A Bayesian Hierarchical Model of Trial-To-Trial Fluctuations in Decision Criterion

Robin Vloeberghs[1], Anne E. Urai[2], Kobe Desender[1,†], and Scott W. Linderman[3,†]

Abstract

Classical decision models assume that the parameters giving rise to choice behavior are stable, yet emerging research suggests these parameters may fluctuate over time. Such fluctuations, observed in neural activity and behavioral strategies, have significant implications for understanding decision-making processes. However, empirical studies on fluctuating human decision-making strategies have been limited due to the extensive data requirements for estimating these fluctuations. Here, we introduce hMFC (Hierarchical Model for Fluctuations in Criterion), a Bayesian framework designed to estimate slow fluctuations in the decision criterion from limited data. We first showcase the importance of considering fluctuations in decision criterion: incorrectly assuming a stable criterion gives rise to apparent history effects and underestimates perceptual sensitivity. We then present a hierarchical estimation procedure capable of reliably recovering the underlying state of the fluctuating decision criterion with as few as 500 trials per participant, offering a robust tool for researchers with typical human datasets. Critically, hMFC does not only accurately recover the state of the underlying decision criterion, it also effectively deals with the confounds caused by criterion fluctuations. Lastly, we provide code and a comprehensive demo to enable widespread application of hMFC in decision-making research.

Keywords: decision-making, signal detection theory, decision criterion, Bayesian modeling, sequential effects

[1] Brain and Cognition, KU Leuven, Belgium
[2] Cognitive Psychology, Leiden University, The Netherlands
[3] Department of Statistics and Wu Tsai Neurosciences Institute, Stanford University, Stanford, CA, USA
[†] Correspondence: `kobe.desender@kuleuven.be` and `scott.linderman@stanford.edu`

Introduction

Every day we make a multitude of decisions: deciding when to cross the street, whether to add a newly discovered song to your playlist, or choosing between coffee or tea. Much effort has gone toward understanding the computational mechanisms underlying perceptual (Gold and Shadlen, 2007), value-based (Krajbich et al., 2010), risky (Peters and D'Esposito, 2020), and social decisions (Chen and Krajbich, 2018). A common theme in these frameworks is that decisions are formed by comparing evidence to a decision criterion. Classical models of decision-making assume that observers generate an internal representation of the relevant information, typically referred to as "evidence" or the "decision variable" (Green et al., 1966). To generate a binary choice, this decision variable is compared to an internal decision criterion. Due to the inherent noisiness of the brain or changes in internal states (e.g., changes in attention), repeated presentation of the same stimulus will not always lead to the same internal representation of that stimulus. Instead, multiple trials of the same stimulus produce a distribution of internal decision variables (fig. 1A). This inherent variability explains why participants make variable responses when repeatedly presented with the same stimulus.

A key component in these models is the decision criterion that explains biases in decision-making. When observers choose one answer more often than another, their behavior can be modelled as the result of a shift in decision criterion, independent of their sensitivity to decision-relevant information (Green et al., 1966). Empirically, the decision criterion depends on internal and external factors: for instance, it shifts in response to external feedback (Haddara and Rahnev, 2022), expectancy (Bang and Rahnev, 2017), task instructions (Kloosterman et al., 2019), monetary rewards (Frithsen et al., 2018), and the base rate of stimuli (Sánchez-Fuenzalida et al., 2023).

Although it is widely recognized that the decision criterion can be flexibly shifted depending on the environment, a common implicit assumption is that in the absence of experimental

manipulations, this (potentially biased) criterion remains constant across a series of experimental trials. Although this assumption allows computational models to be simple and tractable, it may be overly simplistic. Indeed, the idea of trial-by-trial variability in the criterion can be traced back to early work of Wickelgren (1968), and has been noted often in the literature (Benjamin et al., 2009; Ratcliff and Rouder, 1998). However, this line of work in psychophysics and mathematical psychology focuses mostly on the criterion distribution (mean and variance), rather than attempting to estimate the trial-to-trial trajectory of the fluctuating decision criterion.

More recently, the dynamical nature of cognitive processes has received renewed interest in cognitive and computational neuroscience (John et al., 2022). One reason for this trend is the availability of new computational tools to quantify trial-by-trial fluctuations from behavioral and/or neural data. Indeed, increasing evidence suggests that several computational parameters that to underlie decision-making are not stationary, but rather fluctuate over trials (Ashwood et al., 2022; Cowley et al., 2020; Gupta and Brody, 2022; Mochol et al., 2021; Roy et al., 2021). For example, Ashwood et al. (2022) demonstrated that decision behavior in mice can be described by three discrete states persisting for 10 to 100 trials, each with different psychometric parameters. Other work has focused on modeling continuous trial-by-trial changes in the overall bias (i.e., intercept; (Gupta and Brody, 2022)) and other parameters of a logistic regression model (Roy et al., 2021). Combining behavior with neural recordings, Cowley et al. (2020) found that during a multi-hour session, monkeys' hit and false alarm rate slowly drifted, and covaried with neural activity in area V4 and prefrontal cortex. In other work, Mochol et al. (2021) showed that decision behavior in monkeys is influenced by fast (previous tial's response and reward) and slow (moving window of 130 trials) biases, both encoded in the PFC, and were integrated into the current decision.

In this work, we investigate the consequences of fluctuations in decision criterion, a key parameter underlying decision-making. First, we simulate the effects of fluctuations in criterion on common measures of decision-making behavior and highlight the problems that these

fluctuations may cause in interpreting experimental data. We then present a novel hierarchical framework that allows trial-to-trial estimation of fluctuations in the decision-criterion and show that its parameters can be recovered even with a small number of trials (i.e., 500). Lastly, we show that this model solves the confounds caused by unaccounted fluctuations in decision criterion.

Why should we care about fluctuations in decision criterion?

In the next section we simulate the empirical consequences of criterion fluctuations, and especially highlight how these could impact different experimental conclusions. Rather than assuming a fixed decision criterion across time, we model the criterion fluctuations as a first-order autoregressive process (AR(1), see Gupta and Brody (2022) and Roy et al. (2021) for similar approaches). An AR(1) model describes the temporal dynamics of a time series, in this case the decision criterion, via two parameters. First, an autoregressive coefficient a captures the temporal dependency between two successive time points (x_t and x_{t-1}). In general, for an AR(1) model to have stable behavior, a must take values between -1 and +1. When $a = 0$, the resulting criterion is sampled from a normal distribution without across-trial autocorrelation. Second, a zero-mean error term with variance σ^2 causes the random fluctuations. Figure 1B illustrates how variations in both these parameters influence a simulated time series of the decision criterion. Smaller magnitudes of a (bottom left panel) make a time series more mean-reverting, with shorter fluctuations around its mean. In contrast, when a time series is less mean-reverting due to larger values for a (top left panel), it will deviate more strongly from its mean for longer periods. The variance of the error term σ^2 governs the scale of the random fluctuations and thus acts as a scaling parameter, as can be seen by comparing the left and right panels. The left and right panel time series show the exact same dynamics but merely differ in magnitude on the y-axis.

Having established a model to generate fluctuations in decision criterion, we next use simulations to show that these fluctuations can lead to biased and erroneous inferences in various

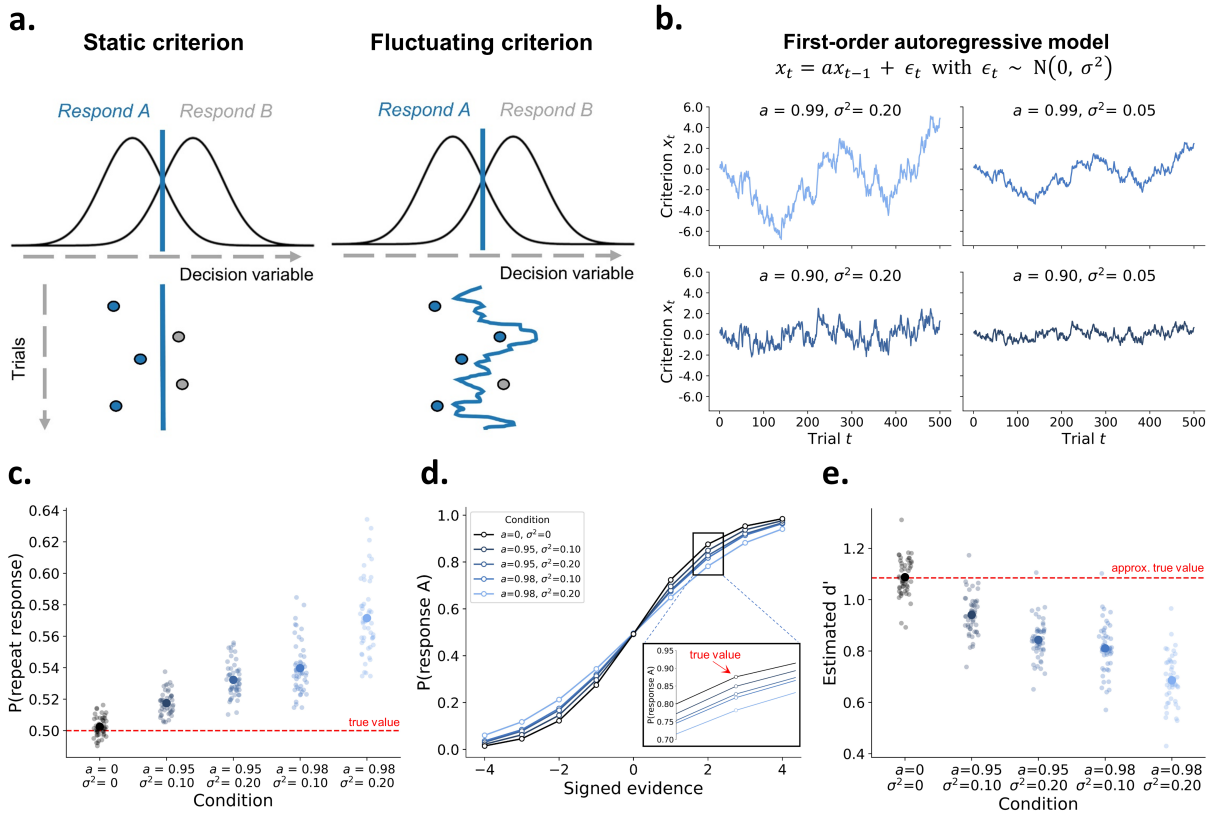

Figure 1: The consequences of fluctuations in decision criterion. **A)** Many theories of decision-making assume that choices are formed by comparing a decision variable to a criterion. According to signal detection theory, an observer will respond A or B, depending on whether the decision variable falls left or right to this criterion, respectively. Whereas this criterion is typically assumed to be fixed across a series of trials (static criterion), here we investigate the consequences of this criterion slowly fluctuating across time (fluctuating criterion). **B)** Fluctuations in criterion can be modelled as a time series using a first-order autoregressive model AR(1). An AR(1) model has two free parameters that control the temporal dependency (a) and the scale of the fluctuations (σ^2). The four panels show how the criterion changes over time under different values of these two parameters. **C)** Simulations show that it is of critical importance to consider fluctuations in decision criterion, because despite the absence of an effect of previous response (true $\beta_{\text{previous response}} = 0$), they give rise to artificial choice history effects. **D)** Likewise, while the generative psychometric slope (true $\beta_{\text{stimulus}} = 1.25$) is constant for all conditions, fluctuations in decision criterion lead to an underestimation of the slope of the psychometric function. **E)** Fluctuations in criterion lead to an underestimation of d', a popular signal-detection measure of sensitivity.

commonly used metrics that (implicitly) assume stable criteria. Below, we discuss three examples highlighting the importance of accounting for a time-varying criterion. First, the presence of criterion fluctuations can mimic a causal influence of previous responses, causing apparent sequential effects in decision-making. Second, criterion fluctuations lead to biased estimates of popular measures of sensitivity. Specifically, both the slope of the psychometric function and d' from SDT are underestimated in the presence of criterion fluctuations. Third, the ability to quantify criterion fluctuations from trial to trial will contribute to a better understanding of decision-making and its neural basis.

Criterion fluctuations can create apparent sequential effects

A wealth of studies have suggested that decisions are not isolated events, but instead show sequential effects: past trials bias the current decision. Such sequential effects have been observed in humans (Braun et al., 2018; Fischer and Whitney, 2014; Fritsche et al., 2020; Fründ et al., 2014; Samaha et al., 2019; Schlunegger and Mast, 2023; Urai et al., 2019; Van Geert et al., 2022; Yu and Cohen, 2008), primates (Lueckmann et al., 2018; Mochol et al., 2021), and rodents (Gupta and Brody, 2022; Lak et al., 2020; Busse et al., 2011; Akrami et al., 2018; Hwang et al., 2017) for a wide range of tasks. In 2AFC tasks, one example of a sequential effect is the tendency to repeat previous responses, also known as choice history bias, repetition bias, serial dependencies, perseveration, or hysteresis (Akaishi et al., 2014; Bosch et al., 2020; Fernberger, 1920; Fischer and Whitney, 2014; Fritsche et al., 2017; Urai et al., 2017, 2019; Urai and Donner, 2022; Van Geert et al., 2022, 2024). These effects are often taken as evidence that observers maintain active representations of their environment or decision history: an influence of previous responses is typically interpreted and modeled as causal effect, where observers update their criterion based on some combination of past stimuli, choices and rewards (Braun et al., 2018; Lak et al., 2020; Urai et al., 2019; Urai and Donner, 2022). For example, Treisman and Williams (1984) suggested that sequential effects can be explained by adjusting the criterion as a function of the observer's previous responses.

However, sequential effects can also arise in absence of any "systematic" adjustments, purely as the consequence of random fluctuations in the decision criterion (see also Gupta and Brody, 2022; Lak et al., 2020; Mendonça et al., 2020). To illustrate this point, we simulate five different datasets consisting of 50 agents with 5000 trials each. Critically, we simulate behavior in which the decision criterion is not stable but instead fluctuates over time according to first-order autoregressive process. To alter the dynamics of the criterion fluctuations, each dataset is generated with different values for a and σ^2. On each trial, the response is drawn from a Bernoulli distribution with the probability governed by stimulus evidence (true $\beta_{\text{stimulus}} = 1.25$), previous response (true $\beta_{\text{previous response}} = 0$), and a fluctuating value over trials representing the fluctuations in criterion. Note that adding this fluctuating value to the Bernoulli probability is mathematically identical to comparing a decision variable to a varying criterion (i.e., fluctuating criterion; fig. 1A). When this value does not fluctuate ($a = 0$ and $\sigma^2 = 0$) this corresponds to a fixed and unbiased criterion (i.e., static criterion; fig. 1A). In fig. 1C, we confirm that simulated observers using a fixed criterion do not display a tendency to repeat the previous response (i.e., no choice history bias). However, a fluctuating decision criterion causes the emergence of apparent choice history bias, despite the absence of any "systematic" criterion shifting in function of previous response (true $\beta_{\text{previous response}} = 0$) (see Lak et al. (2020, figure 2), and Gupta and Brody (2022) for similar simulations). As expected, this apparent choice history bias increases when the random fluctuations in decision criterion become stronger (i.e., increased a and σ^2). In sum, researchers interested in choice history effects cannot ignore the possibility of fluctuations in the decision criterion.

Criterion fluctuations lead to an underestimation of perceptual sensitivity metrics

Many researchers are interested in perceptual sensitivity, which are often measured using one of two popular methods: the psychometric slope and d' in signal detection theory (SDT). Both assess how well individuals can discriminate stimuli, and each has its own method of quantifying sensitivity. The psychometric slope is obtained by presenting stimuli at multiple

levels of intensity and fitting a psychometric function to the data. The slope from this function is often directly used as a measure of sensitivity, or researchers may infer a threshold of interest by evaluating the slope at a particular stimulus intensity strength (Wichmann and Hill, 2001). However, in the presence of criterion fluctuations, the slope of the psychometric function is severely biased (fig. 1D). Specifically, the slope of the psychometric function decreases with increasing strength of criterion fluctuations, despite having the same generative psychometric slope (true $\beta_{\text{stimulus}} = 1.25$). This underestimation occurs because fluctuations in the criterion lead to more variable responding to identical stimuli, which ultimately flattens the psychometric function (Fründ et al., 2014; Hautus et al., 2021; Benjamin et al., 2009). Such distortions may impact scientific conclusions, for example when determining the threshold for conscious access (Del Cul et al., 2007), or a specific level of performance (Fleming et al., 2010). In sum, researchers applying psychometric functions in their work should be aware of criterion fluctuations and its consequences.

Similarly, inspired by signal detection theory researchers often use d', the difference between z-scored hits and z-scored false alarms, as a measure of sensitivity (Hautus et al., 2021). For a given level of true sensitivity, d' remains unchanged in the presence of biased decision criteria (cf. iso-sensitivity curves; (Green et al., 1966)). Although this property holds in the presence of a fixed response bias (i.e., static criterion; fig. 1A), the property does not hold when the criterion fluctuates (Wickelgren, 1968). Specifically, criterion fluctuations lead to underestimation of the true sensitivity, biasing estimates of d'. To demonstrate this point, we again simulate behavior for different observers that all have the same true sensitivity (i.e., same β_{stimulus}), but differ in criterion fluctuation dynamics. The simulation follows the same procedure as before (fig. 1D) but now only includes one level for signed evidence. If we compute d' based on the data without criterion fluctuations, we get an estimated d' of 1.09 (fig. 1E). However, the presence of criterion fluctuations biases d' estimates downwards. In sum, researchers interested in comparing sensitivity, as measured by d', across participants or conditions should account for the possibility of criterion fluctuations.

Criterion fluctuations as a topic of study in human neuroscience

The above examples illustrate how criterion fluctuations can result in various biases. Note that the impact of fluctuations is not limited to only these examples: the concern that non-stationarities might underlie post-error and post-confidence slowing — i.e., slow fluctuation in the decision boundary of the drift diffusion model — has been raised in other domains well (Dutilh et al., 2012; Desender et al., 2019). While these examples stress the importance to control for these fluctuations, it is crucial to highlight that the criterion fluctuations are themselves an intriguing phenomenon worthy of exploration. A growing body of research illustrates that decision-making is characterized by non-stationarities (Ashwood et al., 2022; Cowley et al., 2020; Desender et al., 2019; Dutilh et al., 2012; Gupta and Brody, 2022; Li et al., 2024; Mochol et al., 2021; Roy et al., 2021). Whereas these non-stationarities seem to be a fundamental feature of decision-making, they are still often ignored or considered a nuisance. Recently, a general trend in cognitive neuroscience is emerging towards more dynamical approaches, with more effort spent to capture time-varying internal states (Greene et al., 2023; John et al., 2022). The ability to measure trial-to-trial criterion fluctuations opens up new avenues for research, such as uncovering the neural mechanisms underpinning them, identifying more phenomena in which they play a role, and leveraging this understanding to improve computational models. In sum, studying criterion fluctuations will advance our understanding of decision-making and its underlying mechanisms.

Previous attempts to measure criterion fluctuations

In the previous section, we showcase different scenarios illustrating the importance of considering fluctuations in the decision criterion. In the context of choice history bias (i.e., repetition effects in function of previous response), earlier work attempted to control for criterion fluctuations through a model-free approach (Lak et al., 2020; Mendonça et al., 2020), similar to one developed in the post-error slowing literature (Dutilh et al., 2012). Since the criterion

is assumed to be autocorrelated, it should have a similar influence on temporally adjacent trials. In theory, it should thus be possible to dissociate: (i) a causal influence of the previous response on the current trial, from (ii) an acausal influence of the next response on the current trial, which reflects the value of the criterion that varies only slowly across several trials. Lak et al. (2020) proposed to obtain a drift-free measure of choice history bias by subtracting the acausal from the causal influence. Although this model-free approach sounds promising, Gupta and Brody (2022) demonstrated that this is only effective in the context of a specific subset of updating strategies (i.e., no systematic updating and symmetric win-stay lose-switch). In the presence of other updating strategies, such as win-stay lose-stay or win-stay lose-nothing, their model-free correction resulted in inaccurate and biased estimates. Therefore, this proposed model-free approach is not an appropriate tool to correct for criterion fluctuations. Moreover, it does not allow extracting a trial-by-trial measure of the latent criterion to study its neural and physiological basis.

Instead of the model-free approach, Gupta and Brody (2022) have proposed a model-based approach which explicitly models criterion fluctuations as an AR(1) process in a Bernoulli linear dynamical system. With this approach they were able to correctly recover a wide range of updating strategies. However, there are some key limitations that prevent a widespread usage of their approach, especially in human cognitive neuroscience. First, the authors only tested their model using very large datasets: simulated datasets with 40,000 trials, or rat data with 50.000 trials, obtained over many sessions. Second, their model exhibited inherent trade-offs between parameters. Specifically, their model featured a parameter that weighted the criterion fluctuations in the Bernoulli probability. This generated a trade-off with the error variance (σ^2) as both control the scaling of the criterion fluctuations. Finally, the authors only focused on the recovery of history biases but did not quantitatively investigate the recovery of the criterion trajectory and its underlying computational parameters. Instead, they assumed a fixed autoregressive coefficient a of .9995. Similarly, Roy et al. (2021) developed a method to estimate trial-by-trial fluctuations in the weights of a logistic regression model, which captures

criterion fluctuations through the intercept. However, the weights were assumed to evolve over time according to an AR(1) process with a fixed to 1 (i.e., a random walk). Given that the parameter a plays an important role in capturing the temporal dynamics of a time series (i.e., as shown in fig. 1B), fixing this parameter to a specific value is therefore an overly strong assumption. It is likely that the generative autoregressive coefficient of criterion fluctuations in experimental data deviates from this value, which may bias estimates of trial-to-trial criterion trajectories.

Hierarchical Model for Fluctuations in Criterion (hMFC)

In the current work, we develop the Hierarchical Model for Fluctuations in Criterion (hMFC), which aims to solve these limitations by extending earlier approaches in various respects. First, to solve the issue that the model requires a high number of trials to be fitted, we implement a Bayesian hierarchical estimation approach. This allows to combine data across multiple participants and concurrently models both group level estimates and individual estimates. Therefore, fewer trials are required to obtain reliable parameter estimates. Second, we adjust our model structure compared to Gupta and Brody (2022) to avoid the parameter trade-off discussed earlier. Specifically, in hMFC this trade-off is solved by not additionally weighing the criterion fluctuations in the Bernoulli probability (i.e., setting it to 1). Third, instead of fixing a in the AR(1) to .9995 (Gupta and Brody, 2022) or 1 (Roy et al., 2021), we implement this in our model as a free parameter. This allows to capture a wide range of criterion fluctuation dynamics and ultimately enables researchers not just to control for criterion fluctuations, but also to study these fluctuations and its dynamics as the main topic of investigation (e.g., comparing a over different experimental manipulations or investigating individual differences). Finally, we provide easy to use code and software demos to enable the application of our model by the cognitive neuroscience community.

In sum, in order to model criterion fluctuations in decision criterion we implement a hierarchical

framework within a Bernoulli linear dynamical system, called the Hierarchical Model for Fluctuations in Criterion (hMFC). Next, we elaborate on the technical details of hMFC.

The generative model

First, consider a single session consisting of a sequence of T trials. On each trial $t = 1, \ldots, T$, the agent makes a binary response y_t. Our generative model assumes the response is drawn from a Bernoulli distribution:

$$y_t \sim \text{Bern}(f(\boldsymbol{w}^\top \boldsymbol{u}_t + x_t)) \tag{1}$$

where the probability of a response is governed by the covariates $\boldsymbol{u}_t \in \mathbb{R}^p$, the latent state $x_t \in \mathbb{R}$, and the weights $\boldsymbol{w} \in \mathbb{R}^p$. The sigmoid transformation, $f(a) = 1/(1 + e^{-a})$, constrains the conditional probability to be between 0 and 1. The vector $\boldsymbol{u}_t$ contains a constant of 1 for the model to estimate an intercept, and observed variables or covariates on trial t, such as the stimulus evidence or previous response. The influence of these covariates on the Bernoulli probability is captured by the weights $\boldsymbol{w} \in \mathbb{R}^p$. The latent state $x_t \in \mathbb{R}$ on trial t, which is not directly observable, is assumed to follow a first-order autoregressive process:

$$x_t = a x_{t-1} + \epsilon_t, \qquad \epsilon_t \sim \text{N}(0, \sigma^2) \tag{2}$$

where $a \in [0, 1]$ is the autoregressive coefficient which captures the temporal dynamics of the criterion fluctuations (fig. 1B), and σ^2 controls the variance of the fluctuations (i.e., σ^2 is a scaling parameter). Note that adding a drifting value x_t to the log-odds of the Bernoulli distribution is mathematically equivalent to comparing a decision variable to a fluctuating criterion (fig. 1A). In sum, the model predicts responses based on a combination of the weighted variables u_t and the latent state x_t, as illustrated in fig. 2A. Ultimately, this model allows us to disentangle two sources that influence the response probability: (i) the influence of $\boldsymbol{u}_t$; and (ii) the influence of the latent state x_t.

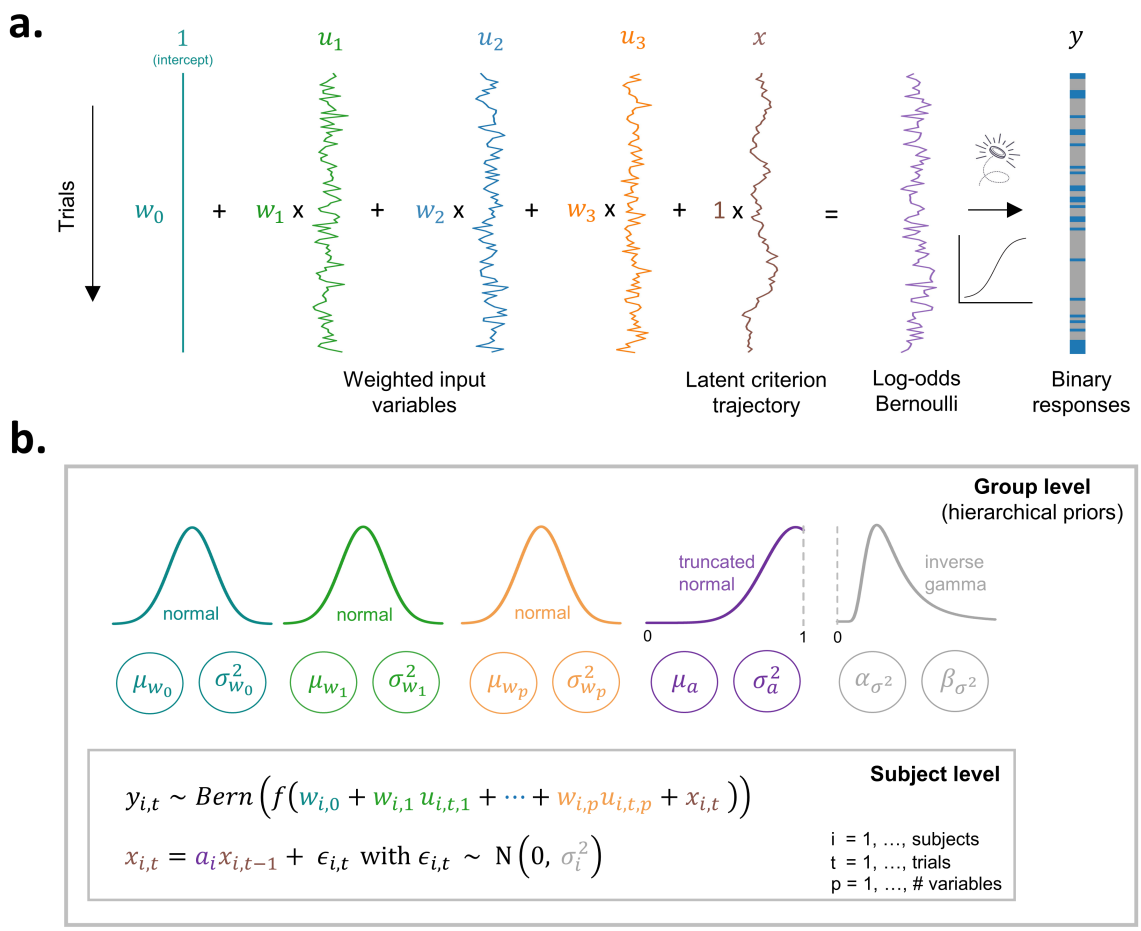

Figure 2: An overview of the generative model and its hierarchical structure. **A)** The figure displays, using 100 simulated trials for one agent, responses sampled from a Bernoulli distribution (i.e., a series of weighted coin flips). The trial-by-trial Bernoulli probabilities are a function of a weighted combination of the observed input variables $\boldsymbol{u}_t$ and the latent criterion fluctuations x_t. The first column consist of a constant to allow the model to estimate an intercept w_0. The next three columns show how input variables, drawn from a standard normal distribution $\mathrm{N}(0, 1)$ evolve randomly over time. Summing the weighted input variables with the criterion trajectory results in the log-odds of the Bernoulli distribution. After applying a sigmoid transformation to transform log-odds into probabilities, we draw from the Bernoulli distribution (i.e., weighted coin flip) to produce binary responses. **B)** For each subject i we infer weights $(w_{i,0}, w_{i,1}, \ldots, w_{i,p})$. We model the latent criterion fluctuations $x_{i,t}$ as an AR(1) process with per-subject autoregressive coefficient a_i and variance σ_i^2. Through the specification of hierarchical priors we allow the sharing of statistical strength across subjects when estimating these individual parameters. Like the individual parameters, the parameters of the hierarchical distributions are iteratively updated during the estimation procedure.

The hierarchical prior: putting the 'h' in hMFC

Having explained the generative model and the behavior it generates (fig. 2), we next focus on the inverse step, namely estimating the underlying latent parameters given observed behavior. In many settings, we need to obtain accurate parameter estimates from only tens to hundreds of trials per subject. Small values of T pose statistical challenges for parameter estimation. However, we can pool information across individuals to aid in inference. To this end, we developed a hierarchical Bayesian model that allows for per-subject parameter estimates, while sharing information across subjects via a global prior distribution. We infer the per-subject parameters, the latent states (i.e., criterion fluctuations) for each subject, as well as the strength of the hierarchical prior using a fully Bayesian inference algorithm.

In the hierarchical model, each subject $i = 1, \ldots, N$ has its own responses $y_{i,t}$, covariates $u_{i,t}$, and latent states, $x_{i,t}$, for $t = 1, \ldots, T_i$. Each subject also has its own parameters $\theta_i = (\boldsymbol{w}_i, a_i, \sigma_i^2)$, where $\boldsymbol{w}_i \in \mathbb{R}^p$ denote the covariate weights, $a_i \in [0, 1]$ is the autoregressive parameter, and σ_i^2 is the variance of the latent criterion fluctuations. We learn a hierarchical prior that captures how the subject-level parameters θ_i vary across the population. The hierarchical prior allows us to share statistical power from other subjects when estimating θ_i. This results in better parameter recovery, even when a relatively small number trials per subject is available (Katahira, 2016).

We specified hierarchical priors for the three per-subject parameters. First, we assume a Gaussian prior on the weights,

$$\boldsymbol{w}_i \sim \mathrm{N}(\boldsymbol{w}_0, \mathrm{diag}(\boldsymbol{\nu}_w^2)), \tag{3}$$

where $\boldsymbol{w}_0 \in \mathbb{R}^p$ and $\boldsymbol{\nu}_w^2 \in \mathbb{R}_+^p$ denote the global mean and variance of the per-subject weights.

Second, we assume a truncated normal prior on the per-subject autoregressive coefficient,

$$a_i \sim \text{TruncNorm}(a_0, \nu_a^2, [0, 1]) \tag{4}$$

where a_0 and $\boldsymbol{\nu}_a^2$ are the global mean and variance, respectively, and the domain is constrained to $[0, 1]$ because negative values of a_i would lead to large jumps in the criterion, and values greater than 1 would lead to unstable dynamics. Third, we specify an inverse gamma (IGa) prior for the variance of the criterion fluctuations,

$$\sigma_i^2 \sim \text{IGa}(\alpha, \beta), \tag{5}$$

with shape parameter $\alpha \in \mathbb{R}_+$ and a scale parameter $\beta \in \mathbb{R}_+$.

Finally, we place weakly informative priors on the parameters of the hierarchical model, $\eta = (\boldsymbol{w}_0, \boldsymbol{\nu}_w^2, a_0, \nu_a^2, \alpha, \beta)$, see the Appendix for complete details. These global parameters will be inferred alongside the per-subject parameters using a fully Bayesian inference algorithm, which we describe next.

A posterior inference algorithm

We develop a fully Bayesian inference algorithm to estimate the the posterior distributions of the per-subject latent states and parameters, as well as the parameters of the hierarchical prior. We use an augmented, blocked Gibbs sampling algorithm to produce samples that are asymptotically drawn from the desired posterior distribution.

The Gibbs sampling algorithm proceeds by iteratively sampling the conditional distribution of one block of variables, holding the rest fixed. However, this approach requires the closed-form expressions for the conditional distributions. This tractability can be achieved if the model is conditionally conjugate. Unfortunately, the Bernoulli likelihood is not conjugate with the Gaussian hierarchical priors for the inputs weights or the AR(1) model for the latent states.

To circumvent this challenge, we employ the Pólya-gamma augmentation trick (Polson et al., 2013). With this augmentation, we can obtain a closed-form expression for the input weights $\boldsymbol{w}_i$ given the response, input variables, criterion fluctuations, and the hierarchical prior.

For the two other per-subject parameters, a_i and σ_i^2, closed-form expressions for the conditional posteriors are available. First, for a_i the truncated normal prior is conjugate with the AR(1) model likelihood. Similarly, the inverse gamma prior for σ_i^2 is conjugate with the AR(1) model of the latent criterion fluctuations.

After Pólya-gamma augmentation, the conditional distribution of the latent criterion fluctuations is a chain-structured Gaussian graphical model. We generate exact samples from this conditional distribution using a forward filtering/backward sampling algorithm (Murphy, 2022). The forward filtering step computes the filtering distributions of the latent criterion fluctuations, given all the responses up to a certain trial. Once the filtering distributions have been computed for each trial, the backward sampling step generates a sequence of latent fluctuations, working backward from the last trial T_i down to the first. This procedure produces a sequence of latent fluctuations $\boldsymbol{x}_i = (x_{i,1}, \ldots, x_{i,T_i})$ that is distributed according to its conditional distribution given the per-subject parameters, covariates, responses, and augmentation variables. Note that after estimating the latent states, we mean-center the states to ensure the intercept $w_{i,0}$ remains identifiable.

Finally, we update the parameters of the hierarchical prior, η, by sampling from their conditional distribution under a weakly informative prior. For the two group-level parameters of the normal prior on weights, $\boldsymbol{w}_0$ and $\boldsymbol{\nu}_w^2$, we have closed-form expressions for the conditional distributions. We update the remaining hierarchical prior parameters — a_0, ν_a^2, α and β — using random-walk Metropolis-Hastings with a symmetric Gaussian proposal distribution.

Together, these updates yield a Markov chain whose stationary distribution is the posterior distribution over the latent per-subject fluctuations, per-subject parameters, and hierarchical prior parameters. Repeatedly applying these updates yields samples that are asymptotically

distributed according to this posterior. We can use these samples to estimate posterior expectations and perform hypothesis tests, as demonstrated in the results below. Complete details of the augmented, blocked Gibbs sampling algorithm can be found in the Appendix.

Accessing the model

Our modeling and inference code, together with comprehensive demonstrations, are available at www.github.com/<tbd>.

Results

To assess the effectiveness of the hMFC parameter estimation, we simulate data with a varying number of trials (500, 1000, 2500 or 5000 trials) per subject, with 50 subjects per dataset and 50 datasets per trial count. For each subject, four weights, $\boldsymbol{w}_i = (w_{i,0}, \ldots, w_{i,3})$, are independently drawn from a normal distribution with mean $\boldsymbol{w}_0 = (0.0, 0.2, -0.3, 0.6)$ variance $\boldsymbol{\nu}_w^2 = (1, 1, 1, 1)$. Similarly, the inputs $\boldsymbol{u}_t \in \mathbb{R}^4$ are also sampled from a standard normal distribution (mean zero, identity covariance), except for $u_{i,0}$, which is a constant of 1 to allow the model to estimate an intercept ($w_{i,0}$). The criterion fluctuations are sampled from an AR(1) model with the autoregressive coefficient a drawn from a truncated normal (mean 0.99, standard deviation 0.025, minimum value 0, and maximum value 1) and σ_i^2 sampled from an inverse gamma distribution ($\alpha = 5.0$, $\beta = 0.45$).

Hierarchical prior (global) parameter recovery

The global parameters ($\eta = (a_0, \nu_a^2, \boldsymbol{w}_0, \boldsymbol{\nu}_w^2, \alpha, \beta)$) are accurately recovered: for most datasets, the generative value lies within the 95% credible interval of the posterior estimates. Figure 4 shows the bias and 95% credible interval of the global parameters in function of the number of trials per subject. Figure 5 shows that the width of the posterior distributions become smaller when more subjects are available.

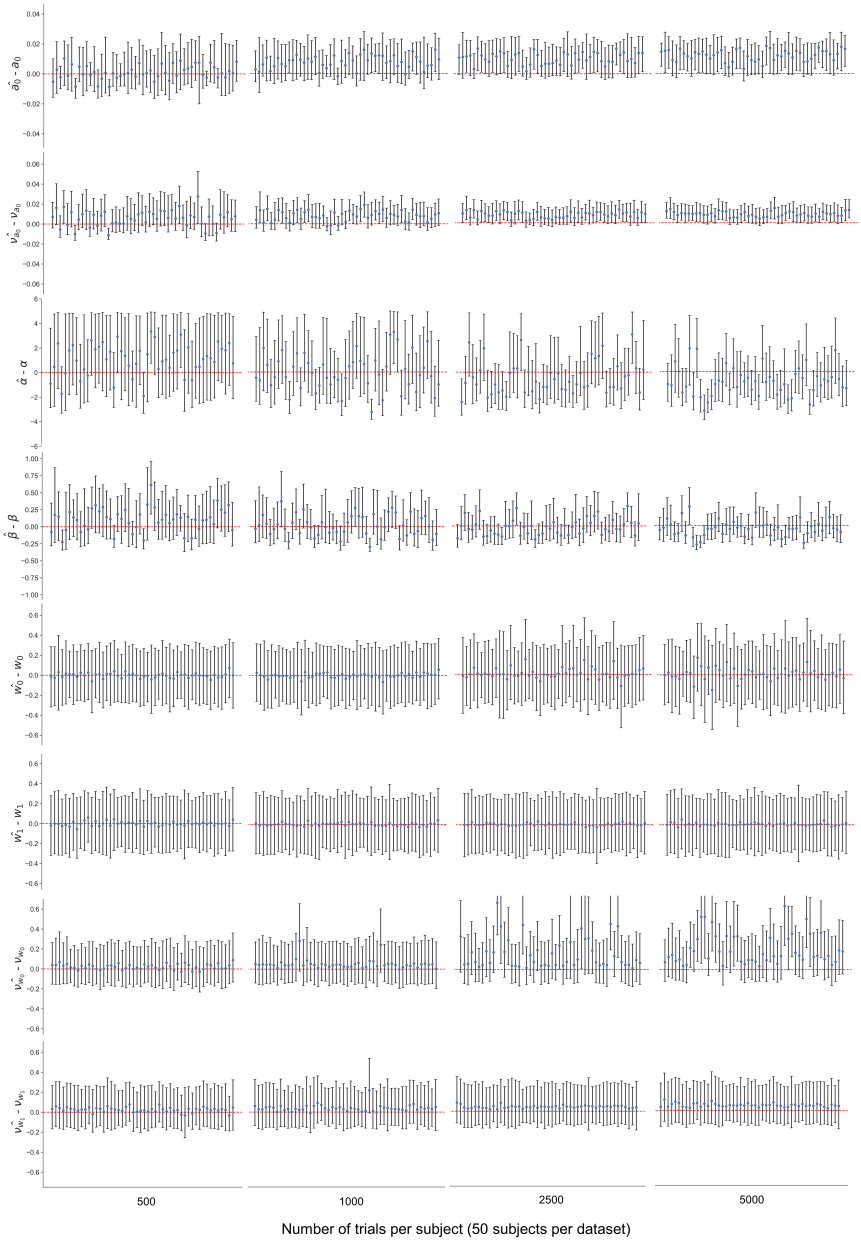

Figure 3: Each panel represents the bias for the global parameter estimates across datasets with varying numbers of trials per subject (500, 1000, 2500, 5000). Each dataset consisted of 50 simulated agents. Dots indicate the estimated bias for each dataset, while the intervals represent the 95% credible intervals. The horizontal red dashed line indicates zero bias.

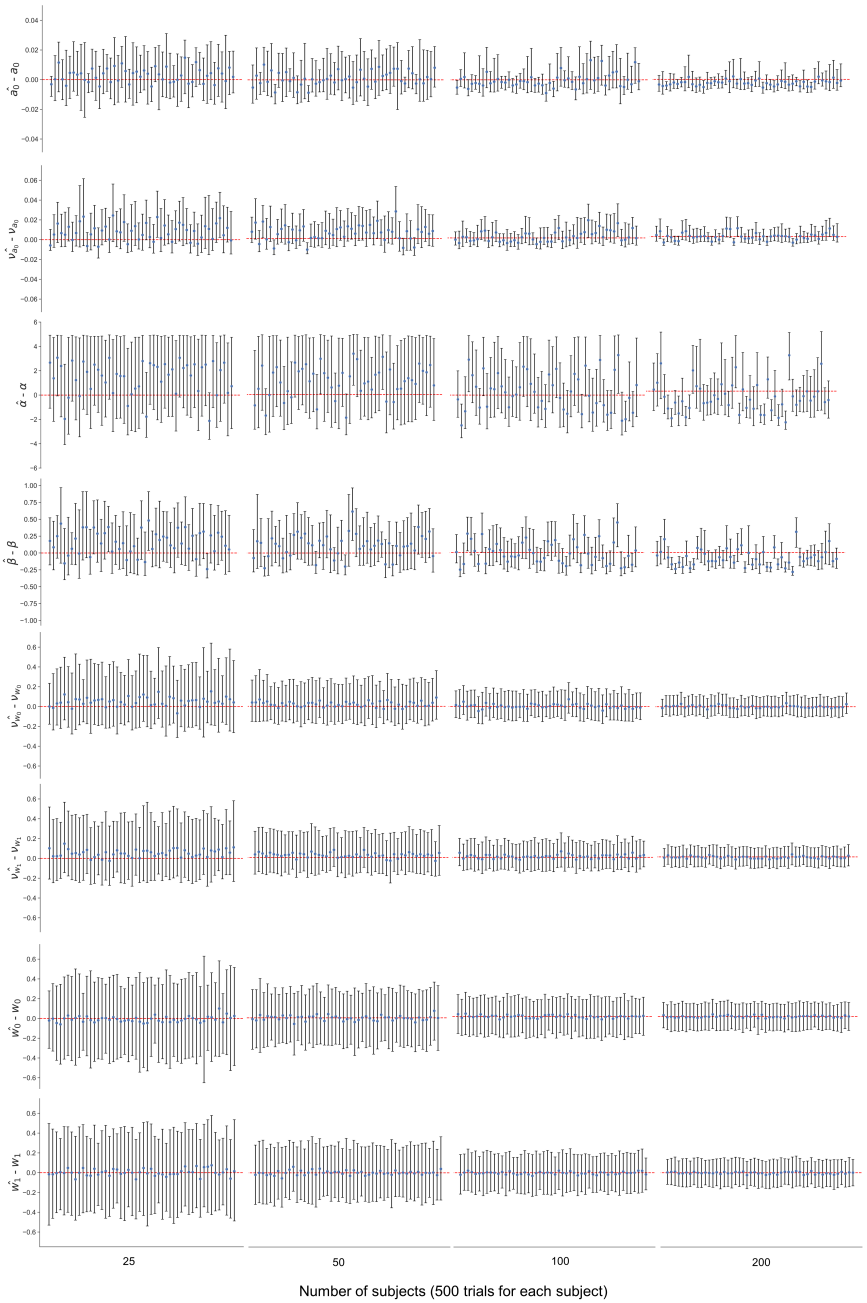

Figure 4: The recovery of the global parameters in function of the number of subjects per dataset (25, 50, 100, 200). The posterior distributions become smaller with more simulated subjects.

Per-subject (local) parameter recovery

The inference algorithm can generally recover all per-subject parameters (Figure 5A). The per-subject weights $w_{i,0}$, $w_{i,1}$, $w_{i,2}$, $w_{i,3}$ show excellent recovery ($r > .98$), even for datasets with only 500 trials per subject (Figure 5D). The autoregressive model parameters a_i and σ_i^2, which describe the criterion fluctuations, show an excellent recovery with many trials per subject ($r = .95$ for a_i, $r = .85$ for σ_i^2, with 5000 trials), but are less accurately estimated when fewer trials per subject are available ($r = .68$ for a_i and $r = .57$ for σ_i^2 with 500 trials) (Figure 5B,C).

Per-subject criterion trajectory recovery

Importantly, although the per-subject parameters a_i and σ_i^2 can be misestimated with low trial numbers, we can still accurately recover the trial-to-trial latent criterion trajectory (fig. 6) crucial for investigations into its neural and physiological bases. The recovery, measured as the correlations between true and inferred latent trajectory, is generally high ($r > .83$), and consistent across datasets (fig. 6A). As shown for one example subject, the true and the inferred trajectory strongly correlate ($r = .85$), even though the model is not able to perfectly recover high-frequency fluctuations in the true latent trajectory (fig. 6C).

At lower trial counts, recovery ranges from X to X across subjects. This variation is partly due to the stochastic criterion fluctuations themselves, but also relates to the combination of parameters governing the individual AR(1) processes. Specifically, trajectory recovery is more difficult for higher values of σ_2 and/or lower values of a.

Computation time

Fitting 50 subjects with 500 trials each for 1000 iterations takes around 57 minutes on a computer with only CPU's. When GPU's are available the computation time is reduced to 4 minutes, making this procedure lightweight enough to run easily on typical human datasets.

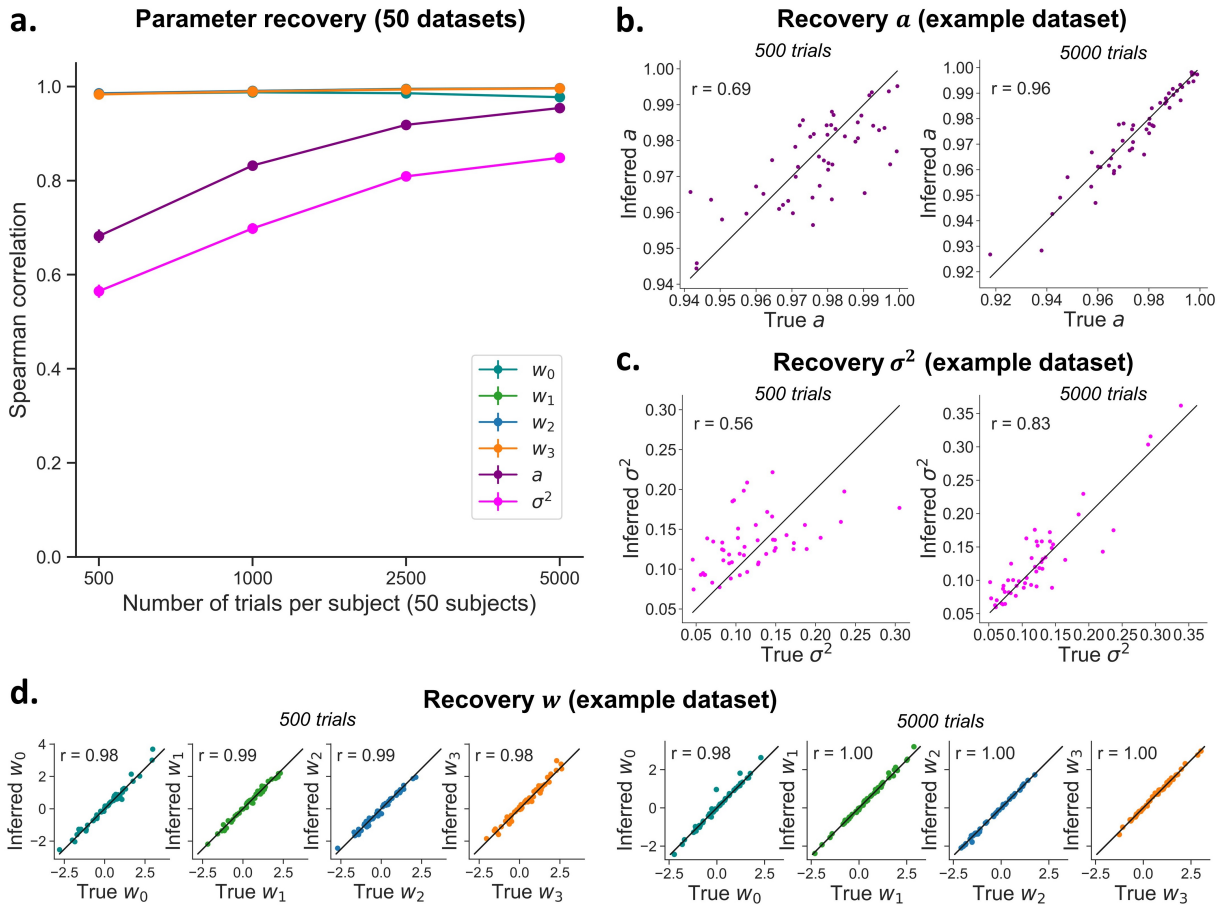

Figure 5: Recovery of the per-subject parameters of the hMFC model. **A)** The covariate weights w show excellent recovery even with low trial numbers (note that the lines for the different weights overlap). The model is able to recover a and σ^2 very well with high trial numbers, but recovery drops with a limited number of data points per subject. Note that error bars (standard error) are shown but they are very small. **B-D)** The recovery for a (B) and σ^2 (C) and input weights w_0, w_1, w_2, w_3 (D) is shown for an example dataset with 500 trials and 5000 trials per subject. Each dot represents one subject. The line shows the diagonal.

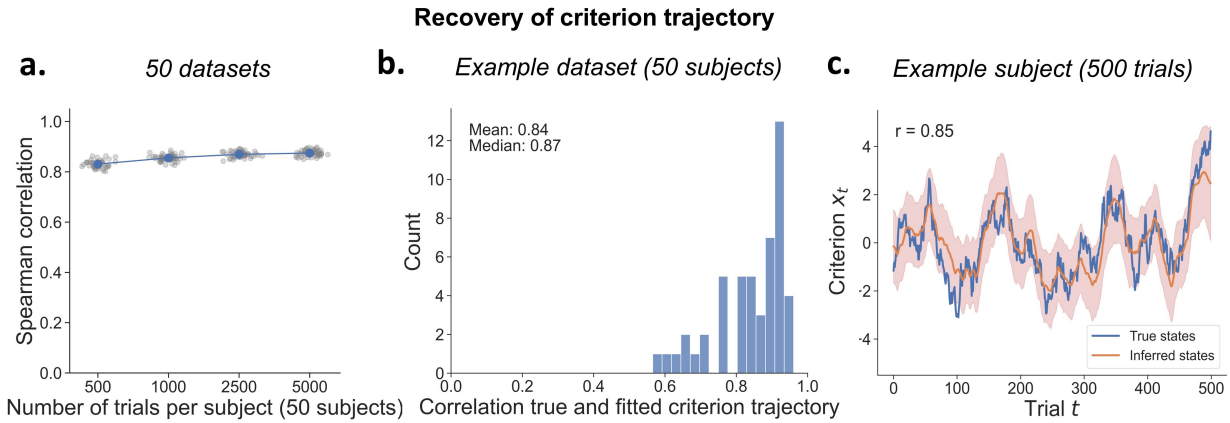

Figure 6: Recovery of the latent criterion trajectory. **A)** Each dot represents the criterion trajectory of one dataset, averaged over subjects. Note that error bars (standard error) are plotted but are very small. **B)** An example dataset of 50 subjects with 500 trials each shows an average correlation between true and inferred criterion fluctuations of $r = .84$, ranging between X-Y. **C)** True and inferred criterion fluctuations for an example session. The shaded area represents the 95% credible interval of the estimated posterior at each trial.

Application hMFC to sequential effects and measures of sensitivity

Having verified that our inference procedure recovers the generative hMFC model, we assess if it also corrects for the spurious effects that can arise from ignoring criterion fluctuations: apparent sequential effects and underestimated d' (as shown in fig. 1). We fit the hMFC model with stimulus evidence and previous response as covariates, to the data shown in fig. 1. Importantly, we ran the fits either with or without estimating criterion fluctuations, allowing us to test if the inclusion of these parameters accurately solves the confounds.

When criterion fluctuations are not estimated, a clear effect of previous response (i.e., tendency to repeat previous response) artificially emerges in the model (fig. 7A): the posterior distributions are shifted further away from zero when the criterion fluctuations become stronger (i.e., replicating the finding from fig. 1C, using our model fit). Fortunately, when criterion fluctuations are estimated, this effect disappears and posteriors center around the generative value ($\beta_{\text{previous response}} = 0$, red line). This shows that our model can accurately distinguish sequential effects caused by slow fluctuations in the decision criterion from sequential effects

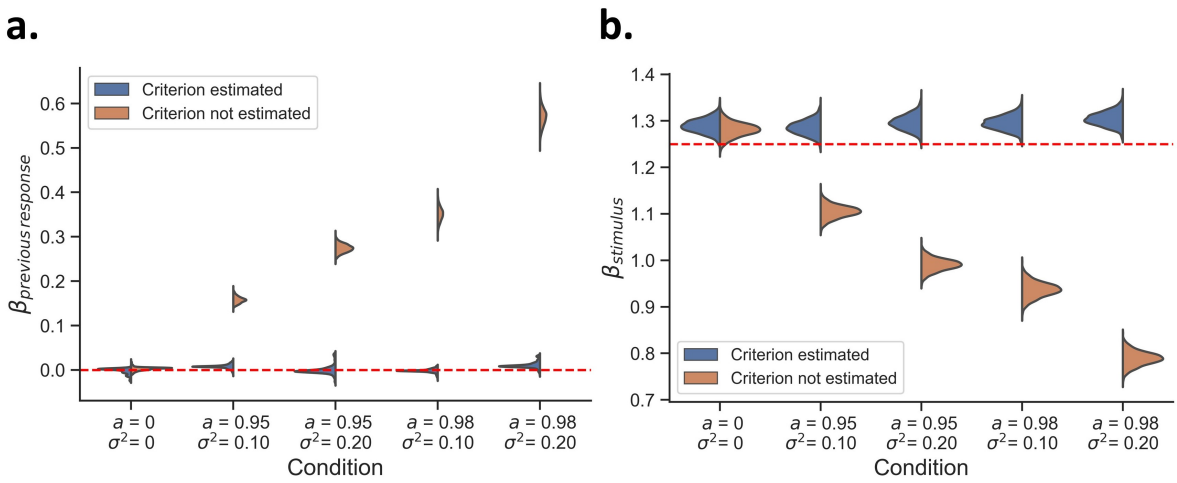

Figure 7: Using hMFC to account for criterion fluctuations. The posterior distributions for β_{stimulus} and $\beta_{\text{previous response}}$ are shown after fitting hMFC to the simulated datasets used in fig. 1C and fig. 1D. **A)** The estimation of criterion fluctuations allows hMFC to correctly recover the absence of any generative criterion updating in function of the response on the previous trial. **B)** Estimating criterion fluctuations helps to address the issue of underestimating the psychometric slope.

caused by systematic updating.

When criterion fluctuations are not estimated, the psychometric slopes are systematically underestimated creffig:fig6B) compared to the generative value ($\beta_{\text{stimulus}} = 1.25$, red line). Fortunately, estimating criterion fluctuations resolves this underestimation. In sum, the hMFC corrects for estimation biases by explicitly estimating the fluctuations in the criterion.

Discussion

Computational models of decision-making typically assume stationarity in the parameters underlying the decision computation. However, an increasing body of research highlights the dynamic nature of decision making and shows that computational variables underlying decision-making may fluctuate over time. One example of these non-stationarities are fluctuations in the decision criterion. Ignoring such latent fluctuations can lead to false interpretations and inaccurate estimates. Specifically, criterion fluctuations can lead to behavior that mimics a

causal relationship in sequential effects, and they can lead to an underestimation of two popular measures of perceptual sensitivity (d' in SDT and the psychometric slope).

Here, we developed the Hierarchical Model for Fluctuations in Criterion (hMFC), a Bayesian hierarchical model to obtain trial-by-trial estimates of the decision criterion. Even with only 500 trials per participant, which is common in studies on human cognitive neuroscience, the model parameters and the latent criterion trajectory itself can be accurately inferred. Importantly, the hMFC can effectively correct for the biases that occur when criterion fluctuations are ignored.

Our model here only allows an AR(1) process to govern the decision criterion, i.e., the intercept of the model. Future work may explore how to extend this hierarchical Bayesian framework to simultaneously infer fluctuations in the weights of other covariates reflecting stimulus sensitivity or trial history (Roy et al., 2021).

We provide open access to the code for hMFC and created an interactive demo for researchers interested in applying this model to their data. Through this work, we hope to enable further research in human cognitive neuroscience aimed at understanding the dynamics of decision-making.

Acknowledgements

AEU is supported by a Veni fellowship from the Netherlands Organisation for Scientific Research. SWL is supported by grants from the NIH BRAIN Initiative (U19NS113201, R01NS131987, R01NS113119, & RF1MH133778), the NSF/NIH CRCNS Program (R01NS130789), and the Sloan, Simons, and McKnight Foundations.

References

Rei Akaishi, Kazumasa Umeda, Asako Nagase, and Katsuyuki Sakai. Autonomous mechanism of internal choice estimate underlies decision inertia. *Neuron*, 81(1):195–206, 2014.

Athena Akrami, Charles D Kopec, Mathew E Diamond, and Carlos D Brody. Posterior parietal cortex represents sensory history and mediates its effects on behaviour. *Nature*, 554(7692): 368–372, 2018.

Zoe C Ashwood, Nicholas A Roy, Iris R Stone, International Brain Laboratory, Anne E Urai, Anne K Churchland, Alexandre Pouget, and Jonathan W Pillow. Mice alternate between discrete strategies during perceptual decision-making. *Nature Neuroscience*, 25(2):201–212, 2022.

Ji Won Bang and Dobromir Rahnev. Stimulus expectation alters decision criterion but not sensory signal in perceptual decision making. *Scientific reports*, 7(1):17072, 2017.

Aaron S Benjamin, Michael Diaz, and Serena Wee. Signal detection with criterion noise: applications to recognition memory. *Psychological review*, 116(1):84, 2009.

Ella Bosch, Matthias Fritsche, Benedikt V Ehinger, and Floris P de Lange. Opposite effects of choice history and evidence history resolve a paradox of sequential choice bias. *Journal of vision*, 20(12):9–9, 2020.

James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL http://github.com/google/jax.

Anke Braun, Anne E Urai, and Tobias H Donner. Adaptive history biases result from confidence-weighted accumulation of past choices. *Journal of Neuroscience*, 38(10):2418–2429, 2018.

Laura Busse, Asli Ayaz, Neel T Dhruv, Steffen Katzner, Aman B Saleem, Marieke L Schölvinck, Andrew D Zaharia, and Matteo Carandini. The detection of visual contrast in the behaving mouse. *Journal of Neuroscience*, 31(31):11351–11361, 2011.

Fadong Chen and Ian Krajbich. Biased sequential sampling underlies the effects of time pressure and delay in social decision making. *Nature communications*, 9(1):3557, 2018.

Benjamin R Cowley, Adam C Snyder, Katerina Acar, Ryan C Williamson, M Yu Byron, and Matthew A Smith. Slow drift of neural activity as a signature of impulsivity in macaque visual and prefrontal cortex. *Neuron*, 108(3):551–567, 2020.

Antoine Del Cul, Sylvain Baillet, and Stanislas Dehaene. Brain dynamics underlying the nonlinear threshold for access to consciousness. *PLoS biology*, 5(10):e260, 2007.

Kobe Desender, Annika Boldt, Tom Verguts, and Tobias H Donner. Confidence predicts speed-accuracy tradeoff for subsequent decisions. *elife*, 8:e43499, 2019.

Gilles Dutilh, Don Van Ravenzwaaij, Sander Nieuwenhuis, Han LJ Van der Maas, Birte U Forstmann, and Eric-Jan Wagenmakers. How to measure post-error slowing: A confound and a simple solution. *Journal of mathematical psychology*, 56(3):208–216, 2012.

Samuel W Fernberger. Interdependence of judgments within the series for the method of constant stimuli. *Journal of Experimental Psychology*, 3(2):126, 1920.

Jason Fischer and David Whitney. Serial dependence in visual perception. *Nature neuroscience*, 17(5):738–743, 2014.

Stephen M Fleming, Rimona S Weil, Zoltan Nagy, Raymond J Dolan, and Geraint Rees. Relating

introspective accuracy to individual differences in brain structure. *Science*, 329(5998):1541–1543, 2010.

Amy Frithsen, Justin Kantner, Brian A Lopez, and Michael B Miller. Cross-task and cross-manipulation stability in shifting the decision criterion. *Memory*, 26(5):653–663, 2018.

Matthias Fritsche, Pim Mostert, and Floris P de Lange. Opposite effects of recent history on perception and decision. *Current Biology*, 27(4):590–595, 2017.

Matthias Fritsche, Eelke Spaak, and Floris P De Lange. A Bayesian and efficient observer model explains concurrent attractive and repulsive history biases in visual perception. *Elife*, 9: e55389, 2020.

Ingo Fründ, Felix A Wichmann, and Jakob H Macke. Quantifying the effect of intertrial dependence on perceptual decisions. *Journal of vision*, 14(7):9–9, 2014.

Joshua I Gold and Michael N Shadlen. The neural basis of decision making. *Annu. Rev. Neurosci.*, 30:535–574, 2007.

David Marvin Green, John A Swets, et al. *Signal detection theory and psychophysics*, volume 1. Wiley New York, 1966.

Abigail S Greene, Corey Horien, Daniel Barson, Dustin Scheinost, and R Todd Constable. Why is everyone talking about brain state? *Trends in Neurosciences*, 2023.

Diksha Gupta and Carlos D Brody. Limitations of a proposed correction for slow drifts in decision criterion. *arXiv preprint arXiv:2205.10912*, 2022.

Nadia Haddara and Dobromir Rahnev. The impact of feedback on perceptual decision-making and metacognition: Reduction in bias but no change in sensitivity. *Psychological Science*, 33 (2):259–275, 2022.

Michael J Hautus, Neil A Macmillan, and C Douglas Creelman. *Detection theory: A user's guide*. Routledge, 2021.

Eun Jung Hwang, Jeffrey E Dahlen, Madan Mukundan, and Takaki Komiyama. History-based action selection bias in posterior parietal cortex. *Nature communications*, 8(1):1242, 2017.

Yohan J John, Kayle S Sawyer, Karthik Srinivasan, Eli J Müller, Brandon R Munn, and James M Shine. It's about time: Linking dynamical systems with human neuroimaging to understand the brain. *Network Neuroscience*, 6(4):960–979, 2022.

Kentaro Katahira. How hierarchical models improve point estimates of model parameters at the individual level. *Journal of Mathematical Psychology*, 73:37–58, 2016.

Niels A Kloosterman, Jan Willem de Gee, Markus Werkle-Bergner, Ulman Lindenberger, Douglas D Garrett, and Johannes Jacobus Fahrenfort. Humans strategically shift decision bias by flexibly adjusting sensory evidence accumulation. *Elife*, 8:e37321, 2019.

Ian Krajbich, Carrie Armel, and Antonio Rangel. Visual fixations and the computation and comparison of value in simple choice. *Nature neuroscience*, 13(10):1292–1298, 2010.

Armin Lak, Emily Hueske, Junya Hirokawa, Paul Masset, Torben Ott, Anne E Urai, Tobias H Donner, Matteo Carandini, Susumu Tonegawa, Naoshige Uchida, et al. Reinforcement biases subsequent perceptual decisions when confidence is low, a widespread behavioral phenomenon. *Elife*, 9:e49834, 2020.

Jing-Jing Li, Chengchun Shi, Lexin Li, and Anne GE Collins. Dynamic noise estimation: A generalized method for modeling noise fluctuations in decision-making. *Journal of Mathematical Psychology*, 119:102842, 2024.

Jan-Matthis Lueckmann, Jakob H Macke, and Hendrikje Nienborg. Can serial dependencies in choices and neural activity explain choice probabilities? *Journal of Neuroscience*, 38(14): 3495–3506, 2018.

André G Mendonça, Jan Drugowitsch, M Inês Vicente, Eric EJ DeWitt, Alexandre Pouget, and Zachary F Mainen. The impact of learning on perceptual decisions and its implication for speed-accuracy tradeoffs. *Nature communications*, 11(1):2757, 2020.

Gabriela Mochol, Roozbeh Kiani, and Rubén Moreno-Bote. Prefrontal cortex represents heuristics that shape choice bias and its integration into future behavior. *Current Biology*, 31(6):1234–1244, 2021.

Kevin P Murphy. *Probabilistic machine learning: Advanced topics*. MIT Press Cambridge, MA, USA, 2022.

Jan Peters and Mark D'Esposito. The drift diffusion model as the choice rule in inter-temporal and risky choice: A case study in medial orbitofrontal cortex lesion patients and controls. *PLoS computational biology*, 16(4):e1007615, 2020.

Nicholas G Polson, James G Scott, and Jesse Windle. Bayesian inference for logistic models using pólya–gamma latent variables. *Journal of the American statistical Association*, 108(504): 1339–1349, 2013.

Roger Ratcliff and Jeffrey N Rouder. Modeling response times for two-choice decisions. *Psychological science*, 9(5):347–356, 1998.

Nicholas A Roy, Ji Hyun Bak, Athena Akrami, Carlos D Brody, and Jonathan W Pillow. Extracting the dynamics of behavior in sensory decision-making experiments. *Neuron*, 109(4):597–610, 2021.

Jason Samaha, Missy Switzky, and Bradley R Postle. Confidence boosts serial dependence in orientation estimation. *Journal of vision*, 19(4):25–25, 2019.

Nicolás Sánchez-Fuenzalida, Simon Van Gaal, Stephen M Fleming, Julia M Haaf, and Johannes J Fahrenfort. Predictions and rewards affect decision-making but not subjective experience. *Proceedings of the National Academy of Sciences*, 120(44):e2220749120, 2023.

Daniel Schlunegger and Fred W Mast. Probabilistic integration of preceding responses explains response bias in perceptual decision making. *Iscience*, 26(7), 2023.

Michel Treisman and Thomas C Williams. A theory of criterion setting with an application to sequential dependencies. *Psychological review*, 91(1):68, 1984.

Anne E Urai and Tobias H Donner. Persistent activity in human parietal cortex mediates perceptual choice repetition bias. *Nature Communications*, 13(1):6015, 2022.

Anne E Urai, Anke Braun, and Tobias H Donner. Pupil-linked arousal is driven by decision uncertainty and alters serial choice bias. *Nature communications*, 8(1):14637, 2017.

Anne E Urai, Jan Willem De Gee, Konstantinos Tsetsos, and Tobias H Donner. Choice history biases subsequent evidence accumulation. *Elife*, 8:e46331, 2019.

Eline Van Geert, Pieter Moors, Julia Haaf, and Johan Wagemans. Same stimulus, same temporal context, different percept? Individual differences in hysteresis and adaptation when perceiving multistable dot lattices. *i-Perception*, 13(4):20416695221109300, 2022.

Eline Van Geert, Tina Ivančir, and Johan Wagemans. An efficient Bayesian observer model of attractive and repulsive temporal context effects when perceiving multistable dot lattices. *Journal of Vision*, 24(4):18–18, 2024.

Felix A Wichmann and N Jeremy Hill. The psychometric function: I. Fitting, sampling, and goodness of fit. *Perception & psychophysics*, 63(8):1293–1313, 2001.

Wayne A Wickelgren. Unidimensional strength theory and component analysis of noise in absolute and comparative judgments. *Journal of Mathematical Psychology*, 5(1):102–122, 1968.

Jesse Windle, Nicholas G Polson, and James G Scott. Sampling pólya-gamma random variates: alternate and approximate techniques. *arXiv preprint arXiv:1405.0506*, 2014.

Angela J Yu and Jonathan D Cohen. Sequential effects: Superstition or rational behavior? *Advances in neural information processing systems*, 21, 2008.

Appendix

Hierarchical Model for Fluctuations in Criterion (hMFC) Details

Notation Let

- N denote the **number of subjects**
- T_i denote the **number of trials** for subject i
- $\eta = (a_0, \nu_a^2, \boldsymbol{w}_0, \boldsymbol{\nu}_w^2, \alpha, \beta)$ denote the **global parameters**.
- $\theta_i = (a_i, \boldsymbol{w}_i, \sigma_i^2)$ denote the **local parameters** for subject i
- $x_{i,t} \in \mathbb{R}$ for $t \in 1, \ldots, T_i$ denote the sequence of **latent states** for subject i.
- $\boldsymbol{u}_{i,t} \in \mathbb{R}^p$ for $t \in 1, \ldots, T_i$ denote the sequence of **inputs** for subject i.
- $y_{i,t} \in \{0, 1\}$ for $t \in 1, \ldots, T_i$ denote the sequence of **binary observations** for subject i.

The Hierarchical Model for Fluctuations in Criterion (hMFC) defines a joint distribution over latent states, observations, and parameters given the inputs,

$$
\begin{aligned}
p(\{\{x_{i,t}, y_{i,t}\}_{t=1}^{T_i}, \theta_i\}_{i=1}^{N}, \eta \mid \{\{\boldsymbol{u}_{i,t}\}_{t=1}^{T_i}\}_{i=1}^{N}) = \\
p(\eta) \times \prod_{i=1}^{N} \Bigg(p(\theta_i \mid \eta) \times \left[\mathrm{N}(x_{i,1} \mid 0, 1) \prod_{t=2}^{T_i} \mathrm{N}(x_{i,t} \mid a_i x_{i,t-1}, \sigma_i^2) \right] \\
\times \left[\prod_{t=1}^{T_i} \mathrm{Bern}\left(y_{i,t} \mid f(x_{i,t} + \boldsymbol{w}_i^\top \boldsymbol{u}_{i,t})\right) \right] \Bigg)
\end{aligned} \tag{6}
$$

where $f(a) = 1/(1 + e^{-a})$ denotes the logistic function and,

$$
p(\theta_i \mid \eta) = \mathrm{TruncNorm}(a_i \mid a_0, \nu_a^2, [0, 1]) \times \mathrm{N}(\boldsymbol{w}_i \mid \boldsymbol{w}_0, \mathrm{diag}(\boldsymbol{\nu}_w^2)) \times \mathrm{IGa}(\sigma_i^2 \mid \alpha, \beta) \tag{7}
$$

is the hierarchical prior. The truncated normal distribution ensures that $a_i \in [0, 1]$, so that

the latent state dynamics for subject i are stable. The hierarchical Gaussian prior on weights $\boldsymbol{w}_i \in \mathbb{R}^p$ allows for sharing of statistical strength across subjects.

We assume weakly informative priors over the global parameters,

$$p(\eta) = \mathrm{Unif}(a_0; [0,1]) \times \mathrm{Unif}(\nu_a; [0, \nu_a^{(\max)}]) \times \mathrm{N}(\boldsymbol{w}_0 \mid 0, \sigma_{w_0}^2 I) \\ \times \mathrm{IGa}(\nu_w^2 \mid \alpha_{\nu_w}, \beta_{\nu_w}) \times \mathrm{Unif}(\alpha; [0,10]) \times \mathrm{Unif}(\beta; [0,10]) \quad (8)$$

with $\sigma_{w_0} \to \infty$, $\alpha_{\nu_w} \to 0$, $\beta_{\nu_w} \to 0$. The uniform priors for a_0, α, and β allow for a wide range of global parameters, while still enforcing marginal stability constraints.

We set $\nu_a^{(\max)} = 0.2$ to prevent the latent states from jumping too far from one trial to the next. Allowing for such flexibility leads to degenerate solutions where the log odds on any trial can be made arbitrarily large using the latent variables $x_{i,t}$.

Posterior Inference

We designed an augmented, blocked Gibbs sampling algorithm to estimate the posterior distribution, $p(\{\{x_{i,t}\}_{t=1}^{T_i}, \theta_i\}_{i=1}^N, \eta \mid \{\{\boldsymbol{u}_{i,t}, y_{i,t}\}_{t=1}^{T_i}\}_{i=1}^N)$.

Gibbs sampling the latent states

The trick is to represent the Bernoulli likelihood as a scale-mean mixture of Gaussians using the Pólya-gamma (PG) augmentation (Polson et al., 2013),

$$\mathrm{Bern}(y \mid \sigma(\psi)) \propto e^{(y-\frac{1}{2})\psi} \int e^{-\frac{1}{2}\omega\psi^2} \, \mathrm{PG}(\omega \mid 1, 0) \, \mathrm{d}\omega, \quad (9)$$

where $\mathrm{PG}(\omega \mid 1, 0)$ denotes the standard Pólya-gamma density on $\omega \in \mathbb{R}_+$. In our case, $\psi_{it} = x_{it} + w_i^\top u_{it}$. After augmenting the model with auxiliary variables ω_{it}, the log likelihood

as a function of x_{it} is,

$$\log p(y_{it} \mid x_{it}, \omega_{it}, \boldsymbol{u}_{it}, \theta_i) = h_{it} x_{it} - \frac{1}{2} \omega_{it} x_{it}^2 \tag{10}$$

where

$$h_{it} = y_{it} - \omega_{it} \boldsymbol{w}_i^\top \boldsymbol{u}_{it} - \frac{1}{2} \tag{11}$$

is the precision-weighted mean, which we recognize as a kind of residual.

After augmentation, the conditional probability of x_i is an exponentiated quadratic,

$$p(x_i \mid y_i, \boldsymbol{u}_i, \omega_i, \theta_i) = \mathrm{N}(x_{i,1} \mid 0, 1) \prod_{t=2}^{T_i} \mathrm{N}(x_{i,t} \mid a_i x_{i,t-1}, \sigma_i^2) \prod_{t=1}^{T} \exp\left\{ h_{it} x_{it} - \frac{1}{2} \omega_{it} x_{it}^2 \right\}. \tag{12}$$

We recognize this as a chain-structured Gaussian graphical model — i.e., a linear Gaussian dynamical system — over x_i. We use information-form message passing algorithms implemented in Dynamax to sample the latent states from their conditional distribution given the emissions, inputs, parameters, and auxiliary PG variables.

Sampling the Pólya-gamma augmentation variables

The conditional distribution over auxiliary variables ω_{it} given the states and parameters is,

$$\begin{aligned} p(\omega_{it} \mid y_{it}, x_{it}, \boldsymbol{u}_{it}, \theta_i) &= \mathrm{PG}(\omega_{it} \mid 1, \psi_{it}) \\ \psi_{it} &= x_{it} + \boldsymbol{w}_i^\top \boldsymbol{u}_{it} \end{aligned} \tag{13}$$

There are highly efficient rejection sampling algorithms for the $\mathrm{PG}(1, \psi)$ distribution (Windle et al., 2014), but we can also use a naïve sampling algorithm based on a representation of the PG distribution as a weighted sum of gamma random variates (Polson et al., 2013).

Note that the augmentation variables (one per subject and trial) are conditionally independent

of one another and thus can be sampled in parallel. Our implementation using the JAX library (Bradbury et al., 2018) takes advantage of this opportunity for fast parallel sampling on GPUs.

Sampling the local parameters

Input weights After augmentation, log likelihood as a function of the weights $\boldsymbol{w}_i$ for subject i is,

$$\log p(y_{it} \mid x_{it}, \omega_{it}, \boldsymbol{u}_{it}, \theta_i) = \boldsymbol{h}_{it}^\top \boldsymbol{w}_i - \frac{1}{2}\boldsymbol{w}_i^\top(\omega_{it}\boldsymbol{u}_i\boldsymbol{u}_i^\top)\boldsymbol{w}_i + c \tag{14}$$

where $\boldsymbol{h}_{it} = (y_{it} - \omega_{it}x_{it} - \frac{1}{2})\boldsymbol{u}_i$ and c is constant w.r.t. $\boldsymbol{w}_i$.

Again, this quadratic log probability is conditionally conjugate with the Gaussian prior on the weights, $\boldsymbol{w}_i \sim \mathrm{N}(\boldsymbol{w}_0, \mathrm{diag}(\boldsymbol{\nu}_w^2))$. The conditional posterior distribution is,

$$\begin{aligned} p(\boldsymbol{w}_i \mid \{y_{it}, x_{it}, \boldsymbol{u}_{it}\}_{t=1}^{T_i}, \theta_i, \eta) &= \mathrm{N}(\boldsymbol{w}_i \mid \boldsymbol{J}^{-1}\boldsymbol{h}, \boldsymbol{J}^{-1}) \\ \boldsymbol{J} &= \mathrm{diag}(\boldsymbol{\nu}_w^{-2}) + \sum_{t=1}^{T_i} \omega_{it}\boldsymbol{u}_i\boldsymbol{u}_i^\top \\ \boldsymbol{h} &= \mathrm{diag}(\boldsymbol{\nu}_w^{-2})\boldsymbol{w}_0 + \sum_{t=1}^{T_i}\left(y_{it} - \omega_{it}x_{it} - \frac{1}{2}\right)\boldsymbol{u}_i \end{aligned} \tag{15}$$

Dynamics coefficient The truncated normal prior on a_i is conjugate with the linear Gaussian dynamics model. Given the latent states,

$$\begin{aligned} p(a_i \mid x_i, \eta) &= \mathrm{TruncNorm}(a_i \mid J^{-1}h, J^{-1}; [0,1]) \\ J &= \nu_a^{-2} + \sigma_i^{-2}\sum_{t=2}^{T_i} x_{i,t-1}^2 \\ h &= \nu_a^{-2}a_0 + \sigma_i^{-2}\sum_{t=2}^{T_i} x_{i,t}x_{i,t-1} \end{aligned} \tag{16}$$

This is a truncated normal distribution with mean $J^{-1}h$ and variance J^{-1}, constrained to $a_i \in [0, 1]$.

Dynamics noise variance The inverse gamma prior is conjugate with the Gaussian noise model. The conditional posterior distribution is,

$$
\begin{aligned}
p(\sigma_i^2 \mid x_i, a_i, \eta) &= \mathrm{IGa}(\alpha', \beta') \\
\alpha' &= \alpha + \frac{T_i - 1}{2} \\
\beta' &= \beta + \frac{1}{2} \sum_{t=2}^{T_i} (x_{i,t} - a_i x_{i,t-1})^2
\end{aligned}
\tag{17}
$$

Sampling the global parameters

The global parameters $\boldsymbol{w}_0$ and $\boldsymbol{\nu}_w^2$ admit closed form Gibbs updates. Under their uninformative priors,

$$
p(w_{0,j} \mid \{w_{i,j}\}_{i=1}^N, \nu_{w,j}^2) = \mathrm{N}\left(w_{0,j} \,\middle|\, \frac{1}{N} \sum_{i=1}^N w_{i,j}, \frac{\nu_{w,j}^2}{N} \right) \tag{18}
$$

and

$$
p(\nu_{w,j}^2 \mid \{w_{i,j}\}_{i=1}^N) = \mathrm{IGa}\left(\nu_{w,j}^2 \,\middle|\, \frac{N}{2}, \frac{1}{2} \sum_{i=1}^N w_{i,j}^2 \right) \tag{19}
$$

for $j = 1, \ldots, p$. Finally, we update the global parameters a_0, ν_a^2, α, and β using random-walk Metropolis-Hastings with symmetric Gaussian proposal distributions. We manually tuned the proposal variance based on trace plots of the log joint probability

Initialization

We initialize the global parameters conservatively, setting $a_0 = 0.95$, $\nu_a = 0.1$, $\boldsymbol{w}_0 = (0, \ldots, 0)^\top$, $\nu_{w,j} = 1.25$ for all j, $\alpha = 4$, and $\beta = 4 \times .4^2$. Thus, at initialization, the global prior is biased

Algorithm 1: Gibbs sampling algorithm for the hierarchical Bernoulli LDS

Input: Observed inputs and choice sequences for N subjects, $\{u_i, y_i\}_{i=1}^N$

Initialize the latent states and parameters as described above.

repeat S times to draw S samples

1. Sample the PG augmentation variables $\omega_{i,t}$ according to eq. (13).
2. Sample the local parameters θ_i according to eq. (15-17).
3. Sample latent states x_i using message passing for eq. (12).
4. Sample the global parameters $\boldsymbol{w}_0$ and $\boldsymbol{\nu}_w^2$ according to eq. (18-19).
5. Sample the global parameters a_0, ν_a^2, α, and β using random-walk Metropolis-Hastings.

end

toward producing per-subject latent states that are reasonably autocorrelated. Then we sample the parameters and states from their conditional distributions given the global prior, without conditioning on the data.

More informative initializations are possible. For example, one could fit a logistic regression (without time-varying latent states) for each subject to obtain initial weights $\boldsymbol{w}_i$, and then initialize the per-subject latent states to $x_i = (0, \ldots, 0)^\top$. based on the initial weights, we can estimate the global weight mean w_0 and variance σ_w^2. Combined with the initializations for a_0, ν_a, α, and β above, this procedure should effectively start the Gibbs sampler near a reasonable basin of the posterior.

Full algorithm

The complete algorithm is given by Algorithm 1.

Supplemental Figures

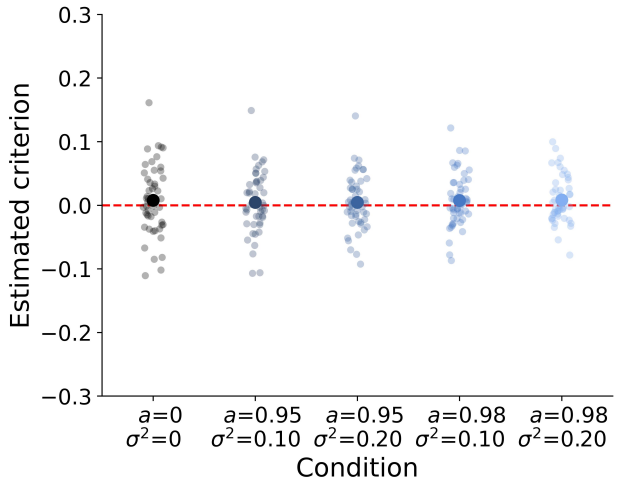

Figure 8: On average the criterion is correctly estimated when using the traditional method assuming a stable criterion.