LETTER **Communicated by Rafal Bogacz**

Evidence Accumulation and Change Rate Inference in Dynamic Environments

Adrian E. Radillo
adrian@math.uh.edu
Department of Mathematics, University of Houston, Houston, TX 77204, U.S.A.

Alan Veliz-Cuba
avelizcuba1@udayton.edu
Department of Mathematics, University of Dayton, Dayton, OH 45469, U.S.A.

Krešimir Josić
josic@math.uh.edu
Department of Mathematics, University of Houston, Houston, TX 77204, U.S.A.;

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 that 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 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 $(n = 1, 2, \ldots)$. We denote each possible choice by H^i, $(i = 1, \ldots, N)$, with H_n being the correct choice at time t_n. The transition rates ^{ij}, $i \neq j$, correspond to the known probabilities that the state changes between two observations: $^{ij} = \mathrm{P}\left(H_n = H^i \mid H_{n-1} = H^j\right)$. The observer makes measurements, ξ_n, at times t_n with known conditional probability densities $f^i(\xi) = \mathrm{P}\left(\xi_n = \xi \mid H_n = \right.$ $\Rightarrow$ $\Rightarrow$. $\Rightarrow$ $\Rightarrow$

Thus, the transition rates, ν_{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 belief subsequently. Such contradictory evidence typically arrives after a change in the environment.

Following Veliz-Cuba et al. (2016), we take logarithms, $x_n^i := \ln P_n(H_n = H^i)$ and denote by $\delta x_n^i := x_n^i - x_{n-1}^i$ the change in log probability due to an observation at time t_n. Finally, we assume the time between observations $\delta t := t_n - t_{n-1}$ is small, and $\nu_t^{ij} = \nu^{ij}\delta t + o(\delta t)$ for $i \neq j$, so that dropping higher-order terms yields

$$\delta x_n^i = \ln f^i{}_t(\xi_n) - \ln\left(1 - \sum_{j\neq i} \nu^{ji}\delta t + \sum_{j\neq i} \nu^{ij} e^{x_{n-1}^j - x_{n-1}^i}\right), \quad i = 1, \ldots, N,$$

where the likelihood function $f^i{}_t(\xi)$ may vary with t. Next, we use the approximation $\ln(1+z) \approx z$ for $|z| \ll 1$ and replace the index n by time, t, to write

$$\delta x_t^i \approx g_t^i \delta t + \overline{\delta W}^i{}_t - \delta t \sum_{j\neq i} \left(\nu^{ij} e^{x_t^j - x_t^i} - \nu^{ji}\right), \quad i = 1, \ldots, N,$$

where the drift $g_t^i = \frac{1}{\delta t} E_\xi\left[\ln f^i{}_t(\xi) \mid H_t\right]$ is the expectation of $\ln f^i{}_t(\xi)$ over ξ, conditioned on the true state of the environment at time t, $H_t = H^1, \ldots, H^N$, and $W_t = (W^1{}_t, \ldots, W^N{}_t)$ follows a multivariate gaussian distribution with mean zero and covariance matrix Σ_t given by

$$\Sigma^{ij}_t = \frac{1}{\delta t}\mathrm{Cov}_\xi\left[\ln f^i{}_t(\xi), \ln f^j{}_t \mid H\right]$$

The nonlinear term in equation 2.2 implies that in the absence of noise, the system 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, ^{ij}, from measurements.

3 Environments with Symmetric Transition Rates

We first derive the ideal observer model when the unknown transition rates are symmetric, ^{ij} constant when j i, and ^{ii} : 1 $(N$ 1) ^{ij}. This simplifies the derivation, since the observer only needs to estimate a single change-point count. The asymmetric case discussed in section 4 follows

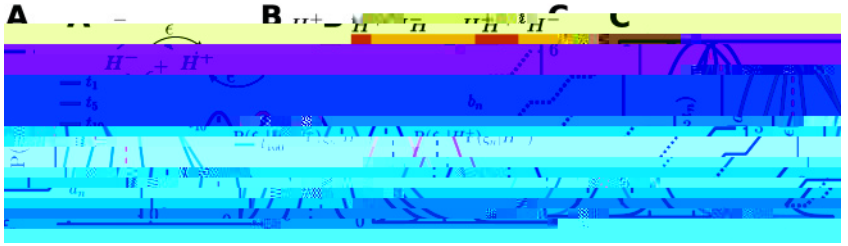

Figure 1: Online inference of the change rate in a dynamic environment. (A) The environment alternates between states H^- and H^+ with transition probabilities , . We analyze the symmetric case (:) in section 3.1 and the asymmetric case ( )

The process $a_{n\ n\ 1}$ is a pure birth process with birth rate . The observer

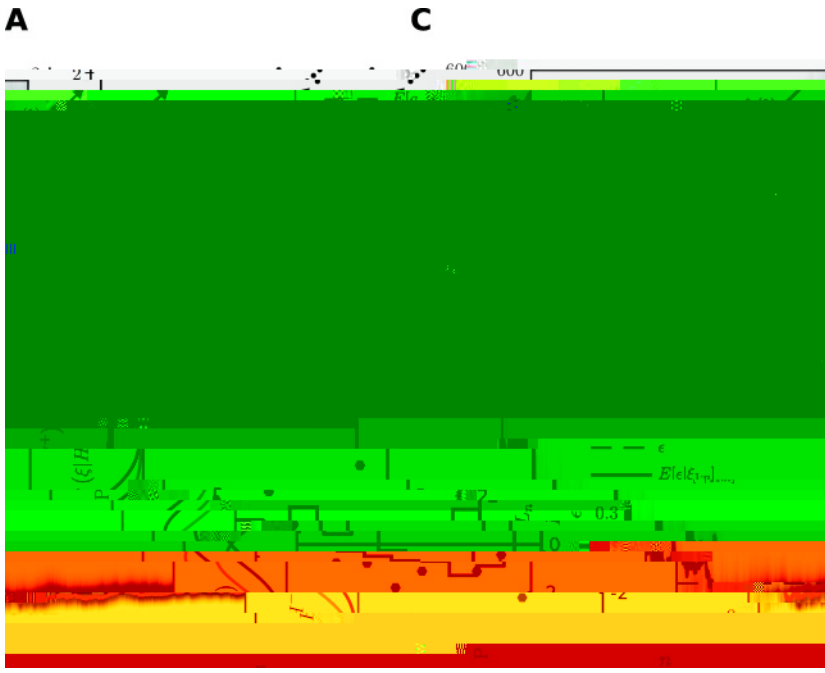

Figure 2: Inference of the states, H , and change rate, . (A) The joint posterior probability, $P_n(H , a)$, is propagated along a directed graph according to equation 3.14. Only paths corresponding to the initial condition (H_1, a_1) $(H , 0)$

and

$$P_n(H, n-1) \quad P(\xi_{1:n}$$

estimate the length of the interval since the last change point. We demonstrate the inference process defined by equation 3.14 in Figure 2.

The observer can compute the posterior odds ratio by marginalizing over the change-point count:

$$R_n : \quad \frac{\mathrm{P}_n\left(H \quad\right)}{\mathrm{P}_n\left(H \quad\right)} \quad \frac{\sum_{a \quad 0}^{n \quad 1} \mathrm{P}_n\left(H \quad, a\right)}{\sum_{a \quad 0}^{n \quad 1} \mathrm{P}_n\left(H \quad, a\right)}. \tag{3.15}$$

Here $\log(R_n) \quad L_n > 0$ implies that $H_n \quad H$ is more likely than $H_n \quad H$ (see Figure 2B). Note that $\mathrm{P}(\xi_{1:n \quad 1})/\mathrm{P}(\xi_{1:n})$ and 1ao81 Tm[(9.56.2401 Tm[())] TJETBT0.000

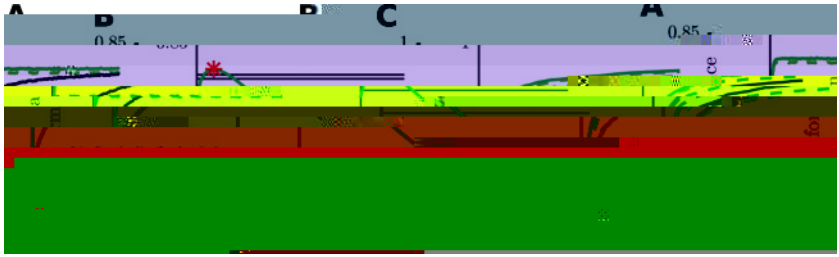

Figure 3: The performance of the inference algorithm. (A) Performance under the interrogation paradigm measured as the percentage of correct responses at the interrogation time. Here and in the next panel, 0.05, and SNR 1. The black curve represents the performance of an ideal observer who infers the change rate from measurements. The green curves represent the performance of observers who assume a fixed change rate (0.3, 0.15, 0.05, 0.03 from darker to lighter; see equation 2.1). The solid green line corresponds to an observer who assumes the true rate and the dashed lines to erroneous rates. (B) The green curve represents the performance at interrogation time t_{300} of an observer who assumes a fixed change rate. The red star marks the maximum of this curve, corresponding to the true change rate 0.05. The horizontal black curves represent the performance at times t_{40}, t_{100}, t_{200}, t_{300} (from bottom to top) of the observer who learns the change rate. (C) The accuracy as a function of the average threshold hitting time in the free response protocol. Here 0.1, and SNR 0.75. See section A.2 for details on numerical simulations. See also Figure 3 in Veliz-Cuba et al. (2016).

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 $n \to$, the posterior over

The observer can also compute the posterior probability $P_n(\ )$ of the transition rate by marginalizing over all states H_n and change-point counts a_n, as in equation 3.16. Furthermore, a point estimate of is given by the mean of the posterior after marginalizing, as in equation 3.17.

4 Environments with Asymmetric Transition Rates

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 in n.

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

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

where the sum is over all $H_{n-1} \in \{H^1, \ldots, H^N\}$ and possible values of the change-point matrix, $\boldsymbol{a}_{n-1}$.

Using P

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

$$P(H_n = H^i, \mathbf{a}_n \mid \epsilon, H_{n-1} = H^j, \mathbf{a}_n$$

Therefore, for $n > 1$, the probability update equation in the case of asymmetric transition rates, equation 4.1, is given by

$$\mathrm{P}_n(H_n \quad H^i, \boldsymbol{a}_n) \\ \frac{\mathrm{P}(\xi_{1:n-1})}{\mathrm{P}(\xi_{1:n})} f^i(\xi_n) \sum_{j=1}^{N} {}^{ij}_{n-1}(\boldsymbol{a}_n \quad \boldsymbol{\delta}^{ij}) \mathrm{P}_{n-1}\left(H_{n-1} \quad H^j, \boldsymbol{a}_n \quad \boldsymbol{\delta}^{ij}\right). \quad (4.6)$$

The point estimates of the transition rates, ${}^{ij}_{n-1}(\boldsymbol{a}_{n-1} \quad \boldsymbol{a}_n \quad \boldsymbol{\delta}^{ij})$, are defined in equation 3.5. As before, $\mathrm{P}_1(H^i, \boldsymbol{a}_1 \quad 0) \quad f^i(\xi_1)\mathrm{P}_0(H^i)/\mathrm{P}(\xi_1)$ and $\mathrm{P}_1(H^i, \boldsymbol{a}_1) \quad 0$ for any $\boldsymbol{a}_1 \quad 0$. At future times, it is only possible to obtain

Equation 4.6 is easier to interpret when $N = 2$. Using equation 4.5, we find

$$\pi^{21}_{n-1}(\boldsymbol{a}_{n-1}) = \frac{a^{21}_{n-1} - 1}{2 - a^{21}_{n-1} - a^{11}_{n-1}}, \quad \pi^{12}_{n-1}(\boldsymbol{a}_{n-1}) = \frac{a^{12}_{n-1} - 1}{2 - a^{12}_{n-1} - a^{22}_{n-1}},$$

and we can express $\pi^{11}_{n-1}(\boldsymbol{a}_{n-1}) = 1 - \pi^{21}_{n-1}(\boldsymbol{a}_{n-1})$ and $\pi^{22}_{n-1}(\boldsymbol{a}_{n-1}) = 1 - \pi^{12}_{n-1}(\boldsymbol{a}_{n-1})$. Expanding the sum in equation 4.6, we have

$$\begin{aligned} \mathrm{P}_n\left(H^1, \boldsymbol{a}_n\right) &= \frac{\mathrm{P}(\xi_{1:n-1})}{\mathrm{P}(\xi_{1:n})} f^1(\xi_n)\left[\pi^{11}_{n-1}(\boldsymbol{a}_n - \delta^{11})\mathrm{P}_{n-1}\left(H^1, \boldsymbol{a}_n - \delta^{11}\right)\right. \\ &\quad \left. \pi^{12}_{n-1}(\boldsymbol{a}_n - \delta^{12})\mathrm{P}_{n-1}\left(H^2, \boldsymbol{a}_n - \delta^{12}\right)\right], \end{aligned} \tag{4.8a}$$

$$\begin{aligned} \mathrm{P}_n\left(H^2, \boldsymbol{a}_n\right) &= \frac{\mathrm{P}(\xi_{1:n-1})}{\mathrm{P}(\xi_{1:n})} f^2(\xi_n)\left[\pi^{22}_{n-1}(\boldsymbol{a}_n - \delta^{22})\mathrm{P}_{n-1}\left(H^2, \boldsymbol{a}_n - \delta^{22}\right)\right. \\ &\quad \left. \pi^{21}_{n-1}(\boldsymbol{a}_n - \delta^{21})\mathrm{P}_{n-1}\left(H^1, \boldsymbol{a}_n - \delta^{21}\right)\right]. \end{aligned} \tag{4.8b}$$

The boundary and initial conditions will be given as above, and the mean inferred transition matrix is given by equation 4.7. Importantly, the inference process described by equations 4.8a and 4.8b allows for both asymmetric change-point matrices, $\boldsymbol{a}_n$,

Figure 5: Schematic showing the reduction of the full inference model, equation 5.7, for a two-state (H) symmetric environment ( ) carried out in sections 5 and 6. (A) Observations ξ_t arrive continuously in time and are used to update the probabilities P_t *(a)* that the environment is in state H after a change points. (B) Red and pink arrows from panels A to B represent, respectively, the summation and averaging of P_t *(a)* over a to obtain equation 5.1b for the zeroth P_t and first A_t moments in section 5.2. Arrows from P_t *(a)* have been omitted for clarity. (C) Moment equations are converted to a neural population model, equation 6.7, by assigning the probabilities to population variables, $P_t \rightarrow u_t$

We can substitute equation 5.1 into equation 3.7 describing the probability of transitions between time t_{n-1}

Finally, note that we can obtain evolution equations for the likelihoods, $P_t(a) \quad P(H_t \quad H, a)$, by applying the change of variables $P_t(a) \quad e^{x_t(a)}$. Itô's change of coordinates rules (Gardiner, 2004) implies that equation 5.6 is equivalent to

$$dP_t(a) \quad P_t(a)\left[\left(g_t \quad \frac{1}{2}\right)dt \quad dW_t\right] \left[\frac{a \quad \alpha}{t \quad \beta}P_t(a \quad 1) \quad \frac{a \quad \alpha}{t \quad \beta}P_t(a)\right]dt, \tag{5.7}$$

where now initial conditions at $t \quad 0$ are simply $P_0(a) \quad P_0(H, a) \quad P_0(H)P_0(a)$

matrix $\boldsymbol{a}$, where we replace the non-change-point counts with dwell times t^j, defined as $a_t^{jj} \quad t^j : \quad a^{jj}{}_t \quad t$ for $t \quad t_n \cdot \quad t_{n\ 1}$. This is necessary due to the divergence of $a^{ii}{}_t$ as $\quad t \to 0$. In the limit $\quad t \to 0$, the modified change-point matrix becomes

$$\boldsymbol{A}_t \quad \begin{pmatrix} t^1 & a^{21} \\ a^{12} & t^2 \end{pmatrix},$$

where $a_t^{ij} \quad \mathbb{Z}\cdot$ is the change-point count from $H^j \to H^i$, while $t^i \quad \mathbb{R}_{\ 0}$ is the dwell time in state H^i. Thus, taking logarithms, linearizing, and taking the limit $\quad t \to 0$, we obtain the following system of stochastic partial differential equations (SPDEs) for the log likelihoods, $x_t^j(\boldsymbol{A}) \quad \ln \mathrm{P}_n(H \ , \boldsymbol{A})$:

$$\mathrm{d}x_t^1(\boldsymbol{A}_t) \quad g_t^1 \mathrm{d}t \quad \mathrm{d}W_t^1 \left(\frac{a_t^{12} \quad \alpha_2 \cdot 1}{t^2 \quad \beta_2} \mathrm{e}^{x_t^2(\boldsymbol{A}_t \quad \delta^{12})} \ x_t^1(\boldsymbol{A}_t) \cdot \frac{a_t^{21} \quad \alpha_1}{t^1 \quad \beta_1} \cdot \frac{\partial x_t^1}{\partial t^1} \right) \mathrm{d}t, \quad (5.9a)$$

$$\mathrm{d}x_t^2(\boldsymbol{A}_t) \quad g_t^2 \mathrm{d}t \quad \mathrm{d}W_t^2 \left(\frac{a_t^{21} \quad \alpha_1 \cdot 1}{t^1 \quad \beta_1} \mathrm{e}^{x_t^1(\boldsymbol{A}_t \quad \delta^{21})} \ x_t^2(\boldsymbol{A}_t) \cdot \frac{a_t^{12} \quad \alpha_2}{t^2 \quad \beta_2} \cdot \frac{\partial x_t^2}{\partial t^2} \right) \mathrm{d}t, \quad (5.9b)$$

where the drift, g_t^i, and noise, W_t^i, are defined as before (for details, see section A.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

the change of variables $\mathrm{P}_t^i(\boldsymbol{A}_t)$ $\mathrm{e}^{x_t^i(\boldsymbol{A}_t)}$, we find

$$\mathrm{d}\mathrm{P}_t^1(\boldsymbol{A}_t) \quad \mathrm{P}_t^1(\boldsymbol{A}_t)\left[\left(g_t^1 \quad \frac{1}{2}\right)\mathrm{d}t \quad \mathrm{d}W_t^1\right]$$
$$\left[\frac{a_t^{12} \quad \alpha_2}{t^2 \quad \beta_2}\mathrm{P}_t^2(\boldsymbol{A}_t \quad \delta^{12}) \quad \frac{a_t^{21} \quad \alpha_1}{t^1 \quad \beta_1}\mathrm{P}_t^1(\boldsymbol{A}_t) \quad \frac{\partial \mathrm{P}_t^1(\boldsymbol{A}_t)}{\partial t^1}\right]\mathrm{d}t, \tag{5.11a}$$

$$\mathrm{d}\mathrm{P}_t^2(\boldsymbol{A}_t) \quad \mathrm{P}_t^2(\boldsymbol{A}_t)\left[\left(g_t^2 \quad \frac{1}{2}\right)\mathrm{d}t \quad \mathrm{d}W_t^2\right]$$
$$\left[\frac{a_t^{21} \quad \alpha_1}{t^1 \quad \beta_1}\mathrm{P}_t^1(\boldsymbol{A}_t \quad \delta^{21}) \quad \frac{a_t^{12} \quad \alpha_2}{t^2 \quad \beta_2}\mathrm{P}_t^2(\boldsymbol{A}_t) \quad \frac{\partial \mathrm{P}_t^2(\boldsymbol{A}_t)}{\partial t^2}\right]\mathrm{d}t, \tag{5.11b}$$

where now initial conditions at t 0 are $\mathrm{P}_0^i(\boldsymbol{A}_t)$ $\mathrm{P}_0(H^i)\mathrm{P}_0(\boldsymbol{A}_t)$.

5.1.3 Multiple States with Symmetric Rates. The continuum limit in the case of N states, $H^1, \ldots, H^N$, with symmetric transition rates can be ded8BT0 TL / F3 1

We begin by summing equation 5.7 over all

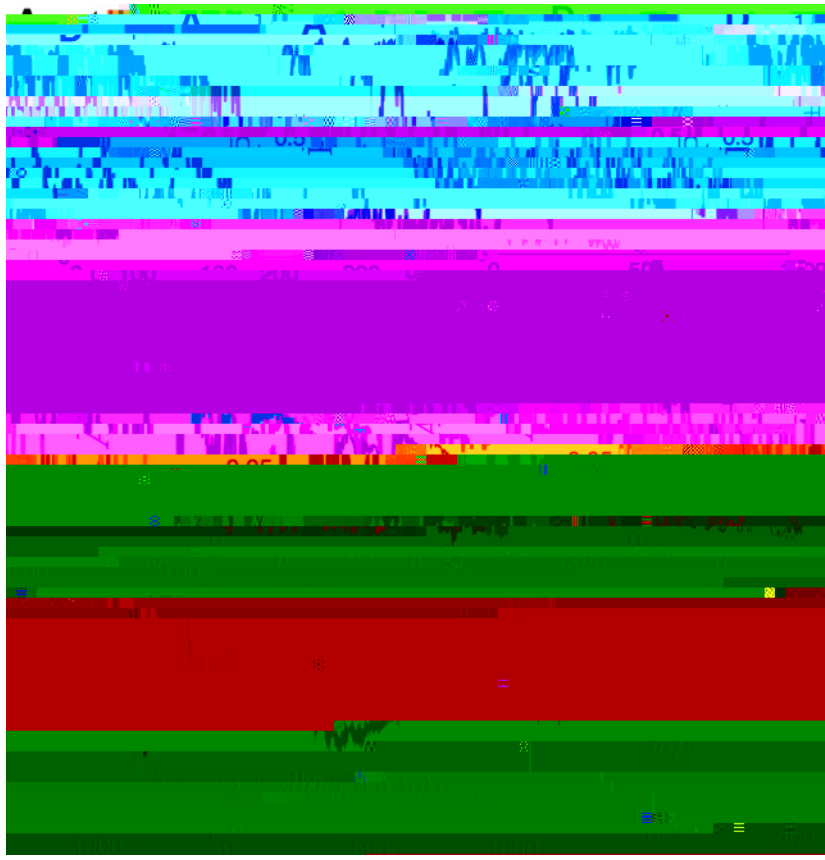

Figure 6: The dynamics of the first two moments as approximated by equa-

We briefly analyze the model, equation 6.1, by considering the limit of no observation noise. That is, we assume g

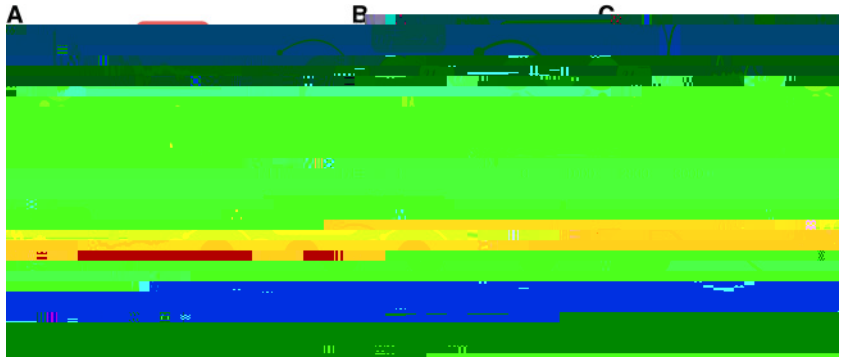

Figure 7: Neural network model with plasticity, inferring the current state H_t and rates of environmental change. (A) Schematic showing the synaptic weight w from neural population $u \to u$ evolving through long-term potentiation (LTP) and long-term depression (LTD) to match the environment's rate of change, : . (B) When the neural populations exchange dominance, their activity levels u are both transiently high. As a result, both synaptic weights, w , increase via LTP. When only one population is active, both weights decay via LTD, as described by equation 6.7b. (C) Inference of the rate, , via long-term plasticity of the weights for 0.01, 0.05, 0.1. Though the signal-to-noise

our model the rate variables, u_t , represent the probability that the environment is in state H . This particular form of the population model leads dynamical equations that are consistent with an accepted rate-correlation-based plasticity rule (Miller, 1994; Pfister & Gerstner, 2006). Using log probabilities would lead to models that contain exponential functions of the rate (Veliz-Cuba et al., 2016), which are less common. In addition, since probabilities can assume a finite range of values, we required that u_t [0, 1]. Using log probabilities would require that we use a semi-infinite range, (, 0] or that we truncate. Note also that the inputs I_t and noise dW_t are gain-modulated using the population rates u_t . Gain-modulating circuits have been identified in many sensory areas (Salinas & Abbott, 1996), and recent studies suggest evidence-accumulating circuits may also modulate input

modifying equation 6.2 for the plasticity agent, so that each C_t decays only when the neural population of origin, u_t , is active. Thus, we obtain the pair of equations

$$\mathrm{d}C_t \quad . \quad H(u_t \; . \; \theta)\left[C_t\right]^2 .$$

Expressing equation 6.8 as a system of equations for the synaptic weights, w_t , yields

$$\mathrm{d}w_t \quad H(u_{t\ \tau} \; . \; \theta)\left[\delta(u_t \; . \; u_t) \; . \; w_t\right]_{|} \, C_t \; \mathrm{d}t. \tag{6.9}$$

Here the function $H(u\ (t \; . \; \tau) \; . \; \theta)$ for $\theta \quad 0$.

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 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 current state, while knowledge of transitions between states is 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 change-point 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 change-point 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

strategies to learn or something close to the normative models we derived here.

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

$$\mathrm{P}(H_n, a_n \backslash H_{n-1}, a_{n-1})$$

$$\begin{cases} 1 - t\dfrac{a_n - \alpha}{t_{n-1} - \beta}, & \& : H_n = H_{n-1} \,\&\, a_n = a_{n-1} \\ t\dfrac{a_n - \alpha - 1}{(} & \end{cases}$$

Furthermore, note that in the limit $t \to$, the $\mathcal{O}(1)$ terms and $\mathcal{O}(t^{-1})$ terms vanish in equation 5.16:

$$\mathrm{d}\overline{P}_t \quad \overline{P}_t \left[\left(g_t \quad \frac{1}{2} \right) \mathrm{d}t \quad \mathrm{d}W \right] \quad \left[A_t^* \quad A_t \right] \mathrm{d}t \tag{A.4a}$$

$$\mathrm{d}A_t \quad A_t \left[\left(g_t \quad \frac{1}{2} \right) \mathrm{d}t \quad \mathrm{d}W \right] \quad \left(A_t^* \quad A_t \right) \left(A_t^* \quad A_t \right) \mathrm{d}t. \tag{A.4b}$$

Therefore, in the event that $A_t \to P_t$ in the long time limit ($t \to$), we find the truncated system, equation A.4, becomes

$$\mathrm{d}P_t \quad P_t \left[\left(g_t \quad \frac{1}{2} \right) \mathrm{d}t \quad \mathrm{d}W \right] \quad ! \left[P_t \quad P_t \right] \mathrm{d}t \tag{A.5a}$$

$$\mathrm{d}P_t \quad P_t \left[\left(g_t \quad \frac{1}{2} \right) \mathrm{d}t \quad \mathrm{d}W \right] \quad ^2 ! \left(P_t \quad P_t \right) \left(P_t \quad P_t \right) \mathrm{d}t. \tag{A.5b}$$

Dividing by and noting that $P_t \quad P_t \quad 1$, equation A.5 becomes

$$\mathrm{d}P_t \quad P_t \left[\left(g_t \quad \frac{1}{2} \right) \mathrm{d}t \quad \mathrm{d}W \right] \quad ! \left[P_t \quad P_t \right] \mathrm{d}t,$$

rules (Gardiner, 2004) imply our population model is equivalent to

$$\mathrm{d}x_t \quad I$$

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, C. K., Romo, R., & Brody, C. D. (2005). Flexible control of mutual inhibition: A neural model of two-interval discrimination. *Science*, *307*(5712), 1121–1124.

McGuire, J. T., Nassar, M. R., Gold, J. I., & Kable, J. W. (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, K. D. (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, S. J., Leong, Y. C., Radulescu, A., & Wilson, R. C. (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, J. M., Heilbronner, S. R., Barack, D. L., Hayden, B. Y., & Platt, M. L. (2011). Posterior cingulate cortex: Adapting behavior to a changing world. *Trends in Cognitive Sciences*, *15*(4), 143–151.

Petit, O., Gautrais, J., Leca, J. B., Theraulaz, G., & Deneubourg, J. L. (2009). Collective decision-making in white-faced capuchin monkeys. *Proceedings of the Royal Society of London B: Biological Sciences*, *276*(1672), 3495–3503.

Pfister, J. P., & 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: Cambridge University Press.

Salinas, E., & Abbott, L. (1996). A model of multiplicative neural responses in parietal cortex. *Proceedings of the National Academy of Sciences*, *93*(21), 11,956–11,961.

Shvartsman, M., Srivastava, V., & Cohen, J. D. (2015). A theory of decision making under dynamic context. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, & R. Garnett (Eds.), *Advances in neural information processing systems*, *19* (pp. 2476–2484). Red Hook, NY: Curran.

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

Socha, L. (2007). *Linearization methods for stochastic dynamic systems*. New York: Springer Science & Business Media.

Sugrue, L. P., Corrado, G. S., & Newsome, W. T. (2004). Matching behavior and the representation of value in the parietal cortex. *Science*, *304*(5678), 1782–1787.

Veliz-Cuba, A., Kilpatrick, Z. P., & Josić, 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. *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)*, *19*, 268–281.

Wilson, R. C., Nassar, M. R., & Gold, J. I. (2010). Bayesian online learning of the hazard rate in change-point problems. *Neural Comput.*, *22*(9), 2452–2476.

Wilson, R. C., & Niv, Y. (2011). Inferring relevance in a changing world. *Frontiers in Human Neuroscience*, *5*, 189.

Wong, K. F., Huk, A. C., Shadlen, M. N., & Wang, X. J. (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

Wyart, V., De Gardelle, V., Scholl, J., & Summerfield, C. (2012). Rhythmic fluctuations