Evidence accumulation and change rate inference in dynamic environments

Adrian E. Radillo¹, Alan Veliz-Cuba², Krešimir Josić^{1, 3, 4*}, and Zachary P. Kilpatrick^{1, 5*}

¹Department of Mathematics, University of Houston, Houston, TX 77204

²Department of Mathematics, University of Dayton, Dayton, OH 45469

³Department of Biology and Biochemistry, University of Houston, Houston, TX 77204

⁴Department of BioSciences, Rice University, Houston, TX 77251

⁵Department of Applied Mathematics, University of Colorado, Boulder, CO 80309

*equal contribution

Keywords: decision making, Bayesian inference, dynamic environment, changepoint, moment closure

Abstract

In a constantly changing world, animals must account for environmental volatility when making decisions. To appropriately discount older, irrelevant information, they need to learn the rate at which the environment changes. We develop an ideal observer model capable of inferring the present state of the environment along with its rate of change. Key to this computation is updating the posterior probability of all possible changepoint counts. This computation can be challenging, as the number of possibilities grows rapidly with time. However, we show how the computations can be simplified in the continuum limit by a moment closure approximation. The resulting low-dimensional system can be used to infer the environmental state and change rate with accuracy comparable to the ideal observer. The approximate computations can be performed by a neural network model via a rate-correlation based plasticity rule. We thus show how optimal observers accumulate evidence in changing environments, and map this computation to reduced models which perform inference using plausible neural mechanisms.

1 Introduction

Animals continuously make decisions in order to find food, identify mates, and avoid predators. However, the world is seldom static. Information that was critical yesterday may be of little value now. Thus, when accumulating evidence to decide on a course of action, animals weight new evidence more strongly than old (Pearson et al, 2011). The rate at which the world changes determines the rate at which an individual should discount previous information (Deneve, 2008; Veliz-Cuba et al, 2016). For instance, when actively tracking prey, a predator may only use visual information obtained within the last second (Olberg et al, 2000; Portugues and Engert, 2009), while social insect colonies integrate evidence that can be hours to days old when deciding on a new home

site (Franks et al, 2002). However, the timescale at which the environment changes is unlikely to be known in advance. To make accurate decisions, animals must also learn how rapidly their environment changes (Wilson et al, 2010).

Evidence accumulators are often used to model decision processes in static and fluctuating environments (Smith and Ratcliff, 2004; Bogacz et al, 2006). These models show how noisy observations can be accumulated to provide a probability that one among multiple alternatives is correct (Gold and Shadlen, 2007; Beck et al, 2008). They explain a variety of behavioral data (Ratcliff and McKoon, 2008; Brunton et al, 2013), and electrophysiological recordings suggest that neural activity can reflect the accumulation of evidence (Huk and Shadlen, 2005; Kira et al, 2015). Since normative evidence accumulation models determine the belief of an ideal observer, they also show the best way to integrate noisy sensory measurements, and can tell us if and how animals fail to use such information optimally (Bogacz et al, 2006; Beck et al, 2008).

Early decision-making models focused on decisions between two choices in a static environment (Wald and Wolfowitz, 1948; Gold and Shadlen, 2007). Recent studies have extended this work to more ecologically relevant situations, including multiple alternatives (Churchland et al, 2008; Krajbich and Rangel, 2011), multidimensional environments (Niv et al, 2015), and cases where the correct choice (McGuire et al, 2014; Glaze et al, 2015), or context (Shvartsman et al, 2015), changes in time. In these cases, normative models are more difficult to derive and analyze (Wilson and Niv, 2011), and their dynamics are more complex. However, methods of sequential and stochastic analysis are still useful in understanding their properties (Wilson et al, 2010; Veliz-Cuba et al, 2016).

We examine the case of a changing environment where an optimal observer discounts prior evidence at a rate determined by environmental volatility. Experiments suggest that humans learn the rate of environmental fluctuations to make choices nearly optimally (Glaze et al, 2015). During dynamic foraging experiments where the choice with the highest reward changes in time, monkeys also appear to use an evidence discounting strategy suited to the environmental change rate (Sugrue et al, 2004).

However, most previous models have assumed that the rate of change of the environment is known ahead of time to the observer (Glaze et al, 2015; Veliz-Cuba et al, 2016). Wilson et al (2010) developed a model of an observer that infers the rate of environmental change from observations. To do so, the observer computes a joint posterior probability of the state of the environment, and the time since the last change in the environment. With more measurements, such observers improve their estimates of the change rate, and are therefore better able to predict the environmental state. Inference of the change rate is most important when an observer makes fairly noisy measurements, and cannot determine the current state from a single observation.

We extend previous accumulator models of decision making to the case of multiple, discrete choices with asymmetric, unknown transition rates between them. We assume that the observer is primarily interested in the current state of the environment, often referred to as the *correct choice* in decision-making models (Bogacz et al, 2006). Therefore, we show how an ideal observer can use sensory evidence to infer the rates at which the environment transitions between states, and simultaneously use these inferred rates to discount old evidence and determine the present environmental state.

Related models have been studied before (Wilson et al, 2010; Adams and MacKay,

2007). However, they relied on the assumption that after a change the new state does not depend on the previous state. This excludes the possibility of a finite number of states: For example, with two choices knowledge of the present state determines with complete certainty the state after a change, and the two are thus not independent. In the case of a finite number of choices our algorithm is simpler than previous ones. The observer only needs to compute a joint probability of the environmental state and the number of changepoints.

The storage needed to implement our algorithms grows rapidly with the number of possible environmental states. However, we show that moment closure methods can be used to decrease the needed storage considerably, albeit at the expense of accuracy and the representation of higher order statistics. Nonetheless, when measurement noise is not too large, these approximations can be used to estimate the most likely transition rate, and the current state of the environment. This motivates a physiologically plausible neural implementation for the present computation: We show that a Hebbian learning rule which shapes interactions between multiple neural populations representing the different choices, allows a network to integrate inputs nearly optimally. Our work therefore links statistical principles for optimal inference with stochastic neural rate models that can adapt to the environmental volatility to make near-optimal decisions in a changing environment.

2 Optimal evidence accumulation for known transition rates

We start by revisiting the problem of inferring the current state of the environment from a sequence of noisy observations. We assume that the number of states is finite, and the state of the environment changes at times unknown to the observer. We first review the case when the rate of these changes is known to the observer. In later sections we will assume that these rates must also be learned. Following Veliz-Cuba et al (2016), we derived a recursive equation for the likelihoods of the different states, and an approximating stochastic differential equation (SDE). Similar derivations were presented for decisions between two choices by Deneve (2008), and Glaze et al (2015).

An ideal observer decides between N choices, based on successive observations at times t_n ($n = 1, 2, \dots$). We denote each possible choice by H^i , ($i = 1, \dots, N$), with H_n being the correct choice at time t_n . The *transition rates* ϵ^{ij} , $i \neq j$, correspond to the *known* probabilities that the state changes between two observations: $\epsilon^{ij} = P(H_n = H^i | H_{n-1} = H^j)$. The observer makes measurements, ξ_n , at times t_n with known conditional probability densities $f^i(\xi) = P(\xi_n = \xi | H_n = H^i)$. Here, and elsewhere, we assume that the observations are conditionally independent. We also abuse notation slightly by using $P(\cdot)$ to denote a probability, or the value of a probability density function, depending on the argument. We use explicit notation for the probability density function when there is a potential for confusion.

We denote by $\xi_{j:n}$ the vector of observations $(\xi_j, \dots, \xi_n)$, and by $P_n(\cdot)$ the conditional probability $P(\cdot | \xi_{1:n})$. To make a decision, the observer can compute the index that maximizes the *posterior probability*, $\hat{i} = \operatorname{argmax}_i P_n(H_n = H^i)$. Therefore $H^{\hat{i}}$ is the most probable state, given the observations $\xi_{1:n}$.

A recursive equation for the update of each of the probabilities $P_n(H_n = H^i)$ after

the n^{th} observation has the form (Veliz-Cuba et al, 2016)

$$P_n(H_n = H^i) \propto f^i(\xi_n) \sum_{j=1}^N \epsilon^{ij} P_{n-1}(H_{n-1} = H^j) \quad (i = 1, \dots, N). \quad (1)$$

Thus, the transition rates, ϵ^{ij} , provide the weights of the previous probabilities in the update equation. Unless transition rates are large or observations very noisy, the probability $P_n(H_n = H^i)$ grows, and can be used to identify the present environmental state. However, with positive transition rates, the posterior probabilities tend to saturate at a value below unity. Strong observational evidence that contradicts an observer’s current belief can cause the observer to change their belief subsequently. Such contradictory evidence typically arrives after a change in the environment.

Taking logarithms, $x_n^i := \ln P_n(H_n = H^i)$, and assuming that evidence from each observation, as well as the time between observations, $\Delta t := t_n - t_{n-1}$, are small, we can approximate the discrete process Eq. (1) with an SDE,

$$d\mathbf{x}(t) = \mathbf{g}(t)dt + \Lambda(t)d\mathbf{W}_t + \mathbf{K}(\mathbf{x}(t))dt, \quad (2)$$

where $\mathbf{K}_i(\mathbf{x}) = \sum_{j \neq i} (\epsilon^{ij} e^{x_j - x_i} - \epsilon^{ji})$, and the drift, $g_i(t) = \frac{1}{\Delta t} E_{\xi} [\ln f^i(\xi) | H(t)]$, depends on the state of the environment at time t , $H(t) \in \{H^1, \dots, H^N\}$. The drift, $g_i(t)$, is largest when the environment is in state H^i . The noise covariance, $\Lambda(t)\Lambda(t)^T = \Sigma(t)$, has entries $\Sigma_{ij} = \frac{1}{\Delta t} \text{Cov}_{\xi} [\ln f^i(\xi), \ln f^j(\xi) | H(t)]$, while $\mathbf{W}_t$ is a vector of independent Wiener processes.

The nonlinear term, $\mathbf{K}_i(\mathbf{x})$, implies that in the absence of noise Eq. (2) has a stable fixed point, and older evidence is discounted. Such continuum models of evidence accumulation are useful because they are amenable to the methods of stochastic analysis (Bogacz et al, 2006). Linearization of the SDE provides insights into the system’s local dynamics (Glaze et al, 2015; Veliz-Cuba et al, 2016), and can be used to implement the inference process in model neural networks (Bogacz et al, 2006; Veliz-Cuba et al, 2016).

We next extend this approach to the case when the observer infers the transition rates, ϵ^{ij} , from measurements.

3 Environments with symmetric transition rates

We first derive the ideal observer model when the unknown transition rates are symmetric, $\epsilon^{ij} \equiv \text{constant}$ when $j \neq i$, and $\epsilon^{ii} := 1 - (N - 1)\epsilon^{ij}$. This simplifies the derivation, since the observer only needs to estimate a single changepoint count. The general, asymmetric case discussed in Section 4 follows the same idea, but the derivation is more involved since the observer must estimate multiple counts.

Our problem differs from previous studies in two key ways (Adams and MacKay, 2007; Wilson et al, 2010): We assume the observer tries to identify the most likely state of the environment at time t_n . To do so the observer computes the joint conditional probability, $P_n(H_n, a_n)$, of the current state, H_n , and the number of environmental changes since beginning the observation, a_n . Previous studies focused on obtaining the predictive distribution, $P_n(H_{n+1})$. The two distributions are closely related, as $P_n(H_{n+1}) = \sum_{H_n} P_n(H_{n+1} | H_n) P_n(H_n)$.

More importantly, Adams and MacKay (2007); Wilson et al (2010) implicitly assumed that only observations since the last changepoint provide information about the current environmental state. That is, if the time since the last changepoint – the current run-length, r_n – is known to the observer, then all observations before that time can be discarded:

$$P(H_n | \xi_{1:n}, r_n) = P(H_n | \xi_{n-r_n:n}).$$

This follows from the assumption that the state after a change is conditionally independent of the state that preceded it. We assume that the number of environmental states is finite. Hence this independence assumption does not hold: Intuitively, if observations prior to a changepoint indicate the true state is H^j , then states $H^i, i \neq j$ are more likely after the changepoint.

Adams and MacKay (2007); Wilson et al (2010) derive a probability update equation for the *run length*, and the number of *changepoints*, and use this equation to obtain the predictive distribution of future observations. We show that it is not necessary to compute run length probabilities when the number of environmental states is finite. Instead we derive a recursive equation for the joint probability of the current state, H_n , and number of changepoints, a_n . As a result, the total number of possible states grows as $N \cdot n$ (linearly in n) where N is the fixed number of environmental states H^i , rather than n^2 (quadratically in n) as in Wilson et al (2010).

3.1 Symmetric 2-state process

We first derive a recursive equation for the probability of two alternatives, $H_n \in \{H^\pm\}$, in a changing environment, where the change process is memoryless, and the change rate, $\epsilon := P(H_n = H^\mp | H_{n-1} = H^\pm)$, is symmetric and initially unknown to the observer (See Fig. 1A). The most probable choice given the observations up to a time, t_n , can be obtained from the log of the posterior odds ratio $L_n = \log \left(\frac{P_n(H_n=H^+)}{P_n(H_n=H^-)} \right)$. The sign of L_n indicates which option is more likely, and its magnitude indicates the strength of this evidence (Bogacz et al, 2006; Gold and Shadlen, 2007). Old evidence should be discounted according to the inferred environmental volatility. Since this is unknown, an ideal observer computes the probability distribution of the change rate, ϵ (See Fig. 1C), along with the probability of environmental states.

Let a_n be the number of changepoints, and $b_n = n - 1 - a_n$ the count of non-changepoints between times t_1 and t_n ($n = 1, 2, \dots$) (See Fig. 1B). The process $\{a_n\}_{n \geq 1}$ is a pure birth process with birth rate ϵ . The observer assumes no changes prior to the start of observation, $P(a_1 = 0) = 1$, and must make at least two observations, ξ_1 and ξ_2 , to detect a change.

To develop an iterative equation for the joint conditional probability density, $P_n(H_n, a_n)$, given the n observations $\xi_{1:n}$, we begin by marginalizing over these quantities at the time of the previous observation, t_{n-1} , for $n > 1$ (See Appendix 7.1 for details):

$$P_n(H_n, a_n) = \frac{P(\xi_{1:n-1})}{P(\xi_{1:n})} P(\xi_n | H_n) \sum_{H_{n-1}=H^\pm} \sum_{a_{n-1}=0}^{n-2} P(H_n, a_n | H_{n-1}, a_{n-1}) P_{n-1}(H_{n-1}, a_{n-1}). \quad (3)$$

Figure 1: Online inference of the change rate in a dynamic environment. (A) The environment alternates between states H^+ and H^- with transition probabilities ϵ^+ , ϵ^- . The state of the environment determines $f^\pm(\xi) = P(\xi|H^\pm)$, which we represent as Gaussian densities. (B) A sample path of the environment (color bar) together with the first ten values of the actual changepoint count, a_n , and non-changepoint count, b_n . (C) Evolution of the conditional probabilities, $P(\epsilon|a_n)$, corresponding to the changepoint count from panel B, until $t_n = t_{100}$. The dashed red line indicates the value of ϵ in the simulation. The densities are scaled so that each equals 1 at the mode.

With two choices we have the following relationships for all $n > 1$:

$$H_n = H_{n-1} \Leftrightarrow a_n = a_{n-1}, \quad \text{and} \quad H_n \neq H_{n-1} \Leftrightarrow a_n = a_{n-1} + 1. \quad (4)$$

The term $P(H_n, a_n | H_{n-1}, a_{n-1})$ in Eq. (3) is therefore nonzero only if either, $H_{n-1} = H_n$, and $a_{n-1} = a_n$, or $H_{n-1} \neq H_n$ and $a_{n-1} = a_n - 1$: If the system is in the joint state (H_{n-1}, a_{n-1}) at t_{n-1} , then at t_n it can either (a) transition to $(H_n \neq H_{n-1}, a_n = a_{n-1} + 1)$ or (b) remain at $(H_n = H_{n-1}, a_n = a_{n-1})$. This observation is central to the message-passing algorithm described in (Adams and MacKay, 2007; Wilson et al, 2010), with probability mass flowing from lower to higher values of a according to a pure birth process (See Fig. 2A). We can thus simplify Eq. (3), leaving only two terms in the double sum. Writing $P_n(H^\pm, a)$ for $P_n(H_n = H^\pm, a_n = a)$, and similarly for any conditional probabilities, we have for $n > 1$:

$$P_n(H^\pm, a) = \frac{P(\xi_{1:n-1})}{P(\xi_{1:n})} f^\pm(\xi_n) [P(H^\pm, a | H_{n-1} = H^\pm, a_{n-1} = a) \cdot P_{n-1}(H^\pm, a) + P(H^\pm, a | H_{n-1} = H^\mp, a_{n-1} = a - 1) \cdot P_{n-1}(H^\mp, a - 1)]. \quad (5)$$

We must also specify *initial conditions* at time t_1 , and *boundary values* when $a \in \{0, n - 1\}$ for these equations. At t_1 we have $P(a_1 = 0) = 1$. Therefore,

$$P_1(H^\pm, 0) = \frac{1}{P(\xi_1)} f^\pm(\xi_1) P_0(H^\pm), \quad (6)$$

and $P_1(H^\pm, a) = 0$ for $a \neq 0$. Here $P_0(H^\pm)$ is the prior over the two choices. The probability $P(\xi_1)$ is unknown to the observer. However, similar to the ratio $\frac{P(\xi_{1:n-1})}{P(\xi_{1:n})}$ in Eq. (5), $P(\xi_1)$ acts as a normalization constant and does not appear in the posterior odds

ratio, R_n (See Eq. (18) below). Finally, at all future times $n > 1$, we have separate equations at the boundaries,

$$P_n(H^\pm, 0) = \frac{P(\xi_{1:n-1})}{P(\xi_{1:n})} f^\pm(\xi_n) P(H^\pm, 0 | H_{n-1} = H^\pm, a_{n-1} = 0) P_{n-1}(H^\pm, 0), \quad (7)$$

and,

$$P_n(H^\pm, n-1) = \frac{P(\xi_{1:n-1})}{P(\xi_{1:n})} f^\pm(\xi_n) P(H^\pm, n-1 | H_{n-1} = H^\mp, a_{n-1} = n-2) \times P_{n-1}(H^\mp, n-2). \quad (8)$$

We next compute $P(H_n, a_n | H_{n-1}, a_{n-1})$ in Eq. (3), with $n > 1$, by marginalizing over all possible transition rates $\epsilon \in [0, 1]$:

$$P(H_n, a_n | H_{n-1}, a_{n-1}) = \int_0^1 P(H_n, a_n | \epsilon, H_{n-1}, a_{n-1}) P(\epsilon | H_{n-1}, a_{n-1}) d\epsilon. \quad (9)$$

Note that $P(\epsilon | H_{n-1}, a_{n-1}) = P(\epsilon | a_{n-1})$, so we need the distribution of ϵ , given a_{n-1} changepoints, for all $n > 1$. We assume that prior to any changepoint observations — that is at time t_1 — the rates follow a Beta distribution with hyperparameters $a_0, b_0 > 0$ (See also Sections 3.1 and 3.2 in Wilson et al (2010)),

$$P_0(\epsilon) = \beta(\epsilon; a_0, b_0) := \frac{\epsilon^{a_0-1} (1-\epsilon)^{b_0-1}}{B(a_0, b_0)},$$

where β denotes the probability density of the associated Beta distribution, and $B(x, y) := \int_0^1 \epsilon^{x-1} (1-\epsilon)^{y-1} d\epsilon$ is the beta function. For any $n > 1$, the random variable $a_n | \epsilon$ follows a Binomial distribution with parameters $(n-1, \epsilon)$, for which the Beta distribution is a conjugate prior. The posterior over the change rate when the changepoint count is known at time $n > 1$ is therefore:

$$\epsilon | a_n \sim \text{Beta}(a_0 + a_n, b_0 + b_n). \quad (10)$$

For simplicity, we assume that prior to any observations, the probability over the transition rates is uniform, $P_0(\epsilon) = 1$ for all $\epsilon \in [0, 1]$, and therefore $a_0 = b_0 = 1$.

We now return to Eq. (9) and use the definition of the transition rate, ϵ , (See Fig. 1) to find:

$$P(H_n, a_n | \epsilon, H_{n-1}, a_{n-1}) = \begin{cases} 1 - \epsilon & H_n = H_{n-1} \text{ \& } a_n = a_{n-1}, \\ \epsilon & H_n \neq H_{n-1} \text{ \& } a_n = a_{n-1} + 1, \\ 0 & \text{otherwise.} \end{cases} \quad (11)$$

Eq. (9) can therefore be rewritten using two integrals, depending on the values of (H_n, a_n) and (H_{n-1}, a_{n-1}) ,

$$P(H^\pm, a | H_{n-1} = H^\pm, a_{n-1} = a) = \int_0^1 (1 - \epsilon) \beta(\epsilon; a_{n-1} + 1, b_{n-1} + 1) d\epsilon, \quad (12)$$

and similarly for $P(H^\pm, a | H_{n-1} = H^\mp, a_{n-1} = a - 1)$.

Figure 2: Inference of the state, H^i , and change rate, ϵ . (A) The joint posterior probability, $P_n(H^\pm, a)$, is propagated along a directed graph according to Eq. (17). Only paths corresponding to the initial condition $(H_1, a_1) = (H^+, 0)$ are shown. (B) A sample sequence of environmental states (color bar, top) together with the first ten observations $\xi_1, \dots, \xi_{10}$ (blue dots), for $\epsilon^+ = \epsilon^- = 0.1$. Superimposed in black (right y-axis) is the log posterior odds ratio L_n as a function of time. (C) Evolution of the posterior over a_n (gray scale). The posterior mean (red) converges to the expected number of change-points $\epsilon(n-1)$ (dashed line). (D) Evolution of the posterior over the change rate ϵ (gray scale). The posterior mean (red) converges to the true value (dashed line) and the variance diminishes with the number of observations.

The mean of the Beta distribution, for $n > 1$, can be expressed in terms of its two parameters:

$$\hat{\epsilon}_{n-1}(a_{n-1}) := \mathbb{E}[\epsilon|a_{n-1}] = \frac{a_{n-1} + 1}{a_{n-1} + b_{n-1} + 2}. \quad (13)$$

We denote this expected value by $\hat{\epsilon}_{n-1}(a_{n-1})$ as it represents a point estimate of the change rate ϵ at time t_{n-1} when the changepoint count is a_{n-1} , $n > 1$. Since $a_{n-1} + b_{n-1} = n - 2$, we have:

$$\hat{\epsilon}_{n-1}(a_{n-1}) = \frac{a_{n-1} + 1}{n}. \quad (14)$$

The expected transition rate, $\hat{\epsilon}_{n-1}(a_{n-1})$, is thus determined by the ratio between the present changepoint count and the number of timesteps, n . Leaving a_0 and b_0 as parameters in the prior gives $\hat{\epsilon}_{n-1}(a_{n-1}) = (a_{n-1} + a_0)/(n - 2 + a_0 + b_0)$. Using the definition in Eq. (14), it follows from Eq. (12) that:

$$P(H^\pm, a|H_{n-1} = H^\pm, a_{n-1} = a) = 1 - \hat{\epsilon}_{n-1}(a), \quad (15)$$

$$P(H^\pm, a|H_{n-1} = H^\mp, a_{n-1} = a - 1) = \hat{\epsilon}_{n-1}(a - 1). \quad (16)$$

Eqs. (15) and (16), which are illustrated in Fig. 2A, can in turn be substituted into Eq. (5) to yield, for all $n > 1$:

$$P_n(H^\pm, a) = \frac{P(\xi_{1:n-1})}{P(\xi_{1:n})} f^\pm(\xi_n) [(1 - \hat{\epsilon}_{n-1}(a)) \cdot P_{n-1}(H^\pm, a) + \hat{\epsilon}_{n-1}(a - 1) \cdot P_{n-1}(H^\mp, a - 1)]. \quad (17)$$

The initial conditions and boundary equations for this recursive probability update have already been described in Eqs. (6–8). Eq. (17) is the equivalent of Eq. (3) in Adams and MacKay (2007), and Eq. (3.7) in Wilson et al (2010). However, here the observer does not need to estimate the length of the interval since the last changepoint. We demonstrate the inference process defined by Eq. (17) in Fig. 2.

The observer can compute the posterior odds ratio by marginalizing over the changepoint count:

$$\begin{aligned} R_n &:= \frac{P_n(H^+)}{P_n(H^-)} = \frac{\sum_{a=0}^{n-1} P_n(H^+, a)}{\sum_{a=0}^{n-1} P_n(H^-, a)} \\ &= \frac{f^+(\xi_n)}{f^-(\xi_n)} \cdot \frac{\sum_{a=0}^{n-1} [(1 - \frac{a+1}{n}) \cdot P_{n-1}(H^+, a) + \frac{a}{n} \cdot P_{n-1}(H^-, a - 1)]}{\sum_{a=0}^{n-1} [(1 - \frac{a+1}{n}) \cdot P_{n-1}(H^-, a) + \frac{a}{n} \cdot P_{n-1}(H^+, a - 1)]}. \end{aligned} \quad (18)$$

Here $\log(R_n) = L_n > 0$ implies that $H_n = H^+$ is more likely than $H_n = H^-$ (See Fig. 2B). Note that $P(\xi_{1:n-1})/P(\xi_{1:n})$ and $1/P(\xi_1)$ need not be known to the observer to obtain the most likely choice.

A posterior distribution of the transition rate ϵ can also be derived from Eq. (17) by marginalizing over (H_n, a_n) ,

$$P_n(\epsilon) = \sum_{s=\pm} \sum_{a=0}^{n-1} P(\epsilon|a_n = a) P_n(H^s, a), \quad (19)$$

Figure 3: The performance of the inference algorithm. (A) Performance under the interrogation paradigm measured as the percentage of correct responses at the interrogation time. Here and in the next panel $\epsilon = 0.05$, and $\text{SNR} = 1$. The black curve represents the performance of an ideal observer who infers the change rate from measurements. The green curves, represent the performance of observers that assume a fixed change rate (0.3, 0.15, 0.05, 0.03 from darker to lighter, see Eq. (1)). The solid green line corresponds to the true rate, dashed lines to erroneous rates. (B) The green curve represents the performance at interrogation time t_{300} of an observer that assumes a fixed change rate. The red star marks the maximum of this curve. The horizontal black curves represent the performance at times $t_{40}, t_{100}, t_{200}, t_{300}$ (from bottom to top) of the observer that learns the change rate. (C) The accuracy as a function of the average threshold hitting time in the free response protocol. Here $\epsilon = 0.1$, and $\text{SNR} = 0.75$. See Appendix 7.2 for details on numerical simulations. See also Fig. 3 in Veliz-Cuba et al (2016).

where $P(\epsilon|a_n)$ is given by the Beta distribution prior Eq. (10). The expected rate is therefore:

$$\bar{\epsilon} := \int_0^1 \epsilon P_n(\epsilon) d\epsilon = \sum_{s=\pm} \sum_{a_n=0}^{n-1} \int_0^1 \epsilon P(\epsilon|a_n) P_n(H^s, a_n) d\epsilon = \sum_{s=\pm} \sum_{a_n=0}^{n-1} \frac{a_n + 1}{n + 1} P_n(H^s, a_n). \quad (20)$$

Explicit knowledge of the transition rate, ϵ , is not used in the inference process described by Eq. (17). However, computing it allows us to evaluate how the observer's estimate converges to the true transition rate (See Fig. 2D). We will also relate this estimate to the coupling strength between neural populations in the model described in Section 6.

We conjecture that when measurements are noisy, the variance of the distribution $P_n(\epsilon)$ does not converge to a point mass at the true rate, ϵ , in the limit of infinitely many observations, $n \rightarrow \infty$, *i.e.* the estimate of ϵ is not consistent. As we have shown, to infer the rate we need to infer the parameter of a Bernoulli variable. It is easy to show that the posterior over this parameter converges to a point mass at the actual rate value if the probability of misclassifying the state is known to the observer (Djuric and Huang, 2000). However, when the misclassification probability is not known, the variance of the posterior remains positive even in the limit of infinitely many observations. In our case, when measurements are noisy, the observer does not know the exact number of change points at finite time. Hence, the observer does not know exactly how to weight

previous observations to make an inference about the current state. As a result, the probability of misclassifying the current state may not be known. We conjecture that this implies that even in the limit $n \rightarrow \infty$ the posterior over ϵ has positive variance (See Fig. 2D).

In Fig. 3 we compare the performance of this algorithm in three cases: when the observer knows the true rate (point mass prior over the true rate ϵ); when the observer assumes a wrong rate (point mass prior over an erroneous ϵ); and when the observer learns the rate from measurements (flat prior over ϵ). We define performance as the probability of a correct decision.

Under the interrogation protocol, the observer infers the state of the environment at a fixed time. As expected, performance increases with interrogation time, and is highest if the observer uses the true rate (See Fig. 3A, also Eq. (1) above). Performance plateaus quickly when the observer assumes a fixed rate, and more slowly if the rate is learned. The performance of observers that learn the rate slowly increases toward that of observers who know the true rate. In panel B, we present the performance of the unknown-rate algorithm at 4 different times ($t_{40}, t_{100}, t_{200}, t_{300}$) and compare it to the asymptotic values with different assumed rates (green curves).

Note, an observer that assumes an incorrect change rate can still perform near optimally (e.g., curve for 0.03 in Fig. 3A), especially when the signal-to-noise ratio (SNR) is quite high. The SNR is the difference in means of the likelihoods divided by their common standard deviation. Change rate inference is more effective at lower SNR values, in which case multiple observations are needed for an accurate estimate of the present state. However, at very low SNR values the observer will not be able to substantially reduce uncertainty about the change rate, resulting in high uncertainty about the state.

In the free response protocol, the observer makes a decision when the log odds ratio reaches a predefined threshold. In Fig. 3C, we present simulation results for this protocol in a format similar to Fig. 3A, with empirical performance as a function of average hitting time. Each performance level corresponds to unique log odds threshold. Similar to the interrogation protocol (Fig. 3A), performance of the free response protocol saturates much more quickly for an observer that fixes their change rate estimate than on that infers this rate over time.

3.2 Symmetric multistate process

We next consider evidence accumulation in an environment with an arbitrary number of states, $\{H^1, H^2, \dots, H^N\}$, with symmetric transition probabilities, $\epsilon^{ij} \equiv \text{constant}$, whenever $i \neq j$. We define $\epsilon := (N-1)\epsilon^{ij}$ for any $i \neq j$, so that the probability of remaining in the same state becomes $\epsilon^{ii} = 1 - \epsilon$, for all $i = 1, \dots, N$. The symmetry in transition rates means that an observer still only needs to track the total number of changepoints, a_n , as in Section 3.1.

Eqs. (3-4) remain valid with N possible choices, $\{H^1, \dots, H^N\}$. When $n > 1$, the

double sum in Eq. (3) simplifies to:

$$P_n(H^i, a) = \frac{P(\xi_{1:n-1})}{P(\xi_{1:n})} f^i(\xi_n) \left[P(H^i, a | H_{n-1} = H^i, a_{n-1} = a) \cdot P_{n-1}(H^i, a) + \sum_{j \neq i} P(H^i, a | H_{n-1} = H^j, a_{n-1} = a - 1) \cdot P_{n-1}(H^j, a - 1) \right]. \quad (21)$$

As in Section 3.1, we have $P_1(H^i, 0) = f^i(\xi_1)P_0(H^i)/P(\xi_1)$ and $P_1(H^i, a_1) = 0$ for $a_1 \neq 0$, where $P_0(H^i)$ describes the observer's belief prior to any observations. At all future times, $n > 1$, we have at the boundaries for all $i = 1, \dots, N$:

$$P_n(H^i, 0) = \frac{P(\xi_{1:n-1})}{P(\xi_{1:n})} f^i(\xi_n) P(H^i, 0 | H_{n-1} = H^i, a_{n-1} = 0) P_{n-1}(H^i, 0),$$

and,

$$P_n(H^i, n - 1) = \frac{P(\xi_{1:n-1})}{P(\xi_{1:n})} f^i(\xi_n) \sum_{j \neq i} P(H^i, n - 1 | H_{n-1} = H^j, a_{n-1} = n - 2) P_{n-1}(H^j, n - 2).$$

Eq. (9) remains unchanged and we still have $P(\epsilon | H_{n-1}, a_{n-1}) = P(\epsilon | a_{n-1})$. Furthermore, assuming a Beta prior on the change rate, Eq. (10) remains valid, and Eq. (11) is replaced by:

$$P(H_n, a_n | \epsilon, H_{n-1}, a_{n-1}) = \begin{cases} 1 - \epsilon & H_n = H_{n-1} \text{ \& } a_n = a_{n-1}, \\ \epsilon / (N - 1) & H_n \neq H_{n-1} \text{ \& } a_n = a_{n-1} + 1, \\ 0 & \text{otherwise.} \end{cases} \quad (22)$$

The integral from Eq. (9) gives, once again, the mean of the Beta distribution, $\hat{\epsilon}_{n-1}(a)$, defined in Eqs. (13-14). As in Section 3.1, $\hat{\epsilon}_{n-1}(a_{n-1})$ is a point estimate of the change rate ϵ at time t_{n-1} when the changepoint count is a_{n-1} . We have,

$$P(H_n, a_n | H_{n-1}, a_{n-1}) = \begin{cases} 1 - \hat{\epsilon}_{n-1}(a_n) & H_n = H_{n-1} \text{ \& } a_n = a_{n-1}, \\ \hat{\epsilon}_{n-1}(a_n - 1) / (N - 1) & H_n \neq H_{n-1} \text{ \& } a_n = a_{n-1} + 1, \\ 0 & \text{otherwise,} \end{cases} \quad (23)$$

and the main probability update equation is now:

$$P_n(H^i, a) = \frac{P(\xi_{1:n-1})}{P(\xi_{1:n})} f^i(\xi_n) \left[(1 - \hat{\epsilon}_{n-1}(a_n)) \cdot P_{n-1}(H^i, a) + \frac{\hat{\epsilon}_{n-1}(a_n - 1)}{N - 1} \sum_{j \neq i} P_{n-1}(H^j, a - 1) \right]. \quad (24)$$

The observer can infer the most likely state of the environments, by computing the index that maximizes the *posterior probability*, marginalizing over all changepoint counts,

$$\hat{i} = \operatorname{argmax}_i P_n(H^i) = \operatorname{argmax}_i \left(\sum_{a=0}^{n-1} P_n(H^i, a) \right).$$

The observer can also compute the posterior probability $P_n(\epsilon)$ of the transition rate, ϵ , by marginalizing over all states H_n and changepoint counts a_n , as in Eq. (19). Furthermore, a point estimate of ϵ is given by the mean of the posterior after marginalizing, as in Eq. (20).

4 Environments with asymmetric transition rates

In this section, we depart from the framework of Adams and MacKay (2007), and Wilson et al (2010), and consider unequal transition rates between states. This includes the possibility that some transitions are not allowed. We consider an arbitrary number, N , of states with unknown transition rates, ϵ^{ij} , between them. The switching process between the states is again memoryless, so that H_n is a stationary, discrete-time Markov chain with finite state space, $\Omega := \{H^1, \dots, H^N\}$. We write the (unknown) transition matrix for this chain as a left stochastic matrix,

$$\epsilon := \begin{pmatrix} \epsilon^{11} & \dots & \epsilon^{1N} \\ \vdots & \ddots & \vdots \\ \epsilon^{N1} & \dots & \epsilon^{NN} \end{pmatrix}, \quad (25)$$

where $\epsilon^{ij} = P(H_n = H^i | H_{n-1} = H^j)$, with $i, j \in \{1, \dots, N\}$. We denote by ϵ^i the i -th column of the matrix ϵ , and similarly for other matrices. Each such column sums to 1. We define the changepoint counts matrix at time t_n as,

$$\mathbf{a}_n := \begin{pmatrix} a_n^{11} & \dots & a_n^{1N} \\ \vdots & \ddots & \vdots \\ a_n^{N1} & \dots & a_n^{NN} \end{pmatrix}, \quad (26)$$

where a_n^{ij} is the number of transitions from state j to state i up to time t_n . There can be a maximum of $n - 1$ transitions at time t_n . For a fixed $n \geq 1$, all entries in $\mathbf{a}_n$ are nonnegative and sum to $n - 1$, i.e. $\sum_{i,j} a_n^{ij} = n - 1$. As in the symmetric case, the changepoint matrix at time t_1 must be the zero matrix, $\mathbf{a}_1 = \mathbf{0}$.

We will show that our inference algorithm assigns positive probability only to changepoint matrices that correspond to possible transition paths between the states $\{H^1, \dots, H^N\}$. Many nonnegative integer matrices with entries that sum to $n - 1$ are not possible changepoint matrices $\mathbf{a}_n$. A combinatorial argument shows that when $N = 2$ state case scales as n^2 , the number of possible pairs, $(H_n, \mathbf{a}_n)$, grows quadratically with the number of steps, n , to leading order. It can also be shown that the growth is polynomial for $N > 2$, although we do not know the growth rate in general (See Fig. 4B). An ideal observer has to assign a probability of each of these states which is much more demanding than in the symmetric rate case where the number of possible states grows linearly.

We next derive an iterative equation for $P_n(H_n, \mathbf{a}_n)$, the joint probability of the state H_n , and an allowable combination of the $N(N - 1)$ changepoint counts (off-diagonal terms of $\mathbf{a}_n$), and N non-changepoint counts (diagonal terms of $\mathbf{a}_n$). The derivation is similar to the symmetric case: For $n > 1$, we first marginalize over H_{n-1} and $\mathbf{a}_{n-1}$,

$$P_n(H_n, \mathbf{a}_n) = \frac{1}{P(\xi_{1:n})} \sum_{H_{n-1}, \mathbf{a}_{n-1}} P(\xi_{1:n} | H_n, H_{n-1}, \mathbf{a}_n, \mathbf{a}_{n-1}) P(H_n, H_{n-1}, \mathbf{a}_n, \mathbf{a}_{n-1}),$$

where the sum is over all $H_{n-1} \in \{H^1, \dots, H^N\}$ and possible values of the changepoint matrix, $\mathbf{a}_{n-1}$.

Using $P(H_n, H_{n-1}, \mathbf{a}_n, \mathbf{a}_{n-1}) = P(H_n, \mathbf{a}_n | H_{n-1}, \mathbf{a}_{n-1})P(H_{n-1}, \mathbf{a}_{n-1})$, and applying Bayes' rule to write

$$P(\xi_{1:n-1} | H_{n-1}, \mathbf{a}_{n-1})P(H_{n-1}, \mathbf{a}_{n-1}) = P(H_{n-1}, \mathbf{a}_{n-1} | \xi_{1:n-1})P(\xi_{1:n-1}),$$

gives

$$P_n(H_n, \mathbf{a}_n) = \frac{P(\xi_{1:n-1})}{P(\xi_{1:n})} P(\xi_n | H_n) \sum_{H_{n-1}, \mathbf{a}_{n-1}} P_{n-1}(H_{n-1}, \mathbf{a}_{n-1}) P(H_n, \mathbf{a}_n | H_{n-1}, \mathbf{a}_{n-1}). \quad (27)$$

We compute the conditional probability $P(H_n, \mathbf{a}_n | H_{n-1}, \mathbf{a}_{n-1})$ by marginalizing over all possible transition matrices, ϵ . To do so, we relate the probabilities of ϵ and $\mathbf{a}$. Note that if the observer assumes the columns ϵ^j are independent prior to any observations, then the exit rates conditioned on the changepoint counts, $\epsilon^j | \mathbf{a}_n^j$, are independent for all states, $j = 1, \dots, N$.

To motivate the derivation we first consider a single state, $j = 1$, and assume that the environmental state has been observed perfectly over $T > 1$ timesteps, but the transition rates are unknown. Therefore, all $\mathbf{a}_n^1$ are known to the observer ($1 \leq n \leq T$), but the ϵ^1 are not. The state of the system at time $n + 1$, given that it was in state H^1 at time n , is a categorical random variable, and $P(H_{n+1} = H^i | H_n = H^1) = \epsilon^{i1}$, for $1 \leq n \leq T - 1$. The observed transitions $H^1 \mapsto H^i$ are independent samples from a categorical distribution with unknown parameters ϵ^1 .

The conjugate prior to the categorical distribution is the Dirichlet distribution, and we therefore use it as a prior on the changepoint probabilities. For simplicity we again assume a flat prior over ϵ^1 , that is $P(\epsilon^1) = \chi_S(\epsilon^1)$, where χ_S is the indicator function on the standard $(N - 1)$ -simplex, S .

Denote by D the sequence of states that the environment transitioned to at time $n + 1$ whenever it was in state H^1 at time n , for all $1 \leq n \leq T - 1$. Therefore D is a sequence of states from the set $\{H^1, \dots, H^N\}$. By definition, $P(D | \epsilon^1) = \prod_{i=1}^N (\epsilon^{i1})^{\sum_{n=1}^{T-1} \chi(H_{n+1}=H^i, H_n=H^1)}$, where $\chi(H_{n+1} = H^i, H_n = H^1)$ is the indicator function, which is unity only when $H_{n+1} = H^i$ and $H_n = H^1$ and zero otherwise. Equivalently, we can write $P(\mathbf{a}_T^1 | \epsilon^1) = \prod_{i=1}^N (\epsilon^{i1})^{\mathbf{a}_T^{i1}}$, since $\mathbf{a}_T^{i1} = \sum_{n=1}^{T-1} \chi(H_{n+1} = H^i, H_n = H^1)$. For general $n > 1$, the posterior distribution for the transition probabilities ϵ^1 given the changepoint vector $\mathbf{a}_n^1$ is then

$$P(\epsilon^1 | \mathbf{a}_n^1) = \frac{\Gamma\left(\sum_{i=1}^N (a_n^{i1} + 1)\right)}{\prod_{i=1}^N \Gamma(a_n^{i1} + 1)} \prod_{i=1}^N (\epsilon^{i1})^{a_n^{i1}} = \text{dir}(\epsilon^1; \mathbf{a}_n^1 + \mathbf{1}).$$

Here $\mathbf{1} = (1, \dots, 1)^T$, so $\mathbf{a}_n^1 + \mathbf{1}$ should be interpreted as the vector with entries $(a_n^{i1} + 1)_{i=1}^N$, $\Gamma(x)$ is the gamma function, and $\text{dir}(\epsilon^1; \mathbf{a}_n^1 + \mathbf{1})$ the probability density function of the N -dimensional Dirichlet distribution, $\text{Dir}(\mathbf{a}_n^1 + \mathbf{1})$.

The same argument applies to all initial states, H^j , $j \in \{1, \dots, N\}$. We assumed that the transition rates are conditionally independent, so that

$$P(\epsilon | \mathbf{a}_n) = \prod_{j=1}^N \text{dir}(\epsilon^j; \mathbf{a}_n^j + 1) = \prod_{j=1}^N \frac{\Gamma\left(\sum_{i=1}^N (a_n^{ij} + 1)\right)}{\prod_{i=1}^N \Gamma(a_n^{ij} + 1)} \prod_{k=1}^N (\epsilon^{kj})^{a_n^{kj}}. \quad (28)$$

Using this observation, the transition probability between two states can be computed by marginalizing over all possible transition matrices, ϵ , conditioned on $\mathbf{a}_{n-1}$,

$$\begin{aligned} P(H_n, \mathbf{a}_n | H_{n-1}, \mathbf{a}_{n-1}) &= \int_{\mathcal{M}} P(H_n, \mathbf{a}_n | \epsilon, H_{n-1}, \mathbf{a}_{n-1}) P(\epsilon | H_{n-1}, \mathbf{a}_{n-1}) d\epsilon \\ &= \int_S \cdots \int_S P(H_n, \mathbf{a}_n | \epsilon^1, \dots, \epsilon^N, H_{n-1}, \mathbf{a}_{n-1}) \\ &\quad \times \text{dir}(\epsilon^1; \mathbf{a}_{n-1}^1 + 1) \times \cdots \times \text{dir}(\epsilon^N; \mathbf{a}_{n-1}^N + 1) d\epsilon^1 \cdots d\epsilon^N, \end{aligned} \quad (29)$$

where $\mathcal{M}$ represents the space of all $N \times N$ left stochastic matrices and S is the $N - 1$ dimensional simplex of $\epsilon^j \in [0, 1]^N$ such that $\sum_{i=1}^N \epsilon^{ij} = 1$.

Let δ^{ij} be the $N \times N$ matrix containing a 1 as its ij -th entry, and 0 everywhere else. For all $i, j \in \{1, \dots, N\}$ we have

$$P(H_n = H^i, \mathbf{a}_n | \epsilon, H_{n-1} = H^j, \mathbf{a}_{n-1}) = \begin{cases} \epsilon^{ij} & \text{if } \mathbf{a}_n = \mathbf{a}_{n-1} + \delta^{ij}, \\ 0 & \text{otherwise.} \end{cases} \quad (30)$$

Implicit in Eq. (30) is the requirement that the environment must have been in the state $H_{n-1} = H^j$ in order for the transition $H^j \mapsto H^i$ to have occurred between t_{n-1} and t_n . This will ensure that the changepoint matrices $\mathbf{a}_n$ that are assigned nonzero probability correspond to admissible paths through the states $\{H^1, \dots, H^N\}$. Applying Eq. (30), we can compute the integrals in Eq. (29) for all pairs (i, j) . We let $\hat{\epsilon}_{n-1}^{ij}(\mathbf{a}_{n-1}) := P(H_n = H^i, \mathbf{a}_n = \mathbf{a}_{n-1} + \delta^{ij} | H_{n-1} = H^j, \mathbf{a}_{n-1})$ to simplify notation, and find

$$\begin{aligned} \hat{\epsilon}_{n-1}^{ij}(\mathbf{a}_{n-1}) &= \int_S \cdots \int_S \epsilon^{ij} \prod_{k=1}^N \text{dir}(\epsilon^k; \mathbf{a}_{n-1}^k + 1) d\epsilon^1 \cdots d\epsilon^N \\ &= \int_S \epsilon^{ij} \text{dir}(\epsilon^j; \mathbf{a}_{n-1}^j + 1) d\epsilon^j \prod_{k \neq j} \int_S \text{dir}(\epsilon^k; \mathbf{a}_{n-1}^k + 1) d\epsilon^k \\ &= \int_S \epsilon^{ij} \text{dir}(\epsilon^j; \mathbf{a}_{n-1}^j + 1) d\epsilon^j = \frac{a_{n-1}^{ij} + 1}{N + \sum_{k=1}^N a_{n-1}^{kj}}. \end{aligned} \quad (31)$$

As in the point estimate of the rate $\hat{\epsilon}_{n-1}(\mathbf{a}_{n-1})$ in Eq. (14), each $\hat{\epsilon}_{n-1}^{ij}(\mathbf{a}_{n-1})$ is a ratio containing the number of $H^j \mapsto H^i$ transitions in the numerator, and the total number of transitions out of the j th state in the denominator. Thus, the estimated transition rate $\hat{\epsilon}_{n-1}^{ij}(\mathbf{a}_{n-1})$ increases with the number of transitions $H^j \mapsto H^i$ in a given interval $\{1, \dots, n\}$. Furthermore, each column sums to unity:

$$\sum_{i=1}^N \hat{\epsilon}_{n-1}^{ij}(\mathbf{a}_{n-1}) = \frac{\sum_{i=1}^N (a_{n-1}^{ij} + 1)}{N + \sum_{k=1}^N a_{n-1}^{kj}} = \frac{N + \sum_{i=1}^N a_{n-1}^{ij}}{N + \sum_{k=1}^N a_{n-1}^{kj}} = 1,$$

Figure 4: Evidence accumulation and change rates inference in a 2-state asymmetric system. (A) Sample path (color bar, top) of the environment between times t_{70} and t_{80} (same simulation as in panels C-E) with corresponding observations (blue dots), and log posterior odds ratio (black step function). Here and in panels C-E, $(\epsilon^+, \epsilon^-) = (0.2, 0.1)$, SNR= 1.4. (B) The number of allowable changepoint matrices as a function of observation number, n , for $N = 2$ (blue circles), and $N = 3$ (blue triangles). (C)-(E) Color plots (gray scale) of the joint density, $P_n(\epsilon^{21}, \epsilon^{12})$, with mean value (red star) approaching the true transition rates (green circle).

so the point estimates $\hat{\epsilon}_{n-1}^{ij}(\mathbf{a}_{n-1})$ for the transition rates out of each state j do provide an empirical probability mass function along each column. However, as in the symmetric case, these estimates are biased toward the interior of the domain. This is a consequence of the hyperparameters we have chosen for our prior density, $\text{dir}(\epsilon; \mathbf{a}_0 + \mathbf{1})$.

Therefore, for $n > 1$, the probability update equation in the case of asymmetric transition rates (Eq. (27)) is given by,

$$P_n(H_n = H^i, \mathbf{a}_n) = \frac{P(\xi_{1:n-1})}{P(\xi_{1:n})} f^i(\xi_n) \sum_{j=1}^N \hat{\epsilon}_{n-1}^{ij}(\mathbf{a}_n - \delta^{ij}) P_{n-1}(H_{n-1} = H^j, \mathbf{a}_n - \delta^{ij}). \quad (32)$$

The point estimates of the transition rates, $\hat{\epsilon}_{n-1}^{ij}(\mathbf{a}_{n-1} = \mathbf{a}_n - \delta^{ij})$, are defined in Eq. (31). As before, $P_1(H^i, \mathbf{a}_1 = \mathbf{0}) = f^i(\xi_1) P_0(H^i) / P(\xi_1)$ and $P_1(H^i, \mathbf{a}_1) = 0$ for any $\mathbf{a}_1 \neq \mathbf{0}$. At future times, it is only possible to obtain changepoint matrices $\mathbf{a}_n$ whose entries sum to $\sum_{i,j} a_n^{ij} = n - 1$, the changepoint matrices $\mathbf{a}_n$ and $\mathbf{a}_{n-1}$ must be related as $\mathbf{a}_n = \mathbf{a}_{n-1} + \delta^{ij}$, as noted in Eq. (30). This considerably reduces the number of terms in the sum in Eq. (32).

The observer can find the most likely state of the environment by maximizing the posterior probability after marginalizing over the changepoint counts $\mathbf{a}_n$,

$$\hat{i} = \operatorname{argmax}_i P_n(H^i) = \operatorname{argmax}_i \left(\sum_{\mathbf{a}_n} P_n(H^i, \mathbf{a}_n) \right).$$

The transition rate matrix can also be computed by marginalizing across all possible states, H_n , and changepoint count matrices, $\mathbf{a}_n$,

$$P_n(\epsilon) = \sum_{s=1}^N \sum_{\mathbf{a}_n} P(\epsilon|\mathbf{a}_n) P_n(H^s, \mathbf{a}_n),$$

where $P(\epsilon|\mathbf{a}_n)$ is the product of probability density functions, $\operatorname{dir}(\epsilon^j; \mathbf{a}_n^j + 1)$, given in Eq. (28). The mean of this distribution is given by

$$\begin{aligned} \bar{\epsilon} &= \int_{\mathcal{M}} \epsilon P_n(\epsilon) d\epsilon = \sum_{s=1}^N \sum_{\mathbf{a}_n} P_n(H^s, \mathbf{a}_n) \int_{\mathcal{M}} \epsilon P(\epsilon|\mathbf{a}_n) d\epsilon \\ &= \sum_{s=1}^N \sum_{\mathbf{a}_n} P_n(H^s, \mathbf{a}_n) \mathbf{E}(\mathbf{a}_n), \end{aligned} \quad (33)$$

where $\mathbf{E}(\mathbf{a}_n)^{ij} = \hat{\epsilon}_n^{ij}(\mathbf{a}_n) = \mathbb{E}[\epsilon^{ij}|\mathbf{a}_n]$ defined in Eq. (31), is a conditional expectation over each possible changepoint matrix $\mathbf{a}_n$.

Eq. (32) is easier to interpret when $N = 2$. Using Eq. (31), we find

$$\hat{\epsilon}_{n-1}^{21}(\mathbf{a}_{n-1}) = \frac{a_{n-1}^{21} + 1}{2 + a_{n-1}^{21} + a_{n-1}^{11}}, \quad \hat{\epsilon}_{n-1}^{12}(\mathbf{a}_{n-1}) = \frac{a_{n-1}^{12} + 1}{2 + a_{n-1}^{12} + a_{n-1}^{22}},$$

and we can express $\hat{\epsilon}_{n-1}^{11}(\mathbf{a}_{n-1}) = 1 - \hat{\epsilon}_{n-1}^{21}(\mathbf{a}_{n-1})$ and $\hat{\epsilon}_{n-1}^{22}(\mathbf{a}_{n-1}) = 1 - \hat{\epsilon}_{n-1}^{12}(\mathbf{a}_{n-1})$. Expanding the sum in Eq. (32), we have

$$P_n(H^1, \mathbf{a}_n) = \frac{P(\xi_{1:n-1})}{P(\xi_{1:n})} f^1(\xi_n) [\hat{\epsilon}_{n-1}^{11}(\mathbf{a}_n - \boldsymbol{\delta}^{11}) P_{n-1}(H^1, \mathbf{a}_n - \boldsymbol{\delta}^{11}) + \hat{\epsilon}_{n-1}^{12}(\mathbf{a}_n - \boldsymbol{\delta}^{12}) P_{n-1}(H^2, \mathbf{a}_n - \boldsymbol{\delta}^{12})], \quad (34a)$$

$$P_n(H^2, \mathbf{a}_n) = \frac{P(\xi_{1:n-1})}{P(\xi_{1:n})} f^2(\xi_n) [\hat{\epsilon}_{n-1}^{22}(\mathbf{a}_n - \boldsymbol{\delta}^{22}) P_{n-1}(H^2, \mathbf{a}_n - \boldsymbol{\delta}^{22}) + \hat{\epsilon}_{n-1}^{21}(\mathbf{a}_n - \boldsymbol{\delta}^{21}) P_{n-1}(H^1, \mathbf{a}_n - \boldsymbol{\delta}^{21})]. \quad (34b)$$

The boundary and initial conditions will be given as above, and the mean inferred transition matrix is given by Eq. (33). Importantly, the inference process described by Eqs. (34) allows for both asymmetric changepoint matrices, $\mathbf{a}_n$, and inferred transition rate matrices $\mathbf{E}(\mathbf{a}_n)$, unlike the process in Eq. (17). However, the variance of the posteriors over the rates will decrease more slowly, as fewer transitions out of each particular state will be observed.

This algorithm can be used to infer unequal transition rates as shown in Fig. 4: Panels C through E show that the mode of the joint posterior distribution, $P_n(\epsilon^{21}, \epsilon^{12})$, approaches the correct rates, while its variance decreases. As in Section 3.1 we conjecture that this joint density does not converge to a point mass at the true rate values unless the SNR is infinite.

5 Continuum limits and stochastic differential equation models

We next derive continuum limits of the discrete probability update equations for the symmetric case discussed in Section 3. We assume that observers make measurements rapidly, so we can derive a stochastic differential equation (SDE) that models the update of an ideal observer's belief (Gold and Shadlen, 2007). SDEs are generally easier to analyze than their discrete counterparts (Gardiner, 2004). For example, response times can be studied by examining mean first passage times of log-likelihood ratios (Bogacz et al, 2006), or log-likelihoods (McMillen and Holmes, 2006), which is much easier done in the continuum limit (Redner, 2001). For simplicity, we begin with an analysis of the two state process, and then extend our results to the multistate case.

5.1 Derivation of the continuum limit

Two-state symmetric process. We first assume that the state of the environment, $\{H_t\}$, is a homogeneous *continuous-time* Markov chain with state space $\{H^+, H^-\}$. The probability of transitions between the two states is symmetric, and given by $P(H_{t+\Delta t} = H^\pm | H_t = H^\mp) = \epsilon \Delta t + o(\Delta t)$, where $0 \leq \epsilon < \infty$. The number of changepoints, a_t , up to time t is a Poisson process with rate ϵ . An observer infers the present state from a sequence of observations, $\xi_{1:n}$, made at equally spaced times, $t_{1:n}$, with $\Delta t = t_j - t_{j-1}$.¹ Each observation, ξ_n , has probability $f_{\Delta t}^\pm(\xi_n) := \Pr(\xi_n | H^\pm)$ (See Veliz-Cuba et al (2016) for more details). We again use the notation $P_n(H^\pm, a) = P(H_{t_n} = H^\pm, a_{t_n} = a | \xi_{1:n})$ where t_n is the time of the n^{th} observation.

As in the previous sections, an estimate of the rate parameter, ϵ , is obtained from the posterior distribution over the changepoint count, a_{t_n} , at the time of the n^{th} observation, t_n . For simplicity, we assume a Gamma prior with parameters α and β over ϵ , so that $\epsilon \sim \text{Gamma}(\alpha, \beta)$. By assumption the changepoint count follows a Poisson distribution with parameter ϵt_n , so that $P(a_{t_n} = a | \epsilon) = (\epsilon t_n)^a e^{-\epsilon t_n} / a!$. Therefore, once a_n changepoints have been observed, we have the posterior distribution $\epsilon | a_n \sim \text{Gamma}(a_n + \alpha, t_n + \beta)$, that is,

$$P(\epsilon | a_n) = \frac{(t_n + \beta)^{a_n + \alpha} \epsilon^{a_n + \alpha - 1} e^{-\epsilon(t_n + \beta)}}{\Gamma(a_n + \alpha)}. \quad (35)$$

We can substitute Eq. (35) into Eq. (9) describing the probability of transitions between time t_{n-1} and t_n to find

$$P(H_n, a_n | H_{n-1}, a_{n-1}) = \int_0^\infty P(H_n, a_n | \epsilon, H_{n-1}, a_{n-1}) \gamma(\epsilon; a_{n-1} + \alpha, t_{n-1} + \beta) d\epsilon, \quad (36)$$

where $\gamma(\epsilon; \alpha, \beta) = \beta^\alpha \epsilon^{\alpha-1} e^{-\epsilon\beta} / \Gamma(\alpha)$ is the density of the Gamma distribution. Using the definition of the transition rate ϵ , we can relate it to the first conditional probability

¹Equal spacing $\Delta t = t_j - t_{j-1}$ is not necessary for all $j = 2, \dots, n$, but it does allow for a more concise derivation of the continuum limit. Irregular spacings would require a more careful selection of the scaling of the log likelihoods $\ln f^\pm(\xi)$.

in the integral of Eq. (36) via

$$P(H_n, a_n | \epsilon, H_{n-1}, a_{n-1}) = \begin{cases} 1 - \epsilon \Delta t & H_n = H_{n-1} \text{ \& } a_n = a_{n-1}, \\ \epsilon \Delta t & H_n \neq H_{n-1} \text{ \& } a_n = a_{n-1} + 1, \\ 0 & \text{otherwise.} \end{cases} \quad (37)$$

We have dropped the $o(\Delta t)$ terms as we are interested in the limit $\Delta t \rightarrow 0$.

Using Eq. (37) and properties of the Gamma distribution we can evaluate the integral in Eq. (36) to obtain

$$P(H_n, a_n | H_{n-1}, a_{n-1}) = \begin{cases} 1 - \Delta t \frac{a_n + \alpha}{t_{n-1} + \beta} & H_n = H_{n-1} \text{ \& } a_n = a_{n-1} \\ \Delta t \frac{a_n + \alpha - 1}{t_{n-1} + \beta} & H_n \neq H_{n-1} \text{ \& } a_n = a_{n-1} + 1 \\ 0 & \text{otherwise.} \end{cases} \quad (38)$$

We can use Eq. (38) in the update equation, Eq. (3), to obtain the probabilities of (H_n, a_n) given observations $\xi_{1:n}$. As before, only terms involving $a_n - 1$ and a_n remain in the sum for $n \geq 1$. Using the same notational convention as in previous sections, we obtain,

$$P_n(H^\pm, a) = \frac{P(\xi_{1:n-1})}{P(\xi_{1:n})} f_{\Delta t}^\pm(\xi_n) \left[\left(1 - \Delta t \frac{a + \alpha}{t_{n-1} + \beta} \right) P_{n-1}(H^\pm, a) + \Delta t \frac{a + \alpha - 1}{t_{n-1} + \beta} P_{n-1}(H^\mp, a - 1) \right]. \quad (39)$$

Note, Eq. (39) is similar to the update Eq. (17) we derived in Section 3, with the time index replaced by $t_{n-1}/\Delta t$ up to the β term. Also, since we have used a Gamma instead of a Beta distribution as a prior, the point estimate of the transition rate is slightly different (See Eq. (14)). As in the discrete time case, a point estimate of the transition rate is required even before the first changepoint can be observed. We therefore cannot use an improper prior, as the rate point estimate would be undefined.

To take the limit of Eq. (39) as $\Delta t \rightarrow 0$ we proceed as in Bogacz et al (2006) and Veliz-Cuba et al (2016), working with logarithms of the probabilities. Dividing Eq. (39) by $P_{n-1}(H^\pm, a)$, taking logarithms of both sides, and using the notation $x_{t_n}^\pm(a) := \ln P_n(H^\pm, a)$, we obtain,²

$$\Delta x_{t_n}^\pm(a) \propto \ln f_{\Delta t}^\pm(\xi_n) + \ln \left[1 - \Delta t \frac{a + \alpha}{t_{n-1} + \beta} + \Delta t \frac{a + \alpha - 1}{t_{n-1} + \beta} e^{x_{t_{n-1}}^\mp(a-1) - x_{t_{n-1}}^\pm(a)} \right]. \quad (40)$$

Using the approximation $\ln(1 + z) \approx z$ for small z yields

$$\Delta x_{t_n}^\pm(a) \propto \ln f_{\Delta t}^\pm(\xi_n) + \Delta t \left(\frac{a + \alpha - 1}{t_{n-1} + \beta} e^{x_{t_{n-1}}^\mp(a-1) - x_{t_{n-1}}^\pm(a)} - \frac{a + \alpha}{t_{n-1} + \beta} \right). \quad (41)$$

²Note, we drop the $\ln [P(\xi_{1:n-1})/P(\xi_{1:n})]$ term below since it is common to all evolution equations. For determining the most likely option, only the relative magnitudes of the log likelihoods are important. In numerical simulations, we normalize to account for this discrepancy.

Since the proportionality constant is equal for all a , we can use the SDE for the log likelihood x_t , (See Veliz-Cuba et al (2016) for the details of the derivation)

$$dx_t^\pm(a) = g^\pm(t)dt + dW_t^\pm + \left(\frac{a + \alpha - 1}{t + \beta} e^{x_t^\mp(a-1) - x_t^\pm(a)} - \frac{a + \alpha}{t + \beta} \right) dt, \quad (42)$$

where $g^\pm(t) = \lim_{\Delta t \rightarrow 0} \frac{1}{\Delta t} E[\ln f_{\Delta t}^\pm(\xi) | H(t)]$ and W^i satisfies $\langle W_t^i W_t^j \rangle = \Sigma^{ij}(t) \cdot t$ with $\Sigma^{ij}(t) = \lim_{\Delta t \rightarrow 0} \frac{1}{\Delta t} \text{Cov}[\ln f_{\Delta t}^i(\xi), \ln f_{\Delta t}^j(\xi) | H(t)]$ for $i, j \in \{+, -\}$.

Note that Eq. (42) is an infinite set of differential equations, one for each pair $(H^\pm, a)$, $a \in \mathbb{Z}_{\geq 0}$. The initial conditions at $t = 0$ are given by $x^\pm(a) = \ln P_0(H^\pm, a)$. To be consistent with the prior over the rate, ϵ , we can choose a Poisson prior over a with mean, α , i.e. $P_0(a) := \frac{\alpha^a e^{-\alpha}}{a!}$. The initial conditions for Eq. (42) are given by $x^\pm = \ln P_0(H^\pm, a) = \ln [P_0(H^\pm) P_0(a)]$. Note also that Eq. (42) at the boundary $a = 0$ is a special case. Since at those values there is no influx of probability from $a - 1$, Eq. (42) reduces to

$$dx_t^\pm(0) = g^\pm(t)dt + dW_t^\pm + \left((\alpha - 1)e^{-x_t^\pm(0)} - \alpha \right) \frac{dt}{t + \beta}.$$

Lastly, note that we can obtain an evolution equations for the the likelihoods, $P_t^\pm(a) = P(H(t) = H^\pm, a)$, by applying the change of variables $P_t^\pm(a) = e^{x_t^\pm(a)}$. Itô's change of coordinates rules (Gardiner, 2004) imply that Eq. (42) is equivalent to

$$dP_t^\pm(a) = P_t^\pm(a) \left[\left(g^\pm(t) + \frac{1}{2} \right) dt + dW_t^\pm \right] + \left[\frac{a + \alpha - 1}{t + \beta} P_t^\mp(a - 1) - \frac{a + \alpha}{t + \beta} P_t^\pm(a) \right] dt, \quad (43)$$

where now initial conditions at $t = 0$ are simply $P_0^\pm(a) = P_0(H^\pm, a) = P_0(H^\pm)P_0(a)$. We will compare the full system, Eq. (43), with an approximation using a moment expansion in subsection 5.2.

Two states with asymmetric rates. Next we consider the case where the state of the environment, $\{H_t\}$, is still a continuous-time Markov chain with state space $\{H^1, H^2\}$, but the probabilities of transition between the two states are asymmetric: $P(H_{t+\Delta t} = H^i | H_t = H^j) = \epsilon^{ij} \Delta t + o(\Delta t)$, $i \neq j$, where $\epsilon^{12} \neq \epsilon^{21}$. Thus, we must separately enumerate changepoints, a_t^{12} and a_t^{21} , to obtain an estimate of the rates ϵ^{12} and ϵ^{21} . In addition, we will rescale the enumeration of non-changepoints so that $a_t^{jj} = a^{jj} \Delta t$, in anticipation of the divergence of a^{jj} as $\Delta t \rightarrow 0$. This will mean the total *dwell time*, a_t^{jj} , will be continuous, while the change point count will be discrete. The quantities a_t^{ij} are then placed into a 2×2 matrix, $\mathbf{A}_t = (a_t^{ij}) \in \mathbb{R}^{2 \times 2}$, where $a^{ij} \in \mathbb{Z}_{\geq 0}$ for $i \neq j$ and $a^{jj} \in \mathbb{R}^*$. Note that if the number of changepoints, a_t^{ij} , and the total dwell time in a state, a_t^{jj} , were known, the escape rate could be estimated as $\epsilon^{ij} = a_t^{ij} / a_t^{jj}$.

As before, we will estimate the rate parameters, ϵ^{ij} , using the posterior probability of the changepoint matrix, $\mathbf{a}_t$. We assume Gamma priors on each rate, so that $\epsilon^{ij} \sim \text{Gamma}(\alpha_j, \beta_j)$. By assumption the changepoint count, a_t^{ij} , follows a Poisson distribution with parameter $\epsilon^{ij} a_t^{jj}$, so that $P(a_t^{ij} = a | \epsilon^{ij} a_t^{jj}) = (\epsilon^{ij} a_t^{jj})^a e^{-\epsilon^{ij} a_t^{jj}} / a!$.

Therefore, once a_t^{ij} changepoints have been observed along with the dwell time a_t^{jj} , we have the posterior distribution $\epsilon^{ij}|(a_t^{ij}, a_t^{jj}) \sim \text{Gamma}(a_t^{ij} + \alpha_j, a_t^{jj} + \beta_j)$, so

$$P(\epsilon^{ij}|a_t^{ij}, a_t^{jj}) = \frac{(a_t^{jj} + \beta_j)^{a_t^{ij} + \alpha_j} (\epsilon^{ij})^{a_t^{ij} + \alpha_j - 1} e^{-\epsilon^{ij}(a_t^{jj} + \beta_j)}}{\Gamma(a_t^{ij} + \alpha_j)}. \quad (44)$$

We now derive the continuum limit of Eq. (34). One key step of the derivation is the application of a change of variables to the changepoint matrix $\mathbf{a}$, where we replace the non-changepoint counts with dwell times t^j , defined as $t_{\Delta t}^i := \Delta t a_{\Delta t}^{ii}$ for $\Delta t = t_n - t_{n-1}$. This is necessary, due to the divergence of $a_{\Delta t}^{ii}$ as $\Delta t \rightarrow 0$. In the limit $\Delta t \rightarrow 0$, the changepoint matrix becomes

$$\mathbf{A} = \begin{pmatrix} t^1 & a^{12} \\ a^{21} & t^2 \end{pmatrix},$$

where $a^{ij} \in \mathbb{Z}^*$ is the changepoint counts from $H^j \mapsto H^i$, while $t^i \in \mathbb{R}^*$ is the dwell time in state H^i . Thus, taking logarithms, linearizing, and taking the limit $\Delta t \rightarrow 0$, we obtain the following system of stochastic partial differential equations (SPDEs) for the log likelihoods, $x_t^j(\mathbf{A}) = \ln P_n(H^\pm, \mathbf{A})$:

$$dx_t^1(\mathbf{A}) = g^1(t)dt + dW_t^1 + \left(\frac{a^{12} + \alpha_2 - 1}{t^2 + \beta_2} e^{x_t^2(\mathbf{A}) - x_t^1(\mathbf{A})} - \frac{a^{21} + \alpha_1}{t^1 + \beta_1} - \frac{\partial x_t^1}{\partial t^1} \right) dt, \quad (45a)$$

$$dx_t^2(\mathbf{A}) = g^2(t)dt + dW_t^2 + \left(\frac{a^{21} + \alpha_1 - 1}{t^1 + \beta_1} e^{x_t^1(\mathbf{A}) - x_t^2(\mathbf{A})} - \frac{a^{12} + \alpha_2}{t^2 + \beta_2} - \frac{\partial x_t^2}{\partial t^2} \right) dt, \quad (45b)$$

where the drift, $g^i(t)$, and noise, $W^i(t)$, are defined as before (for details, see Appendix 7.3). Note that the flux terms, $\frac{\partial x_t^i}{\partial t^i}$, account for the flow of probability to longer dwell times t^i . For example, the SPDE for x_t^1 has a flux term for the linear increase of the dwell time t^1 , since this represents the environment remaining in state H^1 . These flux terms simply propagate the probability densities $e^{x_t^i(\mathbf{A})} = P_t(H^i, \mathbf{A})$ over the space (t^1, t^2) , causing no net change in the probability of residing in either state H^i : $P_t^i = \int_0^\infty \int_0^\infty e^{x_t^i(\mathbf{A})} dt^1 dt^2$.

Eq. (45) generalizes Eq. (34) as an infinite set of SPDEs, indexed by the discrete variables (H^j, a^{12}, a^{21}) where $a^{12}, a^{21} \in \mathbb{Z}_{\geq 0}$. Each SPDE is over the space (t^1, t^2) , and it is always true that $t^1 + t^2 = t$. Initial conditions at $t = 0$ are given by $x^j(\mathbf{A}) = \ln [P_0(H^j) \cdot P_0(\mathbf{A})]$. For consistency with the prior on the rates, ϵ^{ij} , we choose a Poisson prior over the changepoint counts a^{ij} , $i \neq j$, and a Dirac delta distribution prior over the dwell times t^i ,

$$P_0(\mathbf{A}) = \frac{\alpha_1^{a^{21}} e^{-\alpha_1}}{a^{21}!} \frac{\alpha_2^{a^{12}} e^{-\alpha_2}}{a^{12}!} \delta(t^1 - \beta_1) \delta(t^2 - \beta_2). \quad (46)$$

As before, Eq. (45) at the boundaries $a^{12} = 0$ and $a^{21} = 0$ is a special case, since there will be no influx of probability from $a^{12} - 1$ or $a^{21} - 1$.

As in the symmetric case, we can convert Eq. (45) to equations describing the evolution of the likelihoods $P_t^i(\mathbf{A}) = P(H(t) = H^i, \mathbf{A})$. Applying the change of variables $P_t^i(\mathbf{A}) = e^{x_t^i(\mathbf{A})}$, we find

$$dP_t^1(\mathbf{A}) = P_t^1(\mathbf{A}) \left[\left(g^1(t) + \frac{1}{2} \right) dt + dW_t^1 \right] + \left[\frac{a^{12} + \alpha_2 - 1}{t^2 + \beta_2} P_t^2(\mathbf{A} - \delta^{12}) - \frac{a^{21} + \alpha_1}{t^1 + \beta_1} P_t^1(\mathbf{A}) - \frac{\partial P_t^1(\mathbf{A})}{\partial t^1} \right] dt \quad (47a)$$

$$dP_t^2(\mathbf{a}) = P_t^2(\mathbf{a}) \left[\left(g^2(t) + \frac{1}{2} \right) dt + dW_t^2 \right] + \left[\frac{a^{21} + \alpha_1 - 1}{t^1 + \beta_1} P_t^1(\mathbf{A} - \delta^{21}) - \frac{a^{12} + \alpha_2}{t^2 + \beta_2} P_t^2(\mathbf{A}) - \frac{\partial P_t^2(\mathbf{A})}{\partial t^2} \right] dt \quad (47b)$$

where now initial conditions at $t = 0$ are $P_0^i(\mathbf{A}) = P_0(H^i)P_0(\mathbf{A})$.

Multiple states with symmetric rates. The continuum limit in the case of N states, $\{H^1, \dots, H^N\}$, with symmetric transition rates can be derived as with $N = 2$ (See Appendix 7.4 for details). Again, denote the transition probabilities by $P(H_{t+\Delta t} = H^i | H_t = H^j) = \epsilon^{ij} \Delta t + o(\Delta t)$, and the rate of switching from one to any other state by $\epsilon = (N - 1)\epsilon^{ij}$.

Assuming again a Gamma prior on the transition rate, $\epsilon \sim \text{Gamma}(\alpha, \beta)$, and introducing $x_{t_n}^i(a) := \ln P_n(H^i, a)$, we obtain the SDE

$$dx_{t_n}^i(a) = g^i(t)dt + dW_t^i + \left(\frac{a + \alpha - 1}{(N - 1)(t + \beta)} \sum_{j \neq i} e^{x_{t_n}^j(a-1) - x_{t_n}^i(a)} - \frac{a + \alpha}{t + \beta} \right) dt, \quad (48)$$

where $g^i(t) = \lim_{\Delta t \rightarrow 0} \frac{1}{\Delta t} E[\ln f_{\Delta t}^i(\xi) | H(t)]$ and W^i satisfies $\langle W^i W^j \rangle = \Sigma^{ij}(t) \cdot t$ with $\Sigma^{ij}(t) = \lim_{\Delta t \rightarrow 0} \frac{1}{\Delta t} \text{Cov}[\ln f_{\Delta t}^i(\xi), \ln f_{\Delta t}^j(\xi) | H(t)]$.

Eq. (48) is again an infinite set of stochastic differential equations, one for each pair (H^i, a) , $i \in 1, \dots, N$, $a \in \mathbb{Z}_{\geq 0}$. We have some freedom in choosing initial conditions at $t = 0$. For example, since $x^i(a) = \ln P_0(H^i, a)$, we can use the Poisson distribution discussed in the case of two states.

The posterior over the transition rate, ϵ , is

$$P_n(\epsilon) = \sum_{s=1}^N \sum_{a_n=0}^{\infty} P(\epsilon | a_n) P_n(H^s, a_n),$$

where $P(\epsilon | a_n)$ is the Gamma distribution given by Eq. 35. Similar to Eq. 20, the expected rate is

$$\bar{\epsilon} := \int_0^{\infty} \epsilon P_n(\epsilon) d\epsilon = \sum_{s=1}^N \sum_{a_n=0}^{\infty} \int_0^{\infty} \epsilon P(\epsilon | a_n) P_n(H^s, a_n) d\epsilon = \sum_{s=1}^N \sum_{a_n=0}^{\infty} \frac{a_n + \alpha}{t_n + \beta} P_n(H^s, a_n). \quad (49)$$

An equivalent argument can be used to obtain the posterior over the rates in the asymmetric case.

5.2 Moment hierarchy for the 2-state process

In the previous section, we approximated the evolution of the joint probabilities of environmental states and changepoint counts. The result, in the symmetric case, was an infinite set of SDEs, one for each combination of state and changepoint values (H^i, a) . However, an observer is mainly concerned with the current state of the environment. The changepoint count is important for this inference, but may not be of direct interest itself. We next derive simpler, approximate models that do not track the entire joint distribution over all changepoint counts, but only essential aspects of this distribution. We do so by deriving a hierarchy of iterative equations for the moments of the distribution of changepoint counts, $a \in \mathbb{Z}_{\geq 0}$, focusing specifically on the two state symmetric case.

Our goal in deriving moment equations is to have a low-dimensional, and reasonably tractable, system of SDEs. Similar to previous studies of sequential decision making algorithms (Bogacz et al, 2006), such low-dimensional systems can be used to inform neurophysiologically relevant population rate models of the evidence accumulation process. To begin, we consider the infinite system of SDEs given in the two state symmetric case, Eq. (43). Our reduction then proceeds by computing the SDEs associated with the lower order (0th, 1st, and 2nd) moments over the changepoint count a :

$$\bar{P}_t^\pm = \sum_{a \in \mathbb{Z}_{\geq 0}} P_t^\pm(a), \quad \bar{a}_t^\pm = \sum_{a \in \mathbb{Z}_{\geq 0}} (a + \alpha) P_t^\pm(a), \quad \bar{b}_t^\pm = \sum_{a \in \mathbb{Z}_{\geq 0}} (a + \alpha)^2 P_t^\pm(a). \quad (50)$$

We denote the moments using bars ($\bar{b}_t^\pm$). Below, when we discuss cumulants, we will represent them using hats ($\hat{b}_t^\pm$). Note that the “0th” moments are the marginal probabilities of H^+ and H^- .

We begin by summing Eq. (43) over all $a \in \mathbb{Z}_{\geq 0}$ and applying Eq. (50) to find this generates an SDE for the evolution of the moments $\bar{P}_t^\pm$ given

$$d\bar{P}_t^\pm = \bar{P}_t^\pm \left[\left(g^\pm(t) + \frac{1}{2} \right) dt + dW_t^\pm \right] + \frac{1}{t + \beta} [\bar{a}_t^\mp - \bar{a}_t^\pm] dt. \quad (51)$$

where we have used the fact that

$$\overline{(a-1)}_t^\pm = \sum_{a=1}^{\infty} (a + \alpha - 1) P_t^\pm(a-1) = \sum_{a=0}^{\infty} (a + \alpha) P_t^\pm(a) = \bar{a}_t^\pm.$$

The SDE given by Eq. (51) for the zeroth moment, $\bar{P}_t^\pm$, depends on the first moment, $\bar{a}_t^\pm$. We can determine values for the first moment by either obtaining the next SDE in the moment hierarchy, or assuming a reasonable functional form for $\bar{a}_t^\pm$. For instance, if the transition rate ϵ is known we can assume $\bar{a}_t^\pm := (t + \beta)\epsilon\bar{P}_t^\pm + \mathcal{O}(1)$, so that $\bar{a}_t^+ + \bar{a}_t^-$ is approximately the mean of the counting process with rate ϵ . In this case, the continuum limit $t \rightarrow \infty$ of Eq. (51) becomes

$$d\bar{P}_t^\pm = \bar{P}_t^\pm \left[\left(g^\pm(t) + \frac{1}{2} \right) dt + dW_t^\pm \right] + \epsilon \cdot [\bar{P}_t^\mp - \bar{P}_t^\pm] dt. \quad (52)$$

As expected, this is the two state version of Eq. (2), with known rate, ϵ . However, if the observer has no prior knowledge of the rate, ϵ , then $\bar{a}_t^\pm$ should evolve towards $(t + \beta)\epsilon\bar{P}_t^\pm$ at a rate that depends on the noisiness of observations.

To obtain an equation for $\bar{a}_t^\pm$ we multiply Eq. (43) by $(a + \alpha)$ and sum to yield,

$$d\bar{a}_t^\pm = \bar{a}_t^\pm \left[\left(g^\pm(t) + \frac{1}{2} \right) dt + dW_t^\pm \right] + \frac{1}{t + \beta} [\bar{a}_t^\mp + \bar{b}_t^\mp - \bar{b}_t^\pm], \quad (53)$$

This equation relates the zeroth, first, and second moments, $\bar{P}_t^\pm$, $\bar{a}_t^\pm$, and $\bar{b}_t^\pm$. Again, we require an expression for the next moment, $\bar{b}_t^\pm$, to close the system of Eqs. (51, 53). We could obtain an equation for $\bar{b}_t^\pm$ by multiplying Eq. (43) by $(a + \alpha)^2$ and summing. However, as is typical with moment hierarchies, we would not be able to close the system as equations for subsequent moments will depend on successively higher moments (Socha, 2007; Kuehn, 2016). To close the equations for $\bar{P}_t^\pm, \bar{a}_t^\pm, \dots$ we can truncate: One possibility is cumulant-neglect (Whittle, 1957; Socha, 2007), which assumes all cumulants above a given order grow more slowly than the moment itself and can thus be ignored. This allows one to express the highest order moment as a function of the lower order moments, since a moment is an algebraic function of its associated cumulant and lower moments. For instance, neglecting the second cumulant $\hat{b}_t^\pm \approx 0$ allows us to approximate the second moment as $\bar{b}_t^\pm = \hat{b}_t^\pm + (\bar{a}_t^\pm)^2 \approx (\bar{a}_t^\pm)^2$.³

Applying cumulant-neglect to the second moment, $\bar{b}_t^\pm$, in Eqs. (51, 53), using the change of variables, $\bar{A}_t^\pm = \bar{a}_t^\pm / (t + \beta)$, and the fact that $d\bar{A}_t^\pm = [(t + \beta)d\bar{a}_t^\pm - \bar{a}_t^\pm dt] / (t + \beta)^2$, we obtain a closed system of equations for the zeroth and first moments,

$$d\bar{P}_t^\pm = \bar{P}_t^\pm \left[\left(g^\pm(t) + \frac{1}{2} \right) dt + dW_t^\pm \right] + [\bar{A}_t^\mp - \bar{A}_t^\pm] dt \quad (54a)$$

$$d\bar{A}_t^\pm = \bar{A}_t^\pm(t) \left[\left(g^\pm(t) + \frac{1}{2} \right) dt + dW_t^\pm \right] + (\bar{A}_t^\mp - \bar{A}_t^\pm) \left(\frac{1}{t + \beta} + \bar{A}_t^\mp + \bar{A}_t^\pm \right) dt. \quad (54b)$$

Here initial conditions are given by $\bar{P}_0^\pm := P_0(H^\pm)$ and $\bar{A}_0^\pm = \alpha P_0(H^\pm) / \beta$. We show in Appendix 7.5 that Eq. (54) is also consistent with Eq. (2), which holds in the case of two states and known rate ϵ . Trajectories of Eq. (54) are shown in Fig. 5. Note that both $\bar{P}_t^\pm$ and $\bar{A}_t^\pm$ tend to increase when $g^\pm(t)$ is the maximal drift rate, i.e. when $H^\pm$ is the true environmental state. Thus, we expect that when $\bar{P}^\pm$ is high (close to unity) then $\bar{A}_t^\pm$ will tend to be larger than $\bar{A}_t^\mp$.

Immediately after a changepoint (where the maximal drift rate $g^\pm(t)$ changes), there is an additional contribution to the increase of $\bar{A}_t^\pm$ due to the $(\bar{A}_t^\mp - \bar{A}_t^\pm)$ term. It is this brief burst of additional input to the subsequently dominant variable that generates the counting process, which enumerates changepoints. For instance, when a $H^+ \mapsto H^-$ switch occurs, an increase in $\bar{A}_t^-$ will temporarily be driven both by the drift term, $g^-(t)$, and the nonlinear term involving $(\bar{A}_t^+ - \bar{A}_t^-)$. The burst of input generated by the nonlinear term in Eq. (54b) has an amplitude that decays nonautonomously with time. In fact, it can be shown that when the signal-to-noise ratio of the system is quite high, the variables $\bar{A}_t^\pm \approx (m + \alpha) / (t + \beta)$, which is effectively the true changepoint count m divided by elapsed time as modified by the prior.

³Note, we have used a hat to distinguish cumulants ($\hat{b}_t^\pm$), whereas bars still denote moments ($\bar{b}_t^\pm$).

Figure 5: The dynamics of the first two moments, as approximated by Eq. (54). (A) The probabilities $\bar{P}_t^\pm$ track the present state of the environment (bar above), switching with rate $\epsilon = 0.1$, and approach the stationary densities around the equilibria determined by the dichotomous drift terms $g^\pm(t)$. (B) The first moments, $\bar{A}_t^\pm$, also switch with the environmental state and alternate between the neighborhoods of two points. (C) The sum $\tilde{\epsilon}_t := \bar{A}_t^+ + \bar{A}_t^-$ provides a running point estimate of the environmental transition rate, ϵ , as shown in Eq. (55). The estimate is determined by the actual changepoints, and noisily tracks m/t , where m is the actual number of changepoints. (D,E,F) Same as A,B,C, but the moment simulations are compared with numerical simulations of the full system of SDEs given by Eq. (43). (D) Thick red line is $\bar{P}_t^+$ from Eq. (54) and thin black line is $\sum_{a=0}^\infty P_t^+(a)$ using Eq. (43); (E) Thick red line is $\bar{A}_t^+$ from Eq. (54) and thin black line is $\sum_{a=0}^\infty (a + \alpha)P_t^+(a)$ using Eq. (43); (F) Estimates of $\tilde{\epsilon}_t$ using Eq. (54) and Eq. (43). Details about the simulation method, initial conditions, and parameters are provided in Appendix 7.7.

We can also obtain a point-estimate of the transition rate of the environment, which we define as $\tilde{\epsilon}_t := \bar{A}_t^+ + \bar{A}_t^-$, since the following relations hold:

$$\begin{aligned}\bar{A}_t^+ + \bar{A}_t^- &= \frac{1}{t + \beta} \sum_{a \in \mathbb{Z}_{\geq 0}} (a + \alpha) [P(H^+, a | \xi(t)) + P(H^-, a | \xi(t))] \\ &= \sum_{a \in \mathbb{Z}_{\geq 0}} \frac{(a + \alpha)}{t + \beta} P(a | \xi(t)) = \int_0^\infty \epsilon \sum_{a \in \mathbb{Z}_{\geq 0}} P(\epsilon | a) P(a | \xi(t)) d\epsilon, \quad (55)\end{aligned}$$

This estimate is an average over the distribution of possible changepoint counts, a , given the observations, $\xi(t)$. Here $P(\epsilon | a)$ is a Gamma distribution with parameters α and β . In Fig. 5C we compare this approximation, $\tilde{\epsilon}_t$, with the true change rate ϵ and the running estimate m/t , obtained from the actual number of changepoints, m .

In Fig. 5D,E these approximations are compared to Eq. (43), the full SDE giving the distribution over all changepoint counts, a . Notice that the first moments $\bar{A}_t^\pm$ are overestimates of the true average, $\sum_{a=0}^\infty (a + \alpha) P_t^\pm(a) / (t + \beta)$, obtained from Eq. (43). We expect this is due to the fact that the moment equations, Eq. (54), tend to overcount the number of changepoints. Fluctuations lead to an increase in the number of events whereby $\bar{A}_t^+$ and $\bar{A}_t^-$ exchange dominance ($\bar{A}_t^+ = \bar{A}_t^-$), which will lead to a burst of input to one of the variables $\bar{A}_t^\pm$. As a consequence, the transition rate tends to be overestimated by Eq. (54) compared to Eq. (43).

In sum, while the inference approximation given by Eq. (54) does not provide an estimate of the variance, it does provide insight into the computations needed to estimate the changepoint count and transition probability. Transitions increment the running estimate of the changepoint count, and this increment decays over time, inversely with the total observation time t . Similar equations for the moments can be obtained in the case of asymmetric transition rates, or more than two choices using Eq. (47) and Eq. (48) respectively, although we omit their derivation here.

6 Learning transition rate in neural populations with plasticity

Models of decision making often consist of mutually inhibitory neural populations with finely tuned synaptic weights (Machens et al, 2005; McMillen and Holmes, 2006; Wong et al, 2007). For instance, many models of evidence integration in two alternative choice tasks assume that synaptic connectivity is tuned so that the full system exhibits line attractor dynamics in the absence of inputs. Such networks integrate inputs perfectly, and maintain this integrated information in memory after the inputs are removed. However, in changing environments optimal evidence integration should be leaky, since older information becomes irrelevant for the present decision (Deneve, 2008; Glaze et al, 2015).

We previously showed that optimal integration in changing environments can be implemented by mutually excitatory neural populations (Veliz-Cuba et al, 2016). Instead of a line attractor, the resulting dynamical systems contain globally attracting fixed points. Such leaky integrators maintain a limited memory of their inputs on a timescale determined by the frequency of environmental changes. However, in this previous work we assumed that the rates of the environmental changes were known to the observer. Here, we show that when these rates are not known a priori, a plastic neu-

ronal network is capable of learning and implicitly representing them through coupling strengths between neural populations.

Symmetric environment. We begin with Eq. (54), the leading order equations for the likelihood $\bar{P}_t^\pm$ and change rate variables $\bar{A}_t^\pm$ derived in Section 5.2. We interpret the likelihoods as neural population activity variables $u_t^\pm := \bar{P}_t^\pm$, reflecting a common modeling assumption that two populations receive separate streams of input associated with evidence for either choice (Bogacz et al, 2006). Next, we define a new variable $w_t^\pm := \bar{A}_t^\pm / \bar{P}_t^\pm$, which represents the synaptic weight between these neural populations. In particular, $w_t^\pm$ represents the strength of coupling from $u_t^\pm$ to $u_t^\mp$ as well as the local inhibitory coupling within $u_t^\pm$. Applying this change of variables to Eq. (54), we derive a set of equations for the population activities, $u_t^\pm$, and their associated synaptic weights, $w_t^\pm$:

$$du_t^\pm = u_t^\pm \left[\left(g^\pm(t) + \frac{1}{2} \right) dt + dW_t^\pm \right] + [w_t^\mp u_t^\mp - w_t^\pm u_t^\pm] dt \quad (56a)$$

$$dw_t^\pm = - \left[w_t^\pm + w_t^\mp - \frac{w_t^\mp}{u_t^\pm} \right] \left[\frac{1}{t + \beta} + u_t^\mp (w_t^\mp - w_t^\pm) \right] dt. \quad (56b)$$

This is a neural population model with a rate-correlation based plasticity rule (Miller, 1994; Pfister and Gerstner, 2006). We can interpret the non-autonomous term, $1/(t + \beta)$, as modeling the dynamics of a chemical agent involved in the plasticity process whose availability decays over time. Simple chemical degradation kinetics for a concentration C_t yield such a function when

$$dC_t = -C_t^2 dt, \quad C(0) = 1/\beta \quad \Rightarrow \quad C_t = \frac{1}{t + \beta}. \quad (57)$$

We briefly analyze the model, Eq. (56), by considering the limit of no observation-noise. That is, we assume $g^\pm(t) \rightarrow \pm\infty$ when $H_t = H^+$ and $\Sigma^{++} \rightarrow 0$, where $\langle W_t^+ W_t^+ \rangle = \Sigma^{++}(t) \cdot t$, and analogous relations hold when $H_t = H^-$. As a result, when $H_t = H^+$, then $u_t^+ \rightarrow 1$ and $u_t^- \rightarrow 0$, which we demonstrate in Appendix 7.6. Plugging the expressions $u_t^+ = 1$ and $u_t^- = 0$ into Eq. (56b) for w_t^+ , we find

$$dw_t^+ = - \left[\frac{1}{t + \beta} \right] w_t^+ dt. \quad (58)$$

Next, we write Eq. (56b) for w_t^- in the form

$$u_t^- dw_t^- = - [u_t^- w_t^- + u_t^- w_t^+ - w_t^+] \left[\frac{1}{t + \beta} + u_t^+ (w_t^+ - w_t^-) \right] dt,$$

so by plugging in $u_t^+ = 1$ and $u_t^- = 0$, we find $0 = w_t^+ \left[\frac{1}{t + \beta} + w_t^+ - w_t^- \right]$, which, when $w_t^+ \neq 0$, simplifies to

$$w_t^- = w_t^+ + \frac{1}{t + \beta}. \quad (59)$$

An analogous pair of equations holds when $H_t = H^-$ and thus $u_t^- \rightarrow 1$ and $u_t^+ \rightarrow 0$. Solving Eq. (58,59) and their $H_t = H^-$ counterparts iteratively, we find that in the limit of no observation-noise (e.g., $g^\pm(t) \rightarrow \pm\infty$ when $H_t = H^\pm$),

$$w_t^\pm = \frac{w_0\beta + m}{t + \beta}, \quad w_t^\mp = \frac{w_0\beta + m + 1}{t + \beta}, \quad (60)$$

where $H_t = H^\pm$ and $w_0 := w^j(0)$ for $H(0) = H^j$, so w_0 constitutes the initial estimate of the change rate of the environment. Here, m is the number of changepoints in the time series H_t during the time interval $[0, t]$. Eq. (60) can be re-expressed in the form of a rate-based plasticity rule

$$dw_t^\pm = [\delta(u_t^+ - u_t^-) - w_t^\pm] \cdot C_t dt, \quad (61)$$

where $\delta(u)$ is the Dirac delta distribution, along with Eq. (57) for the concentration decay of the agent C_t . Note, that the non-negative term $\delta(u_t^+ - u_t^-)$ in Eq. (61) results in long term potentiation (LTP) of both synaptic weights, $w_t^\pm$, whenever the neural activities, $u_t^\pm$, are both high, *i.e.* when their values cross at $u_t^\pm = 0.5$. During such changes, the weights $w_t^\pm$ are incremented. Outside of these transient switching epochs, there is a constant long term depression (LTD) of the synaptic weights $w_t^\pm$ modeled by the term $(-w_t^\pm)$.

We schematize (Fig. 6A) and simulate (Fig. 6B,C) the resulting plastic neural population network:

$$du_t^\pm = u_t^\pm [I^\pm(t)dt + dW_t^\pm] + [w_t^\mp u_t^\mp - w_t^\pm u_t^\pm] dt \quad (62a)$$

$$dw_t^\pm = [\delta(u_t^+ - u_t^-) - w_t^\pm] \cdot C_t dt. \quad (62b)$$

The constant $1/2$ has been absorbed into the population input, so that $I^\pm(t) = g^\pm(t) + 1/2$. It is important to note that Eq. (62) only performs optimal inference in the limit of no observation-noise. Perturbing away from this case, we expect the performance to be sub-optimal. However, as can be seen in Fig. 6C, the correct change rate is approximated reasonably well.

Eq. (62), thus models evidence accumulation in a symmetrically changing environment when the change rate, ϵ , is not known a priori. The model is based on the recursive equation for the joint probability of the environmental state, $H^\pm$, and change-point count, a , derived in Section 3.1. We obtained a tractable model by first passing to the continuum limit, and then applying a moment closure approximation to reduce the dimension of the resulting equations. Obtaining the low-dimensional approximation in Eq. (54) was crucial to obtaining a neural population model that approximates state inference. We next extend this model to the case of asymmetric rates of change.

Asymmetric environment. The continuum limit of the inference process in an asymmetric environment, Eq. (47), provides several pieces of information we can use to identify an approximate neural population model. First, under the assumption of large signal-to-noise ratios, the synaptic weights should evolve to reflect the number of detected change-points, rescaled by the amount of time spent in each state

$$w_t^+ = \frac{w_0^+\beta^+ + m_t^+}{t^+ + \beta^+}, \quad w_t^- = \frac{w_0^-\beta^- + m_t^-}{t^- + \beta^-}, \quad (63)$$

Figure 6: Neural network model with plasticity, inferring the current state H_t and rates $\epsilon^\pm$ of environmental change. (A) Schematic showing the synaptic weight $w^\pm$ from neural population $u^\pm \mapsto u^\mp$ evolving through long term potentiation (LTP) and long term depression (LTD) to match the environment's rate of change, $\epsilon^\pm := \epsilon$. (B) When the neural populations exchange dominance, their activity levels $u^\pm$ are both transiently high. As a result, both synaptic weights, $w^\pm$, increase via LTP. When only one population is active, both weights decay via LTD, as described by Eq. (62b). (C) Inference of the rate, ϵ , via long term plasticity of the weights for $\epsilon = 0.01, 0.05, 0.1$. Though the signal-to-noise ratio is finite (See Appendix 7.7), the weights in the network described by Eq. (62) converge to the actual rate, ϵ . (D) Schematic showing the evolution of weights $w_t^\pm$ when rates are asymmetric, $\epsilon^+ > \epsilon^-$, so that $w_\infty^+ > w_\infty^-$. The network is described by Eq. (65). (E) Only the weight $w_t^\pm$ decays through LTD when population $u_t^\pm$ is active, and only the weight $w_t^\pm$ is potentiated through LTP when dominance switches from $u_t^\pm$ to $u_t^\mp$, as in Eq. (65). (F) Network weights $w_t^\pm$ converge to the asymmetric rates, $\epsilon^\pm$. See Appendix 7.7 for details about the simulations.

where $w_0^\pm := w^\pm(0)$ are the network's initial estimates of the change rates $\epsilon^\pm$, $m^\pm$ is the true number of changepoints $H^\pm \mapsto H^\mp$ during the time interval $[0, t]$, and $t^\pm$ is the total length of time spent in the state $H^\pm$.⁴ Second, the flux term in Eq. (47) indicates that a memory process is needed to store the estimated time $t^\pm$ spent in each state $H^\pm$. This can be accomplished by modifying Eq. (57) for the plasticity agent, so that each $C_t^\pm$ decays only when the neural population of origin, $u_t^\pm$, is active. Thus we obtain the pair of equations:

$$dC_t^\pm = -H(u_t^\pm - \theta) [C_t^\pm]^2. \quad (64)$$

Expressing Eq. (63) as a system of equations for the synaptic weights, $w_t^\pm$, yields:

$$dw_t^\pm = H(u_{t-\tau}^\pm - \theta) [\delta(u_t^+ - u_t^-) - w_t^\pm] \cdot C_t^\pm dt. \quad (65)$$

Here the function $H(u^\pm(t - \tau) - \theta)$ for $\theta \geq 0.5$, and $\tau > 0$ enforces the requirement that the population $u_t^\pm$ must have a high rate of activity prior to the LTP event. Thus, to learn asymmetric weights, there should be a small delay τ accounting for the time it takes for the presynaptic firing rate to trigger the plasticity process (Gütig et al, 2003). We demonstrate the performance of the network whose weights evolve according to Eq. (65) in Fig. 6D,E. Note, the network with weights evolving according to Eq. (65) can still infer symmetric transition rates $\epsilon^\pm = \epsilon$, but it will do so at half the rate of the network Eq. (62). This is due to the fact that Eq. (65) counts changepoints and dwell times of each state $H^\pm$ separately.

We have thus shown that the recursive update equations for the state probability in a dynamic environment lead to plausible neural network models that approximate the same inference. Previous neural network models of decision making have tended to interpret population rates as a representation of posterior probability (Bogacz et al, 2006; Beck et al, 2008). We have shown that the synaptic weight between populations can represent the change rate of the environment. As a result, standard rate-correlation models of plasticity can be used to implement the change rate inference process.

7 Discussion

Evidence integration models have a long history in neuroscience (Ratcliff and McKoon, 2008). These normative models conform with behavioral observations across species (Brunton et al, 2013), and have been used to explain the neural activity that underpins decisions (Gold and Shadlen, 2007). However, animals make decisions in an environment that is seldom static (Portugues and Engert, 2009). The relevance of available information, the accessibility, and the payoff of different choices can all fluctuate. It is thus important to extend evidence accumulation models to such cases.

We have shown how ideal observers accumulate evidence to make decisions when there are multiple, discrete choices, and the correct choice changes in time. We assumed that the rates of transition between environmental states are initially unknown to the observer. An ideal observer must therefore integrate information from measurements to

⁴We use the notation $H^\pm$ for the two states here, for convenience and consistency with Eq. (62). Similarly, we use $\epsilon^\pm$ and $t^\pm$ rather than the numerical notation of the asymmetric case in Section 5.1.

concurrently estimate both the transition rates and the current state of the environment. Importantly, these two inference processes are coupled: Knowledge of the rate allows the observer to appropriately discount older information to infer the present state, while knowledge of transitions between states is in turn necessary to infer the rate.

Inference when all transition rates are identical is straightforward to implement in resulting models. An ideal observer only needs to track the probability of the environmental state and the total changepoint count, regardless of the states between which the change occurred. However, when the transition rates are asymmetric, the resulting models are more complex. In this case, an ideal observer must estimate a matrix of changepoint counts, distinguished by the starting and ending states. The number of possible matrices grows polynomially with the number of observations. This computation is difficult to implement, and we do not suggest that animals make inferences about environmental variability in this way. However, understanding the ideal inference process allowed us to identify its most important features. In turn, we derived tractable approximations and plausible neural implementations, whose performance compared well with that of an ideal observer (Fig. 5D,E,F). We believe humans and other animals do generally implement approximate strategies when they need to infer such rates (Lange and Dukas, 2009). Ideal observer models allow us to understand what inferences can be made with the available information, which assumptions of the observer are important (e.g., assuming an incorrect transition rate does not always have a large impact on performance), and how such inferences could be approximated in networks of the brain and other biological computers.

In many naturally occurring decisions like foraging, mate selection, and home-site choice, animals simply need to identify the best alternative rather than the rate of environmental change (Johnson et al, 2013). Therefore, rapid approximations, or a guess of the environmental change rate may provide better initial performance than learning the rate, which could be slow. Moreover, it appears that when measurements are noisy, rates cannot be learned precisely even in the limit of infinite observations. Thus, learning the rate may only be useful when noise is too high for single measurements to determine the correct alternative, but sufficiently low to make rate inference possible. There is evidence that humans adjust their rate of evidence-discounting, based on the actual change rate of the environment (Glaze et al, 2015). However, further psychophysical studies are needed to identify whether subjects use heuristic strategies to learn or something close to the normative models we derived here.

A number of related models have been developed previously (Wilson et al, 2010; Adams and MacKay, 2007). The present model is somewhat different, as a finite number of choices implies that the present environmental state is dependent on the previous state. As a result, we found it was more efficient to implement an update equation that estimated the present environmental state and the changepoints, rather than the time in the present state.

Several of the assumptions we have made in this study could be modified to extend our analysis to more general situations. For instance, we have assumed that the observer's eventual choice does not affect the environment. However, in many natural situations changes in the environment are a consequence of the observer's actions (Cisek and Pastor-Bernier, 2014). In more realistic situations it is likely that there is a sequence of actions leading to an ultimate decision, and each action can influence the informa-

tion available to the observer. An animal making a foraging decision in a group collects more evidence once it moves toward a particular food patch, but it may also draw other members with it, changing the subsequent availability of food there (Petit et al, 2009). Thus including a sequence of actions, and their impact on the available information and the environment would be necessary in a realistic model. Another possibility is that changes to the environment are non-Markovian and/or involve multiple timescales. Extending our ideal observer models to estimate such change statistics might require derivation of multi-step update equations. In such cases, we expect the truncations we have applied in this work would be useful for identifying tractable approximations of the optimal inference process.

Optimal models of evidence accumulation are useful both as baselines to compare to performance in psychophysical experiments, and starting points for identifying plausible neuronal network implementations. Our core contribution here has been to present a general model of evidence accumulation in a dynamic environment, when an observer has no prior knowledge of the rate of change. An unavoidable feature of these models is that the number of variables the observer must track grows as more observations are made, and growth is more rapid in asymmetric environments with multiple environmental states. This motivated our development of continuum approximations and low-dimensional moment equations for the optimal models, which suggest more plausible neural computations. We hope this work will foster future theoretical studies that will extend this framework, as well as experiments that could validate the models herein. To fully understand the neural mechanisms of evidence accumulation, we must account for the wide variety of conditions that organisms encounter when making decisions.

Acknowledgments

Funding was provided by NSF-DMS-1517629 (AER, KJ, and ZPK); NSF-DMS-1311755 (ZPK); and NSF/NIGMS-R01GM104974 (KJ).

Appendix

7.1 Two-state system with unknown symmetric rate

We show how to derive Eq. (3) from the main text. Bayes' rule and the law of total probability first yield:

$$P_n(H_n, a_n) = \frac{1}{P(\xi_{1:n})} \sum_{H_{n-1}=H^\pm} \sum_{a_{n-1}=0}^{n-2} P(\xi_{1:n}|H_n, H_{n-1}, a_n, a_{n-1})P(H_n, H_{n-1}, a_n, a_{n-1}).$$

Using the conditional independence of observations,

$$P(\xi_{1:n}|H_n, H_{n-1}, a_n, a_{n-1}) = P(\xi_n|H_n)P(\xi_{1:n-1}|H_{n-1}, a_{n-1}),$$

we find that,

$$P_n(H_n, a_n) = \frac{1}{P(\xi_{1:n})} \sum_{H_{n-1}=H^\pm} \sum_{a_{n-1}=0}^{n-2} P(\xi_n|H_n)P(\xi_{1:n-1}|H_{n-1}, a_{n-1})P(H_n, H_{n-1}, a_n, a_{n-1}).$$

Furthermore, we can use the definition of conditional probability to write,

$$P(H_n, H_{n-1}, a_n, a_{n-1}) = P(H_n, a_n | H_{n-1}, a_{n-1}) P(H_{n-1}, a_{n-1}),$$

and Bayes' rule also implies,

$$P(\xi_{1:n-1} | H_{n-1}, a_{n-1}) P(H_{n-1}, a_{n-1}) = P_{n-1}(H_{n-1}, a_{n-1}) P(\xi_{1:n-1}).$$

Hence, we derive Eq. (3) from the main text,

$$P_n(H_n, a_n) = \frac{P(\xi_{1:n-1})}{P(\xi_{1:n})} P(\xi_n | H_n) \sum_{H_{n-1}=H^\pm} \sum_{a_{n-1}=0}^{n-2} P_{n-1}(H_{n-1}, a_{n-1}) P(H_n, a_n | H_{n-1}, a_{n-1}).$$

7.2 Numerical methods for free response protocol

The free response protocol is simulated by evolving the update Eq. (17) and subsequently computing the log likelihood ratio $L_n := \log(R_n)$ using Eq. (18) at each timestep n . Each point along the curves in Fig. 3C corresponds to an average waiting time and average performance corresponding to a threshold value θ over 100,000 simulations. For each value of θ , the simulation is terminated when $|L_n| > \theta$ and the choice is given by the sign of L_n . To avoid excessively long simulations, we removed any that lasted longer than $n = 5000$, but we found changing this upper bound did not affect averages considerably. There were 400 values of θ chosen, discretizing the interval from $\theta = 0$ to $\theta = 3.89$.

7.3 Continuum limit for two states with asymmetric rates

We begin by considering Eq. (34a), which provides an update of the probability of being in state H^1 after n observations, given the specific changepoint matrix $\mathbf{a}$:

$$P_n(H^1, \mathbf{a}) = \mathcal{F}_{n,\Delta t}^1 \left[\left(1 - \frac{a^{21} + 1}{1 + a^{21} + a^{11}} \right) P_{n-1}(H^1, \mathbf{a} - \delta^{11}) + \frac{a^{12}}{1 + a^{12} + a^{22}} P_{n-1}(H^2, \mathbf{a} - \delta^{12}) \right], \quad (66)$$

where we have defined $\mathcal{F}_{n,\Delta t}^1 = \frac{P(\xi_{1:n-1})}{P(\xi_{1:n})} f_{\Delta t}^1(\xi_n)$. Subsequently, we divide by $P_{n-1}(H^1, \mathbf{a})$ and take the logarithm to find:

$$\Delta x_n^1(\mathbf{a}) = \ln \mathcal{F}_{n,\Delta t}^1 + \ln \left[\left(1 - \frac{a^{21} + 1}{1 + a^{21} + a^{11}} \right) e^{x_{n-1}^1(\mathbf{a} - \delta^{11}) - x_{n-1}^1(\mathbf{a})} + \frac{a^{12}}{1 + a^{12} + a^{22}} e^{x_{n-1}^2(\mathbf{a} - \delta^{12}) - x_{n-1}^1(\mathbf{a})} \right].$$

Now, as we are taking the continuum limit, we consider $P(H_{t+\Delta t} = H^\pm | H_t = H^\mp) = \epsilon^{ij} \Delta t + o(\Delta t)$, where $0 \leq \epsilon^{ij} < \infty$. Given a $\mathbb{R}^{2 \times 2}$ transition rate matrix with off-diagonal entries $\epsilon^{ij} \Delta t + o(\Delta t)$ and diagonal entries $1 - \epsilon^{ij} \Delta t + o(\Delta t)$, the expected changepoint and non-changepoint counts after n timesteps will be $\langle a_{\Delta t}^{ij} \rangle = \Delta t \epsilon^{12} \epsilon^{21} n / (\epsilon^{12} +$

ϵ^{21}) and $\langle a_{\Delta t}^{ii} \rangle = \epsilon^{ij} n / (\epsilon^{12} + \epsilon^{21})$. Thus, while the changepoint counts $a_{\Delta t}^{ij}$ scale with Δt , the non-changepoint counts do not. In the continuum limit, we will choose a time $t := n\Delta t$ and take $\Delta t \rightarrow 0$ while keeping t constant, so the number of timesteps diverges like $n = t/(\Delta t)$ for a fixed time t . While the expected changepoint counts $\langle a_{\Delta t}^{ij} \rangle$ thus remain fixed, the non-changepoint counts $\langle a_{\Delta t}^{ii} \rangle$ will diverge as $(\Delta t)^{-1}$, suggesting we should rescale non-changepoint counts to the absolute dwell time $t_{\Delta t}^i = \Delta t a_{\Delta t}^{ii}$. The expected value of the dwell times is then finite in this limit $\lim_{\Delta t \rightarrow 0} \langle t_{\Delta t}^i \rangle = \epsilon^{ij} t / (\epsilon^{12} + \epsilon^{21})$. Performing this change of variables, we then define the changepoint matrix as involving changepoint counts a^{ij} on the off-diagonal and dwell times along the diagonal: $\mathbf{A} = \begin{pmatrix} t^1 & a^{12} \\ a^{21} & t^2 \end{pmatrix}$ so an increment of non-changepoint count a^{ii} now takes the form $\mathbf{A} + \Delta t \delta^{ii}$. As such, we can now expand:

$$e^{x^i(\mathbf{A} - \delta^{ii})} = e^{x^i(\mathbf{A} - \Delta t \delta^{ii})} = e^{x^i(\mathbf{A})} - \Delta t e^{x^i(\mathbf{A})} \frac{\partial x^i(\mathbf{A})}{\partial t^i} + \mathcal{O}((\Delta t)^2),$$

via application of the chain rule and noting $dt^i/da^{ii} = \Delta t$. Note, we cannot perform such an expansion in Δt to $x^i(\mathbf{A} - \delta^{ij})$, since perturbations to the matrix $\mathbf{A}$ in this case are $\mathcal{O}(1)$. Truncating Eq. (66) to terms of $\mathcal{O}(\Delta t)$ and incorporating the Poisson-delta prior Eq. (46), we find the discrete update equation becomes:

$$\Delta x_n^1(\mathbf{A}) = \ln \mathcal{F}_{n,\Delta t}^1 + \Delta t \cdot \left[\frac{a^{12} + \alpha_2 - 1}{t^2 + \beta_2} e^{x_{n-1}^2(\mathbf{A} - \delta^{12}) - x_{n-1}^1(\mathbf{A})} - \frac{a^{21} + \alpha_1}{t^1 + \beta_1} - \frac{\partial x_{n-1}^1(\mathbf{A})}{\partial t^1} \right].$$

Lastly, upon taking the continuum limit $\Delta t \rightarrow 0$, we find that

$$dx_t^1(\mathbf{A}) = [g^1(t)dt + dW_t^1] + \left[\frac{a^{12} + \alpha_2 - 1}{t^2 + \beta_2} e^{x_t^2(\mathbf{A} - \delta^{12}) - x_t^1(\mathbf{A})} - \frac{a^{21} + \alpha_1}{t^1 + \beta_1} - \frac{\partial x_t^1(\mathbf{A})}{\partial t^1} \right] dt,$$

where the statistics of the drift $g^1(t)$ and noise dW_t^1 are analogous to those given after Eq. (42), only the transition rates of H_t from $H^j \mapsto H^i$ are now ϵ^{ij} . Note, due to the flux term $\frac{\partial x_t^1(\mathbf{A})}{\partial t^1}$ and continuum values for $t^j \in \mathbb{R}^*$, this is a stochastic partial differential equation (SPDE). An analogous SPDE can be derived for $x_t^2(\mathbf{A})$ in the same way.

7.4 Continuum limit with multiple states and symmetric rates

The derivation parallels that with two states. To obtain the continuum limit, we use the generalized version of Eq. (3),

$$P_n(H_n, a_n) = \frac{P(\xi_{1:n-1})}{P(\xi_{1:n})} P(\xi_n | H_n) \sum_{H_{n-1}} \sum_{a_{n-1}=0}^{\infty} P_{n-1}(H_{n-1}, a_{n-1}) P(H_n, a_n | H_{n-1}, a_{n-1}). \quad (67)$$

Assuming again a Gamma prior on the transition rate, $\epsilon \sim \text{Gamma}(\alpha, \beta)$, and following the derivations of Eq. (23) and (38), we obtain

$$P(H_n, a_n | H_{n-1}, a_{n-1}) = \begin{cases} 1 - \Delta t \frac{a_n + \alpha}{t_{n-1} + \beta} & H_n = H_{n-1} \text{ \& } a_n = a_{n-1} \\ \Delta t \frac{a_n + \alpha - 1}{(N-1)(t_{n-1} + \beta)} & H_n \neq H_{n-1} \text{ \& } a_n = a_{n-1} + 1 \\ 0 & \text{otherwise.} \end{cases} \quad (68)$$

Using Eq. (68) in Eq. (67) yields

$$P_n(H^i, a) = \frac{P(\xi_{1:n-1})}{P(\xi_{1:n})} f_{\Delta t}^i(\xi_n) \left[\left(1 - \Delta t \frac{a + \alpha}{t_{n-1} + \beta} \right) P_{n-1}(H^i, a) + \Delta t \frac{a + \alpha - 1}{(N-1)(t_{n-1} + \beta)} \sum_{j \neq i} P_{n-1}(H^j, a-1) \right].$$

Dividing by $P_{n-1}(H^i, a)$, taking logarithms, and denoting $x_{t_n}^i(a) := \ln P_n(H^i, a)$ we obtain

$$\Delta x_{t_n}^i(a) \propto \ln f_{\Delta t}^i(\xi_n) + \ln \left[1 - \Delta t \frac{a + \alpha}{t_{n-1} + \beta} + \Delta t \frac{a + \alpha - 1}{(N-1)(t_{n-1} + \beta)} \sum_{j \neq i} e^{x_{t_{n-1}}^j(a-1) - x_{t_{n-1}}^i(a)} \right].$$

Using the approximation $\ln(1+z) \approx z$ valid for small z yields

$$\Delta x_{t_n}^i(a) \propto \ln f_{\Delta t}^i(\xi_n) + \Delta t \left(\frac{a + \alpha - 1}{(N-1)(t_{n-1} + \beta)} \sum_{j \neq i} e^{x_{t_{n-1}}^j(a-1) - x_{t_{n-1}}^i(a)} - \frac{a + \alpha}{t_{n-1} + \beta} \right).$$

Similar to the $N = 2$ case, we may then take the continuum limit to yield Eq. (48).

7.5 Consistency of the moment hierarchy equations

We begin by taking the SDE given by Eq. (2) for $N = 2$ states, and the known rate ϵ , and changing variables to $\bar{P}_t^\pm = e^{x_t^\pm}$, so

$$d\bar{P}_t^\pm = \bar{P}_t^\pm [g^\pm(t)dt + dW^\pm] + \epsilon [\bar{P}_t^\mp - \bar{P}_t^\pm] dt.$$

Furthermore, note that in the limit $t \rightarrow \infty$, the $\mathcal{O}(1)$ terms and $\mathcal{O}(t^{-1})$ terms vanish in Eq. (54):

$$d\bar{P}_t^\pm = \bar{P}_t^\pm \left[\left(g^\pm(t) + \frac{1}{2} \right) dt + dW^\pm \right] + [\bar{A}_t^\mp - \bar{A}_t^\pm] dt \quad (69a)$$

$$d\bar{A}_t^\pm = \bar{A}_t^\pm(t) \left[\left(g^\pm(t) + \frac{1}{2} \right) dt + dW^\pm \right] + (\bar{A}_t^\mp - \bar{A}_t^\pm) (\bar{A}_t^\mp + \bar{A}_t^\pm) dt. \quad (69b)$$

Therefore, in the event that $\bar{A}_t^\pm \rightarrow \epsilon \bar{P}_t^\pm$ in the long time limit ($t \rightarrow \infty$), we find the truncated system, Eq. (69), becomes

$$d\bar{P}_t^\pm = \bar{P}_t^\pm \left[\left(g^\pm(t) + \frac{1}{2} \right) dt + dW^\pm \right] + \epsilon \cdot [\bar{P}_t^\mp - \bar{P}_t^\pm] dt \quad (70a)$$

$$\epsilon d\bar{P}_t^\pm = \epsilon \bar{P}_t^\pm(t) \left[\left(g^\pm(t) + \frac{1}{2} \right) dt + dW^\pm \right] + \epsilon^2 \cdot (\bar{P}_t^\mp - \bar{P}_t^\pm) (\bar{P}_t^\mp + \bar{P}_t^\pm) dt. \quad (70b)$$

Dividing by ϵ and noting that $\bar{P}_t^+ + \bar{P}_t^- = 1$, Eq. (70b) becomes

$$d\bar{P}_t^\pm = \bar{P}_t^\pm \left[\left(g^\pm(t) + \frac{1}{2} \right) dt + dW^\pm \right] + \epsilon \cdot [\bar{P}_t^\mp - \bar{P}_t^\pm] dt,$$

which is consistent with Eq. (70a), and indicates the truncated moment hierarchy is consistent with the case of known rates and two choices ($N = 2$) in the SDE Eq. (2).

7.6 Noise-free limit of the neural population model

Consider the neural population Eq. (56a) for the evolution of u_t^- in the event of environmental state $H_t = H^+$ and no observation-noise $f^\pm(\xi|H^+) = \delta(\xi - \xi^\pm)$. As a result, the drift terms diverge $g^\pm(t) := \lim_{\Delta t \rightarrow 0} \frac{1}{\Delta t} \mathbb{E} [\ln f_{\Delta t}^\pm(\xi)|H^+] \rightarrow \pm\infty$ and the covariance matrix $\Sigma^{ij}(t) := \lim_{\Delta t \rightarrow 0} \frac{1}{\Delta t} \text{Cov} [\ln f_{\Delta t}^i(\xi), \ln f_{\Delta t}^j(\xi)|H^+] \rightarrow 0$. Thus, the dominant terms on the right hand side of Eq. (56a) for u_t^- come from the input so $du_t^- = -u_t^-|g^-|dt$, and the population activity immediately decays to $u_t^- = 0$. As a result, since $u_t^+ + u_t^- = 1$, we expect $u_t^+ = 1$, when $H_t = H^+$.

7.7 Numerical simulations of SDE models

Stochastic differential equation (SDE) models of evidence accumulation in symmetric environments changing between two states are simulated using a standard Euler-Maruyama integration algorithm (Higham, 2001). Eq. (43) describes the evolution of an infinite number of SDEs over the changepoint vector $a \in \mathbb{Z}_{\geq 0}$, state vector $H_t \in \{H^+, H^-\}$, and time $t \in [0, T]$, so we truncate this space to $a \in \{0, 1, 2, \dots, 1000\}$, which is sufficient for transition rates ϵ and total simulation times T not too large. We compared our results to cases with longer state vectors $a \in \{0, 1, 2, \dots, a_{\max}\}$ and the changes were negligible. Simulations shown in Fig. 5 had a transition rate of $\epsilon = 0.1$ and total run time of $T = 1000$ with timestep $dt = 10^{-3}$. Observations were sampled from a normal distribution $f^\pm(\xi) = \frac{e^{-(\xi \mp \mu)^2/(2\sigma^2)}}{\sqrt{2\pi\sigma^2}}$ with mean $\mu = 0.5$ and variance $\sigma^2 = 1$, so the signal-to-noise ratio was $2\mu/\sigma = 1$. Initial conditions were chosen so that $P_0^\pm = 0.5$ and $P_0(a) = \frac{\alpha^a e^{-\alpha}}{a!}$ where $\alpha = 1$ and $\beta = 5$. A similar approach was used to numerically simulate the neural population model Eq. (62) and its variants to produce Fig. 6, but $\mu = 1$ and $\sigma = 0.1$, so the signal-to-noise ratio was 20.

References

- Adams RP, MacKay DJ (2007) Bayesian online changepoint detection. arXiv preprint arXiv:07103742
- Beck JM, Ma WJ, Kiani R, Hanks T, Churchland AK, Roitman J, Shadlen MN, Latham PE, Pouget A (2008) Probabilistic population codes for bayesian decision making. *Neuron* 60(6):1142–1152
- Bogacz R, Brown E, Moehlis J, Holmes P, Cohen JD (2006) The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol Rev* 113(4):700–65, DOI 10.1037/0033-295X.113.4.700
- Brunton BW, Botvinick MM, Brody CD (2013) Rats and humans can optimally accumulate evidence for decision-making. *Science* 340(6128):95–8, DOI 10.1126/science.1233912
- Churchland AK, Kiani R, Shadlen MN (2008) Decision-making with multiple alternatives. *Nat Neurosci* 11(6):693–702, DOI 10.1038/nn.2123

- Cisek P, Pastor-Bernier A (2014) On the challenges and mechanisms of embodied decisions. *Philosophical transactions of the Royal Society of London Series B, Biological sciences* 369(1655)
- Deneve S (2008) Bayesian spiking neurons i: inference. *Neural Comput* 20(1):91–117, DOI 10.1162/neco.2008.20.1.91
- Djuric PM, Huang Y (2000) Estimation of a bernoulli parameter p from imperfect trials. *IEEE Signal Processing Letters* 7(6):160–163
- Franks NR, Pratt SC, Mallon EB, Britton NF, Sumpter DJ (2002) Information flow, opinion polling and collective intelligence in house-hunting social insects. *Philosophical Transactions of the Royal Society of London B: Biological Sciences* 357(1427):1567–1583
- Gardiner CW (2004) *Handbook of stochastic methods for physics, chemistry, and the natural sciences*, 3rd edn. Springer-Verlag, Berlin
- Glaze CM, Kable JW, Gold JJ (2015) Normative evidence accumulation in unpredictable environments. *Elife* 4:e08,825
- Gold JJ, Shadlen MN (2007) The neural basis of decision making. *Annu Rev Neurosci* 30:535–74, DOI 10.1146/annurev.neuro.29.051605.113038
- Gütig R, Aharonov R, Rotter S, Sompolinsky H (2003) Learning input correlations through nonlinear temporally asymmetric hebbian plasticity. *The Journal of neuroscience* 23(9):3697–3714
- Higham DJ (2001) An algorithmic introduction to numerical simulation of stochastic differential equations. *SIAM review* 43(3):525–546
- Huk AC, Shadlen MN (2005) Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. *J Neurosci* 25(45):10,420–36, DOI 10.1523/JNEUROSCI.4684-04.2005
- Johnson DD, Blumstein DT, Fowler JH, Haselton MG (2013) The evolution of error: Error management, cognitive constraints, and adaptive decision-making biases. *Trends in ecology & evolution* 28(8):474–481
- Kira S, Yang T, Shadlen MN (2015) A neural implementation of wald’s sequential probability ratio test. *Neuron* 85(4):861–873
- Krajibich I, Rangel A (2011) Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. *Proceedings of the National Academy of Sciences* 108(33):13,852–13,857
- Kuehn C (2016) Moment closure—a brief review. In: *Control of Self-Organizing Non-linear Systems*, Springer, pp 253–271

- Lange A, Dukas R (2009) Bayesian approximations and extensions: optimal decisions for small brains and possibly big ones too. *Journal of theoretical biology* 259(3):503–516
- Machens CK, Romo R, Brody CD (2005) Flexible control of mutual inhibition: a neural model of two-interval discrimination. *Science* 307(5712):1121–1124
- McGuire JT, Nassar MR, Gold JJ, Kable JW (2014) Functionally dissociable influences on learning rate in a dynamic environment. *Neuron* 84(4):870–881
- McMillen T, Holmes P (2006) The dynamics of choice among multiple alternatives. *Journal of Mathematical Psychology* 50(1):30–57
- Miller KD (1994) A model for the development of simple cell receptive fields and the ordered arrangement of orientation columns through activity-dependent competition between on-and off-center inputs. *J Neurosci* 14:409–441
- Niv Y, Daniel R, Geana A, Gershman SJ, Leong YC, Radulescu A, Wilson RC (2015) Reinforcement Learning in Multidimensional Environments Relies on Attention Mechanisms. *Journal of Neuroscience* 35(21):8145–8157
- Olberg R, Worthington A, Venator K (2000) Prey pursuit and interception in dragonflies. *Journal of Comparative Physiology A* 186(2):155–162
- Pearson JM, Heilbronner SR, Barack DL, Hayden BY, Platt ML (2011) Posterior cingulate cortex: adapting behavior to a changing world. *Trends in cognitive sciences* 15(4):143–151
- Petit O, Gautrais J, Leca JB, Theraulaz G, Deneubourg JL (2009) Collective decision-making in white-faced capuchin monkeys. *Proceedings of the Royal Society of London B: Biological Sciences* 276(1672):3495–3503
- Pfister JP, Gerstner W (2006) Triplets of spikes in a model of spike timing-dependent plasticity. *J Neurosci* 26(38):9673–9682
- Portugues R, Engert F (2009) The neural basis of visual behaviors in the larval zebrafish. *Current opinion in neurobiology* 19(6):644–647
- Ratcliff R, McKoon G (2008) The diffusion decision model: theory and data for two-choice decision tasks. *Neural computation* 20(4):873–922
- Redner S (2001) A guide to first-passage processes. Cambridge University Press
- Shvartsman M, Srivastava V, Cohen JD (2015) A theory of decision making under dynamic context. In: *Advances in Neural Information Processing Systems*, pp 2476–2484
- Smith PL, Ratcliff R (2004) Psychology and neurobiology of simple decisions. *Trends Neurosci* 27(3):161–8, DOI 10.1016/j.tins.2004.01.006

- Socha L (2007) Linearization methods for stochastic dynamic systems, vol 730. Springer Science & Business Media
- Sugrue LP, Corrado GS, Newsome WT (2004) Matching behavior and the representation of value in the parietal cortex. *science* 304(5678):1782–1787
- Veliz-Cuba A, Kilpatrick ZP, Josic K (2016) Stochastic models of evidence accumulation in changing environments. *SIAM Rev* 58:264–289
- Wald A, Wolfowitz J (1948) Optimum character of the sequential probability ratio test. *The Annals of Mathematical Statistics* 19(3):326–339
- Whittle P (1957) On the use of the normal approximation in the treatment of stochastic processes. *Journal of the Royal Statistical Society Series B (Methodological)* pp 268–281
- Wilson RC, Niv Y (2011) Inferring relevance in a changing world. *Frontiers in human neuroscience* 5:189
- Wilson RC, Nassar MR, Gold JJ (2010) Bayesian online learning of the hazard rate in change-point problems. *Neural Comput* 22(9):2452–76
- Wong KF, Huk AC, Shadlen MN, Wang XJ (2007) Neural circuit dynamics underlying accumulation of time-varying evidence during perceptual decision making. *Front Comput Neurosci* 1:6, DOI 10.3389/neuro.10.006.2007