for sampling with means in the range of [21, 30] were shown in black solid lines, which were vertically rescaled by a factor of 1/4 to aid visualization.

Cross-trial Autocorrelations in Estimates and RTs

Substantial cross-trial autocorrelations are an important, and often unexplained, aspect of human behavior. For example, long-range dependencies in estimates and in RTs known as $1/f$ noise have been observed in many cognitive and perceptual tasks and can explain more variance in behavior than the experimental manipulations (Gilden et al., 1995; Gilden, 2001; Wagenmakers et al., 2004)¹⁴. In these tasks, participants were instructed to repeatedly

¹⁴ $1/f$ noise goes under various names in different literatures (e.g., pink noise and flicker noise). Besides cognition, diverse complex processes have also been found to exhibit this type of

estimate fixed physical quantities (e.g., a 1-sec temporal interval or a 1-inch spatial interval) or to repeatedly choose between two options. The statistical features of the time series produced by participants were analyzed in the frequency domain, with the high-frequency components corresponding to trials that are close together, whereas low-frequency components correspond to trials that are well separated. The power of each of these components for explaining the time series are then calculated (i.e., in a spectral density analysis; Sheu & Ratcliff, 1995; Gilden et al., 1995; Gilden, 2001). Standard statistical processes show different relationships between frequency and power: in a random walk the power falls off with $1/f^2$ noise (i.e., a slope of -2 in log-log power spectra), whereas white noise (also called independent sampling or direct sampling) has a flat power spectrum (i.e., $1/f^0$ noise and a slope of 0). In a time-series containing long-range serial dependence, as has been found in human data, power spectra typically have a slope between -1.5 and -0.5, and are thus categorized as $1/f$ noise. The long-range autocorrelations in $1/f$ noise are not straightforward to produce, generally requiring complex processes to do so (Gardner, 1978).

Further complicating the picture, while time-series of estimates have long-range autocorrelations that are classed as $1/f$ noise (Gilden et al., 1995; Gilden, 2001; Wagenmakers et al., 2004; Zhu et al., 2021), RT time series fluctuate as $1/f$ noise but with a log-log slope that is shallower than that of estimates (Van Orden et al., 2003; Wagenmakers et al., 2004). As shown in Figure 12, the ABS qualitatively reproduces the observed autocorrelations in time series of RTs and estimates. The cross-trial autocorrelation in estimates is predicted by the cross-trial carryover of the sampler's location in the autocorrelated MC³ algorithm (Zhu et al., 2018; Zhu et al., 2022c): the initial location of the sampler for the present trial is the

long-range dependence. For example, MEG and EEG data from human brains, and indeed classical music, all exhibit $1/f$ -like fluctuations (Novikov et al., 1997; Linkenkaer-Hansen et al., 2001)

last sample for the preceding trial. In comparison, the RT time series is predicted to be less autocorrelated because samples generated by the MC³ are accumulated to a threshold to produce the RT; this is a non-linear transformation of autocorrelated samples which “whitens” the power spectrum. Simulations of the full model demonstrate both these effects, and as it is driven by the MC³ algorithm it occurs for all variants except for the direct sampling variant (see Figure 12 and Table 2).

Figure 12. Power spectra for time series of estimates and RT. Dashed lines denote power spectra by a simulated participant with solid lines showing the average. RT time series are colored in red, whereas time series of estimates are in blue. (*Left*) The direct sampling variant predicts independent estimates and RTs and thus exhibits a flat line (i.e., the power spectrum of white noise). (*Right*) The ABS model predicts autocorrelations in estimates and RTs with the latter displaying flatter slopes than the former (i.e., the power spectra of $1/f$ noise).

Summary

Using a consistent set of parameter values, we have shown that the ABS qualitatively captures empirical results and relationships observed across probability judgments, estimates, confidence intervals, choices, confidence judgments, and response times (see Table 2). The wide range of predicted behaviors is based on an internal probabilistic model using a fine-grained set of hypotheses. The process of inferring the posterior probability of the hypotheses is governed by Bayes' rule and approximated using an autocorrelated sampling algorithm. While the autocorrelation in the sampling algorithm is motivated as making the sampling process computationally efficient, it is crucial for explaining many empirical effects such as slow errors, anchoring, and cross-trial autocorrelations. Assuming each sample generated is costly, turning these samples into choices relies on an optional stopping rule that trades the benefits of larger samples against the cost of sampling. In turn, the optional stopping rule helps explain empirical effects such as the repulsion effect, the resolution of confidence, and metacognitive inefficiency. The probabilistic model also learns from trial history, using the adaptive prior. This prior helps explain effects such as conservatism, the conjunction fallacy, partition dependence, and fast errors. In sum, this rational process is fruitful for understanding a wide assortment of human behavior.

Table 2

Which Models Reproduce the Key Empirical Targets

	Empirical Effects	SPRT	Bayesian Sampler (Zhu et al., 2020)	Autocorrelated Bayesian Sampler			
				No prior variant	Direct sampling variant	Fixed sample size variant	Full model
Probability Judgments	Conservatism and probabilistic identities	N/A	✓	✗	✓	✓	✓
	Mean-variance relationship	N/A	✓	✗	✓	✓	✓

Decisions Affecting Later Responses	Explicit subadditivity	N/A	✓	✗	✓	✓	✓
	Conjunction fallacy	N/A	✓	✗	✓	✓	✓
	Implicit subadditivity in typical unpacking	N/A	✗	✓	✗	✓	✓
	Implicit superadditivity in atypical unpacking	N/A	✗	✓	✗	✓	✓
	Partition dependence	N/A	✓	✗	✓	✓	✓
	Anchoring	N/A	N/A	✓	✗	✓	✓
	Repulsion	N/A	N/A	✓	✓	✗	✓
	Speed-accuracy trade- off	✓	N/A	✓	✓	✓	✓
	Slow errors	✗	N/A	✓	✗	✗	✓
	Fast errors	✗	N/A	✗	✓	✗	✓
Accuracy and Response Times	Near-linear relationship of RT quantiles with a fan shape	✓	N/A	✓	✓	✗	✓
	Positive relationship between confidence and stimulus discriminability	✓	N/A	✓	✓	✓	✓
	Resolution of confidence	✗	N/A	✓	✗	✗	✓
	Metacognitive inefficiency	✗	N/A	✓	✓	✓	✓
	Negative (cross-trials) relationship between confidence and RT	✗	N/A	✓	✓	✗	✓
	Positive (cross- conditions) relationship between confidence and RT	✓	N/A	✓	✓	✓	✓
	Strong overconfidence in self-produced confidence intervals	N/A	N/A	✓	✓	✓	✓
	Little or no overconfidence in evaluation of confidence intervals	N/A	N/A	✓	✓	✓	✓
	Large long-range autocorrelation in estimation time series	N/A	N/A	✓	✗	✓	✓

Lesser long-range autocorrelation in RT time series	X	N/A	✓	X	X	✓
---	---	-----	---	---	---	---

Note. N/A denotes that the model is non-applicable to the empirical effect because it does not produce the relevant behavioral measure.

Comparison with Competing Models

There are many models that can produce at least a subset of the empirical effects that the ABS does, and many were briefly mentioned in the text above. Here we compare the ABS first to other models of probability judgments and then to drift-diffusion models of choice, response time, and confidence.

Models of Probability Judgments

Intensive modeling efforts have also been directed at explaining human probabilistic reasoning, spurred on by the identification of biases, particularly those summarized in Table 1, demonstrating that people's judgments do not conform with the rules of probability theory (e.g., Peterson & Beach, 1967; Nilsson et al., 2009; Costello & Watts, 2014; Dasgupta et al., 2017; Zhu et al., 2020; Tversky & Koehler, 1994; Hilbert, 2012). Many models have assumed that probability judgments follow a deterministic process, albeit one that violates the rules of probability theory. For example, one type of model, geared toward accounting for conjunction fallacies, assumes that probability estimates of conjunctions are the weighted average of the probabilities of their constituent events, which produces above-chance conjunction fallacy rates and can reproduce several probabilistic identities (Fantino, Kulik, Stolarz-Fantino, & Wright, 1997; Nilsson, Winman, Juslin, & Hansson, 2009; Nilsson, Juslin, & Winman, 2016). However, these models require additional mechanisms to match the

empirically observed combination of above-chance and below-chance rates of conjunction fallacies that the ABS can produce (Nilsson et al., 2009; Fisk & Pidgeon, 1996).

A different type of deterministic approach that has been used to explain conjunction fallacies is based on quantum probability. Here probabilities are based on projections of event subspaces. If the events are compatible, probability judgments are indistinguishable from classical probability theory, but if the events are incompatible then interference produces probability judgments that deviate from classical probability theory. These deviations are such conjunction and disjunction fallacies will occur at rates above chance, and in this way, quantum probability can produce both above-chance and below-chance conjunction fallacies depending on how the events are represented (Busemeyer et al., 2011). Quantum probability has explained a wide range of probabilistic biases, including some not covered here (Pothos & Busemeyer, 2022). However, there are also probabilistic identities that quantum probability cannot reproduce, that are predicted by sampling-based models, including the ABS (Costello & Watts, 2018; Zhu et al., 2020).

A third deterministic approach is support theory, which was developed to explain subadditivity biases. The core assumption of support theory is that the probability of event descriptions is evaluated rather than the probability of the events themselves and does not incorporate the probabilities of events that are not immediately available (e.g., those not mentioned in the descriptor of events). This approach elegantly explains both a range of implicit subadditivity results, as well as explaining why subadditivity does not occur for mutually exclusive binary events. However, it does not produce the later finding that an atypical unpacking of events produces implicit superadditivity (Sloman et al., 2004), and requires additional mechanism such as an “ignorance prior” which pulls probability judgments towards indifference between the available responses (Fox & Rottenstreich, 2003). These different mechanisms have echoes in the ABS. In the ABS, a hypothesis is “available”

only if it has been sampled, and the event description influences the starting point of the local sampler. Further, the prior on responses is a principled version of the ignorance prior, one that is uncertain about the underlying probabilities because often only a small number of samples is available.

Recent approaches have rejected purely deterministic approaches and explored the alternative possibility that stochastic mechanisms explain the biases in probability judgments. For example, simple unbiased response noise has been shown to produce subadditivity (Bearden, Wallsten, & Fox, 2007; Brenner, 2003). However, unbiased response noise alone does not explain why subadditivity still occurs for median judgments. A more promising alternative is to consider corruptive noise in memory or evidence accumulation, which can produce stronger biases (e.g., Costello & Watts, 2014; Erev et al., 1994; Hilbert, 2012). In a leading stochastic model, Probability Theory plus Noise (PT+N), people are assumed to first draw independent samples from a probabilistic representation, and unbiased “counting noise” is added to individual samples to reflect an error-prone cognitive system (Costello & Watts, 2014). This counting noise pulls probability judgments towards indifference and allows the PT+N to capture empirical results such as explicit subadditivity, the conjunction fallacy, and a wide range of probabilistic identities (Costello & Watts, 2014; 2017; 2018; Costello et al., 2018). The PT+N has impressive empirical coverages, although it has recently been argued that it does not fully reproduce all the mean-variance relationship in probability judgments (Sundh et al., 2021): while it will produce the inverted U-shaped relationship between the mean and variance of judgments, the curve will not be pulled inward as is with the empirical data and as the ABS predicts. This mean-variance relationship also stands as a challenge to deterministic models because it is not easily produced by simply adding response noise to a deterministic model.

Drift-Diffusion Models

One family of models that deserves more extensive discussion is the family of Drift-Diffusion Models (DDMs; Ratcliff et al., 2016; Gold & Shadlen, 2007; Krajbich & Rangel, 2011; Drugowitsch et al., 2012). While there are many members of this family, they all describe decision making as a stochastic process similar to that of a biased random walk (or a biased diffusion process, in continuous time) in which the path of the accumulator is, on average, biased by the drift rate (Ratcliff & McKoon, 2008; Bogacz et al., 2006). For binary choices, this is determined by the difference in the evidence signals supporting for the two alternatives. In line with the stopping rule of the SPRT, the threshold reached in the DDM decides the choice and the time taken to do so the response time. However, unlike the static summation process of log-likelihood ratios in the SPRT, the accumulator in the DDM also diffuses because the accumulator is corrupted by noise, typically white noise. In perceptual tasks, the drift rate is related to which choice is objectively correct (Ratcliff & McKoon, 2008), whereas in high-level cognitive tasks where people are choosing their preferred option the drift rate is assumed to be related to the relative appeal of the alternatives (Krajbich, Armel, & Rangel, 2010; Krajbich & Rangel, 2011).

There are generally strong theoretical links between these models and the normative framework of the SPRT: in the continuous limit, the SPRT converges on the DDM and the drift rate and the corruptive noise in the DDM can jointly mimic calculation of likelihood ratios in the SPRT (Bogacz et al., 2006). However, implementing an optimal statistical decision test in the form of the DDM also generates a number of useful theoretical and empirical insights that were not originally part of the SPRT. First, the psychologically implausible assumption that people are required to have global knowledge of the generative model of the task to calculate the exact likelihoods (e.g., $P(s_t|A)$ or $P(s_t|B)$) is implicitly relaxed by the DDM because the drift rate and diffusion noise are free parameters that are

recovered from fitting to behavioral data. Thus, the DDM does not need to calculate with the exact cumulative differences in evidence as supposed by the SPRT, greatly improving the DDM's computational plausibility, given that the exact likelihood ratios are almost always impractical to compute in real time except in simple toy problems.

Second, extensions of the DDM can also account for empirical features such as those noted above which are not accounted for by the SPRT. For example, slow and fast errors (Ratcliff & Rouder, 1998; Townsend & Ashby, 1983) can be produced by further assuming that model parameters are variable across trials (Laming, 1968; Rouder, 1996; Ratcliff, 1981; Ratcliff & Rouder, 1998). In particular, varying the drift rates trial-by-trial generates slow errors, while varying the starting point of the accumulator predicts fast errors (Laming, 1968; Ratcliff & Rouder, 1998). The ABS model works in a similar fashion, as autocorrelation acts like variability in drift rates and a biased prior of evidence acts like variability in the starting point of the accumulator.

Moreover, the benefits of using the DDM instead of the SPRT also apply to explaining confidence judgments. The SPRT predicts the confidence ratings to be identical between correct and incorrect responses because, for a fixed symmetric threshold, there will always be the same level of evidence difference accumulated favoring the selected option,¹⁵ and the probability that the chosen option is correct is simply read out from the final state of the accumulator. As outlined above, such predictions are contradicted by empirical data in which choice accuracy and confidence are positively related (Baranski & Petrusic, 1998; Dougherty, 2001; Vickers, 1979; Yeung & Summerfield, 2014). To reconcile the confidence data with the SPRT, one kind of DDM introduced another assumption in which the same drift-diffusion process continues to run for a period of times after the choice has been made

¹⁵ Regardless of trial-by-trial variability in drift-rates and/or starting points.

but before the confidence judgments (Pleskac & Busemeyer, 2010). As the accumulator has a bias toward the correct choice, this continued accumulation after the choice and before the confidence judgment, no longer bounded by the fixed threshold of the decision, will drive the confidence ratings towards supporting the correct choice. Hence, with this additional assumption, this DDM can correctly predict that people should report higher confidence in correct responses than in errors, and moreover that the resolution of confidence effect grows with the delay between choosing and reporting confidence. However, since the temporal structure supposed by this assumption is that confidence occurs after the choice, this DDM cannot explain why the relationship between choice accuracy and confidence also appears when confidence judgments are given simultaneously with a decision (e.g., Li & Ma, 2020; Kiani et al., 2014). In explaining these data, researchers have assumed that confidence decreases with response time (Kiani et al., 2014; Calder-Travis et al., 2020), and as occurs in the ABS.

Alternative versions of the DDM, such as the RTCON model, have been developed to capture no-choice confidence rating (Ratcliff & Starn, 2009). RTCON assumes that each confidence rating has an independent diffusion process and the first diffusion process to reach the threshold determines the confidence rating. So, for 7 confidence ratings, there are 7 diffusion processes racing to the threshold. RTCON captures many key empirical relationships between confidence and RT (Ratcliff & Starn, 2009), but appears to have difficulty capturing the interaction between confidence and choice because the model was not initially designed for tasks involving both choices and confidence judgments (Pleskac & Busemeyer, 2010). The approach of RTCON has been generalized to continuous-response paradigms in which there are an infinite number of potential responses: the circular diffusion model (Smith, 2016) and the spatially continuous diffusion model (Ratcliff, 2018). The two

models have recently been integrated into a unified framework where geometric similarity among response options is represented (Kvam & Turner, 2021).

Overall, the family of models encapsulated by the DDM effectively describe behavior even in tasks far different from the perceptual tasks for which it was initially developed (Ratcliff et al., 2016; Bogacz et al., 2006), including tasks in which there is little to no perceptual noise, such as value-based decisions (Busemeyer & Townsend, 1993; Usher & McClelland, 2004; Milosavljevic et al., 2010) and recognition memory tasks (Starns & Ratcliff, 2014; Ratcliff, Gomez, & McKoon, 2004; Ratcliff, Thapar, & McKoon, 2011). Because of its strong normative underpinnings in the SPRT, and its psychologically plausible assumptions, the DDM has been widely used in psychology, economics, and neuroscience (Ratcliff et al., 2016; Drugowitsch et al., 2012; Fehr & Rangel, 2011; Ratcliff & Smith, 2015). Indeed, the DDM has become the default framework in decision-making research. That being said, the broader scope of behavioral responses including choice, RT, confidence, and estimates captured by the ABS have not yet been united within a single implementation of the DDM. This is partly because the task representation needed for choice is different than that needed for estimates or confidence intervals (in particular, the DDM represents an accumulated value – a summary statistics of the sample – but does not retain the sample itself). Thus, while the ABS and the DDM share similar descriptive capabilities, the ABS arguably has the advantage in terms of parsimony given the breadth of behavior covered within its single framework, capturing relationships, such as the combination of anchoring and repulsion effects described above, not explained by current DDM approaches.

Toward Complete Task-Specific Models

The ABS so far has explained the generic judgment and decision-making process. But how might this be applied to specific tasks? Fortunately, we can relate our approach to existing models in the literature that take quite a similar approach to that taken here, and indeed have helped inspire our work. Here we describe two successful existing models, Nosofsky & Palmeri's (1997) exemplar-based random-walk (EBRW) model and Blurton et al.'s (2020) visual attention model, and we outline how using the ABS would involve only minor modifications (such as adding autocorrelations) to them. The result is that we can view these existing models as complete task-specific models in the ABS framework.

The EBRW operates in a hypothesis space of exemplars which represent categories (Nosofsky & Palmeri, 1997). This representational assumption is inherited from the *generalized context model* where each individual exemplar is situated as a point in a multidimensional psychological space, and similarity between exemplars decreases with the distance between points in the space (Nosofsky, 1984; Shepard, 1987). Building on this representation, EBRW further assumes that exemplars are retrieved sequentially as in a random walk process, predicting the time course of categorization and recognition decision-making (Nosofsky & Palmeri, 1997). Similar ideas can be found in the Poisson random walk model of visual attention (Blurton et al., 2020; Bundesen, 1990). In this model, a series of tentative categories (i.e., hypotheses) is proposed and accumulated until one category has accrued enough samples more than any other category (Blurton et al., 2020). The generation of tentative categories is governed by the theory of visual attention (Bundesen, 1990).

Across the two models, there is a common mechanism that integrates over hypotheses for response selection. While neither model was originally motivated from normative principles, they can both be seen as special cases within our framework in which decisions are driven by samples of hypotheses (see Appendix B for detailed comparisons), but with the

addition of features such as autocorrelated samples. More specifically, the exemplar-based representation of hypotheses of the EBRW can be adopted by the ABS when modeling categorization tasks, suggesting how the ABS could be applied to explain the effects of similarity and practice in categorization and RT, which have been captured by the EBRW. When combined with a Bayesian theory of visual attention, the ABS could also be generalized to account for human eye movement and object localization. In addition, the correspondence between the Poisson random walk and the direct sampling variant of the ABS provides a bridging condition that allows the ABS to account for detailed fits in response time distributions. Furthermore, adding autocorrelations and an adaptive prior to both the EBRW and Poisson random walk models generalizes these models to capture a wider range of empirical effects such as those found in confidence judgments. This demonstrates how the ABS can be adapted and work well in specialized tasks.

Discussion

The Autocorrelated Bayesian Sampler is a step towards a unified rational process of human behavior. Through our analysis, we have identified two key ideas that are necessary for such a unified framework: probabilistic models and approximate inference. Approximate inference via accumulating hypothesis samples means that the ABS views response time in a different way from most existing approaches, as primarily determined by time to mentally sample hypotheses, rather than accumulating more sensory data. We discuss this further below, pointing to possible reconciliations on the view of time as well as possible reconciliations between diffusion processes (and noisy probability judgment models) and the ABS. Next, we explore how the ABS could be extended both to multi-alternative tasks and how complex probabilistic representations could be incorporated. Finally, we discuss and comment on the extent to which the ABS has rational, Bayesian, basis.

Contrasting Views on the Role of Time

The existing dominant view on the role of time in decision making is to collect and integrate sensory evidence. In the binary choice example, the likelihood ratio between the two alternatives is represented exactly at any moment and the odds of the correct response are constantly updated in light of new sensory evidence or newly retrieved memories. This view has been adopted by many models including SDT, the SPRT, and by and large DDM approaches.

By contrast, the ABS takes a very different view on the role of time because it sees perception and cognition as emerging from probabilistic representations and computations. Instead of coding a single value of the sensory input, people are assumed to implicitly encode multiple values of the sensory input with their subjective uncertainty about those values. But this posterior is difficult to represent exactly for virtually all cognitive tasks; thus,

approximation is needed to access it. The passage of time is then viewed as being used for generating more samples from the posterior to refine this approximation. Thus, time matters because it allows the computational process of sampling the posterior to unfold, not because additional sensory data must be accumulated. In the limit, the gradual refinement of the posterior belief should lead to a convergence to the optimal choice.

These contrasting perspectives were also studied in more detail in Lengyel et al. (2015) which presents empirical evidence supporting the view of posterior approximation. It is, however, important to note that the two views on the role of time may not be mutually exclusive. One possible reconciliation could be that the brain first conducts evidence integration, and then a posterior based on the sensory evidence can be approximated with sampling, or these processes could be overlapping. Further analyzing the aspects of the two views on time may be an important topic for future research.

Diffusion, Noise, and Sampling

While we have contrasted other models and the ABS, there are also potential links between these approaches. Considering stochastic models of probability judgments, it may be possible to extend a sampling model with noisy counting, such as the PT+N, to explain vast majority of the effects that we explored above. In our previous work, we have shown that the PT+N and the Bayesian Sampler models can mimic one another's predictions of average probability judgments (Zhu et al., 2020). Building on this, it is possible to envision generalizing the PT+N in the same way we generalized the Bayesian Sampler to the ABS. First, rather than using independent samples, the PT+N could instead use a local, autocorrelated sampler such as MCMC or MC³. Indeed, Costello & Watts (2018) have begun to explore this possibility by introducing a type of autocorrelation or carry-over effect allowing earlier samples to influence later probability judgments. Second, when making

decisions, PT+N could also use optional stopping rather than a fixed number of samples to account for RT data. This could be a promising alternative to the ABS, although more work is required to develop this rough sketch into a formal model and determine how well it reproduces human data.

For diffusion models, an interesting starting point is the recent interest in continuous diffusion models (Ratcliff, 2018; Kvam, Marley, & Heathcote, 2022; Kvam, 2019; Kvam & Busemeyer, 2020). The main focus of the continuous diffusion models was to describe the cognitive processes underlying tasks that involve continuous responses such as orientation estimation (Ratcliff, 2018) and pricing (Kvam & Busemeyer, 2020). The accumulator is typically depicted in a two-dimensional space. Without any bias in starting point, the accumulator initializes in state $[0,0]$ (i.e., the origin). The amount of evidence accumulated is described as the distance from the origin, and there is a directional bias towards the option that is most favored at that moment. A continuous absorbing threshold (e.g., a semicircle whose center is the origin) defines the space of possible trajectories; when the threshold is reached, the diffusion process is terminated and a response is triggered. The size of the threshold regions that correspond to each response has been adjusted to make some responses more or less likely (Kvam & Turner, 2021).

This kind of mechanism could link a continuous diffusion process to the ABS. As long as the size of the threshold associated with each discrete fine-grained hypothesis is proportional to its posterior probability, then the continuous diffusion process will effectively be sampling from the posterior on fine-grained hypotheses¹⁶. If the fine-grained hypotheses

¹⁶ Alternatively, without changing the absorbing threshold, we can adjust the drift rate to mimic a sampling algorithm. The idea is that a diffusing particle that moves with a bias in the direction of higher probability can be linked to local sampling algorithms (Rossky, Doll, & Friedman, 1978; Roberts & Stramer, 2002; Ma et al., 2019). Indeed, local samplers move around the space stochastically, which can be described as a diffusion process. Therefore, by carefully designing the biasing method for diffusion, the long-term behaviour will converge to as if drawing samples from a

that are sampled are then processed to produce judgements and decisions as in the ABS, this would formally link the models, through substituting a diffusion-based sampling algorithm for the local sampling algorithm currently used. Thus, the distinction between diffusion and sampling may not be clear cut, and they may potentially be viewed as parts of a larger framework.

Multi-alternative ABS

Asking people to choose among more than two alternatives is often a useful strategy to test the generalizability of computational models developed to explain binary choice data. As previously noted, many aspects of the ABS are readily applicable to such choices: the hypothesis space can be divided into as many regions as required by the query, and the Beta prior generalizes to a Dirichlet distribution when considering more than two options. The optimal stopping rule derived from dynamic programming is, however, more difficult to calculate within realistic times when choosing between more than two alternatives (see Appendix C for detail). As we've noted, the max-minus-next heuristic is again often considered as a good approximation to the optimal stopping rule in multi-alternative choices (Dragalin et al., 1999; 2000) and, indeed, there is computational work suggesting that humans adopt the max-minus-next stopping rule to choose among many alternatives (Brown, Steyvers, & Wagenmakers, 2009). This heuristic stopping rule can also explain the best-known empirical result on the relationship between choice and RT, Hick's Law (Brown et al., 2009): that RT increases logarithmically with the number of alternatives (Hick, 1952; Proctor & Schneider, 2018).

distribution. For instance, a class of Lengevin diffusions has been designed such that the stationary distribution of the diffusion is in proportion to the target distribution of MCMC sampling (Roberts & Stramer, 2002; Ma et al., 2019).

Multi-alternative choice tasks with confidence have also recently been used to argue against the Bayesian confidence hypothesis—that confidence in a choice is the posterior probability of that choice. Li and Ma (2020) found that the best-fitting model for confidence ratings in a three-alternative choice task was *not* the posterior probability of the chosen option, but the difference between the probability of the chosen option and the probability of second most probable option. This result could be reconciled with the Bayesian confidence hypothesis through the ABS, assuming a stopping rule such as the max-minus-next heuristic. More formally, consider a three-alternative choice with options A , B , and C , with respective accumulated evidence i, j , and k . Further assuming that $i > j > k$, then option A is chosen where $i - j = \Delta$ following the max-minus-next stopping rule. The confidence difference between the best and the second-best is thus, $\text{Diff} = \text{Conf}_A - \text{Conf}_B = \frac{i-j}{i+j+k}$, and the total amount of evidence is related to this difference, $i + j + k = \frac{i-j}{\text{Diff}}$. Given that, we can rewrite the confidence for the chosen option A as follows:

$$\text{Conf}_A = \frac{i}{i + j + k} = \text{Diff} \frac{i}{i - j} = \text{Diff} \frac{i}{\Delta} \quad (10)$$

Thus, the best-fitting model in Li & Ma (2020), which was used to argue against the Bayesian confidence hypothesis, is proportional to the predictions of the ABS using the max-minus-next stopping rule.

Extending to Complex Representations

The behaviors discussed here are low-dimensional, with responses being situated within one- or two-dimensional spaces. As a result, the transformation from hypothesis samples to behavior is relatively simple, sometimes just with a linear mapping. But the scope and diversity of human behavior is much broader: many complex human behaviors are both high-dimensional and embedded in hierarchically organized spaces. Drawing, for example, or even copying line-drawings, requires sophisticated mental processes that represent a

description of a drawing's parts (e.g., lines and circles) and higher-order relations (e.g., repetition and hierarchy; Van Sommers, 1984; Tian et al., 2020). The motor system that implements routines and trajectories for turning these rich, structured representations into motor commands to produce drawings is also doing a more complex task than in low-dimensional behavior (for example, complex trajectories may to be segmented into discrete, and hierarchically organized actions). Thus, the hypothesis space from which outputs are selected can be open-ended, and hierarchically organized at a range of levels of abstraction. While it is difficult to see how sensory accumulation models such as the SPRT or the DDM might extend to such cases, it is at least possible in principle to see how a Bayesian, sampling-based approach might operate. For example, many existing models of cognition, perception, and motor control involve Bayesian inference over compositional symbolic representations (e.g., Lake et al., 2017) – and the relevant computations can only be approximated, often using sampling (frequently, using standard MCMC). An interesting direction for future research is how far these Bayesian models can be mapped into fine-grained data relating detailed measures of high-dimensional output (including, for example, accuracy and variability) to fine-grained performance features, such as timing, and autocorrelations across trials. In general, where a Bayesian cognitive model can be defined, a sampling approximation to that model can be created, and compared with detailed process data from experiments. Thus, the ABS provides a possible bridge from simple, but intensively-studied, decisions concerning binary choice or one-dimensional estimation, to models of cognition operating at full scale.

Assessing the Rationality of the ABS

Bayesian models of cognition, pitched at Marr's computational level, combines all available trial information (i.e., prior knowledge of hypotheses, $p(h)$, and the likelihood of the

data presented in the trial given a hypothesis, $p(h|s)$) using the rules of probability theory. In doing so, Bayesian models fully extract trial information with 100% efficiency (Zellner, 2002).

We see the ABS as Bayesian in two ways. In the first way, the ABS can provide a sample-based approximation to an underlying Bayesian representation of a task. Combining the ABS with a Bayesian representation produces a rational process model (Griffiths, Vul, & Sanborn, 2012). That is, it is an algorithmic approximation of a computational-level model, which transforms it into a process model. A rational process model does not align perfectly with the underlying computational-level model because sampling approximations inevitably lead to loss of information. As a result, sampling models predict mistakes, systematic biases, and variability in behavior that are due to using stochastic samples.

While we developed the ABS with a Bayesian perspective in mind, the underlying model does not necessarily have to be Bayesian. The probabilistic representation could be simply the relative frequencies of past events, without an associated model (e.g., Costello & Watts, 2014). Alternatively, the probabilistic representation might not even be described as optimal or rational (Tauber, Navarro, Perfors, & Steyvers, 2017). All that is required is that the representation can be written as a probability distribution, which covers a wide range of representations – any finite set of non-negative numbers can be normalized to become a probability distribution.

The second way in which the ABS is Bayesian is that Bayesian Monte Carlo is used to interpret the samples that are generated. This allows the ABS to take advantage of context-free expectations about the underlying probabilities, which improves probability (and confidence) judgments when a small number of samples give only imprecise information about those probabilities. We assume a conjugate prior on responses, making this process computationally very simple. Our analysis suggests that people do incorporate this prior on

responses in forming behaviors—as we have seen, this assumption explains many classical empirical effects in probability judgments. Simple adaptivity in constructing the prior on responses (e.g., adapt to immediate feedback from the preceding trial) has also been shown to be useful in explaining human data such as fast errors.

This raises the question of whether the ABS makes optimal use of the sampling mechanism it has available. This is a different kind of normative concern than just fully extracting trial information. Here the question is whether the ABS is *resource rational* (Lieder & Griffiths, 2020; Bhui, Lai, & Gershman, 2021). Aside from the Bayesian Monte Carlo prior on responses, which could be argued to be resource rational, there is the smaller-scale temporal tradeoff between either generating another sample which takes time or stopping and making do with the samples collected so far. The optimal stopping rule relies on solving a difficult dynamic programming problem (detailed in Appendix C). Indeed, given the fact that the optimal stopping rule is generally computationally intractable, we assumed another approximation to the optimal stopping rule with the max-minus-next rule, which has been independently suggested to well-approximate the performance of optimal stopping rule in the information theory literature (Dragalin et al., 1999). This optional stopping rule helps explain empirical patterns such as repulsion, slow errors, and resolution of confidence.

In short, we argue that exploiting algorithmic approximations to the optimal solution is the key feature of the ABS that justifies it as a rational process model. Our proposal, nonetheless, does not address the meta-level theoretical question concerning how the mind allocates cognitive resources across these approximation algorithms (e.g., the local sampler, the Bayesian Monte Carlo process, and the approximate dynamic-programming solution for stopping). A fully resource rational analysis would further specify the balance between the times spent on each approximation algorithm and the incentives from the task. Whether an optimal allocation of the limited cognitive resources is at play merits future investigations.

Conclusions

We have outlined a rational process model of human behavior: the ABS. The ABS is rooted in a Bayesian framework, where the cognitive system is presumed to have an internal probabilistic model, which describes how the sensory data is generated in the real world. But rather than representing and computing with probabilities, we assume that the cognitive system uses a tractable approximation of the posterior which is realized via a local sampling algorithm. Distinct aspects of these posterior samples are relevant for different types of query and when appropriately transformed, they provide a unified explanation of probability judgments, estimates, confidence intervals, choices, confidence judgments, and the time course of the posterior sampling accounts for response times.

Our framework shifts the locus of explanations for the accumulation of sensory input to computation (through sampling) over internal hypotheses. Thus, in our framework, the time-course, and variability, of behavior is primarily explained in terms of an internal, noisy, computational process (involved in sampling from the hypothesis space), rather than through the gradual accumulation of noisy sensory data. We demonstrate the usefulness of our theory by reproducing key pair-wise relationships and stylized facts for the six behavioral measures and point the way toward extending the approach to the complex probabilistic models required to describe the richness of human behavior.

References

- Abbott, J. T., & Griffiths, T. L. (2011). Exploring the influence of particle filter parameters on order effects in causal learning. In L. Carlson, C. Hoelscher, & T. F. Shipley (Eds.), *Proceedings of the annual meeting of the cognitive science society*. Austin, TX: Cognitive Science Society.
- Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological review*, 98(3), 409.
- Andrieu, C., De Freitas, N., Doucet, A., & Jordan, M. I. (2003). An introduction to MCMC for machine learning. *Machine learning*, 50(1), 5-43.
- Aragones, E., Gilboa, I., Postlewaite, A., & Schmeidler, D. (2005). Fact-free learning. *American Economic Review*, 95(5), 1355-1368.
- Ariely, D., Tung Au, W., Bender, R. H., Budescu, D. V., Dietz, C. B., Gu, H., ... & Zauberman, G. (2000). The effects of averaging subjective probability estimates between and within judges. *Journal of Experimental Psychology: Applied*, 6(2), 130.
- Attneave, F. (1953). Psychological probability as a function of experienced frequency. *Journal of Experimental Psychology*, 46(2), 81.
- Baker, C. L., Goodman, N. D., & Tenenbaum, J. B. (2008). Theory-based social goal inference. In B. C. Love, K. McRae, & V. M. Sloutsky (Eds.), *Proceedings of the thirtieth annual conference of the cognitive science society* (pp. 1447-1452). Austin, TX: Cognitive Science Society.
- Baker, C. L., Saxe, R., & Tenenbaum, J. B. (2009). Action understanding as inverse planning. *Cognition*, 113(3), 329-349.
- Bardolet, D., Fox, C. R., & Lovallo, D. (2011). Corporate capital allocation: A behavioral perspective. *Strategic Management Journal*, 32(13), 1465-1483.

- Baranski, J. V., & Petrusic, W. M. (1998). Probing the locus of confidence judgments: Experiments on the time to determine confidence. *Journal of Experimental Psychology: Human Perception and Performance*, 24, 929–945.
- Battaglia, P. W., Hamrick, J. B., & Tenenbaum, J. B. (2013). Simulation as an engine of physical scene understanding. *Proceedings of the National Academy of Sciences*, 110(45), 18327-18332.
- Bearden, J. N., Wallsten, T. S., & Fox, C. R. (2007). Contrasting stochastic and support theory accounts of subadditivity. *Journal of Mathematical Psychology*, 51(4), 229-241.
- Blurton, S. P., Kyllingsbæk, S., Nielsen, C. S., & Bundesen, C. (2020). A Poisson random walk model of response times. *Psychological Review*, 127(3), 362.
- 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.
- Bramley, N. R., Dayan, P., Griffiths, T. L., & Lagnado, D. A. (2017). Formalizing Neurath's ship: Approximate algorithms for online causal learning. *Psychological Review*, 124(3), 301.
- Brenner, L. A. (2003). A random support model of the calibration of subjective probabilities. *Organizational Behavior and Human Decision Processes*, 90(1), 87-110.
- Brooks, S., Gelman, A., Jones, G., & Meng, X. L. (Eds.). (2011). *Handbook of Markov chain Monte Carlo*. CRC press.
- Brown, S., Steyvers, M., & Wagenmakers, E. J. (2009). Observing evidence accumulation during multi-alternative decisions. *Journal of Mathematical Psychology*, 53(6), 453-462.

- Bhui, R., Lai, L., & Gershman, S. J. (2021). Resource-rational decision making. *Current Opinion in Behavioral Sciences*, 41, 15-21.
- Bundesen, C. (1990). A theory of visual attention. *Psychological Review*, 97(4), 523.
- Pothos, E. M., & Busemeyer, J. R. (2022). Quantum cognition. *Annual review of psychology*, 73, 749-778.
- Busemeyer, J. R., Pothos, E. M., Franco, R., & Trueblood, J. S. (2011). A quantum theoretical explanation for probability judgment errors. *Psychological Review*, 118(2), 193.
- 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.
- Calder-Travis, J., Charles, L., Bogacz, R., & Yeung, N. (2020). Bayesian confidence in optimal decisions. PsyArXiv.
- Chater, N., & Manning, C. D. (2006). Probabilistic models of language processing and acquisition. *Trends in Cognitive Sciences*, 10(7), 335-344.
- Chater, N., Zhu, J. Q., Spicer, J., Sundh, J., León-Villagrà, P., & Sanborn, A. (2020). Probabilistic biases meet the Bayesian brain. *Current Directions in Psychological Science*, 29(5), 506-512.
- Chen, M. H., Ibrahim, J. G., & Shao, Q. M. (2000). Power prior distributions for generalized linear models. *Journal of Statistical Planning and Inference*, 84(1-2), 121-137.
- Costello, F., & Watts, P. (2014). Surprisingly rational: probability theory plus noise explains biases in judgment. *Psychological Review*, 121(3), 463.
- Costello, F., & Watts, P. (2017). Explaining high conjunction fallacy rates: The probability theory plus noise account. *Journal of Behavioral Decision Making*, 30(2), 304-321.
- Costello, F., & Watts, P. (2018). Invariants in probabilistic reasoning. *Cognitive Psychology*, 100, 1-16.

- Costello, F., Watts, P., & Fisher, C. (2018). Surprising rationality in probability judgment: Assessing two competing models. *Cognition*, 170, 280-297.
- Courville, A. C., & Daw, N. (2007). The rat as particle filter. In J. Platt, D. Koller, Y. Singer, & S. Roweis (Eds.), *Advances in Neural Information Processing Systems*, 20.
- Dasgupta, I., Schulz, E., & Gershman, S. J. (2017). Where do hypotheses come from? *Cognitive Psychology*, 96, 1-25.
- Dayan, P., & Abbott, L. F. (2005). *Theoretical neuroscience: computational and mathematical modeling of neural systems*. MIT press.
- Diaconis, P., & Ylvisaker, D. (1979). Conjugate priors for exponential families. *The Annals of Statistics*, 269-281.
- Diederich, A., & Oswald, P. (2016). Multi-stage sequential sampling models with finite or infinite time horizon and variable boundaries. *Journal of Mathematical Psychology*, 74, 128-145.
- Dougherty, M. R. P. (2001). Integration of the ecological and error models of overconfidence using a multiple-trace memory model. *Journal of Experimental Psychology: General*, 130, 579-599.
- Dragalin, V. P., Tartakovsky, A. G., & Veeravalli, V. V. (1999). Multihypothesis sequential probability ratio tests. I. Asymptotic optimality. *IEEE Transactions on Information Theory*, 45(7), 2448-2461.
- Dragalin, V. P., Tartakovsky, A. G., & Veeravalli, V. V. (2000). Multihypothesis sequential probability ratio tests. II. Accurate asymptotic expansions for the expected sample size. *IEEE Transactions on Information Theory*, 46(4), 1366-1383.
- Drugowitsch, J., Mendonça, A. G., Mainen, Z. F., & Pouget, A. (2019). Learning optimal decisions with confidence. *Proceedings of the National Academy of Sciences*, 116(49), 24872-24880.

Drugowitsch, J., Moreno-Bote, R., Churchland, A. K., Shadlen, M. N., & Pouget, A. (2012).

The cost of accumulating evidence in perceptual decision making. *Journal of Neuroscience*, 32(11), 3612-3628.

Drugowitsch, J., Moreno-Bote, R., Churchland, A. K., Shadlen, M. N., & Pouget, A. (2012).

The cost of accumulating evidence in perceptual decision making. *Journal of Neuroscience*, 32(11), 3612-3628.

Drugowitsch, J., Wyart, V., Devauchelle, A. D., & Koechlin, E. (2016). Computational precision of mental inference as critical source of human choice suboptimality. *Neuron*, 92(6), 1398-1411.

Edwards, W. (1965). Optimal strategies for seeking information—Models for statistics, choice reaction-times, and human information-processing. *Journal of Mathematical Psychology*, 2, 312–329.

Epley, N., & Gilovich, T. (2006). The anchoring-and-adjustment heuristic: Why the adjustments are insufficient. *Psychological Science*, 17(4), 311-318.

Erev, I., Wallsten, T. S., & Budescu, D. V. (1994). Simultaneous over-and underconfidence: The role of error in judgment processes. *Psychological Review*, 101(3), 519.

Fantino, E., Kulik, J., Stolarz-Fantino, S., & Wright, W. (1997). The conjunction fallacy: A test of averaging hypotheses. *Psychonomic Bulletin & Review*, 4(1), 96-101.

Fehr, E., & Rangel, A. (2011). Neuroeconomic Foundations of Economic Choice--Recent Advances. *Journal of Economic Perspectives*, 25(4), 3-30.

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

Fiedler, K. (1991). Heuristics and biases in theory formation: On the cognitive processes of those concerned with cognitive processes. *Theory & Psychology*, 1(4), 407-430.

- Findling, C., & Wyart, V. (2021). Computation noise in human learning and decision-making: origin, impact, function. *Current Opinion in Behavioral Sciences*, 38, 124-132.
- Fiser, J., Berkes, P., Orbán, G., & Lengyel, M. (2010). Statistically optimal perception and learning: from behavior to neural representations. *Trends in Cognitive Sciences*, 14(3), 119-130.
- Fischhoff, B., Slovic, P., & Lichtenstein, S. (1978). Fault trees: Sensitivity of estimated failure probabilities to problem representation. *Journal of Experimental Psychology: Human Perception and Performance*, 4(2), 330.
- Fisk, J. E., & Pidgeon, N. (1996). Component probabilities and the conjunction fallacy: Resolving signed summation and the low component model in a contingent approach. *Acta Psychologica*, 94(1), 1-20.
- Fleming, S. M., & Lau, H. C. (2014). How to measure metacognition. *Frontiers in Human Neuroscience*, 8, 443.
- Fox, C. R. & Rottenstreich, Y. S. (2006). Between ignorance and truth: Partition dependence and learning in judgment under uncertainty. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 32(6), 1385.
- Fox, C. R., & Rottenstreich, Y. (2003). Partition priming in judgment under uncertainty. *Psychological Science*, 14(3), 195-200.
- Fränken, J. P., Theodoropoulos, N. C., & Bramley, N. R. (2022) Inductive Inference as Adaptive Search using Markov Chain Monte Carlo. PsyArXiv.
- Gardner, M. (1978). Mathematical games. *Scientific American*, 239(2), 18-25.
- Garrett, H. E. (1922). A study of the relation of accuracy to speed. *Archives of Psychology*, 56, 1-105.
- Gelman, A., Stern, H. S., Carlin, J. B., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2013). *Bayesian data analysis*. London, UK: Chapman and Hall.

- Gershman, S. J., & Beck, J. M. (2017). Complex probabilistic inference. *Computational models of brain and behavior*, 453.
- Gershman, S. J., Blei, D. M., & Niv, Y. (2010). Context, learning, and extinction. *Psychological Review*, 117(1), 197.
- Geyer, C. (1991). Markov chain Monte Carlo maximum likelihood. In *Computing science and statistics: Proceedings of 23rd Symposium on the Interface*, Interface Foundation, Fairfax Station, 1991 (pp. 156-163).
- Gilden, D. L. (2001). Cognitive emissions of $1/f$ noise. *Psychological Review*, 108(1), 33.
- Gilden, D. L., Thornton, T., & Mallon, M. W. (1995). $1/f$ noise in human cognition. *Science*, 267(5205), 1837-1839.
- Gold, J. I., & Shadlen, M. N. (2002). Banburismus and the brain: Decoding the relationship between sensory stimuli, decisions, and reward. *Neuron*, 36(2), 299–308.
- Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*.
- Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics*. Oxford, England: Wiley.
- Griffiths, T. L., Vul, E., & Sanborn, A. N. (2012). Bridging levels of analysis for probabilistic models of cognition. *Current Directions in Psychological Science*, 21(4), 263-268.
- Haefner, R. M., Berkes, P., & Fiser, J. (2016). Perceptual decision-making as probabilistic inference by neural sampling. *Neuron*, 90(3), 649-660.
- Hamrick, J. B., Smith, K. A., Griffiths, T. L., & Vul, E. (2015). Think again? The amount of mental simulation tracks uncertainty in the outcome. In D. C. Noelle, R. Dale, A. S. Warlaumont, J. Yoshimi, T. Matlock, C. D. Jennings, & P. P. Maglio (Eds.), *Proceedings of the thirtieth annual conference of the cognitive science society* (pp. 866--871). Austin, TX: Cognitive Science Society.

- Hawkins, G. E., & Heathcote, A. (2021). Racing against the clock: Evidence- based vs. time- based decisions. *Psychological Review*, 128(2), 222–263.
- Henmon, V. A. C. (1911). The relation of the time of a judgment to its accuracy. *Psychological Review*, 18(3), 186.
- Hick, W. E. (1952). On the rate of gain of information. *Quarterly Journal of Experimental Psychology*, 4(1), 11-26.
- Hilbert, M. (2012). Toward a synthesis of cognitive biases: how noisy information processing can bias human decision making. *Psychological Bulletin*, 138(2), 211.
- Hoyer, P. O., & Hyvärinen, A. (2003). Interpreting neural response variability as Monte Carlo sampling of the posterior. In *Advances in neural information processing systems* (pp. 293-300).
- Ibrahim, J. G., Chen, M. H., Gwon, Y., & Chen, F. (2015). The power prior: theory and applications. *Statistics in Medicine*, 34(28), 3724-3749.
- Irwin, F. W., Smith, W. A. S., & Mayfield, J. F. (1956). Tests of two theories of decision in an "expanded judgment" situation. *Journal of Experimental Psychology*, 51(4), 261.
- Jazayeri, M., & Movshon, J. A. (2007). A new perceptual illusion reveals mechanisms of sensory decoding. *Nature*, 446(7138), 912-915.
- Johnson, D. M. (1939). Confidence and speed in the two-category judgment. *Archives of Psychology*, 34, 1–53.
- Juslin, P., & Olsson, H. (1997). Thurstonian and Brunswikian origins of uncertainty in judgment: A sampling model of confidence in sensory discrimination. *Psychological Review*, 104(2), 344.
- Juslin, P., & Persson, M. (2002). PROBABILITIES from EXemplars (PROBEX): A “lazy” algorithm for probabilistic inference from generic knowledge. *Cognitive Science*, 26(5), 563-607.

- Juslin, P., Winman, A., & Hansson, P. (2007). The naïve intuitive statistician: A naïve sampling model of intuitive confidence intervals. *Psychological Review*, 114(3), 678.
- Juslin, P., Winman, A., & Olsson, H. (2003). Calibration, additivity, and source independence of probability judgments in general knowledge and sensory discrimination tasks. *Organizational Behavior and Human Decision Processes*, 92, 34–51.
- Kellen, D., Winiger, S., Dunn, J. C., & Singmann, H. (2021). Testing the foundations of signal detection theory in recognition memory. *Psychological Review*.
- Kepecs, A., & Mainen, Z. F. (2012). A computational framework for the study of confidence in humans and animals. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 367(1594), 1322-1337.
- Kiani, R., Corthell, L. & Shadlen, M. N. (2014). Choice certainty is informed by both evidence and decision time. *Neuron*, 84(6), 1329–1342.
- Körding, K. P., & Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature*, 427(6971), 244-247.
- Krajbich, I., & Rangel, A. (2011). Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. *Proceedings of the National Academy of Sciences*, 108(33), 13852-13857.
- Krajbich, I., Armel, C., & Rangel, A. (2010). Visual fixations and the computation and comparison of value in simple choice. *Nature Neuroscience*, 13(10), 1292-1298.
- Kvam, P. D. (2019). A geometric framework for modeling dynamic decisions among arbitrarily many alternatives. *Journal of Mathematical Psychology*, 91, 14–37.
- Kvam, P. D., & Busemeyer, J. R. (2020). A distributional and dynamic theory of pricing and preference. *Psychological Review*.
- Kvam, P. D., & Turner, B. M. (2021). Reconciling similarity across models of continuous selections. *Psychological Review*.

- Kvam, P. D., Marley, A., & Heathcote, A., (2022). A unified theory of discrete and continuous responding. *Psychological Review*.
- Kwisthout, J., Wareham, T., & Van Rooij, I. (2011). Bayesian intractability is not an ailment that approximation can cure. *Cognitive Science*, 35(5), 779-784.
- Lake, B. M., Ullman, T. D., Tenenbaum, J. B., & Gershman, S. J. (2017). Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40, e253.
- Laming, D. R. J. (1968). *Information theory of choice-reaction times*. New York, NY: Academic Press.
- Levy, R., Reali, F., & Griffiths, T. (2008). Modeling the effects of memory on human online sentence processing with particle filters. In D. Koller, D. Schuurmans, Y. Bengio, & L. Bottou (Eds.), *Advances in neural information processing systems*, 21.
- Li, H. H., & Ma, W. J. (2020). Confidence reports in decision-making with multiple alternatives violate the Bayesian confidence hypothesis. *Nature Communications*, 11(1), 1-11.
- Lieder, F., Griffiths, T. L., Huys, Q. J., & Goodman, N. D. (2018). The anchoring bias reflects rational use of cognitive resources. *Psychonomic Bulletin & Review*, 25(1), 322-349.
- Lieder, F., & Griffiths, T. L. (2020). Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources. *Behavioral and Brain Sciences*, 43.
- Link, S. W., & Heath, R. A. (1975). A sequential theory of psychological discrimination. *Psychometrika*, 40(1), 77-105.
- Linkenkaer-Hansen, K., Nikouline, V. V., Palva, J. M., & Ilmoniemi, R. J. (2001). Long-range temporal correlations and scaling behavior in human brain oscillations. *Journal of Neuroscience*, 21(4), 1370-1377.

- Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review*, 95(4), 492.
- Luce, R. D. (1986). *Response times: Their role in inferring elementary mental organization* (No. 8). Oxford University Press on Demand.
- Luu, L., & Stocker, A. A. (2018). Post-decision biases reveal a self-consistency principle in perceptual inference. *eLife*, 7, e33334.
- Ma, Y. A., Chen, Y., Jin, C., Flammarion, N., & Jordan, M. I. (2019). Sampling can be faster than optimization. *Proceedings of the National Academy of Sciences*, 116(42), 20881-20885.
- Mamassian, P. (2016). Visual confidence. *Annual Review of Vision Science*, 2, 459-481.
- Maniscalco, B., & Lau, H. (2012). A signal detection theoretic approach for estimating metacognitive sensitivity from confidence ratings. *Consciousness and Cognition*, 21(1), 422-430.
- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. San Francisco: W. H. Freeman.
- McCarley, J. S., & Benjamin, A. S. (2013). Bayesian and signal detection models. In *The Oxford Handbook of Cognitive Engineering*.
- Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review*, 85(3), 207.
- Mehle, T., Gettys, C. F., Manning, C., Baca, S., & Fisher, S. (1981). The availability explanation of excessive plausibility assessments. *Acta Psychologica*, 49(2), 127-140.
- Milosavljevic, M., Malmaud, J., Huth, A., Koch, C., & Rangel, A. (2010). The drift diffusion model can account for value-based choice response times under high and low time pressure. *Judgment and Decision Making*, 5(6), 437-449.

- Moran, R. (2015). Optimal decision making in heterogeneous and biased environments. *Psychonomic Bulletin & Review*, 22(1), 38-53.
- Nosofsky, R. M. (1984). Choice, similarity, and the context theory of classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 10(1), 104.
- Nosofsky, R. M. (1986). Attention, similarity, and the identification–categorization relationship. *Journal of Experimental Psychology: General*, 115(1), 39.
- Nosofsky, R. M., & Palmeri, T. J. (1997). An exemplar-based random walk model of speeded classification. *Psychological Review*, 104(2), 266.
- Novikov, E., Novikov, A., Shannahoff-Khalsa, D., Schwartz, B., & Wright, J. (1997). Scale-similar activity in the brain. *Physical Review E*, 56(3), R2387.
- Nilsson, H., Juslin, P., & Winman, A. (2016). Heuristics can produce surprisingly rational probability estimates: Comment on Costello and Watts (2014). *Psychological Review*, 123(1), 103-111.
- Nilsson, H., Winman, A., Juslin, P., & Hansson, G. (2009). Linda is not a bearded lady: configural weighting and adding as the cause of extension errors. *Journal of Experimental Psychology: General*, 138(4), 517.
- Olson, C. L. (1976). Some apparent violations of the representativeness heuristic in human judgment. *Journal of Experimental Psychology: Human Perception and Performance*, 2(4), 599.
- Orbán, G., Berkes, P., Fiser, J., & Lengyel, M. (2016). Neural variability and sampling-based probabilistic representations in the visual cortex. *Neuron*, 92(2), 530-543.
- Pachella, R. G. (1974). The interpretation of reaction time in information-processing research. In B. H. Kantowitz (Ed.), *Human information processing: Tutorials in performance and cognition* (pp. 41–82). Hillsdale, NJ: Erlbaum.

- Peterson, C. R., & Beach, L. R. (1967). Man as an intuitive statistician. *Psychological Bulletin*, 68(1), 29.
- Pleskac, T. J., & Busemeyer, J. R. (2010). Two-stage dynamic signal detection: a theory of choice, decision time, and confidence. *Psychological Review*, 117(3), 864.
- Pouget, A., Drugowitsch, J., & Kepecs, A. (2016). Confidence and certainty: distinct probabilistic quantities for different goals. *Nature Neuroscience*, 19(3), 366-374.
- Proctor, R. W., & Schneider, D. W. (2018). Hick's law for choice reaction time: A review. *Quarterly Journal of Experimental Psychology*, 71(6), 1281-1299.
- Rasmussen, C. E., & Ghahramani, Z. (2002). Bayesian Monte Carlo. In S. Becker, S. Thrun, & K. Obermayer (Eds.), *Advances in neural information processing systems*, 505-512.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85(2), 59.
- Ratcliff, R. (1981). A theory of order relations in perceptual matching. *Psychological Review*, 88, 552-572.
- Ratcliff, R. (2018). Decision making on spatially continuous scales. *Psychological Review*, 125(6), 888.
- Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. *Neural Computation*, 20(4), 873-922.
- Ratcliff, R., & Rouder, J. N. (1998). Modeling response times for two-choice decisions. *Psychological Science*, 9(5), 347-356.
- Ratcliff, R., & Smith, P. (2015). Modeling simple decisions and applications using a diffusion model. In J. R. Busemeyer, Z. Wang, J. T. Townsend, & A. Eidels (Eds.), *The Oxford handbook of computational and mathematical psychology* (pp. 35-62). Oxford University Press.
- Ratcliff, R., & Starns, J. J. (2009). Modeling confidence and response time in recognition memory. *Psychological Review*, 116(1), 59.

- Ratcliff, R., Gomez, P., & McKoon, G. (2004). A diffusion model account of the lexical decision task. *Psychological Review*, 111(1), 159.
- Ratcliff, R., Smith, P. L., Brown, S. D., & McKoon, G. (2016). Diffusion decision model: Current issues and history. *Trends in Cognitive Sciences*, 20(4), 260-281.
- 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.
- Ratcliff, R., Thapar, A., & McKoon, G. (2011). Effects of aging and IQ on item and associative memory. *Journal of Experimental Psychology: General*, 140(3), 464.
- Robert, C. P., & Casella, G. (2004). *Monte Carlo statistical methods* (Vol. 2). New York: Springer.
- Roberts, G. O., & Stramer, O. (2002). Langevin diffusions and Metropolis-Hastings algorithms. *Methodology and Computing in Applied Probability*, 4(4), 337-357.
- Rosky, P. J., Doll, J. D., & Friedman, H. L. (1978). Brownian dynamics as smart Monte Carlo simulation. *The Journal of Chemical Physics*, 69(10), 4628-4633.
- Rotello, C. M. (2018). Signal detection theories of recognition memory. In J. H. Byrne & J. T. Wixted (Eds.), *Learning and memory: A comprehensive reference*, Vol. 2: Cognitive Psychology of Memory (2nd ed., pp. 201– 226). Elsevier.
- Roth, D. (1996). On the hardness of approximate reasoning. *Artificial Intelligence*, 82(1-2), 273-302.
- Rouder, J.N. (1996). Premature sampling in random walks. *Journal of Mathematical Psychology*, 40, 289–296.
- Russo, J. E., & Kolzow, K. J. (1994). Where is the fault in fault trees? *Journal of Experimental Psychology: Human Perception and Performance*, 20(1), 17.

- Sanborn A.N. & Beierholm U.R. (2016). Fast and accurate learning when making discrete numerical estimates. *PLoS Computational Biology* 12(4), e1004859.
- Sanborn, A. N., & Chater, N. (2016). Bayesian brains without probabilities. *Trends in Cognitive Sciences*, 20(12), 883-893.
- Sanborn, A. N., Griffiths, T. L., & Navarro, D. J. (2010). Rational approximations to rational models: alternative algorithms for category learning. *Psychological Review*, 117(4), 1144.
- Sanborn, A. N., Mansinghka, V. K., & Griffiths, T. L. (2013). Reconciling intuitive physics and Newtonian mechanics for colliding objects. *Psychological Review*, 120(2), 411.
- Shekhar, M., & Rahnev, D. (2021a). The nature of metacognitive inefficiency in perceptual decision making. *Psychological Review*, 128(1), 45.
- Shekhar, M., & Rahnev, D. (2021b). Sources of metacognitive inefficiency. *Trends in Cognitive Sciences*, 25(1), 12-23.
- Shepard, R. N. (1987). Toward a universal law of generalization for psychological science. *Science*, 237(4820), 1317-1323.
- Sheu, C. F., & Ratcliff, R. (1995). The application of Fourier deconvolution to reaction time data: A cautionary note. *Psychological Bulletin*, 118(2), 285.
- Sloman, S., Rottenstreich, Y., Wisniewski, E., Hadjichristidis, C., & Fox, C. R. (2004). Typical versus atypical unpacking and superadditive probability judgment. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 30, 573–582.
- Smith, P. L. (2016). Diffusion theory of decision making in continuous report. *Psychological Review*, 123(4), 425.
- Spicer, J., Zhu, J-Q., Chater, N., & Sanborn, A. N. (2022). Perceptual and cognitive judgments show both anchoring and repulsion. *Psychological Science*.

- Starns, J. J., & Ratcliff, R. (2014). Validating the unequal-variance assumption in recognition memory using response time distributions instead of ROC functions: A diffusion model analysis. *Journal of Memory and Language*, 70, 36-52.
- Stengård, E., & Van den Berg, R. (2019). Imperfect Bayesian inference in visual perception. *PLoS Computational Biology*, 15(4), e1006465.
- Stone, M. (1960). Models for choice-reaction time. *Psychometrika*, 25(3), 251-260.
- Strack, F., & Mussweiler, T. (1997). Explaining the enigmatic anchoring effect: Mechanisms of selective accessibility. *Journal of Personality and Social Psychology*, 73(3), 437.
- Sundh, J., Zhu, J.-Q., Chater, N., & Sanborn, A. N. (2021). The mean-variance signature of Bayesian probability judgment. <https://doi.org/10.31234/osf.io/yuhaz>
- Swensson, R. G. (1972). The elusive tradeoff: Speed vs accuracy in visual discrimination tasks. *Perception & Psychophysics*, 12(1), 16-32.
- Tauber, S., Navarro, D. J., Perfors, A., & Steyvers, M. (2017). Bayesian models of cognition revisited: Setting optimality aside and letting data drive psychological theory. *Psychological Review*, 124(4), 410-441.
- Tentori, K., Crupi, V., & Russo, S. (2013). On the determinants of the conjunction fallacy: probability versus inductive confirmation. *Journal of Experimental Psychology: General*, 142(1), 235.
- Teigen, K. H. (1974). Overestimation of subjective probabilities. *Scandinavian Journal of Psychology*, 15(1), 56-62.
- Tian, L., Ellis, K., Kryven, M., & Tenenbaum, J. (2020). Learning abstract structure for drawing by efficient motor program induction. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, & H. Lin (Eds.), *Advances in Neural Information Processing Systems*, 33, 2686-2697.

- Tickle, H., Tsetsos, K., Speekenbrink, M., & Summerfield, C. (2021). Human optional stopping in a heteroscedastic world. *Psychological Review*.
- Townsend, J. T., & Ashby, F. G. (1983). *Stochastic modeling of elementary psychological processes*. CUP Archive.
- Trippas, D., Kellen, D., Singmann, H., Pennycook, G., Koehler, D. J., Fugelsang, J. A., & Dubé, C. (2018). Characterizing belief bias in syllogistic reasoning: A hierarchical-Bayesian meta-analysis of roc data. *Psychonomic Bulletin & Review*, 25, 2141–2174.
- Tversky, A., & Kahneman, D. (1971). Belief in the law of small numbers. *Psychological Bulletin*, 76(2), 105.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157), 1124-1131.
- Tversky, A., & Koehler, D. J. (1994). Support theory: a nonextensional representation of subjective probability. *Psychological Review*, 101(4), 547.
- Usher, M., & McClelland, J. L. (2004). Loss aversion and inhibition in dynamical models of multialternative choice. *Psychological Review*, 111(3), 757.
- Van Orden, G. C., Holden, J. G., & Turvey, M. T. (2003). Self-organization of cognitive performance. *Journal of Experimental Psychology: General*, 132(3), 331.
- Van Rooij, I. (2008). The tractable cognition thesis. *Cognitive Science*, 32(6), 939-984.
- Van Sommers, P. (1984). *Drawing and cognition: Descriptive and experimental studies of graphic production processes*. Cambridge University Press.
- Van Zandt, T. (2000). ROC curves and confidence judgments in recognition memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26(3), 582.
- Varey, C. A., Mellers, B. A., & Birnbaum, M. H. (1990). Judgments of proportions. *Journal of Experimental Psychology: Human Perception and Performance*, 16(3), 613.

- Vickers, D. (1979). *Decision processes in visual perception*. New York, NY: Academic Press.
- Vickers, D., & Packer, J. (1982). Effects of alternating set for speed or accuracy on response time, accuracy and confidence in a unidimensional discrimination task. *Acta Psychologica*, 50(2), 179-197.
- Vickers, D., & Pietsch, A. (2001). Decision making and memory: A critique of Juslin and Olsson's (1997) sampling model of sensory discrimination. *Psychological Review*, 108(4), 789–804.
- Vul, E., Goodman, N., Griffiths, T. L., & Tenenbaum, J. B. (2014). One and done? Optimal decisions from very few samples. *Cognitive Science*, 38(4), 599-637.
- Wagenmakers, E. J., Farrell, S., & Ratcliff, R. (2004). Estimation and interpretation of $1/f^{\alpha}$ noise in human cognition. *Psychonomic Bulletin & Review*, 11(4), 579-615.
- Wald, A. (1949). Statistical decision functions. *The Annals of Mathematical Statistics*, 165-205.
- Wald, A., & Wolfowitz, J. (1948). Optimum character of the sequential probability ratio test. *The Annals of Mathematical Statistics*, 326-339.
- Wickelgren, W. A. (1977). Speed–accuracy tradeoff and information-processing dynamics. *Acta Psychologica*, 41, 67–85.
- Wyart, V., & Koechlin, E. (2016). Choice variability and suboptimality in uncertain environments. *Current Opinion in Behavioral Sciences*, 11, 109-115.
- Yeung, N., & Summerfield, C. (2014). Shared mechanisms for confidence judgements and error detection in human decision making. In *The Cognitive Neuroscience of Metacognition* (pp. 147-167). Springer, Berlin, Heidelberg.
- Yuille, A., & Kersten, D. (2006). Vision as Bayesian inference: analysis by synthesis? *Trends in Cognitive Sciences*, 10(7), 301-308.

- Zamboni, E., Ledgeway, T., McGraw, P. V., & Schluppeck, D. (2016). Do perceptual biases emerge early or late in visual processing? Decision-biases in motion perception. *Proceedings of the Royal Society B: Biological Sciences*, 283(1833), 20160263.
- Zellner, A. (2002). Information processing and Bayesian analysis. *Journal of Econometrics*, 107(1-2), 41-50.
- Zhu, J.-Q., León-Villagrà, P., Chater, N., & Sanborn, A. N. (2022c). Understanding the structure of cognitive noise. *PLOS Computational Biology*, 18(8).
- Zhu, J.-Q., Sanborn, A. N., & Chater, N. (2018). Mental sampling in multimodal representations. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, & R. Garnett (Eds.), *Advances in Neural Information Processing Systems*, Montréal, Canada.
- Zhu, J.-Q., Sanborn, A. N., & Chater, N. (2020). The Bayesian sampler: Generic Bayesian inference causes incoherence in human probability judgments. *Psychological Review*, 127(5), 719-748. <http://dx.doi.org/10.1037/rev0000190>
- Zhu, J.-Q., Sanborn, A., & Chater, N. (2019). Why decisions bias perception: an amortised sequential sampling account. In A. K. Goel, C. M. Seifert, & C. Freksa (Eds.), *Proceedings of Annual Meeting of the Cognitive Science Society*, Montreal, Canada.
- Zhu, J.-Q., Spicer, J., Sanborn, A. N., & Chater, N. (2021). Cognitive variability matches speculative price dynamics. PsyArXiv. <https://doi.org/10.31234/osf.io/gfjvs>
- Zhu, J.-Q., Newall, P. W., Sundh, J., Chater, N., & Sanborn, A. N. (2022b). Clarifying the relationship between coherence and accuracy in probability judgments. *Cognition*, 223, 105022.

Appendix A

Algorithm and Simulation Details

Throughout the paper, we used a consistent set of parameters in simulating the Autocorrelated Bayesian Sampler. All simulations were repeated 2^{14} times. The target posterior distributions (i.e., $P(h|s)$ in Equation 3) for sampling were Gaussian distributions: $N(\mu = 0, \sigma = 1)$ for choice tasks and $N(\mu = \mu_{target}, \sigma = 4)$ for estimation tasks. The stimuli discriminability was manipulated at 5 different values by fully randomizing the decision boundaries between $\{-0.2, 0.2\}$ (for the least discriminable stimuli), $\{-0.5, 0.5\}$, $\{-0.7, 0.7\}$, $\{-1, 1\}$, and $\{-2, 2\}$ (for the most discriminable stimuli). When accuracy was emphasized, we set a larger value for decision threshold $\Delta = 6$, whereas we set $\Delta = 2$ when speed was emphasized. To qualitatively reflect that increasing visual contrasts reduces task difficulty, as shown in Figure 9B, we simulated Gaussian target distributions with means all equal to 0 but standard deviations of 1.5 (contrast 1), 1.3 (contrast 2), and 1 (contrast 3).

We simulated a local sampling algorithm, MC³, which draws samples from the unnormalized posterior (Geyer, 1991). MC³ runs parallel Markov chains at different temperatures and allows individual chains to exchange information. We fixed the number of parallel chains at $C = 6$ and varied the temperature between the chains where the c -th chain has a temperature of $T_c = \frac{1}{1+4(c-1)}$. The higher temperature chains will draw samples from flattened target distributions, thus prefer making long distance moves and assist lower temperature chains to make non-local jumps through the swapping step (see Algorithm for details). This means that the first chain is also the cold chain (temperature of 1) such that its target distribution is not heated up and thus undistorted. The widths of proposal distribution were always in smaller proportions to the width of target distribution, where the ratios were set at 1:100 for choice tasks and 1:1.5 for estimation tasks. The average speed of generating a sample $\lambda = 0.1 \text{ sec}^{-1}$.

Algorithm Metropolis-coupled Markov chain Monte Carlo (MC³)

```

Choose a starting hypothesis  $h_1$ , target unnormalized posterior  $\pi(h) = P^*(h|s)$ 
for  $n=1$  to  $N$                                 ▷ generate  $N$  samples
    for  $c=1$  to  $C$                                 ▷ update all  $C$  chains
        Draw a candidate hypothesis  $h' \sim \mathcal{N}(h_{n-1}^c, \sigma^2)$     ▷ Gaussian proposal
         $A^c = \min\{1, [\frac{\pi(h')}{\pi(h_{n-1}^c)}]^{1/T_c}\}$                 ▷ acceptance probability
        Sample  $u \sim U[0,1]$                                 ▷ acceptance threshold
        if  $u < A^c$  then  $h_n^c = h'$  else  $h_n^c = h_{n-1}^c$  end if    ▷ decide acceptance
    end for
    repeat floor( $C/2$ ) times                                ▷ chain swapping
        Randomly select two chains  $i, j$  without repetition
         $A^{swap} = \min\{1, \frac{\pi(h_n^j)^{1/T_i} \pi(h_n^i)^{1/T_j}}{\pi(h_n^i)^{1/T_i} \pi(h_n^j)^{1/T_j}}\}$     ▷ swapping probability
        Sample  $u \sim U[0,1]$                                 ▷ swapping threshold
        if  $u < A^{swap}$  then swap( $h_n^i, h_n^j$ ) end if        ▷ decide swapping
    end repeat
end for

```

Appendix B**Reinterpretations of Nosofsky & Palmeri's (1997) Exemplar-based Random Walk****Model and Blurton et al.'s (2020) Poisson Random Walk Model**

In this section, we provide detailed reinterpretations of Nosofsky & Palmeri's (1997) EBRW model and Blurton et al.'s (2020) Poisson random walk model. This analysis helps shed light on the relationship between the ABS and these existing computational models. In short, the direct sampling variant of the ABS parallels the model mechanisms of both the EBRW and Poisson random walk models.

Nosofsky & Palmeri's (1997) Exemplar-based Random Walk Model

The EBRW model, a generalization of Logan's instance-based model (Logan, 1988), aims to explain the relationship between categorization and response times by integrating exemplar-based models of categorization and response-time models of evidence accumulation (Nosofsky & Palmeri, 1997). Exemplar models assume that people store many

instances (“exemplars”) of events in memory, and evaluate new events by activating stored exemplars according to their similarity (Medin & Schaffer, 1978; Nosofsky, 1986). This idea has been formalized in the Generalized Context Model (Nosofsky, 1986): individual exemplars are assumed to be represented as points in a multidimensional psychological space where similarity between exemplars decreases with the distance between objects in the space (Shepard, 1987). For a new observation x and a set of K stored exemplars $X^* = \{x_1^*, x_2^*, \dots, x_K^*\}$, all exemplars are activated in proportion to their similarity with the new observation $s(x, x^*)$. The predictive value for the new observation is:

$$\tilde{f}(x) = \frac{\sum_{i=1}^K f(x_i) s(x, x_i)}{\sum_{i=1}^K s(x, x_i)} \quad (\text{S1})$$

where $f(x_i)$ is the information associated with the exemplar x_i . This equation is a general expression for many cognitive applications of the exemplar model. For example, to identify x within the set of exemplars X^* (i.e., $\tilde{f}(x) = p(x^*|x)$), we can set the function $f(x_i) =$

$\begin{cases} 1 & \text{if } x_i = x \\ 0 & \text{otherwise} \end{cases}$; similarly, to judge the probability of x belonging to a category c (i.e., $\tilde{f}(x) =$

$p(c|x)$), $f(x_i) = \begin{cases} 1 & \text{if } c_i = c \\ 0 & \text{otherwise} \end{cases}$.

When the function $f(x_i)$ is an indicator function as it was used in the categorization tasks, the resultant $\tilde{f}(x)$ is a discrete probability distribution within a continuous psychological space. The probability associated with the i -th exemplar equals $\frac{s(x, x_i)}{\sum_{i=1}^K s(x, x_i)}$ if its indicator function takes a value of 1. In this case, the exemplar-based model provides one way to construct the posterior probability of hypotheses by incorporating similarity-weighted information from stored exemplars. EBRW then supposes that people sequentially draw samples according to $\tilde{f}(x)$ until a threshold number of samples in favor of a response is reached. The waiting times between samples is exponentially distributed, as in the ABS, though each waiting time also includes an additive constant which is not currently present in our model. Thus, the EBRW is nearly equivalent to the direct sampling variant of the ABS in

which there is no interdependence between samples of exemplars. What the full ABS can add to this model are autocorrelations between samples and an adaptive prior, which potentially could allow it to account for a wider range of behavior.

Blurton et al.'s (2020) Poisson Random Walk Model

While closely related to the EBRW, Blurton et al.'s (2020) Poisson Random Walk (PRW) model instead focuses on assigning confusable visual stimuli to perceptual categories (e.g., whether a Gabor patch is oriented to the left or right of vertical) and is based on the Theory of Visual Attention (Bundesen, 1990). The TVA's contribution can be captured by Equation S1, with the additional assumption of PRW being that the inter-step time between two successive tentative categories is exponentially distributed with a rate C . Thus, for N samples generated in this way, the response time should follow an Erlang distribution (the sum of N exponentially distributed intervals):

$$f(t|N, C) = \text{Erlang}(N, C) \quad (\text{S2})$$

where, in the binary case, the average speed of generating a tentative category $C = v_A + v_B$

where v_A, v_B represent the strengths of category A, B respectively. In other words, the

probability of a tentative category being A is $p_A = \frac{v_A}{v_A + v_B}$ and that of being B is $p_B = 1 -$

$p_A = \frac{v_B}{v_A + v_B}$. If we consider the category A as the upper bound and the category B as the lower

bound, the random walk of the accumulator has an increment of +1 with probability p_A and -

1 with probability p_B . Coupling the Poisson process of tentative category generation with an

optional stopping rule whose accumulator starts at z ($0 < z < a$) and evidence thresholds

$k_A = a, k_B = 0$, we can derive the cumulative distribution function of first passage time to

state k_B (Equation A5 of Blurton et al., 2020):

$$F_B(t|v_A, v_B, a, z) = \sum_{n=z}^{\infty} F(t|n, v_A + v_B) f_B(n|\frac{v_A}{v_A + v_B}, a, z) \quad (\text{S3})$$

where $F(t|n, v_A + v_B)$ is the cumulative distribution function of the Erlang distribution above. Because of the fixed thresholds $(0, a)$ and starting point (z) , the number of samples (or number of increments) in the random walk to reach one of the thresholds is a random variable. n can be as small as z (for reaching threshold B) or $a - z$ (for reaching threshold A) if all the tentative categories generated in the process are for the same category. But due to the randomness in generating tentative categories, n can also be far greater than those numbers and the probability of n has a closed form solution (Feller, 1968; Blurton et al., 2020 equation A1). Marginalizing over all possible n s produces the CDF of response times.

This basic version of the PRW is mathematically very close to both the decision process of the EBRW (as noted in Blurton et al., 2020) and a restricted version of the ABS that uses direct sampling and no adaptive prior. First, we can interpret p_A, p_B as posterior probabilities of category A and B respectively and they also normalize to 1. Second, the Erlang distribution of response times given N samples is also assumed in the ABS (compare Equation 6 and S2). Third, the random walk process with an increment of ± 1 is equivalent to the heuristic stopping rule used in the ABS where the differences in number of samples is tracked. While p_A, p_B, z are free parameters for the Poisson random walk model, our reinterpretation will constrain their values. p_A, p_B as the posterior probability should be further constrained by the mental representation of stimuli.

Beyond this basic formulation, the PRW also includes extensions to improve its fit to the data and its scope. To produce fast or slow average errors (in comparison to average correct responses), the PRW incorporates, following the approach used in DDMs, trial-by-trial variability in category strengths and starting points. To explain the variability in the leading edge of response time distributions, the PRW allows for processing rates to vary within a trial. The full version of the ABS produces fast or slow errors because of an adaptive prior and autocorrelated sampling respectively. While we did not explore the ABS's

predictions for variability in the leading edge of response time distributions, it would be interesting to see if autocorrelated sampling – effectively varying processing rates within a trial – produces this effect as well, or whether other mechanisms, such as a variable non-decision time, are needed.

The PRW also has been extended beyond two-alternative choice tasks to any number of alternatives by assuming that evidence for one response suppresses evidence for all other responses. The suppression is strong enough that only one counter has non-zero counts at any one time. This extension differs from our proposal in the Discussion of using a max-minus-next stopping rule for multi-alternative choice as there is no inhibition, but future work is needed to explore what contrasting predictions these mechanisms make.

Appendix C

Optimal Stopping Rule is Approximated by the Heuristic Stopping Rule

The optimal solution to terminate the sampling algorithm is to solve the following dynamic programming problem. To illustrate, we demonstrate the optimal stopping rule in binary choice. Let p_{00} denote the prior probability of one alternative being true (i.e., the prior on responses) and p_{ij} the posterior probability after i 0s and j 1s have been accumulated. The posterior probability that the other alternative is true is simply $1 - p_{ij}$. The binary decision task is equivalent to a sequential statistical test on whether $p_{ij} < 1/2$, and the optimal stopping rule for this test can be formally described using dynamic programming as follows (Wald, 1949; Zhu, Sanborn, & Chater, 2019):

$$F(i, j) = \min \{F_0(i, j), c + p_{ij}F(i, j + 1) + (1 - p_{ij})F(i + 1, j)\} \quad (S4)$$

where $F(i, j)$ and $F_0(i, j)$ are respectively expected cost of sampling and expected cost of termination after i samples against and j samples in favour have been observed, and c is the opportunity cost of generating a sample. Suppose that $p_{00} = \text{Beta}(0,0)$, then the posterior of

evidence is also Beta distributed $p_{ij} = \text{Beta}(i, j)$. If the punishment for an incorrect decision is one unit of utility, the expected cost of termination should be the expected cost of incorrectly choosing an alternative when posterior of that alternative is p_{ij} :

$$F_0(i, j) = \min \left\{ \frac{i}{i+j}, \frac{j}{i+j} \right\} \quad (\text{S5})$$

The expected cost of drawing another sample is the sum of (i) the cost of generating one sample, c , (ii) the expected cost if the new sample turns out to be in favor, $p_{ij}F(i, j+1)$, and (iii) the expected cost if the new sample turns out to be against, $(1 - p_{ij})F(i+1, j)$. The sampling process should stop whenever the expected cost of sampling exceeds that of termination: $F(i, j) \geq F_0(i, j)$. We solved the dynamic programming and visualised its stopping rules in Figure S1 by setting the cost of collecting one sample, $c = 0.006$, and the prior probability of evidence to Beta(1,1). While the optimal decision threshold is collapsing over time, it can be approximated by a heuristic max-minus-next stopping rule. Indeed, in multi-alternative settings, it has been shown that implementing the max-minus-next rule closely matches the performance of an optimal stopping rule (Dragalin et al., 1999; 2000).

Figure S1. Decision boundaries obtained from solving the dynamic programming problem.

The sampling algorithm should be terminated once the accumulator reaches the yellow terminating regions. In this illustration, the accumulator started at the top-left corner $\{i=0, j=0\}$ (white triangle) and terminated at the state $\{i=3, j=5\}$ (white circle).

Appendix D

Confidence From Binomial Sample Counts Can Produce Metacognitive Inefficiency

In this section, we demonstrate the problem that when confidence judgments are assumed to be tally of discrete sample counts (especially when the sample size is small), the so-called “metacognitive inefficiency” effect demonstrated by Shekhar & Rahnev (2021a) can also be present even without any loss of information (i.e., noise) in forming confidence judgments. To recap, metacognitive inefficiency describes a relationship between an informativeness measure ($meta_d'/d'$) and the confidence criterion which classifies

confidence judgments as either high or low (Shekhar & Rahnev, 2021a). Both $meta_d'$ and d' are model-based measures of informativeness that crucially assume Gaussian distributions (Maniscalco & Lau, 2012; Fleming & Lau, 2014). More specifically, the inverse Gaussian cumulative distribution function (CDF) is used to calculate $meta_d'$ and d' :

$$meta_d'(c) = \phi^{-1}(\text{hit rate}|c) - \phi^{-1}(\text{false alarm rate}|c)$$

where ϕ^{-1} is the inverse Gaussian CDF and c is the confidence criterion for which individual confidence judgments that are greater (smaller) than c are considered high (low). Note that when $c = 50\%$ (i.e., when choosing either option always receives the same level of confidence), $d' = meta_d'$.

The empirical results uncovered using this approach indicate that as the confidence criterion (c) increases, $meta_d'/d'$ decreases (Shekhar & Rahnev, 2021a). As d' is fixed in this calculation, that means that as c moves further into the upper tails, $meta_d'(c)$ becomes smaller. This result has been used to argue that a certain amount of information about a decision was lost in the process of forming metacognitive judgments and that more extreme confidence judgments are noisier (Shekhar & Rahnev, 2021a; 2021b).

Here, we do not challenge the empirical validity of the effect. Instead, we show that assuming a Gaussian distribution when calculating the informativeness of metacognitive judgment can lead to a decrease in the measured metacognitive efficiency when there is no loss of information but the data-generating distribution is not Gaussian. To illustrate, we present an example when the data is generated by binomial distributions (see Figure S2). In this example, the cumulative Gaussian distributions that respectively intersect with the hit and false alarm distributions (which are cumulative binomial) become closer together as the confidence criterion becomes more extreme: $meta_d'(x) < meta_d'(y)$ and $x > y \geq Np$. The quantitative difference between the two $meta_d'$ of 0.16 is small but noticeable in this

example (see Figure S2). However, this bias becomes smaller when N is large because the data-generating binomial distribution is better approximated by a Gaussian distribution.

Figure S2. An illustration of metacognitive inefficiency arising not from loss of information in confidence judgment, but from incorrectly assuming Gaussian generating distributions. Two confidence criteria were shown (x, y , where $x > y$). Hit rates were calculated based on the binomial distribution $\text{Bin}(N, p)$ (black circles) given its intersection with a confidence criterion, whereas similarly false alarm rates were calculated using the symmetric binomial distribution $\text{Bin}(N, 1 - p)$ (black squares). In this illustrative example, $N = 12$ and $p = 0.8$. Using a Gaussian to compute $meta_d'$ (difference in the horizontal positions of the solid curves) will lead to a decrease in value when the confidence criterion increases: $meta_d'(x) < meta_d'(y)$.

Appendix E

ABS Predictions When Estimates are the Sample Average Instead of the Last Sample

Here we consider an alternative process of making estimates using sample averages instead of taking the last sample as the estimate. As shown in Figure S3, with this new process of producing estimates, the qualitative model behavior of the ABS was reproduced. Using the sample average as the estimate, however, does make the anchoring effect stronger, while repulsion and the degree of autocorrelation become weaker. This is because, compared to taking the last sample, averaging of all samples increased the impact of the earlier samples generated by the autocorrelated sampler.

Figure S3. Using the sample average, rather than the last sample, as the estimate does not qualitatively change model behaviors. **(A)** The model predicts co-occurrence of repulsion and anchoring as shown in Figure 11. **(B)** The model produces $1/f$ noise as shown in Figure 12. The sample size used in the simulations was fixed at 5.

Appendix F

The Process and Time for Producing Probability Judgment

In this section, we provide further justification for the ABS process of probability judgment in which (i) a fixed number of samples is collected and (ii) the waiting time between two consecutive samples are exponentially distributed. The fixed-sample-size rule was assumed both for simplicity and because of the lack of an obvious target for deriving an optimal stopping rule for probability judgments. As shown in the main text, this process should produce Gamma-distributed RTs (i.e., continuous generalization of the Erlang distribution), which are unimodal and positively skewed. In Figure S4, this prediction is confirmed by the RT data collected from three experiments (Zhu et al., 2020; Sundh et al., 2021). The best-fitting shape and scale parameters for experiment 1, 2, and 3 are Gamma(2.36, 3.43), Gamma(2.37, 3.97), Gamma(2.22, 3.68) respectively. The RT distributions resemble each other despite being based on different groups of participants and different probability judgments, suggesting that people may indeed use a fixed-sample-size rule to produce probability judgments.

Alternatively, one could replace the fixed-sample-size rule with an optional stopping rule (e.g., the max-minus-next stopping rule) or a fixed-time stopping rule (e.g., stop after thinking about the task for 10 sec). A max-minus-next stopping rule would predict a mixture of Erlang distributions but would also predict no indifferent probability judgments because sampling continues under this rule until probability judgments are no longer indifferent. This mismatches empirical data as discussed earlier. On the other hand, a fixed-time rule predicts a Gaussian distribution of RTs if we further assume Gaussian noise in the estimation of elapsed time. This prediction should be rejected given the skewness in the RT distributions.

Figure S4. Histograms of response-time data collected from the three experiments of probability judgment shown in the legend. Across the three experiments, RTs for probability judgments are unimodal and positively skewed. Colored lines are best-fitting Gamma distributions. RTs greater than 60 sec were excluded from analysis.