Psychological Review

Grounding Computational Cognitive Models

Casimir J. H. Ludwig, Erik Stuchlý, and Gaurav Malhotra

Online First Publication, January 6, 2025. https://dx.doi.org/10.1037/rev0000533

AMERICAN PSYCHOLOGICAL ASSOCIATION

Psychological Review

ISSN: 0033-295X

https://doi.org/10.1037/rev0000533

THEORETICAL NOTE

Grounding Computational Cognitive Models

Casimir J. H. Ludwig[1], Erik Stuchlý[2], and Gaurav Malhotra[3]
[1] School of Psychological Science, University of Bristol
[2] Institute of Psychology, University of Hamburg
[3] Department of Psychology, State University of New York at Albany

Cognitive scientists and neuroscientists are increasingly deploying computational models to develop testable theories of psychological functions and make quantitative predictions about cognition, brain activity, and behavior. Computational models are used to explain target phenomena such as experimental effects, individual, and/or population differences. They do so by relating these phenomena to the underlying components of the model that map onto distinct cognitive mechanisms. These components make up a "cognitive state space," where different positions correspond to different cognitive states that produce variation in behavior. We examine the rationale and practice of such model-based inferences and argue that model-based explanations typically miss a key ingredient: They fail to explain *why* and *how* agents occupy specific positions in this space. A critical insight is that the agent's position in the state space is not fixed, but that the behavior they produce is the result of a *trajectory*. Therefore, we discuss (a) the constraints that limit movement in the state space; (b) the reasons for moving around at all (i.e., agents' objectives); and (c) the information and cognitive mechanisms that guide these movements. We review existing research practices, from experimental design to the model-based analysis of data, and through simulations we demonstrate some of the inferential pitfalls that arise when we ignore these dynamics. By bringing the agent's perspective into sharp focus, we stand to gain better and more complete explanations of the variation in cognition and behavior over time, between different environmental conditions, and between different populations or individuals.

Keywords: computational model, individual differences, model-based inference, sampling, temporal dynamics

A core aim of cognitive science is to develop testable theories of psychological functions that make quantitative predictions about behavior. To this end, a theory may be cast as a computational model (a formal mathematical model or a computer simulation) that instantiates the psychological mechanisms and processes assumed by the theory. A computational model embodies the core assumptions of a psychological theory, along with auxiliary assumptions that are needed to connect the theory with empirical observations. Such a model may be regarded as a representation of a *target system* (Suárez & Pero, 2019), that is, the "ground truth" model that operates in an agent's brain. Many authors have written about the importance of computational modeling for theory development and testing in psychology (e.g., Farrell & Lewandowsky, 2018; Guest & Martin, 2021; Lee & Wagenmakers, 2014; Oberauer & Lewandowsky, 2019; Roberts & Pashler, 2000; Robinaugh et al., 2021; Simon, 1992; van Rooij & Baggio, 2021). To highlight just a few benefits, models allow us: (a) to generate predictions from a theory and compare different theories against data; (b) to test "proof-of-principle" explanations of an empirical phenomenon; and (c) to identify latent psychological mechanisms and processes that underlie some cognitive capacity (e.g., decision making, object recognition, visual working memory). With advances in computing power and computing software, computational modeling is increasingly widely adopted in psychology, neuroscience and

Hongjing Lu served as action editor.
Casimir J. H. Ludwig https://orcid.org/0000-0003-4656-0751
Erik Stuchlý https://orcid.org/0000-0002-9174-8860
Gaurav Malhotra https://orcid.org/0000-0002-6868-9655
Code for generating the simulation results in Figures 4, 5, and C1 is available from https://github.com/CasLudwig/Grounding-computational-cognitive-models. The author-formatted version of the article is publicly available from https://osf.io/preprints/psyarxiv/vur6t. This work was not preregistered. Elements of this work were presented at the annual meeting of the Society for Mathematical Psychology, Amsterdam 2023. The authors have no conflicts of interest to disclose. Casimir J. H. Ludwig was supported by a Research Fellowship from the Leverhulme Trust (Grant RF-2021-168).

Casimir J. H. Ludwig played a lead role in conceptualization, data curation, formal analysis, funding acquisition, investigation, methodology, project administration, resources, software, supervision, visualization, writing–original draft, and writing–review and editing. Erik Stuchlý played a supporting role in conceptualization, data curation, formal analysis, investigation, methodology, resources, software, writing–original draft, and writing–review and editing. Gaurav Malhotra played a supporting role in conceptualization, investigation, methodology, supervision, writing–original draft, and writing–review and editing.
Correspondence concerning this article should be addressed to Casimir J. H. Ludwig, School of Psychological Science, University of Bristol, 12a Priory Road, Bristol BS8 1TU, United Kingdom. Email: c.ludwig@bristol.ac.uk

psychiatry (Huys et al., 2016; Jarecki et al., 2020; Kriegeskorte & Douglas, 2018; Robinaugh et al., 2021).

Across a wide variety of domains, researchers have taken great strides in constructing models that can explain and predict behavior (and link this behavior with neural mechanisms; e.g., Forstmann et al., 2011; Love, 2015; Turner et al., 2019). These models frequently take the form of *process models* that describe how information flows through the cognitive system (Jarecki et al., 2020), at what Marr (1982), dubbed the "algorithmic level" of analysis. Crucially, most models contain parameters that allow them to be flexible (Roberts & Pashler, 2000): The same model can be used to fit data from different conditions or individuals by tuning these parameters. By fitting parameters to data, we often seek to identify "the underlying mechanism of …" some target phenomenon (e.g., a behavioral effect, cognitive capacity, neural activation). The model then provides a mechanistic explanation of this phenomenon—an explanation in terms of cognitive mechanisms and the parameters that govern their operation (Bechtel & Abrahamsen, 2005, 2010; Kaplan & Craver, 2011; Simon, 1992). This is frequently where the explanation of the target phenomenon ends: We point to a relation between the target and the estimated model parameters, and these model parameters are mapped onto meaningful psychological constructs. What more could we want?

Take, for example, evidence accumulation models of decision making (illustrated in Figure 1 and explained in more detail in the section "Examples of model-based analysis"). These models are used widely to account for reaction time (RT) and accuracy data and decompose these observed data into meaningful, latent psychological variables (i.e., model parameters) such as processing speed and response caution (for reviews, see e.g., Donkin & Brown, 2018; Gold & Shadlen, 2001; Smith & Ratcliff, 2004). An empirical phenomenon might be that older adults are slower to respond than younger adults. Fitting the behavioral data from both groups with an evidence accumulation model suggests that older adults are slower to respond, not because they have a lower processing speed, but because their decision threshold is higher: They need more evidence to respond (e.g., Ratcliff et al., 2004). The increase in the decision threshold here is a mechanistic explanation for the difference in behavior between young and older participants. Such model-based explanations are rife in cognitive science, and we could easily have picked examples from our own back catalog (e.g., Farrell et al., 2010; Ludwig, Butler, et al., 2009; Ludwig, Farrell, et al., 2009; Perez Santangelo et al., 2022).

However, such explanations are incomplete: what is critically missing is *why* and *how* parameters come to take on the values that best account for the data. In other words, the hypothesized mechanisms themselves beget explanation: Why was the decision threshold higher? What prevented older adults from adopting a lower, more appropriate threshold (given the objective of the task)? How did they come to settle on the (high) value of the decision threshold they adopted? Although such questions may be addressed in the Discussion sections, modelers generally consider them beyond the scope of the model (see Starns & Ratcliff, 2010, for an attempt to address such questions for this particular case).

We recognize that the explanatory scope of models has to be restricted for them to be useful. However, why and how agents "settle" on some combination of parameters are important psychological questions in their own right that must be answered if we are to produce better and more complete explanations of cognitive, neural and behavioral phenomena. Addressing these questions involves adopting the agent's perspective and considering the constraints, information, and cognitive mechanisms at play while the agent attempts to navigate a latent "cognitive parameter space" over the course of a task. We argue that this perspective is lacking in many applications of model-based analysis and failing to adopt this perspective has several consequences. First, models and their parameter estimates typically give a static representation of the "average latent state" that gave rise to behavior. This average state may not be representative of the agent's state at any one point in time. Second, models are overly flexible in their predictions, because they do not adequately capture the constraints that the agent operates under. Both these factors can

Figure 1

Relation Between the Ground Truth Model, the Data Generated From the Ground Truth Model, and the Model(s) That Are Fit to These Data

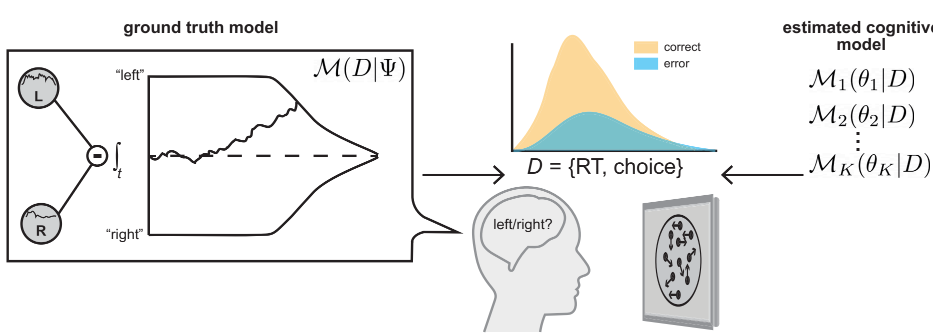

Note. An agent performing the classic random dot motion discrimination task. They are making decisions by accumulating the net evidence in favor of one or the other motion direction toward a decision boundary. In this case, the decision boundaries "collapse" over time. This ground truth model generates data, D, conditional on its parameter values, Ψ. The data in this example consist of RTs and choice outcomes, summarized by RT distributions for correct and error decisions. These data may be fit with a variety of cognitive models, $k = 1, \ldots, K$, each with their own set of parameters, θ_k. The challenge for the modeler is to estimate the parameters of the model, conditioned on the observed data, and to select the model that best approximates the ground truth. RT = reaction time. See the online article for the color version of this figure.

lead to erroneous parameter estimates and model-based inferences. Third, by focusing on this average state, modelers often overlook important and more general questions about the flexibility with which cognition and behavior is adapted, and the mechanisms underlying this flexibility. In this article, we outline a program of *grounding cognitive model parameters* that puts the agent's perspective center stage, with the aim of expanding and improving explanations in the cognitive and brain sciences.

We start by examining the rationale and practice of *model-based analysis*, with examples from a variety of cognitive domains. We then operationalize the research program by identifying three focused research questions that are typically overlooked in such analyses, but that must be addressed in order to understand why and how cognitive model parameters take on the values they do. First, we need to know what constraints are acting on the latent states that generated the empirical data. Second, we need to establish the agent's objectives (such as maximizing speed, accuracy, reward) that drive the change in parameter values. Third, we need to work out what mechanisms and information are available to the agent to achieve their objective(s). Along the way, we also address the analytic problem of how, as cognitive scientists, we can estimate the change in parameter values over time, and the inferential pitfalls of not doing so. The overall argument we pursue here is that grounding cognitive model parameters is critical for understanding variation in cognition and behavior over time, between different environmental conditions and between different populations or individuals.

Model-Based Analysis

To help illustrate the logic of model-based analysis, we consider the broad class of evidence accumulation models of rapid decision making. We will use evidence accumulation models as a running example throughout this article, because these models are a particularly popular choice for model-based analyses. We will give a brief description of this model class (readers familiar with these models can skip the next subsection). We then provide examples of model-based inferences, drawn from evidence accumulation models, but also from a selection of other models from other cognitive domains. These examples illustrate how models and their parameters are used to explain empirical phenomena such as experimental effects, neural signals, and individual and/or population differences. Importantly, these examples underscore that our arguments apply to cognitive models in any domain of psychological science.

Evidence Accumulation Models

Models of rapid choice come in many different flavors (e.g., Brown & Heathcote, 2008; Heathcote & Matzke, 2022; LaBerge, 1962; Link & Heath, 1975; Ratcliff, 1978; Usher & McClelland, 2001; Vickers, 1970), but they all share the basic idea that evidence in favor of the decision alternatives accumulates over time to a critical threshold. There is competition in this accumulation process between the choice options, and the process is corrupted by one or more sources of noise. Once the evidence for one of the alternatives reaches the critical threshold, the motor response associated with the chosen option is initiated after some "nondecision" time. The decision/drift diffusion model (DDM; Ratcliff, 1978; Ratcliff et al., 2016; Ratcliff & Rouder, 1998; Smith & Ratcliff, 2004) for two-choice tasks is probably the most frequently used model from this class. It assumes that the evidence in favor of one option is subtracted from the evidence in favor of the other option. The net evidence then drifts toward an upper or lower decision boundary, which represent the two decision alternatives.

The agent illustrated in Figure 1 makes decisions in a classic motion discrimination task in line with the DDM, in this case augmented with time-varying decision boundaries (Hawkins et al., 2015; Ludwig, 2009; Smith, 2000). Suppose this model is the *ground truth*. The agent monitors the responses of sensory channels tuned to the different motion directions and subtracts the response of, say, the rightward channel from the leftward channel. Positive net evidence then points toward the motion stimulus moving left; negative net evidence points toward the motion stimulus moving right. The decision boundaries correspond to the difference in evidence that would be needed in order to commit to one or the other choice. An error occurs when, due to noise in evidence accumulation process, the integrated evidence hits the incorrect boundary (e.g., hitting the boundary for a rightward response when the pattern is moving left). Better quality sensory evidence (e.g., higher motion coherence) results in a faster drift toward the corresponding boundary. As a result, the decision is made more quickly and more accurately. The probability of making an error can be reduced by increasing the separation between the decision boundaries, but at the cost of an increase in RTs. Other parameters of the model (typically) include the mean rate of evidence accumulation, the internal noise in the evidence accumulation process, the nondecision time, the prior bias of an agent toward one decision and the intertrial variability in the starting position and accumulation process.

Model-Based Inferences

Given the ground truth model, an agent will produce observed behavioral data (in this example: RTs and choices), D, conditioned on a set of parameter values Ψ. As cognitive scientists, we want to know what $\mathcal{M}(D|\Psi)$ is. However, the true model and its parameters cannot be known. We can only try to approximate it by fitting the data with one or several models. For example, many different versions of the DDM have been proposed that include mechanisms such as: between-trial noise in the starting point; between-trial noise in the drift rate (Ratcliff & Rouder, 1998); urgency signal (Cisek et al., 2009); collapsing decision boundaries (Ditterich, 2006; Smith, 2000); attentional biases in the drift rate (Krajbich et al., 2010). This variety of mechanisms, and more generally the variety of available models within the broader class, presents the modeler with a plethora of choices when fitting the observed data (Dutilh et al., 2019). These choices give rise to the K alternative models, $\mathcal{M}_1(\theta_1|D) \ldots \mathcal{M}_k(\theta_k|D)$, on the right-hand side of Figure 1. Note the reverse dependence of the parameters on the data: Our challenge is to estimate the parameter values (for each different model), given some observed data. Where several models are fit to the data, we (somehow) need to select the model that comes closest to the true model (Burnham & Anderson, 2002; Navarro, 2019; Pitt et al., 2002; Shiffrin et al., 2008).

The methods used to estimate model parameters and to select between competing models are highly active research areas in their own right, which we will not deal with in detail here (for general overviews, see Farrell & Lewandowsky, 2018; Lee & Wagenmakers, 2014). Broadly speaking, cognitive modelers seek two types of inferences: inferences through model-selection or through parameter comparison. Inferences through model-selection involve setting up models that represent competing explanations for some target phenomenon, and then selecting between these models through some procedure that trades off the quality of the fit and model complexity (e.g., various information criteria, Bayes factors, cross-validation). Inferences through parameter comparison involve fitting the data with different sets of parameters to account for the target phenomenon. We then adopt some statistical procedure to decide which model parameters relate most strongly to this target.

For example, we may have two explanations for the slowing of RTs with age (e.g., Ratcliff et al., 2004, as discussed in the introduction): Older people are slower to accumulate evidence (lower drift rate) or they may need more evidence before committing to a response (increased decision boundary separation). Inference through model selection would involve fitting the data from both groups with a model in which the drift rate varies between the two groups but the decision boundary remains constant, and a model in which the drift rate is constant but the boundary separation may vary. Inference through parameter comparison involves fitting the data from both groups independently with their own sets of parameters, and then performing statistical tests to assess which parameters differentiate the groups. Note that modelers often adopt both or a mixture of the model-selection and parameter comparison approaches. For instance, they may use both approaches independently as a way to provide converging evidence for a particular mechanistic signature. Alternatively, model-selection may be used to decide on a particular instantiation of the broader model class (e.g., selecting between different flavors of evidence accumulation models, such as the DDM, linear ballistic accumulator, leaky competing accumulator, etc.). The target phenomenon is then explained through a parameter comparison for the selected model instantiation.

In Table 1, we have compiled a (relatively arbitrary) selection of examples of model-based inferences, regardless of whether these inferences were obtained through model-selection or parameter comparisons. These examples come from a variety of models/cognitive domains: decision making, memory, learning, (overt) attention, perception, and categorization.[1] Within each domain, we list model-based explanations for different classes of target phenomena: experimental effects, individual/population differences and (where possible) neural correlates. No doubt these examples are biased by our own knowledge of the field and the reader can probably think of examples from their own area of expertise.

In each of these cases, a model-based analysis decomposes the observed data into a small number of meaningful psychological dimensions, as illustrated in Figure 2. We then look for differences along those dimensions between different experimental conditions, individuals or populations, or we try to identify neural signals that correlate with the variation along these dimensions (van Maanen & Miletić, 2021). For instance, in Figure 2, the data consist of a sequence of RT and accuracy measurements. When we fit a model to these data, such as the DDM, these data are effectively projected into a low dimensional space (downward dashed arrow, estimated cognitive model). For ease of illustration, we have shown just three dimensions, corresponding to key mechanisms that may induce differences in RTs (and accuracy) between individuals and groups (the full model has more than three parameters). Each individual may be represented as a single point within this space. For two hypothetical groups of participants (e.g., young and older people, represented by blue and orange points, respectively), there is some degree of individual variation along all three dimensions. In addition, the two groups differ systematically in their position along the "decision threshold" axis, suggesting that this mechanism specifically is responsible for group differences in the behavioral data. We can now ask why and how different people or groups occupy different locations in this space. In the remainder of this article, we break this problem down into three focused research questions that must be addressed to understand why and how cognitive model parameters take on the value they do.

What Are the Constraints on the Model Parameters?

Consider again the agent illustrated in Figure 1. They perform a perceptual decision-making task using one particular instantiation of the broad class of evidence accumulation models. This system is characterized by a set of parameters, Ψ. The true values for Ψ may be represented as a point in a *cognitive parameter or state space*, where the dimensions of the space represent the parameters of the ground truth model. Figure 2 (bottom-left, ground truth model) illustrates the agent's position in a three-dimensional (sub)space. The data generated by the agent depends on their (ground truth) position in this space (upward solid arrow).

It may be tempting to explain why an agent occupies a particular position in this space by identifying why this position is special in some way. For example, we could justify an observed value of decision threshold by comparing it to the value that maximizes some objective in an experiment (e.g., Balci et al., 2011; Bogacz et al., 2006; Hawkins et al., 2012; Starns & Ratcliff, 2010; Zacksenhouse et al., 2010). Indeed, we discuss the importance of objectives, and individual variation therein, in the next section ("What is the agent's objective?"). For now, suffice to say that agents may attempt to achieve some objective by moving about strategically in the cognitive parameter space.

Almost all cognitive models allow for some degree of strategic control over (a subset of) model parameters. Take each of the model classes represented in Table 1 in turn. *Evidence accumulation models* assume that agents can trade-off accuracy against speed by controlling the decision threshold (Ratcliff & Rouder, 1998; Voss et al., 2004). Parameters such as the drift rate and nondecision time are under much less strategic control, but may still be modulated to some extent through, for example, selective attention and arousal. *Multinomial processing tree models* of memory typically contain at least one or two thresholds for carving up the decision space (Batchelder & Riefer, 1990; Bayen et al., 1996), and these thresholds are to some degree under strategic control (e.g., influenced by expectations in Bayen et al., 2000). *Flexible resource models of visual working memory* (e.g., Bays & Husain, 2008) assume that there is a limited capacity resource that may be configured adaptively according to the demands of the task, which in turn determines the precision with which items are represented. *Reinforcement learning models* feature learning rates that may be adjusted to alter the sensitivity to recent reward feedback and volatility in the environment

[1] Readers may disagree about whether (some of) these examples really constitute *explanatory*, *mechanistic* cognitive models (Jarecki et al., 2020; Oberauer & Lin, 2017). However, we note that these models are often given a mechanistic interpretation or implementation (e.g., "discrete slot" vs. "continuous resource" models of visual working memory give rise to different flavors of mixture models Bays et al., 2009; Van den Berg et al., 2012; W. Zhang & Luck, 2008).

Table 1
Examples of Model-Based Analyses in Multiple Cognitive Domains

Domain	Target phenomenon	Model-based inference
Rapid decision making (evidence accumulation models)	Participants are able to trade-off speed with accuracy based on instructions.	*Decision boundaries* are higher under accuracy instructions (Voss et al., 2004).
	Older adults are slower to respond in a variety of cognitive tasks.	Older adults set *decision boundaries* too high and also have longer *nondecision times* (Ratcliff et al., 2004).
	Individual variation in responsiveness to speed–accuracy instructions.	Flexibility in *decision boundary* adjustment is linked to connectivity between presupplementary motor area and striatum (Forstmann et al., 2010).
Episodic memory (multinomial processing tree models)	Better source memory for information that is consistent with prior expectations.	*Guessing parameter* is biased by expectations about the source derived from schematic knowledge (Bayen et al., 2000).
	Memory deficits in Alzheimer's disease.	Greater deterioration in *immediate retrieval primacy* compared to *immediate retrieval recency* with disease progression (Lee et al., 2020).
	Activity in posterior parietal cortex (PPC) is correlated with item and source memory.	PPC activity linked to item and source memory *guessing biases* (Pergolizzi & Chua, 2016).
Visual working memory (VWM; mixture models)	VWM performance decreases with set size.	The decrease in performance with memory load largely reflects a decrease of the *probability that the item is represented in memory* (W. Zhang & Luck, 2008).
	VWM deteriorates with age.	The *precision of VWM representations* decreases with age; the *probability of reporting a nontarget feature* increases with age (Peich et al., 2013).
	VWM load and delay related activity in a network of frontal, parietal and occipital regions.	Load-dependent increases in *precision variability* is linked to the quality of neural representations in the superior intraparietal sulcus (Galeano Weber et al., 2016).
Learning (reinforcement learning models)	"Positivity bias" in learning: people often learn more from positive compared to negative prediction errors.	*Learning rates* depend on outcome valence, but in opposite ways for factual and counter-factual learning (Palminteri et al., 2017).
	Impaired performance in reversal learning in participants with more obsessive compulsive disorder symptoms.	Higher obsessive–compulsive symptoms linked to increased (subjective) *transition uncertainty* (Fradkin et al., 2020).
	Exploration–exploitation dilemma in reward-based learning.	*Inverse temperature* is related to Locus Coeruleus—Norepinephrine system, suggesting this system is involved in controlling choice strategy (Jepma & Nieuwenhuis, 2011).
Active vision (dynamic eye movement control models)	Oculomotor control is modulated by reading text in different layouts (e.g., normal vs. inverted), suggesting cognitive control over eye movements.	Text layout manipulations affect the *perceptual span* and the *scale of an autonomous timer* mechanism (among others; Rabe et al., 2021).
	Individual differences in scan paths during scene viewing.	Differences in *attentional window size* parameters (among others) mediate variation in saccade amplitudes between individuals (and tasks; Schwetlick et al., 2023).
	Development of reading ability manifests itself in changes in fixation duration and saccade amplitudes.	Primary difference between children and adults lies in the *rate of lexical processing* (Reichle et al., 2013).
Perceptual learning (template matching models)	Orientation discrimination in noise improves with training (thresholds decrease over time).	Learning is mediated by narrowing of *filter tuning* (exclusion of external noise) and suppression of *additive internal noise* (Dosher & Lu, 1999).
	Deficits in (spatial) vision in adults with amblyopia.	Visual performance deficits in amblyopia are mediated by increased *internal noise* and deficient *perceptual templates* (Xu et al., 2006).
	Perceptual learning may be mediated by changes in early sensory encoding or by the changes in the readout of the early sensory code.	Learning results in changes in the *perceptual weighting* of the input and is reflected in anterior cingulate cortex activity, suggesting a higher order, nonsensory locus of perceptual learning (Kahnt et al., 2011).
Categorization (category learning models)	Category learning is more impaired by a concurrent task when learning simple, explicit rules compared to more complex implicit rules that require integration across multiple dimensions.	A deficit in learning explicit rules is best captured by a decline in the *ability to select new rules*; this decline leaves implicit rule learning unaffected (Waldron & Ashby, 2001).
	Intact category learning despite impaired recognition for individual items in amnesic patients.	The *threshold for forming a new representational cluster* is different between amnesic patients and controls (Love & Gureckis, 2007).
	Prototype and exemplar models generate similar predictions for behavioral categorization performance.	Trial-by-trial activation of lateral occipital cortex and posterior parietal cortex correlates with a measure of *representational match* derived from exemplar models (Mack et al., 2013).

Note. The examples were chosen to span a range of model classes, cognitive capacities and empirical phenomena. In particular, for each domain/model class, we selected a basic experimental effect, a population or individual difference, and (where possible) a neural correlate. We have focused generally on relatively well-established models and generally highly cited papers. Verbal labels for model parameters (or model-based metrics) are marked in italics.

Figure 2

Cognitive Modeling Involves a Mapping Between High-Dimensional Data and a Low Dimensional, Psychologically Meaningful Parameter Space

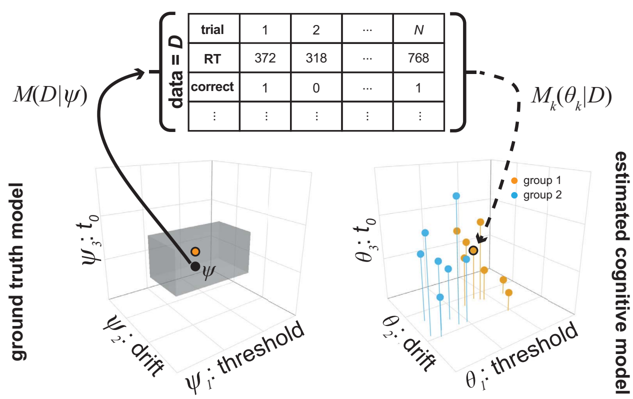

Note. The matrix at the top shows some hypothetical data for a given participant, which consists of a sequence of reaction times and accuracy measurements, but could include other dependent variables (e.g., skin conductance, EEG power in a particular spectral band). These data are generated by a ground truth model (bottom-left) of the type shown in Figure 1. For the sake of illustration, we only show a three-dimensional subspace of parameters: the height of the decision threshold, the mean drift rate, and the nondecision time, t_0. Suppose these ground truth model parameters are subject to constraints that limit the room for maneuver in the cognitive parameter space. These constraints are represented schematically by the shaded cuboid. An agent occupies a point in this space, which results in a noisy sample of observed data (upward solid arrow from the black point; $\mathcal{M}(D|\Psi)$). The modeler takes the observed data and fits one or more models to these data by tuning the free parameters (downward dashed arrow; $\mathcal{M}_k(\theta_k|D)$). The data from a participant (N tuples of the dependent variables) are then mapped to a point in this low dimensional space (estimated cognitive model, bottom right). Different participants will occupy different positions in this space and, in this example, there is also a systematic difference between two groups of participants along the decision threshold axis. The participant who generated the data is highlighted with a black outline (and re-plotted in the ground truth model). Note that in this illustration we assume that the estimated cognitive model is a good approximation of the ground truth model; that is, the two models share these underlying dimensions. Nevertheless, the agent's position in the cognitive parameter space is not recovered exactly in the parameter estimates: The estimated parameter vector (orange) is not aligned with the parameters Ψ in the ground truth model that generated the data. There may be a number of reasons for this discrepancy, such as: (a) The structure of the estimated cognitive model may not align exactly with that of the ground truth model (i.e., the two models may not have exactly the same dimensions of variation); (b) even if the estimated cognitive model matches the (unknown) ground truth model, the data may not be sufficiently diagnostic for identifying all the parameters of the ground truth model accurately (e.g., due to methodological limitations, and the data being a noisy and finite sample from the ground truth model); and (c) there may be error in the estimation procedure and/or parameter trade-offs. RT = reaction time; EEG = electroencephalography. See the online article for the color version of this figure.

(Behrens et al., 2007), as well as parameters that control the degree of exploration (Daw et al., 2006; Wilson et al., 2014). *Models of eye movement control assume* that observers have some control over their perceptual span (the region from which perceptual information is taken to select the next fixation location) and the rate of an internal timer that paces the eye movements (Engbert et al., 2005; Nuthmann et al., 2010). *Perceptual template matching models* typically assume that observers can tune the properties of their perceptual template, such as its shape and its sensitivity (Lu & Dosher, 2008). *Models of category learning* may include a threshold for the abstraction of exemplars to a new category representation (e.g., cluster, prototype or rule; Love & Gureckis, 2007).

However, before we attribute an agent's position in the cognitive parameter space to strategic control, we should consider the various constraints that limit the space that is accessible to the agent. That is, the more constrained the range for a particular parameter is, the less likely it is that the participant settled on a particular value through strategic control. Rather, in that case, the explanation for the agent's position lies in these constraints and they should form part of the model. In this section, we highlight three broad classes of constraints—biophysical, environmental, and cognitive—and then discuss how such constraints may be included in our models.

Constraints on Strategic Control

Biophysical Constraints

Some parameters will be subject to "hard-wired" constraints imposed by physical or physiological limits. For instance, afferent and efferent delays (making up nondecision time in models of evidence accumulation) cannot be reduced below a certain minimum time that is determined by the speed of communication along the neural pathways for sensory signaling and movement production (e.g., Bullier, 2001; Munoz & Wurtz, 1995; Pruszynski et al., 2010). Metabolic costs associated with neural firing and synaptic transmission may limit representational capacity (Attwell & Laughlin, 2001; Levy & Baxter, 1996). Indeed, such costs may contribute to the spatial and/or temporal filtering characteristics of the sensory apparatus, which constrain the fidelity with which an input may be represented (Vincent & Baddeley, 2003). Moreover, noise is inherent in neural processing and is often correlated between neurons, so that it cannot be eliminated completely by spatial or temporal pooling (Averbeck et al., 2006; Shadlen et al., 1996). However, in our modeling, we often ignore these factors. That is, nondecision times are often left to vary over a wide range, sometimes resulting in implausibly large estimates (e.g., over half the RT in Ludwig, 2009). Likewise, the mechanisms that encode the inputs are typically left unspecified and the noise parameters in the model are often used as "scaling" parameters (i.e., fixed to some arbitrary constant to make the model identifiable), or otherwise simply tuned to account for the variability observed in the data.

Environmental Constraints

Properties of the environment may directly constrain some of the model parameters to be within a certain range. One such property is the information available to perform a particular task. Information is often incomplete and uncertain (at least in "large-world" problems; Gigerenzer & Gaissmaier, 2011; Simon, 1955) and, even if all the relevant information is available, there may not be enough time to use all that information. Even in simple experiments, information is often noisy and feedback about performance is often either withheld or stochastic. For example, in the classic perceptual decision-making task illustrated in Figure 1, the stimulus itself is noisy (and sometimes nonstationary; e.g., Holmes et al., 2016; Ludwig & Evens, 2017). The quality of this information affects the rate at which evidence is accumulated: The drift rate for a low coherence random dot motion stimulus is lower than that for a high coherence motion pattern. However, other than constraining the model to respect this ordinal relation, drift rate parameters are typically left free to accommodate variation in behavior introduced by variation in the quality of information.

Cognitive Constraints

The computations that can be performed on the available information in the available time may be limited further by constraints on cognitive capacities (e.g., processing speed, attention, working memory capacity, etc.) and costs associated with using these capacities (Anderson, 1990; Howes et al., 2009; Lieder & Griffiths, 2020; Simon, 1976). For example, in evidence accumulation models of multi-attribute choice, attention cannot be deployed to all different attributes of the different choice options simultaneously (Busemeyer et al., 2019; Busemeyer & Townsend, 1993). Rather, comparisons are made on an attribute-wise basis with attention switching sequentially between attributes and options. This pattern of attending to different attributes at different points in time, produces temporal fluctuation in the drift rate during a choice epoch, even when the environment (i.e., information about the choice options) remains constant. Although we may equip a model with such attention switching mechanisms (Busemeyer & Townsend, 1993; Krajbich et al., 2010), the drift rates associated with different attributes and options, and the temporal profile of attention switches are left free to vary over a wide range in order to account for the observed behavior.

Identifying and Modeling Constraints

A useful way to think about these different constraints is that they limit the room for maneuver in the cognitive parameter space. It is likely that along any one dimension, a (weighted) mixture of constraints limits the range of parameter values that may be adopted (and/or the resolution with which parameters can be varied). Moreover, the range will vary between parameters: Some parameters will be constrained more strongly than others. In Figure 2 (ground truth model, bottom-left), the grey shaded cuboid represents schematically the constraints that limit the possible positions that an agent may adopt. This volume represents the part of the cognitive parameter space that is accessible to the agent: This is the space in which the agent can exert strategic control over the parameters.

When we fit a model to data, it would clearly be helpful to incorporate these constraints in our parameter estimates. Consider what would happen if we ignored the constraints in Figure 2. That is, from the perspective of the modeler, the shaded cuboid does not exist. This absence renders the space of possible parameter estimates wide open, making the model more flexible than it should be. That is, the parameters are free to vary over too wide a range, producing greater freedom in the model predictions (Roberts & Pashler, 2000). As a result, there is a good chance that our parameter estimates fall outside the plausible region given by the properties of the environment, and the cognitive and neural systems involved in the task (e.g., implausibly large estimates of nondecision time in Ludwig, 2009). For example, the estimated parameters shown in Figure 2 (bottom right), display greater variation along the nondecision time axis than the constraints should allow for (indeed, asobserved by Bompas et al., 2024). Such erroneous estimates will have knock-on effects on other parameter estimates: too long a nondecision time implies too short a decision time, which can only be produced by increasing the mean drift rate and/or lowering the decision threshold.

Of course, often modelers *do* impose constraints on parameters when fitting models, for instance, by imposing hierarchical structure in models and priors on the parameters (Lee & Wagenmakers, 2014). However, common practice is to leave such priors relatively vague and letting the data "speak for themselves." Even if priors are more informed, they are typically determined by the modeler's experience with what typical parameter estimates look like for the model in question, rather than mechanistic considerations. For the purpose of grounding cognitive parameters, constraints on parameters should be informed by such mechanistic considerations.

For example, the drift rate is thought to be determined primarily by the input (i.e., environmental constraints) and the transduction of that input into an internal representation (biophysical constraints). Therefore, if we know the nature of the inputs an agent encounters and have a model of the basic (sensory) machinery by which those inputs are represented, we can estimate the drift rate for any one specific input. At least for perceptual tasks, we can use our knowledge of early visual coding in order to turn an input image (sequence) into a time-varying drift rate (Ludwig, 2009; Smith, 1995; Zylberberg et al., 2012). With the emergence of convolutional neural networks, it may be possible to use such a system as a front-end, in order to generate internal visual representations of more complex inputs (e.g., medical images in Holmes et al., 2020). Augmenting the basic evidence accumulation model with a perceptual front-end obviously introduces additional complexity to the model. However, basic perceptual mechanisms have been well-characterized psychophysically and/or physiologically, and deep neural networks are typically trained independently on a different task and different data sets. To the extent that this front-end can be specified independently, these additional mechanisms can actually *reduce* the degrees of freedom of the model, by removing the freedom for the drift rate to take on any value. Note that this approach is not limited to perceptual tasks or, indeed, evidence accumulation models (see Sanders & Nosofsky, 2020; Zou & Bhatia, 2021, for examples of this approach in category learning with naturalistic visual or linguistic inputs, respectively).

To summarize, cognition and behavior are dependent on the position of an agent in a cognitive parameter space. This state space is formed by the dimensions (i.e., parameters) of the ground truth model that the agent brings to bear on a particular task. We have reviewed the different sources of constraints that limit the room for maneuver in this multidimensional space. A (weighted) mixture of constraints will act on each dimension, and this weighted mixture may vary for the different dimensions. As a result, movement along some dimensions will be more constrained than movement along other dimensions. This variation in freedom of movement along different dimensions corresponds to variation in the degree to which parameters may be controlled strategically.[2] Grounding cognitive parameters involves

[2] In this article, we often refer to "strategic control parameters," which gives the impression that some parameters are under strategic control (e.g., decision threshold) and others are not (e.g., nondecision time). This terminology is merely useful shorthand and, as illustrated in Figure 2, this dichotomy is false. Exerting strategic control is a matter of moving about in a multidimensional space. It is possible that an agent simplifies the problem by ignoring those dimensions that are under less control (thereby reducing the dimensionality of the space they have to navigate). Nevertheless, even then the remaining parameters may vary in the degree to which they can be controlled strategically.

identifying and, if possible, modeling these constraints. When these constraints are not adequately incorporated, models are too flexible in that they assume people can "go anywhere" in the parameter space. As a result, we are likely to make incorrect inferences about the agent's position and movement in the state space. Of course, a key question is why an agent needs to move in this space at all: why are they exerting strategic control? In order to address this question, we need to understand what it is that they are trying to achieve, that is, their objective.

What Is the Agent's Objective?

A satisfying answer to the *why* question—why does an agent adopt a particular position in the cognitive parameter space in Figure 2—would be a *rational explanation* of this position (Anderson, 1990; Lieder & Griffiths, 2020; Oaksford & Chater, 2007). Is the position instrumental in achieving some objective? Different positions in this space will generate different behaviors and consequently lead to different outcomes. The utility of a particular outcome will depend on an agent's objective. In a rational analysis a single, normative objective is typically assumed. The question is then how the agent might achieve this objective. Unfortunately, many (if not most) psychological experiments are not designed with the goal of inferring a participant's objective and linking it to the estimated model parameters (see below for examples of exceptions). As a result, objectives are free to vary and there may be a great deal of variability between participants and populations in their chosen objectives. This variability in objectives is a source of individual or population differences in cognition and behavior. In this section, we will discuss the importance of setting a clear task objective, but also the importance of recognizing individual differences in the objectives that are *actually* adopted by agents.

Variable Objectives = Variable Behavior

Figure 3A illustrates schematically the relation between a strategic control parameter Ψ_j and two different objectives. For example, the objectives might be reward rate and accuracy, with the decision threshold as the control parameter. The relation between the height of the decision threshold and reward rate is nonmonotonic: set the decision threshold too low and the agent will make too many errors; set the decision threshold too high and they will spend too long on any one decision when they should be moving on to the next trial (i.e., reward opportunity). In contrast, the relation between decision threshold and accuracy is monotonic, with accuracy improving up to a ceiling as the decision threshold is increased. Maximizing the objective involves varying the relevant strategic control parameter(s) and finding the peak of the function (we discuss this search process in more detail in "What mechanisms and information are available to navigate the cognitive parameter space?"). As such, different (groups of) participants who have adopted different objectives are likely to settle on different values of this control parameter (regardless of whether they succeed in finding the optimal parameter value).

In most cognitive modeling endeavors, the participant's objective is not linked to the parameter estimates and, in any event, the objective in many psychological experiments is either vaguely specified or not at all. Take, for example, the common instruction to "try to respond as quickly and as accurately as possible." It is left for participants to figure out what these instructions mean for them.

Figure 3
Objectives as a Function of a Single Strategic Control Parameter

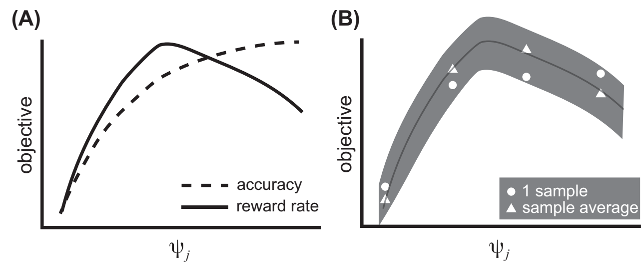

Note. (A) Different objectives may depend on the value of a strategic control parameter in different ways. The solid line shows a nonmonotonic relation between an objective (e.g., reward rate) and ground truth parameter, Ψ_j (e.g., decision threshold). The dashed line shows a different objective (e.g., accuracy). Agents looking to maximize these different objectives will settle on different values of Ψ_j. For the sake of illustration, the curves are drawn schematically and scaled arbitrarily. (B) The agent's estimates of the objective will often be uncertain, illustrated with the error ribbon around the reward rate objective from Panel A. The agent samples the objective at four different values of control parameter Ψ_j. The circles illustrate the noisy objective estimates that might result from taking just a single sample. The triangles illustrate the objective estimates that result from averaging over multiple samples taken at each point in the parameter space. Clearly, integrating the objective estimates over multiple samples results in more accurate estimates, but at a greater temporal cost. The thin black line shows the true (mean) objective.

Some participants may only want to achieve a minimum acceptable level of accuracy and want to leave the lab as quickly as possible (Hawkins et al., 2012). These participants may adopt a low decision threshold so that they respond quickly, but at a cost of a higher error rate. Others may care more about accuracy than response speed and set their threshold higher (e.g., Bohil & Maddox, 2003; Starns & Ratcliff, 2010). Still others may be aiming for a certain level of confidence in their decisions (Lee et al., 2014; Vickers & Lee, 1998) or the maximum reward rate (Bogacz et al., 2006). If participants are left to make up their own objectives, the variation therein will be a source of uncontrolled variance in the data and in any parameter estimates derived from those data.

An obvious solution to counteract this uncontrolled variance is to formulate a *task* objective explicitly, for instance, by specifying a clear and fully transparent incentive structure. This approach is standard in experimental economics (for reviews, see Camerer & Hogarth, 1999; Hertwig & Ortmann, 2001; Houser & McCabe, 2014) and has found some traction in psychology as well. For example, Malhotra et al. (2017) designed a decision-making experiment as a game where participants had a limited time to collect as many reward points as possible and had to make correct decisions to obtain a reward. Therefore, participants were incentivized to balance their speed and accuracy in order to maximize their reward rate (see also Evans et al., 2017; Starns & Ratcliff, 2012). In unspeeded decision paradigms, participants may be incentivized to maximize their accuracy and confidence judgements jointly, through the use of Brier scores (or variations thereof; cf. Bang et al., 2017; Brier, 1950; Yu et al., 2015). In experimental economics, even when incentives do not shift behavior on average (compared to some baseline that involves no or a different kind or level of incentives), they tend to make behavior less variable (Camerer & Hogarth, 1999). Presumably, this reduction in variability is at least partly down to more people adopting the same, experimenter intended objective, and having (learnt) some knowledge of how to aim for that objective given the task environment.

In addition to specifying an explicit objective, participants also need knowledge of how this objective relates to the task environment, such as the number of trials, the penalties associated with incorrect decisions, dependence on other agents' actions, and so forth. These task features may be conveyed through instructions and/or through experience with the task (learning). More generally, people may vary in how they represent the task environment due to variations in background knowledge, cognitive capacities, and how they perceive and respond to feedback from the environment (Szollosi et al., 2023). To the extent to which this variability is left unconstrained, it will result in variation in cognition and behavior.

When a clear task objective has been defined, and the task environment is structured so as to encourage maximizing this objective, agents may nevertheless fail to meet the objective. This failure may have various reasons, including limitations on cognitive and/or temporal resources (Anderson, 1990; Howes et al., 2009; Lieder & Griffiths, 2020; Simon, 1976). If agents replace the specified task objective with their own personal objective, what appears to be suboptimal performance according to the standards defined by the task, may actually be adaptive with regard to the agent's personal objective(s; Rahnev & Denison, 2018). For instance, in reward rate experiments with blocks of a fixed duration, participants may favor accuracy over reward rate (Balci et al., 2011), perhaps because it is easier to track performance accuracy than it is to estimate reward rate. As a result, they end up setting their decision thresholds too high with respect to the defined task objective (e.g., Balci et al., 2011; Bogacz et al., 2010; Bohil & Maddox, 2003; Starns & Ratcliff, 2010). Alternatively, agents may adopt *robust* strategies that incorporate their uncertainty about the variables that determine the task objective (e.g., maximize the minimum reward rate, given uncertainty about intertrial intervals; Zacksenhouse et al., 2010). Therefore, the *actual*, variable objectives adopted by participants, and their relation to estimated model parameters, should be investigated.

Investigating Variable Objectives

We consider three possible approaches to investigate such individual (or population) differences in objectives: (a) assess which objective function best explains behavior; (b) elicit subjective reports; and (c) take independent measurements of task-related variables. First, we suggest that researchers run model comparisons to assess how different objective functions relate to the strategic control parameters under investigation. When two populations or individuals differ in estimated parameter values (e.g., as illustrated in Figure 2), this analysis can help identify whether a difference in objectives can account for observed variability in parameter estimates. For example, in experiments where participants are asked to maximize reward rate, we might find that a subset of participants are unable to find the optimal value for the decision threshold and set their threshold too high (Bogacz et al., 2010; Starns & Ratcliff, 2012; Zacksenhouse et al., 2010). We may then assess whether the data from these "suboptimal" participants are better modeled by assuming that they were aiming for a different objective, such as maximizing accuracy (see Figure 3A; Balci et al., 2011), minimizing time spent in the lab without making too many errors (Hawkins et al., 2012), or maximizing minimum reward rate (Zacksenhouse et al., 2010). This approach may be generalized to settings where there is no specific task objective, in which case the variability in adopted objectives is likely to be greater.

Our second suggestion is to investigate objectives more directly. For instance, where possible, we can ask participants to report their objectives or elicit ratings of how much they care about different possible objectives. Even in simple choice experiments of the type discussed throughout this article, we might elicit ratings of how much participants cared about speed, accuracy, reward and so forth. Where objectives are less accessible for subjective report, we may be able to infer them by giving participants a choice between different tasks that embody different objectives. For example, Kool and Botvinick (2014) let participants choose (and switch) between a high-effort, high-reward task and a low-effort, low-reward task. Presumably, participants who spent more time pursuing the high-effort task cared more about the reward than participants who chose the low-effort task, who likely preferred to save energy. It is also plausible that this type of variability is linked to stable personality traits such as "need for cognition" (Cacioppo et al., 1996; Gheza et al., 2023; Sandra & Otto, 2018). Either way, a critical step is to relate the individual differences in subjective reports or inferred goals to the variability in estimated parameters.

Finally, differences in estimated parameters between participants may not stem from a difference in their objectives, but from their ability to estimate these objectives. Therefore, our third suggestion is to take independent measurements of task-related capacities that are involved in estimating the objective. For example, for participants to estimate an objective like reward rate, they need to be able to keep track of the amount of time that has elapsed (Zacksenhouse et al., 2010). Measuring individual differences in timing ability may then offer insights into why participants do not adopt optimal decision thresholds (Zacksenhouse et al., 2010) or collapsing boundaries (Miletić & van Maanen, 2019). Again, the variability in these task-related variables needs to be linked to the variability in the estimated parameters. We have not encountered many examples of this last, critical step and we believe this is fertile territory for a program of grounding cognitive model parameters.

In summary, task objectives are not fully specified in most psychological studies and participants can bring their own objectives to bear on the task. This variation in objectives can lead people to different positions in the cognitive parameter space, resulting in variation in behavior. Understanding these individual and/or population differences is key to understanding why cognitive model parameters take on the values they do. However, even when different agents have adopted the same objective, they may still end up in different locations. To understand this source of variation, we need to examine the information and cognitive mechanisms available to agents to navigate the cognitive parameter space. Therefore, we now take a closer look at the dynamics of strategic control.

What Mechanisms and Information Are Available to Navigate the Cognitive Parameter Space?

Suppose a participant has adopted a certain objective—either the one prescribed by the experimenter or one they have defined for themselves. They now have to try to achieve this objective by varying the strategic control parameters of the ground truth model. Understanding *how* they vary these control parameters is a critical ingredient of a more complete explanation of the participant's behavior. In essence, this is "simply" a standard optimization problem: The agent has to move around in the cognitive parameter space in order to find the peak of the objective function.[3] There are many algorithms available to solve such problems, such as gradient descent, the Nelder–Mead simplex method, genetic algorithms, simulated annealing and so forth. As such, a natural hypothesis is that human participants try to approximate this kind of algorithm (Busemeyer & Myung, 1992), either during performance of some specific task or over the course of their lifetime (or, indeed, across generations over the course of evolution). Although few articles make such a claim explicitly, the assumption that participants perform some kind of optimization over the short timescales of a specific task is often implicit (including in our own work; e.g., Malhotra et al., 2017). In this section, we examine this idea critically. In particular, we highlight how this problem is much more complex for an agent than it appears at first sight, due to limitations in the cognitive resources and information available for a task. This complexity is likely to bear on the mechanism(s) adopted by agents to navigate the parameter space in many realistic environments.

Uncertainty in Objective Estimates

To appreciate the difficulty of the problem faced by the agent, consider again the simple, one-dimensional objective function from Figure 3A. This objective function is set by the environment (or the experimenter in a cognitive task), and the agent is unlikely to have much prior knowledge of global shape of the objective function that they can use to guide their task performance. The best they can probably do is to *sample* the objective at various different points along the strategic control parameter axis (x-axis in Figure 3A) and experience *local* feedback about the objective value at those points. In the absence of prior knowledge or a rich internal model of the global function, the agent cannot sample the objective through mental simulation, but actually has to interact with the environment. That is, for any given value of the control parameter, the agent generates an action, observes the outcome and updates their estimate of the objective function at that location. They then choose another point along the axis, observe another outcome and form another local estimate. The goal of the agent is to estimate simultaneously the underlying objective function and to move to a point along this axis that maximizes this objective function.

In many realistic situations, estimating even the local value of the objective function is not straightforward because information gained about the objective is likely to be uncertain (Lee et al., 2014; Mikhael & Bogacz, 2016). We distinguish between two sources of uncertainty. *Environmental uncertainty* is induced when the environment provides noisy, delayed, or no feedback. *Representational uncertainty* refers to the internal representation of the feedback and/or the internal (mis)estimation of the objective. Clearly, both environmental and representational uncertainty compound each other. Therefore, in what follows we assume that estimates of the objective function are uncertain, without specifying the source of that uncertainty.

We illustrate this uncertainty in our estimates of the objective by the grey ribbon around the objective in Figure 3B. The ribbon shows the scale of a distribution from which the objective estimates are sampled when varying strategic control parameter Ψ_j. Suppose the agent samples the objective at four parameter values. For each value, they generate an action (e.g., response in an experimental trial), observe the outcome and estimate the objective. They then select the location with the highest objective estimate. The circles are possible draws from the distribution when the agent just takes a single sample of the objective. Clearly, these samples are quite noisy: The highest estimate is given by the right-most parameter value, even though this location is suboptimal if the agent's objective is to maximize reward rate.

If an agent wants to gain more precise estimates of the objective, they must take several actions at each value of the strategic control parameter and integrate multiple samples of the objective at each location (e.g., compute the average objective value). As a result, the estimates will be closer to the true values (illustrated by the triangles). If the agent now picked the parameter value with the highest estimated objective value (the third point from the left), they would get quite close to the peak of the function. However, estimating the objective over a larger number of samples comes with an obvious temporal cost of generating repeated actions at each control parameter value. If time (or the total number of trials in an experiment) is limited, integrating samples over multiple actions (i.e., trials in an experiment) limits the number of different control parameter values that may be selected and the extent to which the cognitive state space may be explored.

Costs of Exploring the Cognitive Parameter Space

The temporal cost associated with sampling the objective (i.e., information gathering) creates an exploration–exploitation dilemma (Cohen et al., 2007; Sutton & Barto, 1998). That is, any time spent *exploring* a (potentially) low-value region of the cognitive parameter space is time that could have been spent *exploiting* a previously sampled location with a higher objective value—an opportunity cost. A rational agent would explore if the cost of doing so is outweighed by the long-run benefit of being able to find and exploit a region where the objective value is higher. In other words, the utility of exploration depends on the available time-horizon (Wilson et al., 2014). However, in most psychological experiments, this time horizon is relatively brief. As a result, exploration of the objective may be limited to only a small set of strategic control parameter values, rather than an exhaustive search for the peak.

In addition to this opportunity cost, sampling the objective also comes with cognitive costs. Take the simple strategy described in relation to Figure 3B: The agent samples a small number of points in the cognitive parameter space and then settles for the spot at which the objective estimate was the highest. At the very least, this agent would need to remember the location of the best objective estimate so far and the estimated value itself. They then need to update these when a better value is sampled in order to be able to return to and exploit this location once they have taken a sufficient number of samples. Moreover, they need decision mechanisms for selecting the strategic control parameters to sample and the number of samples to take. Therefore, even quite a "simplistic" sampling algorithm already requires some (working) memory resources and decision mechanisms.

[3] This function is also often referred to as a "cost" or "loss" function (e.g., in statistics or economics) or "fitness" function (e.g., in theoretical biology). Note that where the objective is framed as a cost or loss, the function will need to be minimized, but flipping of the objective makes no difference to our arguments.

Note that these demands come on top of the cognitive demands imposed by the primary task itself.

The opportunity and cognitive costs associated with sampling the objective are amplified when there are multiple strategic control parameters, as in the three-dimensional cognitive parameter space shown in Figure 2. Even if each parameter is sampled at only a few different values, their combined number scales nonlinearly with the number of dimensions ("the curse of dimensionality"; Bellman, 1957). The question is how participants navigate this multidimensional space in pursuit of their objective, considering that they can only acquire noisy, local information about the objective. Addressing this question is core to the grounding of cognitive model parameters. We outline several ways in which agents might respond to this challenge.

First, as noted in the previous section, agents may choose an objective that is easier to achieve and/or estimate. For example, instead of maximizing reward rate, participants may choose to maximize accuracy (Balci et al., 2011) because it may be easier to track accuracy. We have already argued for the importance of probing the objectives that are actually adopted by different agents, recognizing that they may differ from the one set by the experimenter.

Second, given an objective, agents may control only one or two key parameters that have the biggest effect on this objective, thereby reducing the dimensionality of the cognitive parameter space to be searched through. For example, in an effort to maximize reward rate, participants may only modulate their decision threshold and not bother trying to control the more constrained drift rate and nondecision time parameters. As a result, they need not navigate a three-dimensional state space (Figure 2), but only need to move along one axis (e.g., Ψ_1 in Figure 2).

Third, instead of searching for the parameters that yield the maximum objective value, the agent may aim for an objective value that is "good enough"—satisficing instead of maximizing (Simon, 1955). All these approaches can potentially simplify the search problem, but they do not eliminate it. Even if an agent has adopted an easy-to-estimate objective and controls only a small set of parameters in order to achieve a satisfactory value for this objective, they still need to adjust these control parameters in order to find this minimum satisfactory level. Therefore, any attempt to understand why participants adopted a particular set of parameters also needs to account for the way these parameters values have been arrived at.

Sampling in Cognitive Parameter Space

Given the constraints outlined, highly sophisticated forms of statistical optimization (e.g., that rely on derivatives or function learning) are cognitively implausible.[4] In this section, we argue that navigating the cognitive parameter space under these constraints is a problem that is well-suited to a popular class of sampling algorithms for approximate probabilistic inference. After discussing the advantages of this framework, we will demonstrate the behavior of an agent who adopts such an algorithm for getting around the cognitive parameter space.

In statistics and machine learning, many sampling algorithms have been developed to allow for approximate inference in complex, multidimensional search (or hypothesis) spaces, where inference would be intractable otherwise ("Monte Carlo" algorithms such as importance sampling, Metropolis–Hastings, Gibbs sampling, and sequential methods such as particle filtering; Andrieu et al., 2003; Doucet et al., 2001; MacKay, 2003). Consider Markov Chain Monte Carlo (MCMC) techniques that are widely used for parameter estimation in statistical and, indeed, cognitive modeling (Andrieu et al., 2003; Van Ravenzwaaij et al., 2018). Broadly speaking, these algorithms involve a set of "particles" in the (multidimensional) search space, with each particle representing a parameter vector for which the objective value is computed (in modeling applications, the objective is typically the likelihood or posterior density of the parameters, given some data). The algorithm performs a form of local, stochastic, and greedy search where, for each particle, a *proposal* value is drawn randomly from a transition distribution centered on the current position. If this proposed position improves the objective value, the proposal is *accepted* and the particle moves to the proposed location; otherwise, the proposal is likely to be *rejected*, in which case the particle does not move. In this way, a *chain* of samples is produced that eventually will be concentrated in a high-value region.

In the past 2 decades, various authors have proposed that such sampling algorithms may also describe the way humans explore complex, open-ended, internal hypothesis spaces in a wide variety of cognitive domains, such as (among others): categorization (Sanborn et al., 2010), causal learning (Bonawitz, Denison, Gopnik, & Griffiths, 2014; Bramley & Xu, 2023; Bramley et al., 2017), probabilistic inference (Dasgupta et al., 2017, 2018), concept learning (Goodman et al., 2008; Ullman et al., 2012; Zhao et al., 2024), perception (Bill et al., 2022; Gershman et al., 2012; Haefner et al., 2016), judgment, and decision making (Lieder et al., 2018). Sampling algorithms are often presented as a cognitively plausible way for the brain to approximate rational (Bayesian) inference with lower computational overheads compared to the full normative solution (for reviews, see Bonawitz, Denison, Griffiths, & Gopnik, 2014; Bramley et al., 2023; Chater et al., 2020; Fiser et al., 2010; Sanborn, 2017; Shi et al., 2010; Suchow et al., 2017; Vul et al., 2014).

Aside from their rational appeal, sampling algorithms have several features that make them attractive and plausible candidates for how an agent might navigate the cognitive parameter space under the constraints discussed previously. First, they do not require (much) prior knowledge of or assumptions about the global objective function. For sure, it is easier for these algorithms to find high-value regions when the underlying objective is smooth and parametrically simple to describe, but the algorithm does not need to "know" this in order to reach these regions. Second, MCMC algorithms search locally by generating samples randomly around the current position of the particles. Local search is consistent with the "stickiness" with which humans update their beliefs in probabilistic inference (Dasgupta et al., 2017) or causal learning (Bramley et al., 2017). In this way, MCMC embodies a very natural trade-off between exploration and exploitation. If the algorithm has found a high-value region, samples tend to be concentrated in that region (exploitation), although there is some probability of accepting a move to a lower value region and, via that route, potentially finding a different mode (exploration). Third, the

[4] We do not claim that function learning itself is implausible; clearly it is not (e.g., Brehmer, 1976; Kalish et al., 2004; Lucas et al., 2015). Indeed, people can exploit spatial correlations in complex multimodal, multidimensional functions in order to find high-value regions, with their behavior being well described as a form of Gaussian process regression (Wu et al., 2018). The key question here is whether these (effortful) cognitive capacities are directed to the "metatask" of finding good strategic control parameters that govern performance in a cognitively demanding primary task.

memory requirements of these algorithms are very limited. A basic MCMC sampler only needs to be able to compare a proposal sample with the current value. It does not use information from other particles or past iterations in order to extract information about the global function (of course, more sophisticated MCMC samplers do use this kind of knowledge and perform better as a result; e.g., Differential Evolution MCMC; Heathcote et al., 2019; ter Braak & Vrugt, 2008). As such, an agent navigating cognitive parameter space may be conceptualized as a single particle (chain) moving around the state space through local search (Bramley et al., 2017; Vul et al., 2014).

In typical MCMC applications, proposals are evaluated without actually moving the particles (indeed, any move is contingent on the proposal being accepted after comparison with the current objective value). However, we suggest that human agents frequently do not have that luxury: They may not be able to simulate mentally what will happen when moving to a different location in the state space. Therefore, they can only evaluate a "proposal" location after actually moving there, interacting with the environment and getting feedback. Note the contrast with how sampling is often conceived of in other domains (e.g., judgment and decision making; Lieder et al., 2018; Stewart et al., 2006). In such scenarios, samples may be drawn mentally before committing to an overt response (even then, sampling takes time and incurs a cost; Lieder et al., 2018). Moreover, as noted earlier, a single interaction at a given location in the state space will often not be sufficient, due to uncertainty in the objective estimates. Depending on the level of uncertainty, the agent may have to remain in one place for some time and integrate feedback from several interactions (e.g., trials).[5]

Figure 4 illustrates an agent who uses a sampling approach to navigate the cognitive parameter space, looking for the peak of an objective function under the constraints discussed above. Panel A shows a two-dimensional cognitive parameter space, with the true objective value (unknown to the agent) shown as the color temperature for each combination of strategic control parameters (Ψ_1, Ψ_2). This particular objective function is arbitrary and chosen for illustration only. The peak of the objective is shown by the plus sign. Panel B shows an example search trajectory of an agent who can only sample noisy estimates of the underlying objective function and who has very limited memory of their search history in this space. Specifically, at any one point in the space, the agent samples the objective function over a small window of trials and averages the estimates from that window (as described in the context of Figure 3B). If the average objective is better than that from the previously visited location, then the search continues from here; otherwise, the agent returns to the previous location. The algorithm is described in more detail in the figure caption and in Appendix A. We assume the agent has some (coarse) information about the objective value scale and is able to perform at least an ordinal comparison between the current location and the immediately preceding one (i.e., the memory load is minimal). In the trajectory illustrated in Figure 4B, the agent starts in a relatively low-value location, explores some other low-value regions, but ultimately ends up in a good spot, quite close to the peak objective value.

We do not claim that this algorithm is *the* algorithm that agents use to navigate the cognitive parameter space. It is entirely possible that the agent uses some other sampling algorithm or heuristic strategy. Indeed, the precise nature of the sampling algorithm should be an empirical question of interest. Sampling algorithms, like the one discussed here and others developed in studies of perception and cognition, provide a general framework for exploring this question and can accommodate a variety of strategies (Bramley et al., 2017). Our aim was to demonstrate the utility of the sampling approach as a cognitively minimal search algorithm. A secondary aim was to use this algorithm to generate trajectories in the cognitive parameter space, so that we could assess the consequences of these dynamics in the context of model-based analysis. We now turn to these consequences.

Estimating the Trajectory in Cognitive Parameter Space

At this point, it should be clear that in many psychological paradigms, there is no single point in the cognitive parameter space that is responsible for the observed data. Rather, as the agent tries to achieve an objective, they will adjust the cognitive parameters that are under strategic control, thereby moving around in this state space, as illustrated in Figure 4B. In other words, the ground truth parameters change over time. In turn, this change produces behavior (i.e., data) that changes over time: $\mathcal{M}(D(t)|\psi(t))$. We now switch to the perspective of the cognitive scientist who, through model-based analysis, is looking to infer the latent mechanisms that gave rise to the observed behavior.

Static Models Are Often Misspecified

For the sake of illustration, consider a cognitive parameter space defined by the mean and standard deviation of a Gaussian distribution (we deliberately frame this example in a generic, abstract format; we present a cognitively more realistic example in the following subsection "Capturing Parameter Dynamics"). Figure 4C shows a time series of data, generated by the agent's movements illustrated in Panel B, with each data point $y_t \sim \mathcal{N}(\mu_t, \sigma_t)$. The agent initially generates observations from a tight distribution centered around a value of −2, but ends in a spot where they generate much more variable data with a positive mean (hence the upward trend and "fanning out" of the data). The modeler only has access to these data.

Standard practice in model-based analysis is to treat the data-generating process as stationary and to estimate a single set of parameters of a cognitive model, θ, from either a set of summary statistics derived from the data or the time-unordered distribution of data. Figure 4C shows the marginal distribution of the data on the right-hand side. Suppose the modeler fits a Gaussian distribution to these data. The posterior estimates for the parameters are shown in Figure 4D as the violin plots on the right (see Appendix B for further details). The means of these distributions are represented by the filled triangle in Panel B. This simple demonstration highlights several important points. First, note that the functional form of the estimated model actually matches the ground truth *at any one point*. However, it is clear that this model will not be a good fit when applied to the overall distribution: The data are generated by a mixture of Gaussians, with weights proportional to the number of trials spent in each location of the parameter space. As a result, the observed distribution (Figure 4C) has fatter tails, rightward skew

[5] This situation is more analogous to a modeler who can only generate model predictions through simulation and does not have access to an analytic likelihood. A standard solution in "approximate Bayesian computation" is to estimate a synthetic likelihood, where multiple noisy estimates of the likelihood are averaged for a given set of parameters (Hartig et al., 2011; Palestro et al., 2018).

Figure 4

Navigating a 2D Cognitive Parameter Space With Noisy Objective Estimates

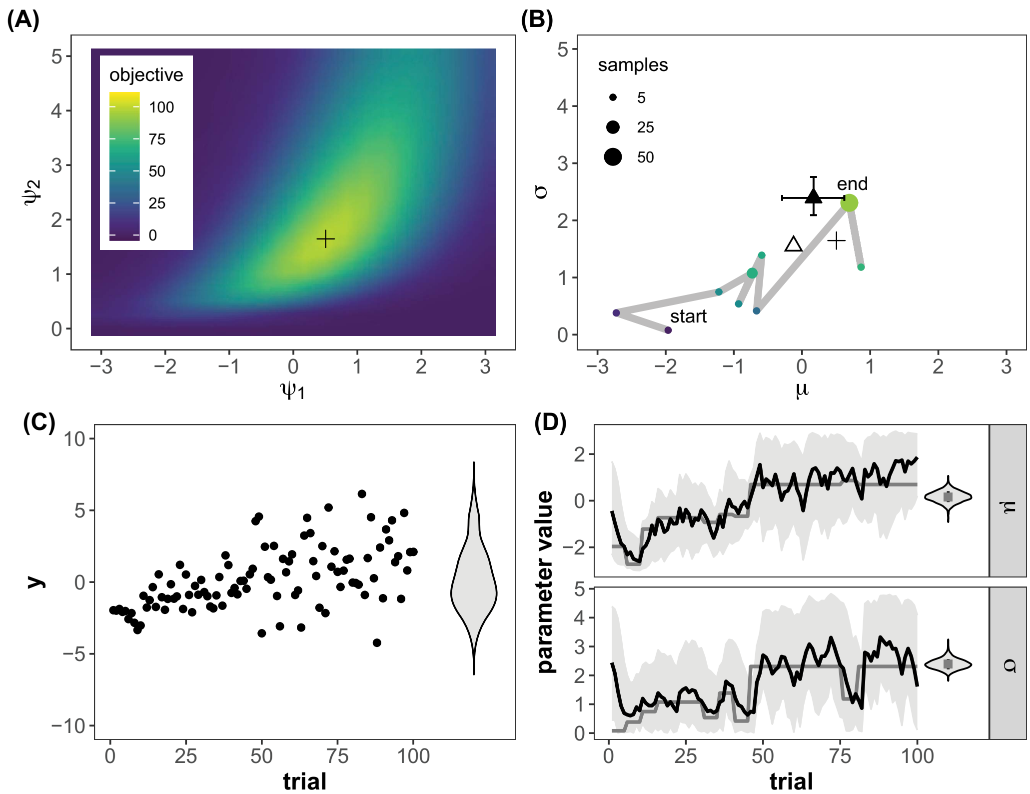

Note. (A) An arbitrary 2D objective function. The peak of the function is marked with a +. The function corresponds to the mean objective value at a given location (Ψ_1, Ψ_2). An agent sampling a given location will obtain a noisy estimate of the objective around this mean. (B) An example trajectory from an agent following a simple sampling algorithm over the course of *N* trials. For the sake of illustration, the parameter space is defined by the parameters of a Gaussian distribution. At each location, the agent generates a data point drawn from $\mathcal{N}(\mu(t), \sigma(t))$ and obtains a noisy objective estimate around the mean value shown in Panel A. These estimates are averaged over a small window of *W* trials. If the agent jumps, the value at the new location is compared to the value at the previous location and the new location is "accepted" if the objective estimate is better. The probability of a jump, π is inversely related to the estimated objective value, $\bar{v}$, through: $\pi(\bar{v}) = 1 - \frac{\bar{v}}{\gamma}, 0 \leq \pi \leq 1$, where γ is the threshold objective value at which the jump probability drops to 0 (ensuring search terminates when the estimated objective reaches some satisfactory value). The parameters of this simulation were: $N = 100$; $W = 5$; $\sigma_{\text{obj}} = 25$ (standard deviation around the objective); $\sigma_{\text{jump}} = 1$ (standard deviation in 2D of a bivariate Gaussian jump distribution); $\gamma = 100$. The true maximum of the objective is shown by the plus sign. A weighted average of all the locations visited is shown by the open triangle. The filled triangle (with 95% credible intervals) shows the posterior parameter estimates based on the fit of a Gaussian distribution to the data generated by the agent. The color temperature of the visited locations indicates the true, underlying mean objective value (not the noisy estimate) on the same color scale as the objective in Panel A. (C) The time series of observations generated by the agent, along with the marginal distribution. (D) Posterior parameter estimates for μ and σ from a particle filter (solid black line; ribbon shows the 95% credible interval). The true parameter values are shown by the dark grey solid line underneath. The posterior densities from the static model fit are shown on the right-hand side in each panel (the posterior means, shown by the dark grey squares, correspond to the filled triangle in Panel B). The vertical black line inside the densities represents the interquartile range (which is hard to see due to the tight posteriors). 2D = two-dimensional. See the online article for the color version of this figure.

and even a hint of bimodality (in the right tail). In other words, exploration of the cognitive parameter space can give rise to data that appears to come from a different functional form. Second, the overall parameter estimates reflect this mixture, in that they are some kind of average of all the locations visited by the agent. However, the posterior mean or modal parameters were not actually adopted at any one point (i.e., the filled triangle in Panel B does not coincide with any of the visited locations). In this case, the estimated parameters come reasonably close to the ones that were adopted most of the time by the agent (in half the trials), but of course that need not be the case if the agent explored more extensively. Third, it might be argued that Bayesian parameter estimates will reflect the temporal structure in the data implicitly in their posterior densities. That is, the posterior density of a parameter might be expected to scale with the number of trials in which that parameter was adopted by the agent. However, that appears not to be the case. For instance, the agent spends half the trials at $\mu \approx 0.07$, but the posterior density here is really low. Moreover, the point of highest density corresponds to a location where *no* time was spent. Therefore, the analyst would be hard-pressed to infer the mixture weights from the posterior densities.

Consider how the analyst might respond to this situation. We know that behavior can change over time as participants become familiar with a task (e.g., Heathcote et al., 2000; Logan, 1992; Newell & Rosenbloom, 1981) and, therefore, we often include a (brief) practice phase in our studies. The hope (and it is often just that) is that by the time actual data collection starts, participants have settled on a point in the cognitive parameter space and their performance has stabilized. So the analyst may discard the initial data, which in this particular example would work well if they had decided (in advance) that the practice block should contain about 50 trials. However, we suspect that in many instances decisions about practice trials are not that well informed and, in any event, there is of course no guarantee that the agent will eventually settle on some stable point in the cognitive parameter space (as more or less happens in this example). Indeed, in many standard paradigms the variation in behavior (e.g., RT) over the course of a block or the entire experiment can be much larger than the effect of any experimental manipulation(s; e.g., Dutilh et al., 2009; Gunawan et al., 2022). So setting aside the inclusion of practice trials, and dealing with the data as they are in Figure 4C, the modeler has fit a simple Gaussian model to the observed data and notes that the fit is poor. A likely response is that they would look for a more complex functional form that can accommodate the data better, for instance, an ex-Gaussian or some other skewed distribution. This model would undoubtedly produce a better fit, but is clearly mismatched to the model that actually generated the data. As a result, inferences drawn from the estimated model and its parameters are not to be trusted.

Basically, we are dealing with a form of model misspecification: The analyst tries to approximate the dynamic ground truth, $\mathcal{M}(D(t)|\psi(t))$, with a static model, $\mathcal{M}(\theta|D')$, where D' is a time-unordered or summary statistics representation of the actual time series of behavior $D(t)$. In our example, the modeler initially actually adopted the correct "basis function," but attributed the poor fit of the model to its functional form, rather than the nonstationarity of the underlying data-generating process. As a result, they were misled into adopting the wrong functional form (e.g., an ex-Gaussian). Of course, it could be argued that the modeler in this scenario should have explored their data better and realized that behavior was nonstationary. Nevertheless, this scenario is representative of common practice in model-based analysis. For instance, evidence accumulation models are typically fit to the marginal distributions of choice and RT data over trials and modelers (including some of the present authors) often do not consider temporal structure in the data (we highlight exceptions below). Variance in the data introduced by parameter dynamics will then have to be captured somehow, for instance, through variation in the model architecture (e.g., adding noise components, such as between-trial noise in drift rate) or in the parameter values themselves (e.g., inflating existing noise components and/or pushing parameters to "compromise" values that were never adopted). In this way, a good model fit might be obtained, but at the cost of taking the modeler further away from the ground truth.[6]

Capturing Parameter Dynamics

To capture the variation in the data-generating process (and the resulting behavioral measurements) over time, we need to (a) consider the temporal structure in the data, the fact that the data form a *time series* of observations and (b) fit the time series with models that allow for temporal variation in the parameters. Of course, there are many models in cognitive science that produce or account for sequential behavior, such as reinforcement learning models (Sutton & Barto, 1998). However, these models typically do so with a single static set of parameters; rather, the variation in predicted behavior stems from the dynamics of and/or noise in the input (e.g., nonstationarity in the reward structure; Behrens et al., 2007). Our agenda here is to promote assessment of the dynamics in the model parameters themselves. For this purpose, it is useful to distinguish between the "core" cognitive model that generates or predicts an observation at a particular point in time, and a "meta" level of control that governs the way core model parameters change over time (for early examples of this idea in cognitive modeling, see Busemeyer & Myung, 1992; Vickers & Lee, 1998). For instance, in the example of Figure 4, the core or point-wise observation model is a simple Gaussian distribution, the parameters of which are controlled by a MCMC-like *transition model*. The question is then how to make this transition model visible (Schumacher et al., 2023).

The challenge for the analyst is to identify the sequence of hidden states from the overall collection of observations. State space models refer to the general class of statistical techniques for solving this problem (Durbin & Koopman, 2012). Of most relevance here are flexible methods that apply to non-Gaussian and nonlinear systems, such as particle filtering (Doucet et al., 2001; Gordon et al., 1993; MacKay, 2003; Speekenbrink, 2016). In particle filtering, the assumption is that an observation y_t is produced by the latent state at this time, x_t, and that the current latent state only depends on the previous state (i.e., we can define a transition model that takes us from the previous state to a new state).[7] The current state is then estimated sequentially with each incoming observation. A large population of particles represents possible latent states that generated the data. Particles associated with a high(er) likelihood for the new data are propagated, whereas those that do not capture the new data well are eliminated. In this way, the distribution of particles can adapt when the data suggest a change in the latent state. Given that we conceptualized the agent's movements in

[6] As discussed in the previous sections and illustrated in Figures 3 and 4, parameter dynamics may stem from participants exerting strategic control (i.e., exploring the cognitive parameter space in order to achieve an objective). However, other parameters that are under less strategic control may also change with time (Dutilh et al., 2009; J. Zhang & Rowe, 2014), for example, through perceptual learning (Dosher & Lu, 2017; Watanabe & Sasaki, 2015) or fatigue (Ratcliff & Van Dongen, 2011). The grounding of cognitive model parameters (and, indeed, the approaches discussed next) is not restricted to capturing only the movements due to strategic control, but can (and should) be applied to capture nonstrategic changes as well.

[7] The latent state x_t may or may not include all parameters of the cognitive model we are trying to estimate. In many particle filtering applications, some of the parameters are allowed to evolve over time, whereas others remain constant (Liu & West, 2001). For our purposes, we consider the scenario where the modeler allows for movement in the parameter space along all dimensions, that is, $x_t = \theta_t$. Nevertheless, a strategy of allowing some parameters to evolve and fixing others may be a good approach for identifying the major dimensions of variation (or strategic control) over time. We revisit this issue in the General Discussion section.

cognitive parameter space as a single particle Markov Chain, adopting a large population of particles to approximate the trajectory of the agent is a natural fit.

Figure 4D illustrates the sequential parameter estimates from a basic bootstrap particle filter (Gordon et al., 1993; Speekenbrink, 2016). Details for this method are given in Appendix B. The posterior means for the parameter estimates are shown by the black solid line (with 95% credible intervals). The initial estimates are some way off, because at this point there is not yet enough data and the estimates are driven by the prior (arbitrarily centered on $\mu = 0$, $\sigma = 2.5$). However, after about 10 trials, the particle filter estimates track the underlying trajectory very well, even when there are sudden larger jumps in the state space (e.g., for σ just before trial 50). The root-mean-squared error (RMSE) between the mean particle filter estimates and the underlying true values is much lower than that between the static posterior mean and the true values (μ: 0.42 vs. 0.89; σ: 0.46 vs. 0.84).

The mixture-of-Gaussians example from Figure 4 tells a cautionary tale of the inferential pitfalls associated with fitting a static model to data from an agent who dynamically explores the cognitive parameter space. However, most computational cognitive models have more complex architectures, contain nonlinearities (e.g., decision thresholds), behave nonlinearly in their parameters, and may have a mixture of static and dynamic parameters. Therefore, it is important to assess whether our observations generalize to a cognitively more realistic setting. We have previously used an expanded judgment paradigm to test whether and when people collapse their decision boundaries in evidence accumulation models (Malhotra et al., 2017, 2018). In this setup, a discrete sequence of binary evidence samples points probabilistically toward one of two decision alternatives, and the agent is set the objective of maximizing their reward rate. The theoretically optimal "policy" (i.e., combination of decision threshold height and gradient) may be derived for this task environment using dynamic programming (Bellman, 1957; Malhotra et al., 2018; see also Drugowitsch et al., 2012; Moran, 2015). This optimal dynamic programming solution is a normative strategy—it does not consider the constraints an agent actually faces when trying to adjust their decision threshold *online* as they interact with the environment and receive feedback. This paradigm is therefore a good test bed for generating and identifying trajectories in cognitive parameter space during a realistic (experimental) task.

Details of the task environment and simulated agents are given in Appendix C. After each binary evidence sample, the agent makes a probabilistic decision to wait for more evidence or commit to one of the two alternatives. If they commit and their choice is correct, they are rewarded; if they make an error they are penalized. The overall amount of time is limited, so the agent will want to fit as many correct trials into this limited period as possible—they have to balance speed and accuracy. To do so, they can modulate the height and gradient of their decision boundary, after estimating the reward rate over a (small) window of trials. That is, the model assumes a "meta" level of control that adjusts these parameters in pursuit of an objective function.

Figure 5A illustrates three different simulated agents: a "stationary" agent and two "nonstationary" agents who adopted the algorithm described in "Sampling in cognitive parameter space" (see also Appendix A). These trajectories may or may not be realistic (that is an empirical question we leave for future investigation). For the nonstationary agents, the "failure" to converge to the optimal policy may be down to a number of reasons: The objective function allows for a broad range of policies that generate reasonable reward rates; noise in the reward rate estimates can lead the agents astray; the sampling algorithm itself is stochastic. For the present purposes, the main thing that matters is that we have three agents who vary in the extent to which they explored the cognitive parameter space.

The core model has three parameters: threshold height α, threshold slope β, and decision noise η. Decision noise allows for a realistic amount of stochasticity in behavior—it allows for different actions given the same amount of accumulated evidence at the same time. We fixed this parameter in our simulations to a constant value (over time and between agents), to create a scenario in which the ground truth model contained a mix of dynamic and static parameters. Of course, the modeler does not know a priori which parameters evolve and which ones do not. This situation is a common occurrence in many cognitive modeling applications. Figure 5B illustrates the estimates for each of the three model parameters. On the right-hand side of each panel, the violin plots show the posterior densities of the parameters estimated from a static model (i.e., a single set of $\{\alpha, \beta, \eta\}$ for the entire run of trials). When the agent is stationary, these estimates are unbiased and precise. For the nonstationary agents, as in our earlier example (Figure 4D), the static estimates to some extent reflect a weighted average of the threshold parameters that were adopted by the agent. However, the noise parameter is greatly overestimated. The reason for this error is obvious: Movement in the latent state space introduces variability in behavior and the static model can only capture this variability by increasing the noise component.

Now consider the particle filter estimates. For the stationary agent, the estimates are concentrated around the true values, albeit with much greater variance compared to the static estimates. The nonstationary agents are tracked very well with the particle filter. There are some periods where there seems to be some parameter trade-off (e.g., for the limited agent toward the end, higher intercepts are compensated for by steeper gradient). Such trade-offs are common in many computational cognitive models (Moran, 2016; Spektor & Kellen, 2018) and can also occur for the static fits (e.g., for the extensive agent, the overall intercept estimate seems to be underestimated and the gradient overestimated). More importantly, the particle filter adequately captures the level and constancy of the noise. Figure 5C summarizes the results by plotting the RMSE for the three agents and parameters. Unsurprisingly, the advantage of the particle filter is more pronounced the greater the movement in the state space. Overall, there appears to be an asymmetry in the consequences of model misspecification: when the data are generated by a stationary agent, the cost of allowing parameters to vary over time is not great (although the estimates are clearly much less precise); when the data are generated by a nonstationary agent, the cost of a static approximation can be severe. In summary, allowing for parameter dynamics not only enables identification of the agent's journey in cognitive parameter space (or the near absence of one), it also guards against erroneous inferences that can arise when a static model has to capture variability in behavior that stems from nonstationary latent states.

General Discussion

Computational cognitive modeling is increasingly ubiquitous in psychology, neuroscience and psychiatry. A model-based analysis of behavior is used to provide mechanistic explanations of empirical

Figure 5

Trajectories of Three Simulated Agents in an Expanded Judgment Paradigm, Along With Static and Dynamic Parameter Estimates

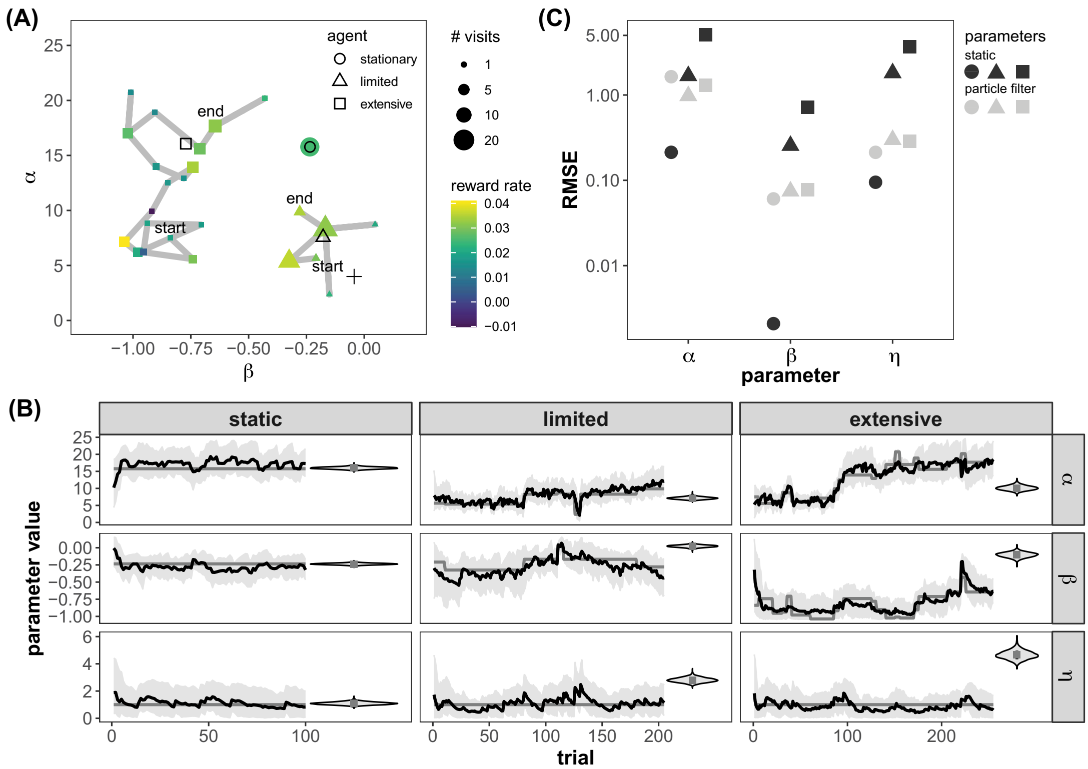

Note. (A) Three agents who vary in the extent to which they explore the parameter space. The *stationary* agent does not move at all (by design). The two nonstationary agents adopted exactly the same sampling algorithm, but started off in different places and ended up with quite different trajectories (*limited* and *extensive* exploration). The reward rate scale corresponds to the mean reward rate averaged over all visits (although the agent only ever estimates the reward rate over a window $W = 5$ trials during the experiment). The mean weighted position is shown for each agent by the open symbol. The optimal policy is indicated by the +. (B) Parameter estimates derived from a static model fit (violins on the right in each panel) and from a particle filter. Conventions as in Figure 4D. (C) RMSE for the different parameters and model fits. The symbol shapes correspond to the different agents, as in Panel A (from left to right for each parameter: stationary—limited—extensive). Errors are shown on a log-scale to visualize the differences better, given that parameters are specified on different scales. For the same reason, only within-parameter comparisons of the error are meaningful. RMSE = root-mean-squared error. See the online article for the color version of this figure.

phenomena such as experimental effects, neural activation and individual or population differences. This type of analysis can generate important insights, such as those listed in Table 1. However, we have argued that more complete explanations should also address *how* and *why* model parameters take on the values that they do. To address these questions, we need to consider the cognitive parameter space—a latent state space formed by the parameters of the (ground truth) model that generates the data—from the perspective of the agent *while they are engaged in a task*. The agent may be trying to achieve some objective (either one they have formulated themselves or one that is imposed by the experimenter), but face several difficulties. First, agents are unlikely to have a good representation of how their position within the cognitive parameter space maps to the objective value. Second, the (local) information they obtain about the objective is probably highly uncertain. Third, their room for maneuver in the state space is limited by various biophysical, environmental, and cognitive constraints. Under these circumstances, the agent may be feeling their way around the task, trying to find a region in the state space where they meet their objective. We have set out a program of *grounding cognitive model parameters* that involves a mix of empirical and computational work to identify the constraints that the agent operates under, the objective(s) they adopt, the information available to them to estimate their objective online over the course of task performance, and the mechanisms by which they move around the state space to achieve their objective.

Part of this endeavor involves an assessment of the agent's behavior from one moment (trial) to the next, taking into account the feedback they receive from the environment. This perspective is often missing from a model-based analysis, where a single set of parameters is typically estimated on the basis of a set of summary statistics or time-unordered representation of the data. At best, these estimated parameters provide a good reflection of the average position of the agent over the course of the task. However, in most cognitive models, the relation between parameter values and predicted behavior is often highly nonlinear, and there is no guarantee that the average position of the agent may be recovered from the time-unordered data. Moreover, the movement of the agent in the state space introduces variability in behavior over time that will have to be absorbed somehow in the static model parameters. As our simulations show (see Figures 4 and 5), this variability can lead to poor model fits, seriously biased parameter estimates and inferential errors. A poor fit might lead the modeler to consider alternative static models (with additional mechanisms or more complex functional form) that may do a better job in capturing the average behavior, but that nevertheless take the analyst further away from the nonstationary ground truth.

Related Approaches

We are by no means the first to discuss these topics. The recency of many citations in the sections above suggests that a good deal of relevant work has been and is being conducted. For instance, there is burgeoning interest in capturing the dynamics of behavior by modeling the time-dependence of parameters of either descriptive, statistical or cognitive models (e.g., Gunawan et al., 2022; Kunkel et al., 2021; Miletić et al., 2021; Schumacher et al., 2023). Hidden Markov Models are being used to identify a small set of latent states that gave rise to the observed data—with participants switching between a limited number of cognitive “regimes” (e.g., on/off-task; Gunawan et al., 2022; Kunkel et al., 2021; Visser, 2011). Alternatively, when parameters change smoothly over time (e.g., as a result of practice or fatigue), the temporal trend may be captured by some simple parametric form (e.g., polynomial; Gunawan et al., 2022). The challenge is then to estimate the parameters of this functional form that, in turn, controls the core model parameters at any one point in time. Part of our contribution is to place recent work like this in the context of a wider research program of grounding cognitive model parameters and open up avenues for future research. In that regard, estimating these trajectories is just a starting point—they still need to be related to the substantive psychological questions about the constraints, objectives, and mechanisms involved in navigating cognitive parameter space.

In our simulations, we adopted a particle filter approach for several reasons. First, most computational cognitive models are highly nonlinear and non-Gaussian, which is the domain particle filters are designed for. Second, particle filtering is an inherently flexible approach, so that we need not commit to a fixed number of latent states in advance (as in Hidden Markov Models) or to a particular fixed functional form of trajectory through the state space (e.g., polynomial). Third, and most importantly, we can think of the agent as a single particle MCMC chain in cognitive parameter space. Using a population of particles to estimate their trajectory therefore seems like a natural choice.

Nevertheless, ideally we would include the *cognitive mechanisms* that underlie the trajectory in the state space into our models directly. One example of this approach is to have a subset of parameters controlled by some learning algorithm that responds to feedback in the environment and updates parameters accordingly (e.g., drift rates or boundary separation in evidence accumulation models; Fontanesi et al., 2019; Lee et al., 2014; Ludwig et al., 2012; Miletić et al., 2020, 2021; Pedersen et al., 2017). Similarly, it may be possible to extend our cognitive model with samplinglike mechanisms for exploring the cognitive parameter space. For instance, the local sampling and comparison routine we outlined as a cognitively minimal search algorithm, is characterized by a number of parameters (e.g., the scale of the jump distribution, the window of trials over which the objective is estimated). Recently, Schumacher et al. (2023) developed a general framework for estimating the parameters of a transition model along with the core model parameters that determine the response at a given moment in time. This framework is applicable to a wide variety of possible transition models, and it is conceivable that we might specify more cognitively plausible transition models (and test between them). We see this challenge as an important and exciting problem to be addressed in future work.

There are also long-standing efforts to ground model parameters through a *rational* or *computational* analysis of cognitive capacities and the world they operate in (Anderson, 1990; Lieder & Griffiths, 2020; Oaksford & Chater, 2007). Some of our arguments echo these approaches. Specifically, a computational or rational analysis requires that we consider the task objective and the information available in the environment for the agent to pursue that objective. Importantly, there is a strong normative commitment to adaptive objectives and the rational use of information—a focus on what the agent *should* do given the task, the environment in which they operate and, potentially, their cognitive constraints (Anderson, 1990; Lieder & Griffiths, 2020). Our analysis encourages researchers to recognize the variety of objectives that might be adopted, to investigate the *actual* objectives that were adopted, and to model the cognitive mechanisms by which these (variable) objectives were pursued (see also Rahnev & Denison, 2018, for similar arguments in the specific domain of perceptual decision making). The ability of humans to make up their own objectives, and the mechanisms by which they pursue those objectives, are fundamental to their psychological makeup. Endowing cognitive models with the capacity to generate this variability is a key step toward a better understanding of variation in cognition and behavior between different individuals and over time.

Challenges and Pragmatic Considerations

Grounding-computational cognitive models involves a form of “zooming out”: going beyond the core model and interrogating the systems that provide input to this core and/or control the way the core model adapts over time (e.g., to achieve some objective). Zooming out is likely to expand the model with additional mechanisms (e.g., a perceptual front-end; sampling mechanisms), a “metalevel” of input and control. Such expansion raises a number of questions, obstacles and objections. Additional mechanisms could make models more flexible and more difficult to falsify. Indeed, it can already be difficult to estimate parameters accurately and reliably for the core model, and expanding this model will make it even harder. A natural question is then how far to take this model expansion.

Model expansion brings with it major challenges in parameter estimation and model selection. First, there is the bias–variance trade-off, clearly visible in Figure 5B. When the ground truth is nonstationary, estimating a single, static set of parameters yields more precise estimates: The posterior densities tend to be narrower than the 95% credible intervals around the time-varying estimates. However, the static parameter estimates are precise, but wrong: They can be a poor reflection of the system at any one point in time and/or display a large amount of bias (e.g., inflated noise estimates). The time-varying estimates track the temporal evolution in the latent states, so that this variation is not accommodated through biasing other model parameters. When the ground truth is stationary, the difference in estimation precision is even more pronounced, although in this case both the simple static and the time-varying estimates show little bias. What level of bias and variance is acceptable depends on the modeling goals. For instance, bias may be acceptable if the primary aim is the detection of some experimental effects and the model is simply a vehicle for extracting a more sensitive dependent measure from multidimensional data (van Ravenzwaaij et al., 2017). However, when the aim is to understand the underlying mechanisms that generated the data better, the loss in estimation precision may be a price worth paying for avoiding bias and gaining insight into the way participants move through the state space.

Second, many cognitive models have a "scaling constraint" in that parameters can only be estimated up to a scale factor (i.e., may be multiplied by a common factor to produce equivalent predictions; Donkin et al., 2009; van Maanen & Miletić, 2021). To ensure the model and its parameters are identifiable, we typically fix one (or more) parameter to a constant, which sets the scale for the remaining parameters. For instance, in the DDM, the within-trial noise is often fixed arbitrarily. In a model-based analysis where we seek to explain some target phenomenon (e.g., experimental effect), the choice of scaling parameter can influence which of the free parameters covary with the target phenomenon (van Maanen & Miletić, 2021). The same problem may arise in the context of estimating time-varying parameters. In particular, suppose there is true variation in the parameter we use to fix the scale of the model—this variation will then have to be absorbed by one or more parameters that *are* allowed to vary over time. As a result, we would introduce spurious dynamics in model parameters that may look like strategic exploration of the cognitive parameter space.

Preliminary simulations (reported in Appendix C—scaling constraint) of the expanded judgment paradigm we used here, confirm that such spurious dynamics can occur: When data are generated with a time-varying scaling parameter, and this parameter is fixed when fitting the model, its temporal variation is partly reflected in other model parameters. To some extent, this problem may be regarded as one of model selection. Any model-based analysis requires decisions to be made about which parameters are free to vary, which parameters are estimated but held constant, and which parameters are fixed as scaling constants. These decisions are no different when we consider strategic (or otherwise) adaptation of model parameters over time. Different configurations of time-varying, constant-but-free and fixed parameters instantiate different models that may be compared (e.g., in their predictive accuracy; see Appendix C—scaling constraint for an example). However, this issue is not *just* a model-selection problem. Deciding on the set of model instantiations (i.e., parameter configurations) to select between, is ultimately a theoretical judgment call. A choice of scaling parameter has to be justified theoretically, ideally with independent empirical support. Considering the cognitive parameter space from the perspective of the agent may be helpful in this regard: If a parameter is to be fixed at all, choose one that is strongly constrained and over which the agent has little strategic control.

Third, if we take movement in cognitive parameter space seriously, it is clear that a wide variety of trajectories are possible (e.g., the two nonstationary agents in Figure 5A). It is likely that participants will start off at different points in the state space, as a result of their own motivations, background knowledge, history with broadly similar tasks, and their own conceptualization of what the task is (Szollosi et al., 2023). Once you take into account the possibility of different transition models, noise in the objective estimates, and stochasticity in the transition model, it will be extremely difficult to disentangle these many sources of variability. There is no single approach to deal with these challenges, but our article suggests a number of possible empirical and computational avenues to explore (e.g., setting clear task objectives, identifying objectives that were actually adopted). At least at a practical level, there are various tools for estimating individual trajectories, as demonstrated by our own simulations (see Figures 4 and 5), and by other authors recently (Gunawan et al., 2022; Schumacher et al., 2023). The estimated trajectories may then become the data of interest for the development and tuning of more cognitively principled search algorithms. For instance, if an estimated trajectory suggests that a participant tends to step in the same direction after a previous move improved their objective value, that behavior would suggest some form of hill-climbing (Bramley et al., 2017; Busemeyer & Myung, 1992); however, if they took very large steps and behavior was uncorrelated from one epoch to the next, that might suggest a much more random search process, such as independent sampling. In any event, just because identifying individual trajectories in cognitive parameter space is a formidable challenge, pretending that they do not exist is likely to lead us down inferential blind alleys.

Aside from the challenges in estimating trajectories in the state space, a reasonable question is how we stop model complexity spiraling out of control. If we expand a core model with higher level mechanisms that control the ones lower down, those control mechanisms themselves may have parameters in need of explanation. How many layers of control do we build? In our view, the correct level of explanation should be determined by the questions that a researcher wishes to answer. Some metalevel model will be required to answer questions about fine-grained behavior, such as how participants navigate cognitive parameter space in order to meet some objective. It is entirely conceivable that interesting questions can be asked about the mechanisms at this metalevel, such as the generality of the transition model across different task domains, the flexible tuning of its parameters (e.g., scale of the jump distribution), and so forth. We certainly do not advocate an extensive regress of meta-meta-… -level models and complexity for complexity's sake. At some (meta-meta-… -)level, the higher level questions may no longer be interesting, or we may simply lack sufficiently diagnostic data that speak to these questions. However, we hope to have persuaded the reader that there is great scope for expanding model-based analyses to incorporate at least one layer of metalevel processes.

Finally, trying to understand the trajectories in cognitive parameter space implies a shift of focus in the development of our theories and models. Current model-based practice typically aims to capture some average behavior over the course of a single, specific task. We invite the modeler to take the view of the participant navigating that task, from one moment to the next. This perspective entails a focus on the mechanisms by which cognition and behavior evolve over time and adapt to a set of task demands (Bramley et al., 2023; Donkin et al., 2022). Typically, participants are feeling their way through a task, with little knowledge of how their position in the cognitive state space influences the objective they are looking to achieve. As a result, different participants will generate very different trajectories in the state space and produce very different (average) behaviors. Returning to the two agents illustrated in Figure 5A, the static parameter estimates suggest very different strategies (policies) for solving the task. At the same time, at the higher control level, the agents' strategies *were exactly the* *same* (by design in this instance). They had exactly the same objective and used exactly the same sampling algorithm to explore that objective. As cognitive scientists, we are often interested in robust invariances that generalize across environments, tasks, and people. We speculate that it is at the metalevel of control that many interesting cognitive invariances will be found.

References

Anderson, J. R. (1990). *The adaptive character of thought*. Psychology Press.

Andrieu, C., De Freitas, N., Doucet, A., & Jordan, M. I. (2003). An introduction to MCMC for machine learning. *Machine Learning*, *50*(1), 5–43. https://doi.org/10.1023/A:1020281327116

Attwell, D., & Laughlin, S. B. (2001). An energy budget for signaling in the grey matter of the brain. *Journal of Cerebral Blood Flow & Metabolism*, *21*(10), 1133–1145. https://doi.org/10.1097/00004647-200110000-00001

Averbeck, B. B., Latham, P. E., & Pouget, A. (2006). Neural correlations, population coding and computation. *Nature Reviews Neuroscience*, *7*(5), 358–366. https://doi.org/10.1038/nrn1888

Balci, F., Simen, P., Niyogi, R., Saxe, A., Hughes, J. A., Holmes, P., & Cohen, J. D. (2011). Acquisition of decision making criteria: Reward rate ultimately beats accuracy. *Attention, Perception, & Psychophysics*, *73*(2), 640–657. https://doi.org/10.3758/s13414-010-0049-7

Bang, D., Aitchison, L., Moran, R., Herce Castanon, S., Rafiee, B., Mahmoodi, A., Lau, J. Y., Latham, P. E., Bahrami, B., & Summerfield, C. (2017). Confidence matching in group decision-making. *Nature Human Behaviour*, *1*(6), Article 0117. https://doi.org/10.1038/s41562-017-0117

Batchelder, W. H., & Riefer, D. M. (1990). Multinomial processing models of source monitoring. *Psychological Review*, *97*(4), 548–564. https://doi.org/10.1037/0033-295X.97.4.548

Bayen, U. J., Murnane, K., & Erdfelder, E. (1996). Source discrimination, item detection, and multinomial models of source monitoring. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*(1), 197–215. https://doi.org/10.1037/0278-7393.22.1.197

Bayen, U. J., Nakamura, G. V., Dupuis, S. E., & Yang, C.-L. (2000). The use of schematic knowledge about sources in source monitoring. *Memory & Cognition*, *28*(3), 480–500. https://doi.org/10.3758/BF03198562

Bays, P. M., Catalao, R. F., & Husain, M. (2009). The precision of visual working memory is set by allocation of a shared resource. *Journal of Vision*, *9*(10), Article 7. https://doi.org/10.1167/9.10.7

Bays, P. M., & Husain, M. (2008). Dynamic shifts of limited working memory resources in human vision. *Science*, *321*(5890), 851–854. https://doi.org/10.1126/science.1158023

Bechtel, W., & Abrahamsen, A. (2005). Explanation: A mechanist alternative. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, *36*(2), 421–441. https://doi.org/10.1016/j.shpsc.2005.03.010

Bechtel, W., & Abrahamsen, A. (2010). Dynamic mechanistic explanation: Computational modeling of circadian rhythms as an exemplar for cognitive science. *Studies in History and Philosophy of Science Part a*, *41*(3), 321–333. https://doi.org/10.1016/j.shpsa.2010.07.003

Behrens, T. E., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. (2007). Learning the value of information in an uncertain world. *Nature Neuroscience*, *10*(9), 1214–1221. https://doi.org/10.1038/nn1954

Bellman, R. (1957). *Dynamic programming*. Princeton University Press.

Bill, J., Gershman, S. J., & Drugowitsch, J. (2022). Visual motion perception as online hierarchical inference. *Nature Communications*, *13*(1), Article 7403. https://doi.org/10.1038/s41467-022-34805-5

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (2006). The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review*, *113*(4), 700–765. https://doi.org/10.1037/0033-295X.113.4.700

Bogacz, R., Hu, P. T., Holmes, P. J., & Cohen, J. D. (2010). Do humans produce the speed–accuracy trade-off that maximizes reward rate? *Quarterly Journal of Experimental Psychology*, *63*(5), 863–891. https://doi.org/10.1080/17470210903091643

Bohil, C. J., & Maddox, W. T. (2003). On the generality of optimal versus objective classifier feedback effects on decision criterion learning in perceptual categorization. *Memory & Cognition*, *31*(2), 181–198. https://doi.org/10.3758/BF03194378

Bompas, A., Sumner, P., & Hedge, C. (2024). Non-decision time: The higgs boson of decision. *Psychological Review*. Advance online publication. https://doi.org/10.1037/rev0000487

Bonawitz, E., Denison, S., Gopnik, A., & Griffiths, T. L. (2014). Win-stay, lose-sample: A simple sequential algorithm for approximating Bayesian inference. *Cognitive Psychology*, *74*, 35–65. https://doi.org/10.1016/j.cogpsych.2014.06.003

Bonawitz, E., Denison, S., Griffiths, T. L., & Gopnik, A. (2014). Probabilistic models, learning algorithms, and response variability: Sampling in cognitive development. *Trends in Cognitive Sciences*, *18*(10), 497–500. https://doi.org/10.1016/j.tics.2014.06.006

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), 301–338. https://doi.org/10.1037/rev0000061

Bramley, N. R., & Xu, F. (2023). Active inductive inference in children and adults: A constructivist perspective. *Cognition*, *238*, Article 105471. https://doi.org/10.1016/j.cognition.2023.105471

Bramley, N. R., Zhao, B., Quillien, T., & Lucas, C. G. (2023). Local search and the evolution of world models. *Topics in Cognitive Science*. Advance online publication. https://doi.org/10.1111/tops.12703

Brehmer, B. (1976). Subjects' ability to find the parameters of functional rules in probabilistic inference tasks. *Organizational Behavior and Human Performance*, *17*(2), 388–397. https://doi.org/10.1016/0030-5073(76)90074-X

Brier, G. W. (1950). Verification of forecasts expressed in terms of probability. *Monthly Weather Review*, *78*(1), 1–3. https://doi.org/10.1175/1520-0493(1950)078<0001:VOFEIT>2.0.CO;2

Brown, S. D., & Heathcote, A. (2008). The simplest complete model of choice response time: Linear ballistic accumulation. *Cognitive Psychology*, *57*(3), 153–178. https://doi.org/10.1016/j.cogpsych.2007.12.002

Bullier, J. (2001). Integrated model of visual processing. *Brain Research Reviews*, *36*(2–3), 96–107. https://doi.org/10.1016/S0165-0173(01)00085-6

Burnham, K. P., & Anderson, D. R. (2002). *Model selection and inference: A practical information-theoretic approach*. Springer.

Busemeyer, J. R., Gluth, S., Rieskamp, J., & Turner, B. M. (2019). Cognitive and neural bases of multi-attribute, multi-alternative, value-based decisions. *Trends in Cognitive Sciences*, *23*(3), 251–263. https://doi.org/10.1016/j.tics.2018.12.003

Busemeyer, J. R., & Myung, I. J. (1992). An adaptive approach to human decision making: Learning theory, decision theory, and human performance. *Journal of Experimental Psychology: General*, *121*(2), 177–194. https://doi.org/10.1037/0096-3445.121.2.177

Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, *100*(3), 432–459. https://doi.org/10.1037/0033-295X.100.3.432

Cacioppo, J. T., Petty, R. E., Feinstein, J. A., & Jarvis, W. B. G. (1996). Dispositional differences in cognitive motivation: The life and times of individuals varying in need for cognition. *Psychological Bulletin*, *119*(2), 197–253. https://doi.org/10.1037/0033-2909.119.2.197

Camerer, C. F., & Hogarth, R. M. (1999). The effects of financial incentives in experiments: A review and capital-labor-production framework. *Journal of Risk and Uncertainty*, *19*, 7–42. https://doi.org/10.1023/A:1007850605129

Chater, N., Zhu, J.-Q., Spicer, J., Sundh, J., León-Villagrá, P., & Sanborn, A. (2020). Probabilistic biases meet the Bayesian brain. *Current Directions in Psychological Science*, *29*(5), 506–512. https://doi.org/10.1177/0963721420954801

Cisek, P., Puskas, G. A., & El-Murr, S. (2009). Decisions in changing conditions: The urgency-gating model. *Journal of Neuroscience*, *29*(37), 11560–11571. https://doi.org/10.1523/JNEUROSCI.1844-09.2009

Cohen, J. D., McClure, S. M., & Yu, A. J. (2007). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *362*(1481), 933–942. https://doi.org/10.1098/rstb.2007.2098

Dasgupta, I., Schulz, E., & Gershman, S. J. (2017). Where do hypotheses come from? *Cognitive Psychology*, *96*, 1–25. https://doi.org/10.1016/j.cogpsych.2017.05.001

Dasgupta, I., Schulz, E., Goodman, N. D., & Gershman, S. J. (2018). Remembrance of inferences past: Amortization in human hypothesis generation. *Cognition*, *178*, 67–81. https://doi.org/10.1016/j.cognition.2018.04.017

Daw, N. D., O'doherty, J. P., Dayan, P., Seymour, B., & Dolan, R. J. (2006). Cortical substrates for exploratory decisions in humans. *Nature*, *441*(7095), 876–879. https://doi.org/10.1038/nature04766

Ditterich, J. (2006). Stochastic models of decisions about motion direction: Behavior and physiology. *Neural Networks*, *19*(8), 981–1012. https://doi.org/10.1016/j.neunet.2006.05.042

Donkin, C. & Brown, S. D. (2018). Response times and decision-making. In J. T. Wixted (Ed.), *Stevens' handbook of experimental psychology and cognitive neuroscience* (Vol. 5, pp. 349–377). Wiley. https://doi.org/10.1002/9781119170174.epcn509

Donkin, C., Brown, S. D., & Heathcote, A. (2009). The overconstraint of response time models: Rethinking the scaling problem. *Psychonomic Bulletin & Review*, *16*, 1129–1135. https://doi.org/10.3758/PBR.16.6.1129

Donkin, C., Szollosi, A., & Bramley, N. R. (2022). Observing effects in various contexts won't give us general psychological theories. *Behavioral and Brain Sciences*, *45*, Article e13. https://doi.org/10.1017/S0140525X21000479

Dosher, B. A., & Lu, Z.-L. (1999). Mechanisms of perceptual learning. *Vision Research*, *39*(19), 3197–3221. https://doi.org/10.1016/S0042-6989(99)00059-0

Dosher, B. A., & Lu, Z.-L. (2017). Visual perceptual learning and models. *Annual Review of Vision Science*, *3*, 343–363. https://doi.org/10.1146/annurev-vision-102016-061249

Doucet, A., De Freitas, N., & Gordon, N. (2001). An introduction to sequential Monte Carlo methods. In A. Doucet, N. Freitas, & N. Gordon (Eds.), *Sequential Monte Carlo methods in practice* (pp. 3–14). Springer. https://doi.org/10.1007/978-1-4757-3437-9_1

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. https://doi.org/10.1523/JNEUROSCI.4010-11.2012

Durbin, J., & Koopman, S. J. (2012). *Time series analysis by state space methods* (Vol. 38). Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199641178.001.0001

Dutilh, G., Annis, J., Brown, S. D., Cassey, P., Evans, N. J., Grasman, R. P., Hawkins, G. E., Heathcote, A., Holmes, W. R., Krypotos, A.-M., Kupitz, C. N., Leite, F. P., Lerche, V., Lin, Y.-S., Logan, G. D., Palmeri, T. J., Starns, J. J., Trueblood, J. S., van Maanen, L., … Donkin, C. (2019). The quality of response time data inference: A blinded, collaborative assessment of the validity of cognitive models. *Psychonomic Bulletin & Review*, *26*(4), 1051–1069. https://doi.org/10.3758/s13423-017-1417-2

Dutilh, G., Vandekerckhove, J., Tuerlinckx, F., & Wagenmakers, E.-J. (2009). A diffusion model decomposition of the practice effect. *Psychonomic Bulletin & Review*, *16*, 1026–1036. https://doi.org/10.3758/16.6.1026

Engbert, R., Nuthmann, A., Richter, E. M., & Kliegl, R. (2005). SWIFT: A dynamical model of saccade generation during reading. *Psychological Review*, *112*(4), 777–813. https://doi.org/10.1037/0033-295X.112.4.777

Evans, N. J., & Brown, S. D. (2017). People adopt optimal policies in simple decision-making, after practice and guidance. *Psychonomic Bulletin & Review*, *24*(2), 597–606. https://doi.org/10.3758/s13423-016-1135-1

Evans, N. J., Rae, B., Bushmakin, M., Rubin, M., & Brown, S. D. (2017). Need for closure is associated with urgency in perceptual decision-making. *Memory & Cognition*, *45*(7), 1193–1205. https://doi.org/10.3758/s13421-017-0718-z

Farrell, S., & Lewandowsky, S. (2018). *Computational modeling of cognition and behavior*. Cambridge University Press. https://doi.org/10.1017/CBO9781316272503

Farrell, S., Ludwig, C. J., Ellis, L. A., & Gilchrist, I. D. (2010). Influence of environmental statistics on inhibition of saccadic return. *Proceedings of the National Academy of Sciences*, *107*(2), 929–934. https://doi.org/10.1073/pnas.0906845107

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. https://doi.org/10.1016/j.tics.2010.01.003

Fontanesi, L., Gluth, S., Spektor, M. S., & Rieskamp, J. (2019). A reinforcement learning diffusion decision model for value-based decisions. *Psychonomic Bulletin & Review*, *26*(4), 1099–1121. https://doi.org/10.3758/s13423-018-1554-2

Forstmann, B. U., Anwander, A., Schäfer, A., Neumann, J., Brown, S., Wagenmakers, E.-J., Bogacz, R., & Turner, R. (2010). Cortico-striatal connections predict control over speed and accuracy in perceptual decision making. *Proceedings of the National Academy of Sciences of the United States of America*, *107*(36), 15916–15920. https://doi.org/10.1073/pnas.1004932107

Forstmann, B. U., Wagenmakers, E.-J., Eichele, T., Brown, S., & Serences, J. T. (2011). Reciprocal relations between cognitive neuroscience and formal cognitive models: Opposites attract? *Trends in Cognitive Sciences*, *15*(6), 272–279. https://doi.org/10.1016/j.tics.2011.04.002

Fradkin, I., Ludwig, C., Eldar, E., & Huppert, J. D. (2020). Doubting what you already know: Uncertainty regarding state transitions is associated with obsessive compulsive symptoms. *PLOS Computational Biology*, *16*(2), Article e1007634. https://doi.org/10.1371/journal.pcbi.1007634

Galeano Weber, E. M., Peters, B., Hahn, T., Bledowski, C., & Fiebach, C. J. (2016). Superior intraparietal sulcus controls the variability of visual working memory precision. *Journal of Neuroscience*, *36*(20), 5623–5635. https://doi.org/10.1523/jneurosci.1596-15.2016

Gershman, S. J., Vul, E., & Tenenbaum, J. B. (2012). Multistability and perceptual inference. *Neural Computation*, *24*(1), 1–24. https://doi.org/10.1162/NECO_a_00226

Gheza, D., Kool, W., & Pourtois, G. (2023). Need for cognition moderates the relief of avoiding cognitive effort. *PLOS One*, *18*(11), Article e0287954. https://doi.org/10.1371/journal.pone.0287954

Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. *Annual Review of Psychology*, *62*(1), 451–482. https://doi.org/10.1146/annurev-psych-120709-145346

Gold, J. I., & Shadlen, M. N. (2001). Neural computations that underlie decisions about sensory stimuli. *Trends in Cognitive Sciences*, *5*(1), 10–16. https://doi.org/10.1016/S1364-6613(00)01567-9

Goodman, N. D., Tenenbaum, J. B., Feldman, J., & Griffiths, T. L. (2008). A rational analysis of rule-based concept learning. *Cognitive Science*, *32*(1), 108–154. https://doi.org/10.1080/03640210701802071

Gordon, N. J., Salmond, D. J., & Smith, A. F. (1993). Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Proceedings F (Radar and Signal Processing)*, *140*(2), 107–113. https://doi.org/10.1049/ip-f-2.1993.0015

Guest, O., & Martin, A. E. (2021). How computational modeling can force theory building in psychological science. *Perspectives on Psychological Science*, *16*(4), 789–802. https://doi.org/10.1177/1745691620970585

Gunawan, D., Hawkins, G. E., Kohn, R., Tran, M.-N., & Brown, S. D. (2022). Time-evolving psychological processes over repeated decisions. *Psychological Review*, *129*(3), 438–456. https://doi.org/10.1037/rev0000351

Haefner, R. M., Berkes, P., & Fiser, J. (2016). Perceptual decision-making as probabilistic inference by neural sampling. *Neuron*, *90*(3), 649–660. https://doi.org/10.1016/j.neuron.2016.03.020

Hartig, F., Calabrese, J. M., Reineking, B., Wiegand, T., & Huth, A. (2011). Statistical inference for stochastic simulation models—Theory and application. *Ecology Letters*, *14*(8), 816–827. https://doi.org/10.1111/j.1461-0248.2011.01640.x

Hawkins, G. E., Brown, S. D., Steyvers, M., & Wagenmakers, E.-J. (2012). An optimal adjustment procedure to minimize experiment time in decisions with multiple alternatives. *Psychonomic Bulletin & Review*, *19*(2), 339–348. https://doi.org/10.3758/s13423-012-0216-z

Hawkins, G. E., Forstmann, B. U., Wagenmakers, E.-J., Ratcliff, R., & Brown, S. D. (2015). Revisiting the evidence for collapsing boundaries and urgency signals in perceptual decision-making. *Journal of Neuroscience*, *35*(6), 2476–2484. https://doi.org/10.1523/JNEUROSCI.2410-14.2015

Hawkins, G. E., & Heathcote, A. (2021). Racing against the clock: Evidence-based versus time-based decisions. *Psychological Review*, *128*(2), 222–263. https://doi.org/10.1037/rev0000259

Heathcote, A., Brown, S., & Mewhort, D. J. (2000). The power law repealed: The case for an exponential law of practice. *Psychonomic Bulletin & Review*, *7*(2), 185–207. https://doi.org/10.3758/BF03212979

Heathcote, A., Lin, Y.-S., Reynolds, A., Strickland, L., Gretton, M., & Matzke, D. (2019). Dynamic models of choice. *Behavior Research Methods*, *51*, 961–985. https://doi.org/10.3758/s13428-018-1067-y

Heathcote, A., & Matzke, D. (2022). Winner takes all! What are race models, and why and how should psychologists use them? *Current Directions in Psychological Science*, *31*(5), 383–394. https://doi.org/10.1177/09637214221095852

Hertwig, R., & Ortmann, A. (2001). Experimental practices in economics: A methodological challenge for psychologists? *Behavioral and Brain Sciences*, *24*(3), 383–403. https://doi.org/10.1017/S0140525X01004149

Holmes, W. R., O'Daniels, P., & Trueblood, J. S. (2020). A joint deep neural network and evidence accumulation modeling approach to human decision-making with naturalistic images. *Computational Brain & Behavior*, *3*(1), 1–12. https://doi.org/10.1007/s42113-019-00042-1

Holmes, W. R., Trueblood, J. S., & Heathcote, A. (2016). A new framework for modeling decisions about changing information: The piecewise linear ballistic accumulator model. *Cognitive Psychology*, *85*, 1–29. https://doi.org/10.1016/j.cogpsych.2015.11.002

Houser, D., & McCabe, K. (2014). Experimental economics and experimental game theory. In P. W. Glimcher & E. Fehr (Eds.), *Neuroeconomics* (pp. 19–34). Elsevier.

Howes, A., Lewis, R. L., & Vera, A. (2009). Rational adaptation under task and processing constraints: Implications for testing theories of cognition and action. *Psychological Review*, *116*(4), 717–751. https://doi.org/10.1037/a0017187

Huys, Q. J., Maia, T. V., & Frank, M. J. (2016). Computational psychiatry as a bridge from neuroscience to clinical applications. *Nature Neuroscience*, *19*(3), 404–413. https://doi.org/10.1038/nn.4238

Jarecki, J. B., Tan, J. H., & Jenny, M. A. (2020). A framework for building cognitive process models. *Psychonomic Bulletin & Review*, *27*(6), 1218–1229. https://doi.org/10.3758/s13423-020-01747-2

Jepma, M., & Nieuwenhuis, S. (2011). Pupil diameter predicts changes in the exploration–exploitation trade-off: Evidence for the adaptive gain theory. *Journal of Cognitive Neuroscience*, *23*(7), 1587–1596. https://doi.org/10.1162/jocn.2010.21548

Kahnt, T., Grueschow, M., Speck, O., & Haynes, J.-D. (2011). Perceptual learning and decision-making in human medial frontal cortex. *Neuron*, *70*(3), 549–559. https://doi.org/10.1016/j.neuron.2011.02.054

Kalish, M. L., Lewandowsky, S., & Kruschke, J. K. (2004). Population of linear experts: Knowledge partitioning and function learning. *Psychological Review*, *111*(4), 1072–1099. https://doi.org/10.1037/0033-295X.111.4.1072

Kaplan, D. M., & Craver, C. F. (2011). The explanatory force of dynamical and mathematical models in neuroscience: A mechanistic perspective. *Philosophy of Science*, *78*(4), 601–627. https://doi.org/10.1086/661755

Kool, W., & Botvinick, M. (2014). A labor/leisure tradeoff in cognitive control. *Journal of Experimental Psychology: General*, *143*(1), 131–141. https://doi.org/10.1037/a0031048

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. https://doi.org/10.1038/nn.2635

Kriegeskorte, N., & Douglas, P. K. (2018). Cognitive computational neuroscience. *Nature Neuroscience*, *21*(9), 1148–1160. https://doi.org/10.1038/s41593-018-0210-5

Kunkel, D., Yan, Z., Craigmile, P. F., Peruggia, M., & Van Zandt, T. (2021). Hierarchical hidden Markov models for response time data. *Computational Brain & Behavior*, *4*, 70–86. https://doi.org/10.1007/s42113-020-00076-w

LaBerge, D. (1962). A recruitment theory of simple behavior. *Psychometrika*, *27*(4), 375–396. https://doi.org/10.1007/BF02289645

Lee, M. D., Bock, J. R., Cushman, I., & Shankle, W. R. (2020). An application of multinomial processing tree models and Bayesian methods to understanding memory impairment. *Journal of Mathematical Psychology*, *95*, Article 102328. https://doi.org/10.1016/j.jmp.2020.102328

Lee, M. D., Newell, B. R., & Vandekerckhove, J. (2014). Modeling the adaptation of search termination in human decision making. *Decision*, *1*(4), 223–251. https://doi.org/10.1037/dec0000019

Lee, M. D., & Wagenmakers, E.-J. (2014). *Bayesian cognitive modeling: A practical course*. Cambridge University Press. https://doi.org/10.1017/CBO9781139087759

Levy, W. B., & Baxter, R. A. (1996). Energy efficient neural codes. *Neural Computation*, *8*(3), 531–543. https://doi.org/10.1162/neco.1996.8.3.531

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*, Article e1. https://doi.org/10.1017/S0140525X1900061X

Lieder, F., Griffiths, T. L., Huys, Q. J. M., & Goodman, N. D. (2018). The anchoring bias reflects rational use of cognitive resources. *Psychonomic Bulletin & Review*, *25*(1), 322–349. https://doi.org/10.3758/s13423-017-1286-8

Link, S., & Heath, R. (1975). A sequential theory of psychological discrimination. *Psychometrika*, *40*(1), 77–105. https://doi.org/10.1007/BF02291481

Liu, J., & West, M. (2001). Combined parameter and state estimation in simulation-based filtering. In A. Doucet, N. de Freitas, & N. Gordon (Eds.), *Sequential Monte Carlo methods in practice* (pp. 197–223). Springer. https://doi.org/10.1007/978-1-4757-3437-9_10

Logan, G. D. (1992). Shapes of reaction-time distributions and shapes of learning curves: A test of the instance theory of automaticity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *18*(5), 883–914. https://doi.org/10.1037/0278-7393.18.5.883

Love, B. C. (2015). The algorithmic level is the bridge between computation and brain. *Topics in Cognitive Science*, *7*(2), 230–242. https://doi.org/10.1111/tops.12131

Love, B. C., & Gureckis, T. M. (2007). Models in search of a brain. *Cognitive, Affective, & Behavioral Neuroscience*, *7*(2), 90–108. https://doi.org/10.3758/CABN.7.2.90

Lu, Z.-L., & Dosher, B. A. (2008). Characterizing observers using external noise and observer models: Assessing internal representations with external noise. *Psychological Review*, *115*(1), 44–82. https://doi.org/10.1037/0033-295X.115.1.44

Lucas, C. G., Griffiths, T. L., Williams, J. J., & Kalish, M. L. (2015). A rational model of function learning. *Psychonomic Bulletin & Review*, *22*(5), 1193–1215. https://doi.org/10.3758/s13423-015-0808-5

Ludwig, C. J. (2009). Temporal integration of sensory evidence for saccade target selection. *Vision Research*, *49*(23), 2764–2773. https://doi.org/10.1016/j.visres.2009.08.012

Ludwig, C. J., Butler, S. H., Rossit, S., Harvey, M., & Gilchrist, I. D. (2009). Modelling contralesional movement slowing after unilateral brain damage. *Neuroscience Letters*, *452*(1), 1–4. https://doi.org/10.1016/j.neulet.2009.01.033

Ludwig, C. J., & Evens, D. R. (2017). Information foraging for perceptual decisions. *Journal of Experimental Psychology: Human Perception and Performance*, *43*(2), 245–264. https://doi.org/10.1037/xhp0000299

Ludwig, C. J., Farrell, S., Ellis, L. A., & Gilchrist, I. D. (2009). The mechanism underlying inhibition of saccadic return. *Cognitive Psychology*, *59*(2), 180–202. https://doi.org/10.1016/j.cogpsych.2009.04.002

Ludwig, C. J., Farrell, S., Ellis, L. A., Hardwicke, T. E., & Gilchrist, I. D. (2012). Context-gated statistical learning and its role in visual-saccadic decisions. *Journal of Experimental Psychology: General*, *141*(1), 150–169. https://doi.org/10.1037/a0024916

Mack, M. L., Preston, A. R., & Love, B. C. (2013). Decoding the brain's algorithm for categorization from its neural implementation. *Current Biology*, *23*(20), 2023–2027. https://doi.org/10.1016/j.cub.2013.08.035

MacKay, D. J. (2003). *Information theory, inference and learning algorithms*. Cambridge University Press.

Malhotra, G., Leslie, D. S., Ludwig, C. J., & Bogacz, R. (2017). Overcoming indecision by changing the decision boundary. *Journal of Experimental Psychology: General*, *146*(6), 776–805. https://doi.org/10.1037/xge0000286

Malhotra, G., Leslie, D. S., Ludwig, C. J., & Bogacz, R. (2018). Time-varying decision boundaries: Insights from optimality analysis. *Psychonomic Bulletin & Review*, *25*(3), 971–996. https://doi.org/10.3758/s13423-017-1340-6

Marr, D. (1982). *Vision: A computational investigation into the human representation and processing*. W. H. Freeman.

Mikhael, J. G., & Bogacz, R. (2016). Learning reward uncertainty in the basal ganglia. *PLOS Computational Biology*, *12*(9), Article e1005062. https://doi.org/10.1371/journal.pcbi.1005062

Miletić, S., Boag, R. J., & Forstmann, B. U. (2020). Mutual benefits: Combining reinforcement learning with sequential sampling models. *Neuropsychologia*, *136*, Article 107261. https://doi.org/10.1016/j.neuropsychologia.2019.107261

Miletić, S., Boag, R. J., Trutti, A. C., Stevenson, N., Forstmann, B. U., & Heathcote, A. (2021). A new model of decision processing in instrumental learning tasks. *eLife*, *10*, Article e63055. https://doi.org/10.7554/eLife.63055

Miletić, S., & van Maanen, L. (2019). Caution in decision-making under time pressure is mediated by timing ability. *Cognitive Psychology*, *110*, 16–29. https://doi.org/10.1016/j.cogpsych.2019.01.002

Moran, R. (2015). Optimal decision making in heterogeneous and biased environments. *Psychonomic Bulletin & Review*, *22*(1), 38–53. https://doi.org/10.3758/s13423-014-0669-3

Moran, R. (2016). Thou shalt identify! The identifiability of two high-threshold models in confidence-rating recognition (and super-recognition) paradigms. *Journal of Mathematical Psychology*, *73*, 1–11. https://doi.org/10.1016/j.jmp.2016.03.002

Munoz, D. P., & Wurtz, R. H. (1995). Saccade-related activity in monkey superior colliculus. II. Spread of activity during saccades. *Journal of Neurophysiology*, *73*(6), 2334–2348. https://doi.org/10.1152/jn.1995.73.6.2334

Navarro, D. J. (2019). Between the devil and the deep blue sea: Tensions between scientific judgement and statistical model selection. *Computational Brain & Behavior*, *2*(1), 28–34. https://doi.org/10.1007/s42113-018-0019-z

Newell, A., & Rosenbloom, P. S. (1981). Mechanisms of skill acquisition and the law of practice. In J. R. Anderson (Ed.), *Cognitive skills and their acquisition* (pp. 1–55). Lawrence Erlbaum.

Nuthmann, A., Smith, T. J., Engbert, R., & Henderson, J. M. (2010). CRISP: A computational model of fixation durations in scene viewing. *Psychological Review*, *117*(2), 382–405. https://doi.org/10.1037/a0018924

Oaksford, M., & Chater, N. (2007). *Bayesian rationality: The probabilistic approach to human reasoning*. Oxford University Press.

Oberauer, K., & Lewandowsky, S. (2019). Addressing the theory crisis in psychology. *Psychonomic Bulletin & Review*, *26*(5), 1596–1618. https://doi.org/10.3758/s13423-019-01645-2

Oberauer, K., & Lin, H.-Y. (2017). An interference model of visual working memory. *Psychological Review*, *124*(1), 21–59. https://doi.org/10.1037/rev0000044

Palestro, J. J., Sederberg, P. B., Osth, A. F., Van Zandt, T., & Turner, B. M. (2018). *Likelihood-free methods for cognitive science*. Springer. https://doi.org/10.1007/978-3-319-72425-6

Palminteri, S., Lefebvre, G., Kilford, E. J., & Blakemore, S.-J. (2017). Confirmation bias in human reinforcement learning: Evidence from counterfactual feedback processing. *PLOS Computational Biology*, *13*(8), Article e1005684. https://doi.org/10.1371/journal.pcbi.1005684

Pedersen, M. L., Frank, M. J., & Biele, G. (2017). The drift diffusion model as the choice rule in reinforcement learning. *Psychonomic Bulletin & Review*, *24*(4), 1234–1251. https://doi.org/10.3758/s13423-016-1199-y

Peich, M.-C., Husain, M., & Bays, P. M. (2013). Age-related decline of precision and binding in visual working memory. *Psychology and Aging*, *28*(3), 729–743. https://doi.org/10.1037/a0033236

Perez Santangelo, A., Ludwig, C. J., Navajas, J., Sigman, M., & Leone, M. J. (2022). Background music changes the policy of human decision-making: Evidence from experimental and drift-diffusion model-based approaches on different decision tasks. *Journal of Experimental Psychology: General*, *151*(9), 2222–2236. https://doi.org/10.1037/xge0001189

Pergolizzi, D., & Chua, E. F. (2016). Transcranial direct current stimulation over the parietal cortex alters bias in item and source memory tasks. *Brain and Cognition*, *108*, 56–65. https://doi.org/10.1016/j.bandc.2016.06.009

Pitt, M. A., Myung, I. J., & Zhang, S. (2002). Toward a method of selecting among computational models of cognition. *Psychological Review*, *109*(3), 472–491. https://doi.org/10.1037/0033-295X.109.3.472

Pruszynski, J. A., King, G. L., Boisse, L., Scott, S. H., Flanagan, J. R., & Munoz, D. P. (2010). Stimulus-locked responses on human arm muscles reveal a rapid neural pathway linking visual input to arm motor output. *European Journal of Neuroscience*, *32*(6), 1049–1057. https://doi.org/10.1111/j.1460-9568.2010.07380.x

Rabe, M. M., Chandra, J., Krügel, A., Seelig, S. A., Vasishth, S., & Engbert, R. (2021). A Bayesian approach to dynamical modeling of eye-movement control in reading of normal, mirrored, and scrambled texts. *Psychological Review*, *128*(5), 803–823. https://doi.org/10.1037/rev0000268

Rahnev, D., & Denison, R. N. (2018). Suboptimality in perceptual decision making. *Behavioral and Brain Sciences*, *41*, Article e223. https://doi.org/10.1017/S0140525X18000936

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*(2), 59–108. https://doi.org/10.1037/0033-295X.85.2.59

Ratcliff, R., & Rouder, J. N. (1998). Modeling response times for two-choice decisions. *Psychological Science*, *9*(5), 347–356. https://doi.org/10.1111/1467-9280.00067

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. https://doi.org/10.1016/j.tics.2016.01.007

Ratcliff, R., Thapar, A., Gomez, P., & McKoon, G. (2004). A diffusion model analysis of the effects of aging in the lexical-decision task. *Psychology and Aging*, *19*(2), 278–289. https://doi.org/10.1037/0882-7974.19.2.278

Ratcliff, R., & Van Dongen, H. P. (2011). Diffusion model for one-choice reaction-time tasks and the cognitive effects of sleep deprivation. *Proceedings of the National Academy of Sciences of the United States of America*, *108*(27), 11285–11290. https://doi.org/10.1073/pnas.1100483108

Reichle, E. D., Liversedge, S. P., Drieghe, D., Blythe, H. I., Joseph, H. S., White, S. J., & Rayner, K. (2013). Using E–Z reader to examine the concurrent development of eye-movement control and reading skill. *Developmental Review*, *33*(2), 110–149. https://doi.org/10.1016/j.dr.2013.03.001

Roberts, S., & Pashler, H. (2000). How persuasive is a good fit? A comment on theory testing. *Psychological Review*, *107*(2), 358–367. https://doi.org/10.1037/0033-295X.107.2.358

Robinaugh, D. J., Haslbeck, J. M., Ryan, O., Fried, E. I., & Waldorp, L. J. (2021). Invisible hands and fine calipers: A call to use formal theory as a toolkit for theory construction. *Perspectives on Psychological Science*, *16*(4), 725–743. https://doi.org/10.1177/1745691620974697

Sanborn, A. N. (2017). Types of approximation for probabilistic cognition: Sampling and variational. *Brain and Cognition*, *112*, 98–101. https://doi.org/10.1016/j.bandc.2015.06.008

Sanborn, A. N., Griffiths, T. L., & Navarro, D. J. (2010). Rational approximations to rational models: Alternative algorithms for category learning. *Psychological Review*, *117*(4), 1144–1167. https://doi.org/10.1037/a0020511

Sanders, C. A., & Nosofsky, R. M. (2020). Training deep networks to construct a psychological feature space for a natural-object category domain. *Computational Brain & Behavior*, *3*(3), 229–251. https://doi.org/10.1007/s42113-020-00073-z

Sandra, D. A., & Otto, A. R. (2018). Cognitive capacity limitations and need for cognition differentially predict reward-induced cognitive effort expenditure. *Cognition*, *172*, 101–106. https://doi.org/10.1016/j.cognition.2017.12.004

Schumacher, L., Bürkner, P.-C., Voss, A., Köthe, U., & Radev, S. T. (2023). Neural superstatistics for Bayesian estimation of dynamic cognitive models. *Scientific Reports*, *13*(1), Article 13778. https://doi.org/10.1038/s41598-023-40278-3

Schwetlick, L., Backhaus, D., & Engbert, R. (2023). A dynamical scan-path model for task-dependence during scene viewing. *Psychological Review*, *130*(3), 807–840. https://doi.org/10.1037/rev0000379

Shadlen, M. N., Britten, K. H., Newsome, W. T., & Movshon, J. A. (1996). A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *Journal of Neuroscience*, *16*(4), 1486–1510. https://doi.org/10.1523/JNEUROSCI.16-04-01486.1996

Shi, L., Griffiths, T. L., Feldman, N. H., & Sanborn, A. N. (2010). Exemplar models as a mechanism for performing Bayesian inference. *Psychonomic Bulletin & Review*, *17*(4), 443–464. https://doi.org/10.3758/PBR.17.4.443

Shiffrin, R. M., Lee, M. D., Kim, W., & Wagenmakers, E.-J. (2008). A survey of model evaluation approaches with a tutorial on hierarchical Bayesian methods. *Cognitive Science*, *32*(8), 1248–1284. https://doi.org/10.1080/03640210802414826

Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, *69*(1), 99–118. https://doi.org/10.2307/1884852

Simon, H. A. (1976). From substantive to procedural rationality. In T. Kastelein, S. Kuipers, W. Nijenhuis, & G. Wagenaar (Eds.), *25 years of economic theory* (pp. 65–86). Springer. https://doi.org/10.1007/978-1-4613-4367-7_6

Simon, H. A. (1992). What is an "explanation" of behavior? *Psychological Science*, *3*(3), 150–161. https://doi.org/10.1111/j.1467-9280.1992.tb00017.x

Smith, P. L. (1995). Psychophysically principled models of visual simple reaction time. *Psychological Review*, *102*(3), 567–593. https://doi.org/10.1037/0033-295X.102.3.567

Smith, P. L. (2000). Stochastic dynamic models of response time and accuracy: A foundational primer. *Journal of Mathematical Psychology*, *44*(3), 408–463. https://doi.org/10.1006/jmps.1999.1260

Smith, P. L., & Ratcliff, R. (2004). Psychology and neurobiology of simple decisions. *Trends in Neurosciences*, *27*(3), 161–168. https://doi.org/10.1016/j.tins.2004.01.006

Speekenbrink, M. (2016). A tutorial on particle filters. *Journal of Mathematical Psychology*, *73*, 140–152. https://doi.org/10.1016/j.jmp.2016.05.006

Spektor, M. S., & Kellen, D. (2018). The relative merit of empirical priors in non-identifiable and sloppy models: Applications to models of learning and decision-making: Empirical priors. *Psychonomic Bulletin & Review*, *25*(6), 2047–2068. https://doi.org/10.3758/s13423-018-1446-5

Stan Development Team. (2024a). *RStan: The R interface to Stan* (R package Version 2.32.6) [Computer software]. https://mc-stan.org/

Stan Development Team. (2024b). *Stan modeling language users guide and reference manual* (Version 2.35) [Computer software]. https://mc-stan.org/docs/

Starns, J. J., & Ratcliff, R. (2010). The effects of aging on the speed–accuracy compromise: Boundary optimality in the diffusion model. *Psychology and Aging*, *25*(2), 377–390. https://doi.org/10.1037/a0018022

Starns, J. J., & Ratcliff, R. (2012). Age-related differences in diffusion model boundary optimality with both trial-limited and time-limited tasks. *Psychonomic Bulletin & Review*, *19*, 139–145. https://doi.org/10.3758/s13423-011-0189-3

Stewart, N., Chater, N., & Brown, G. D. (2006). Decision by sampling. *Cognitive Psychology*, *53*(1), 1–26. https://doi.org/10.1016/j.cogpsych.2005.10.003

Suárez, M., & Pero, F. (2019). The representational semantic conception. *Philosophy of Science*, *86*(2), 344–365. https://doi.org/10.1086/702029

Suchow, J. W., Bourgin, D. D., & Griffiths, T. L. (2017). Evolution in mind: Evolutionary dynamics, cognitive processes, and Bayesian inference. *Trends in Cognitive Sciences*, *21*(7), 522–530. https://doi.org/10.1016/j.tics.2017.04.005

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. MIT Press.

Szollosi, A., Donkin, C., & Newell, B. R. (2023). Toward nonprobabilistic explanations of learning and decision-making. *Psychological Review*, *130*(2), 546–568. https://doi.org/10.1037/rev0000355

ter Braak, C. J., & Vrugt, J. A. (2008). Differential evolution Markov chain with snooker updater and fewer chains. *Statistics and Computing*, *18*, 435–446. https://doi.org/10.1007/s11222-008-9104-9

Turner, B. M., Forstmann, B. U., & Steyvers, M. (2019). *Joint models of neural and behavioral data*. Springer. https://doi.org/10.1007/978-3-030-03688-1

Ullman, T. D., Goodman, N. D., & Tenenbaum, J. B. (2012). Theory learning as stochastic search in the language of thought. *Cognitive Development*, *27*(4), 455–480. https://doi.org/10.1016/j.cogdev.2012.07.005

Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review*, *108*(3), 550–592. https://doi.org/10.1037/0033-295X.108.3.550

Van den Berg, R., Shin, H., Chou, W.-C., George, R., & Ma, W. J. (2012). Variability in encoding precision accounts for visual short-term memory limitations. *Proceedings of the National Academy of Sciences of the United States of America*, *109*(22), 8780–8785. https://doi.org/10.1073/pnas.1117465109

van Maanen, L., & Miletić, S. (2021). The interpretation of behavior-model correlations in unidentified cognitive models. *Psychonomic Bulletin & Review*, *28*(2), 374–383. https://doi.org/10.3758/s13423-020-01783-y

Van Ravenzwaaij, D., Cassey, P., & Brown, S. D. (2018). A simple introduction to Markov chain Monte–Carlo sampling. *Psychonomic Bulletin & Review*, *25*(1), 143–154. https://doi.org/10.3758/s13423-016-1015-8

van Ravenzwaaij, D., Donkin, C., & Vandekerckhove, J. (2017). The EZ diffusion model provides a powerful test of simple empirical effects. *Psychonomic Bulletin & Review*, *24*, 547–556. https://doi.org/10.3758/s13423-016-1081-y

van Rooij, I., & Baggio, G. (2021). Theory before the test: How to build high-verisimilitude explanatory theories in psychological science. *Perspectives on Psychological Science*, *16*(4), 682–697. https://doi.org/10.1177/1745691620970604

Vickers, D. (1970). Evidence for an accumulator model of psychophysical discrimination. *Ergonomics*, *13*(1), 37–58. https://doi.org/10.1080/00140137008931117

Vickers, D., & Lee, M. D. (1998). Dynamic models of simple judgments: I. Properties of a self-regulating accumulator module. *Nonlinear Dynamics,* *Psychology, and Life Sciences*, *2*(3), 169–194. https://doi.org/10.1023/A:1022371901259

Vincent, B. T., & Baddeley, R. J. (2003). Synaptic energy efficiency in retinal processing. *Vision Research*, *43*(11), 1285–1292. https://doi.org/10.1016/S0042-6989(03)00096-8

Visser, I. (2011). Seven things to remember about hidden Markov models: A tutorial on Markovian models for time series. *Journal of Mathematical Psychology*, *55*(6), 403–415. https://doi.org/10.1016/j.jmp.2011.08.002

Voss, A., Rothermund, K., & Voss, J. (2004). Interpreting the parameters of the diffusion model: An empirical validation. *Memory & Cognition*, *32*(7), 1206–1220. https://doi.org/10.3758/BF03196893

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. https://doi.org/10.1111/cogs.12101

Waldron, E. M., & Ashby, F. G. (2001). The effects of concurrent task interference on category learning: Evidence for multiple category learning systems. *Psychonomic Bulletin & Review*, *8*(1), 168–176. https://doi.org/10.3758/BF03196154

Watanabe, T., & Sasaki, Y. (2015). Perceptual learning: Toward a comprehensive theory. *Annual Review of Psychology*, *66*, 197–221. https://doi.org/10.1146/annurev-psych-010814-015214

Wilson, R. C., Geana, A., White, J. M., Ludvig, E. A., & Cohen, J. D. (2014). Humans use directed and random exploration to solve the explore–exploit dilemma. *Journal of Experimental Psychology: General*, *143*(6), 2074–2081. https://doi.org/10.1037/a0038199

Wu, C. M., Schulz, E., Speekenbrink, M., Nelson, J. D., & Meder, B. (2018). Generalization guides human exploration in vast decision spaces. *Nature Human Behaviour*, *2*(12), 915–924. https://doi.org/10.1038/s41562-018-0467-4

Xu, P., Lu, Z.-L., Qiu, Z., & Zhou, Y. (2006). Identify mechanisms of amblyopia in gabor orientation identification with external noise. *Vision Research*, *46*(21), 3748–3760. https://doi.org/10.1016/j.visres.2006.06.013

Yu, S., Pleskac, T. J., & Zeigenfuse, M. D. (2015). Dynamics of postdecisional processing of confidence. *Journal of Experimental Psychology: General*, *144*(2), 489–510. https://doi.org/10.1037/xge0000062

Zacksenhouse, M., Bogacz, R., & Holmes, P. (2010). Robust versus optimal strategies for two-alternative forced choice tasks. *Journal of Mathematical Psychology*, *54*(2), 230–246. https://doi.org/10.1016/j.jmp.2009.12.004

Zhang, J., & Rowe, J. B. (2014). Dissociable mechanisms of speed–accuracy tradeoff during visual perceptual learning are revealed by a hierarchical drift-diffusion model. *Frontiers in Neuroscience*, *8*, Article 69. https://doi.org/10.3389/fnins.2014.00069

Zhang, W., & Luck, S. J. (2008). Discrete fixed-resolution representations in visual working memory. *Nature*, *453*(7192), 233–235. https://doi.org/10.1038/nature06860

Zhao, B., Lucas, C. G., & Bramley, N. R. (2024). A model of conceptual bootstrapping in human cognition. *Nature Human Behaviour*, *8*(1), 125–136. https://doi.org/10.1038/s41562-023-01719-1

Zou, W., & Bhatia, S. (2021). Learning new categories for natural objects. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *43*, 2575–2581. https://escholarship.org/uc/item/89b6z6ns

Zylberberg, A., Barttfeld, P., & Sigman, M. (2012). The construction of confidence in a perceptual decision. *Frontiers in Integrative Neuroscience*, *6*, Article 79. https://doi.org/10.3389/fnint.2012.00079

(*Appendices follow*)

Appendix A

Local Sampling and Comparison Algorithm

Here, we provide a more detailed description of the local sampling and comparison algorithm used to generate trajectories in the cognitive parameter space (Figures 4 and 5 in the main text). The objective function used for Figure 4 is arbitrary and defined as:

$$f(\Psi_1, \Psi_2) = A \exp\left(-\frac{1}{2(1-\rho^2)} \times \frac{(\Psi_1 - \mu_{\Psi_1})^2}{\sigma^2_{\Psi_1}} - 2\rho \times \frac{(\Psi_1 - \mu_{\Psi_1})(\ln \Psi_2 - \mu_{\Psi_2})}{\sigma_{\Psi_1}\sigma_{\Psi_2}} + \frac{(\ln \Psi_2 - \mu_{\Psi_2})^2}{\sigma^2_{\Psi_2}}\right), \quad \text{(A1)}$$

with the following constants: $A = 100$; $\mu_{\Psi_1} = \mu_{\Psi_2} = 0.5$; $\sigma_{\Psi_1} = 1$; $\sigma_{\Psi_2} = 1.5$; $\rho = 0.75$. The objective function for the expanded judgment paradigm simulated for Figure 5, is determined by the parameters of the behavioral task, given below in Appendix C.

Agents are initialized at random points in the cognitive parameter space. At any one point, the agent interacts with the environment and estimates the objective value V with some uncertainty, that is, $p(V = v)$. To reduce the uncertainty, the agent integrates the objective estimates over a window of W trials, to obtain $\bar{v}$. With the total number of trials, N, being limited, the sampling window should be small enough to allow for some exploration, that is, $W \ll N$.

If the estimated objective value, $\bar{v}$, after W trials is low, the agent is likely to jump to a different location in the space. If the value is high, the agent is likely to stay in the current location and perform another set of W trials. Specifically, the probability of a jump, π depends on the estimated objective value through: $\pi(\bar{v}) = 1 - \frac{\bar{v}}{\gamma}$, $0 \leq \pi \leq 1$, where γ is the threshold objective value at which search can be stopped. The precise form of this function is not too important and the agent does not need precise knowledge of the maximum attainable objective value. What matters is that they have some crude sense of whether a location is good or bad, and that this estimate bears an ordinal relation with the true objective value. Whether the agent jumps is a Bernoulli random sample with $p = \pi(\bar{v})$.

If the agent jumps, a "proposal" location is drawn from a symmetrical transition distribution around their current location; in other words, the jump is essentially random, but likely to be close by. We formalized the transition distribution as a multivariate Gaussian, but again the precise implementation is probably not too important. We could have chosen a Lévy flight, a uniform distribution or some other form; the scale of the distribution is much more important, so that the jumps are of an appropriate size. After a transition, the agent samples the objective for W trials at the new location. If the estimated objective is better than the estimate from the previous location, the proposal location is "accepted" and search continues from here. If the estimated objective at the new location is worse than the estimate from the previous location, the proposal location is "rejected" and the agent reverts to the previous location. After an accept/reject decision, search continues as before: jumping with a probability that depends inversely on the estimated objective value at the new location (if the proposal was accepted) or the previous location (if the proposal was rejected).

One way an agent may simplify the search process is through satisficing. This behavior is controlled by the stopping criterion γ, which controls when the jump probability goes down to 0. Satisficing then involves setting γ to a value below the peak of the objective function. Again, the precise implementation details do not matter a great deal; the key ingredient is that agents have some way of sensing that a particular objective value is good enough and that this sensation drives down the probability of jumping. In our simulations, γ was always set to the true maximum of the objective (i.e., 100 for the arbitrary objective from Figure 4A; the peak reward rate derived from the dynamic programming solution in Figure 5). In other words, our agents were maximizers. Note that as a result of noise in the objective estimates, it can (quite easily) happen that the estimated value at a given location exceeds the stopping criterion. In this case, the jump probability is set to 0 and the agent does not move. However, the next window of trials at the same location might easily generate an objective estimate lower than γ, in which case there is some nonzero probability of a transition. Therefore, noise in the objective values may result in the agent moving away from a good (or even optimal) spot.

For the example trajectory shown in Figure 4B, the agent arrives at a good spot on the surface and is able to spend just over half their time there. However, this outcome is not guaranteed: every run of the algorithm will produce a different result, as demonstrated in Figure 5 (top row). The performance of the algorithm will depend on its parameters (W, σ_{jump}, γ), the nature of the objective surface and the amount of noise in the objective estimates. Variation in some or all of these components will produce variation in behavior over time (as search progresses) and variation between individuals.

The algorithm described here is similar to standard MCMC sampling algorithms such as Metropolis–Hastings, except that: (a) it operates on estimates derived from several (i.e., W) trials rather than just a single evaluation of the objective (target); (b) a proposal is generated and evaluated with a (jump) probability that is inversely related to the estimated value of the objective; and (c) the agent simply rejects proposals with a lower estimated objective value than the previous position (rather than accepts them with some probability). The first modification ensures that the agent has a more reliable estimate of the objective (analogous to the synthetic likelihood approach in approximate Bayesian computation; Hartig et al., 2011; Palestro et al., 2018). The second modification is necessary, because we assume that the agent actually has to move to a new location and interact with the environment in order to gain information about the objective at that location (i.e., they cannot estimate the objective elsewhere through, say, mental simulation). Such a move is potentially costly, because the new location may have a lower objective value and the agent will have to spend W trials there to find out. Therefore, rather than always generating proposals (as in standard MCMC), we want to do so adaptively, depending on how good or bad the current location is. The third modification is appropriate in this setting, because the agent's goal is not to sample the full objective function with a probability that is proportional to its value—the goal is simply to find the best (or good enough) location and spend as much time there as possible.

(*Appendices continue*)

Appendix B

Static and Dynamic Parameter Estimates for the Mixture-of-Gaussians Example

For the simple model used to generate Figure 4 in the main text, we used the following procedures to estimate parameters.

Static Parameter Estimates

For the static model, we simply need to estimate the mean and standard deviation of a Gaussian distribution, that is, $y \sim \mathcal{N}(\mu, \sigma)$. We adopted the following priors for the parameters μ and σ:

$$\mu \sim \mathcal{N}(0, 1)$$
$$\sigma^2 \sim \text{inv} - \chi^2(5, 5), \tag{B1}$$

we approximated the posteriors through sampling, using the Stan probabilistic programming language (Stan Development Team, 2024b) and the RStan interface (Stan Development Team, 2024a). Parameters were bound to the intervals $\mu \in [-3, 3]$ and $\sigma \in$ [1e-6, 5]. We used four chains with 1k postwarm-up iterations, giving 4k posterior samples for each parameter.

Particle Filter

Latent states correspond to the time-evolving parameters of the model: For each trial, we assume there is an underlying state $\theta_t = \{\mu_t, \sigma_t\}$. We implemented a basic bootstrap particle filter to approximate the posterior distribution of latent states at each time point (Durbin & Koopman, 2012; Gordon et al., 1993; Speekenbrink, 2016). We used $J = 2{,}000$ particles. The initial locations of the particles are drawn from a bivariate Gaussian distribution with prior mean $\mu_0 = \{0, 2.5\}$ and unit variances $\Sigma_0 = \mathbb{I}$, where $\mathbb{I}$ denotes the identity matrix. Particles are assigned weights of $w_t^{(j)} = J^{-1}$. For each incoming observation y_t, with $t = 1, \ldots, T$, the distribution of particles then evolves as follows:

- Propagate particles through a *transition distribution*, $p(\theta_t^{(j)}|\theta_{t-1}^{(j)})$. We chose a bivariate Gaussian with $\Sigma_{\text{trans}} = \text{diag}(0.5, 0.5)$.
- For $j = 1, \ldots, J$, compute the likelihood of the observation, that is, $p(y_t|\theta_t^{(j)})$.
- Weight update: $\tilde{w}_t^{(j)} = w_{t-1}^{(j)} p(y_t|\theta_t^{(j)})$. This step assigns higher weights to particles that are more consistent with the new observation. Normalize the weights: $w_t^{(j)} = \frac{\tilde{w}_t^{(j)}}{\sum_{j=1}^{J} \tilde{w}_t^{(j)}}$.
- To avoid weight degeneracy, resample the weights if they are concentrated on too few particles, that is, if $\frac{1}{\sum_{j=1}^{J} (w_t^{(j)})^2} < 0.5J$. Resampling eliminates particles with low weights and replicates particles with higher weights. If the resampling step was performed, set $w_t^{(j)} = J^{-1}$.

Particles were constrained to the intervals $\mu_t \in [-3, 3]$ and $\sigma_t \in$ [1e-6, 5]. At each iteration, the collection of particles and their associated weights may be used to compute various (weighted) quantities of interest. In Figure 4D, we show the means, along with the 2.5th and 97.5th percentiles, of the marginal posteriors.

Appendix C

Expanded Judgment Task Simulations

Task Environment

The agent is presented with a sequence of binary evidence samples, x with $x_i \in \{-1, 1\}$. That is, each evidence sample points to one or the other decision alternative, with a consistent bias during any one trial, that is, $x_i \sim \text{Bernoulli}(0.5 \pm \varepsilon)$. The agent terminates the evidence stream by committing to a decision alternative. The accuracy of that decision determines whether they are rewarded or penalized. The next trial then starts after a delay. If there is time remaining in the task, the next trial is then presented. For the agents simulated in the current article, we adopted the task parameters listed in Table C1.

Decision Model

Malhotra et al. (2018) described the optimal decision policy for this task, derived through dynamic programming (Bellman, 1957). This policy prescribes what the agent should do for each possible combination of time (number of evidence samples observed) and accumulated evidence. A given decision policy may be represented as a tripartite division of the (time, evidence)-space: a region where the agent should go for option A+, a region where the agent should go for option A−, and a region in between where the agent should wait for further evidence. The boundaries separating the "go" from the "wait" regions are symmetric and may be summarized conveniently by an intercept and slope. Let α and β represent, respectively, the intercept (in evidence units) and gradient (in radians) of the positive decision boundary. The long-run expected reward rate can then be computed for each (α, β) combination (Malhotra et al., 2017). This 2D objective function, given the current task parameters, is illustrated in Figure C1A.

Our simulated agents essentially adopt this model for making their decisions. On a particular trial t, their (positive) evidence boundary at the kth sample is given by: $b^{+}(k) = \max(\alpha_t + k \tan \beta_t, 0)$;

Table C1
Parameters of the Simulated Expanded Judgment Task

Parameter	Value
Interstimulus interval	0.2 s
Intertrial interval (correct)	3 s
Intertrial interval (error)	3 s
Monetary reward (correct)	1
Monetary penalty (error)	−2
Evidence bias, ε	0.2
Within-trial response deadline	10 s
Task duration	480 s

(*Appendices continue*)

Figure C1

Reward Rates in a Simulated Expanded Judgment Paradigm

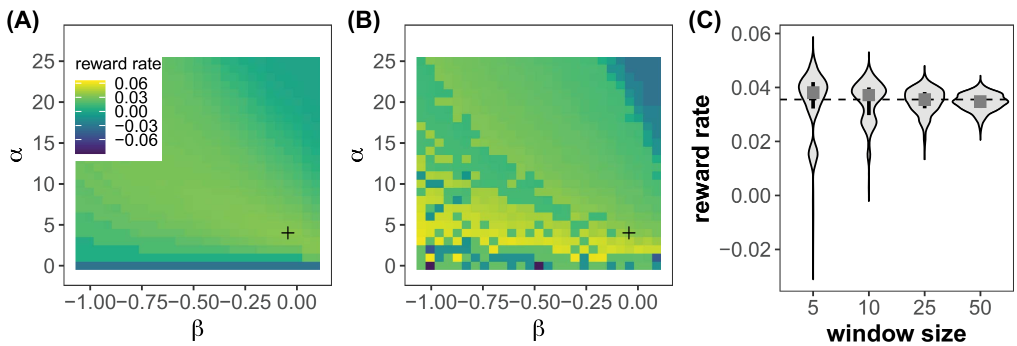

Note. (A) Long-run reward rate from dynamic programming. (B) Noisy reward rates estimated over a window of five simulated trials. At each point in the state space, agents were presented with the same (initial) sequence of evidence samples. Due to differences in the policies, these samples will be distributed differently across trials. For instance, agents who sample more (e.g., due to a high intercept) will see evidence samples that were never seen by agents who sampled much less (e.g., due to a low intercept). (C) Distribution of reward rates over 1,000 windows of varying sizes. See the online article for the color version of this figure.

the negative boundary is simply $b^{-}(k) = -b^{+}(k)$. Within a trial, the agent integrates the evidence perfectly, so that the decision variable after k samples is: $z_k = \sum_{i=1}^{k} x_i$. However, to introduce a realistic amount of variability in behavior, we assume that the agent makes a probabilistic "wait"/"go" decision after each sample. That is, they draw a sample from a Gaussian distribution centered on z_k: $z^* \sim \mathcal{N}(z_k, \eta)$, where η corresponds to the decision noise. If $z^* \geq b^{+}(k)$, the agent goes for option A+; if $z^* \leq b^{-}(k)$, the agent goes for option A−; if $b^{+}(k) = b^{-}(k) = 0$, a random choice is made between A+ and A−; otherwise, the agent waits for the next evidence sample. The decision noise parameter generates variability in choice and decision times—it ensures that the agent does not always make the same choice when presented with the same accumulated evidence at the same point in time. This simple decision model is characterized by three parameters: $\theta = \{\alpha, \beta, \eta\}$.

Navigating Cognitive Parameter Space

Our simulated agents have two parameters under strategic control: the height and gradient of their decision boundary. Decision noise is constant and set to $\eta = 1$ throughout. Figure C1A shows the long-run expected objective values across this 2D cognitive parameter space for an agent without any decision noise. There is a broad ridge in this space where good reward rates may be obtained: Higher boundaries (which would slow people down) may be compensated for by having them collapse more steeply (which forces a decision by effectively imposing a deadline; Hawkins & Heathcote, 2021). Each agent is initialized in a random position within the cognitive parameter space, with $\alpha \in [0, 25]$ and $\beta \in [-1.047, 0.087]$ radians. However, they *never really* experience the long-run expected reward rate for a given threshold height and gradient. Aside from the particular decision policy adopted, the experienced reward rate depends strongly on the stochastic evidence sequences that are presented, the amount of decision noise and the window over which the agent estimates the reward rate.

Figure C1B and C1C give an indication of the variability in reward rate an agent might actually experience when estimating the reward rate from a window of W trials. Panel B shows the experienced reward rate across the cognitive parameter space with $W = 5$ trials. Clearly, the globally "true" optimal policy (marked with a +) is not necessarily optimal for a particular window of trials. Panel C shows the distribution of reward rates experienced for repeated windows of W trials at the same location (in this case, the globally optimal location). The long-run expected reward rate is indicated by the horizontal dashed line. We show these distributions for a range of window sizes. As expected, larger windows yield estimated reward rates that are more tightly clustered around the long-run expected value. However, these simulated agents generally performed between 100 and 300 trials over the course of the experiment, so a window of 50 trials would not allow for much exploration. The variability in experienced reward rate (particularly for smaller, more realistic window sizes) is very pronounced and will make it extremely difficult to find the optimal policy in a limited duration (Evans & Brown, 2017; Malhotra et al., 2017).

We are not so much concerned with whether agents have sufficient information to find the optimal position in the cognitive parameter space. Rather, our aim was simply to generate trajectories in the state space and assess the resulting parameter estimates. To generate the trajectories for the nonstationary agents shown in Figure 5, we followed the local sampling and comparison routine described in Appendix A. Specifically, we adopted the following parameters for the algorithm: $W = 5$; the standard deviation of the jump distribution was set to 10% of the permissible range of the parameters, that is, $\Sigma_{\text{jump}} = \text{diag}(6.25, 0.013)$ (for α and β, respectively; note these are diagonal variances); $\gamma = 0.036$ (close to the true optimal reward rate). The stationary agent was given a random (initial) position, just like the nonstationary agents, but was simply not allowed to move.

Parameter Estimation

Static Parameter Estimates

Given the decision model described above and given a particular decision boundary and accumulated evidence after sample k, the

(*Appendices continue*)

probability for the three possible actions $a \in \{A+, A-, \text{wait}\}$ are:

$$p(a = A+) = 1 - \Phi(b^+(k), z_k, \eta)$$
$$p(a = A-) = \Phi(b^-(k), z_k, \eta)$$
$$p(a = \text{wait}) = 1 - p(a = A+) - p(a = A-), \quad (C1)$$

where Φ denotes the cumulative Gaussian distribution function. The likelihood of an observed sequence of K actions on trial t is now simply: $\mathcal{L}_t = \prod_{k=1}^{K_t} p(a = a_k^* | \theta)$, where a_k^* indicates the chosen action after sample k. Across the entire set of trials $t = 1, \ldots, T$, likelihoods may combined through: $\mathcal{L} = \prod_{t=1}^{T} \mathcal{L}_t$. As usual, it is more convenient to compute the log-likelihood: $\ln \mathcal{L} = \sum_{t=1}^{T} \sum_{k=1}^{K_t} \ln p(a = a_k^* | \theta)$.

We adopted the following priors and parameter transformations:

$$\alpha^*, \beta^*, \eta^* \sim \mathcal{N}(0, 1)$$
$$\alpha = \exp(\alpha^* + 1.5)$$
$$\beta = 0.5\beta^*$$
$$\eta = \exp(\eta^* + 1). \quad (C2)$$

For α and η, the prior and subsequent transformation result in right-skewed distributions, bound at 0 from below and with most of their mass between 0 and 10 (with a longer tail for α). For β, the prior is Gaussian and centered on a gradient of 0, that is, a flat decision boundary. Posterior densities for the parameters were again approximated using Stan, using four chains with 5k postwarm-up iterations. After thinning the chains by a factor of 5, we were left with 4k posterior samples for each parameter.

Particle Filter

We adopted the perspective of a modeler who does not know which parameters change over time and which ones do not. So all parameters were allowed to evolve: $\theta_t = \{\alpha_t, \beta_t, \eta_t\}$. We used the bootstrap particle filter described in Appendix B, again with $J = 2{,}000$ particles. Initial particle locations were drawn from a multivariate Gaussian prior with $\mu_0 = \{5, 0, 2\}$ and variances $\Sigma_0 = \text{diag}(25, 0.16, 9)$. The prior needs to be wide enough so that there will be some particles that account for the early data. The transition distribution was also a multivariate Gaussian, with variances $\Sigma_{\text{trans}} = \text{diag}(1, 0.0025, 0.04)$. The scale of this transition distribution determines how quickly particles can adapt when there has been a change in the latent state (i.e., if particles can only move a small amount, it will take a long time to adapt to a large jump in the state space). Parameters were constrained to the intervals $\alpha_t \in [0.01, 25]$, $\beta_t \in [-1.05, 0.175]$, and $\eta_t \in [0.01, 5]$.

Scaling Constraint

As discussed in the main text, many cognitive models can only be identified up to a scaling constraint. Here, we consider a situation in which the parameter we use to fix the scale of the model, actually varies over time. At first sight, this problem might appear less pressing for the expanded judgment paradigm and model we have used as our test bed. The reason is that the model operates on the *actual* evidence samples presented to the agent (i.e., binary samples, $x_i \in \{-1, 1\}$), which anchor the scale of the parameters. However, the assumption that the agent integrates the evidence as is, effectively implies that they apply a gain of $\lambda = 1$ when translating the external evidence to an internal representation. Alternative values for the gain may be adopted without altering the model predictions, by scaling the other model parameters that can be represented in evidence units (i.e., the threshold height and noise).

We simulated a scenario in which data are generated with a time-varying gain (periodic switches between $\lambda = 1$ and $\lambda = 2$). This temporal pattern was unrelated to strategic exploration of the cognitive parameter space. We simply aimed to have a strong, recognizable temporal signature to detect in the parameter estimates. Figure C2A illustrates particle filter estimates for two models that differ in the choice of constant scaling parameter, fit to the data from the same simulated agent. In both cases, the scaling parameter is set to a value that does not match the ground truth. The left column shows the estimates when we fixed the gain as a scaling parameter to a value of $\lambda = 0.5$. The low adopted value of the gain scales the estimated threshold height and the noise: the effects of a larger gain in the data generation (i.e., faster evidence accumulation) can only be accommodated by lowering the distance to the threshold and reducing the noise. Moreover, the fluctuations in the ground truth gain seem to be compensated for by fluctuations in the threshold height (most clearly visible from about trial 175 onward). Note that in other situations (models, tasks), the temporal variation may be absorbed by a combination of parameters, which would inflate their covariance. The right column of Figure C2A shows a fit in which we have selected the decision noise as a scaling parameter, with $\eta = 0.5$. Here, we recover the temporal fluctuation in the gain accurately, which in turn reduces the periodic fluctuation in the threshold height. The close fit with the ground truth should not be overemphasized; with different (misspecified) values for η, the model misses the ground truth. However, importantly, the parameter estimates remained closely correlated with it.

Figure C2B shows a measure of the predictive accuracy that comes “for free” with the particle filter: the posterior state after observation y_{t-1} is used to predict the current state at time t, and the likelihood of y_t under that predicted current state is therefore a *predictive density*. We show the log expected predictive densities, averaged over trials, for a random sample of 12 simulated agents. In all cases, the predictive densities were larger for the model in which decision noise was (correctly) treated as constant. These simulation results demonstrate that, in principle, a model-selection approach may be viable when different choices regarding time-varying and fixed scaling parameters result in different models (i.e., that do not mimic each other perfectly). No doubt more sophisticated parameter estimation and model-selection approaches may give better results, and may be necessary for more complex models and task environments. We expect this topic will be an important avenue for future research.

(*Appendices continue*)

Figure C2

Alternative Model Specifications for Fitting Data Generated With a Time-Varying Gain

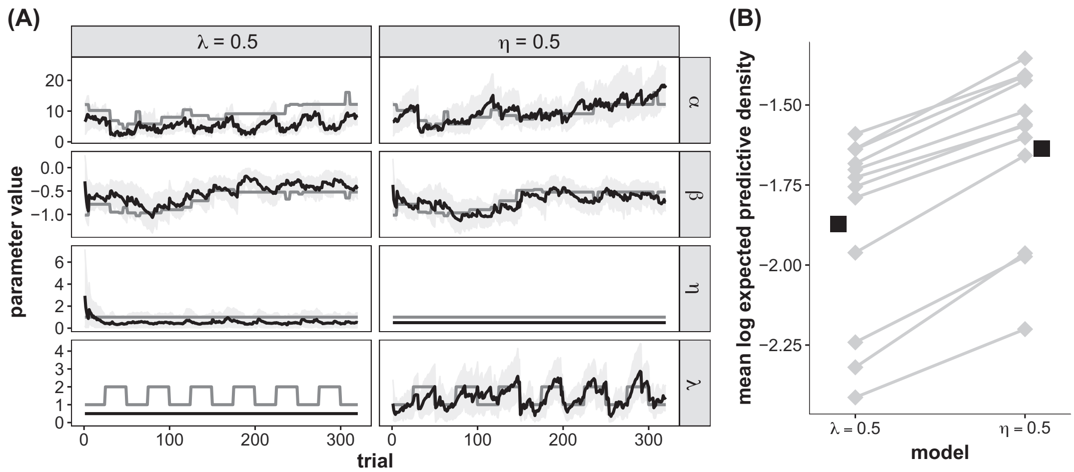

Note. The agent's trajectory in the two-dimensional (α, β) space was generated using the same method as described above. The only difference was that the incoming evidence was multiplied with a gain that switched between values of 1 and 2 periodically (every 25 trials). (A) Left column: Particle filter estimates for a three-parameter model, where the gain is treated as a scaling parameter and set to a value of 0.5. Right column: Particle filter estimates for a three-parameter model, where the decision noise is treated as a scaling parameter and set to a value of 0.5. The solid grey lines show the ground truth values; the solid black lines show the estimated values (and the fixed values for the scaling parameters). For the $\eta = 0.5$ model, the gain estimates were constrained to the interval $\lambda \in [0.01, 5]$. Initial values for this parameter were drawn from $\mathcal{N}(0.5, 1)$; the transition distribution around the current gain values had a standard deviation of 0.2. The particle filter was allowed to explore a wider range of values for $\alpha \in [0.01, 50]$, because higher gains might require higher decision thresholds to capture the data. (B) Log expected predictive density (over particles), averaged across trials, for 12 simulated agents. Black squares show the sample means.

Received July 19, 2023
Revision received October 24, 2024
Accepted November 11, 2024 ■