

Stochastic models of evidence accumulation in changing environments

Alan Veliz-Cuba,* Zachary P. Kilpatrick[✉], and Krešimir Josić[✉]
Department of Mathematics, University of Houston, Houston, TX 77204
 (Dated: May 15, 2015)

Organisms and ecological groups accumulate evidence to make decisions. Classic experiments and theoretical studies have explored decisions between alternatives when the correct choice is fixed during a trial. However, the natural world constantly changes. Using sequential analysis we derive equations for the likelihood of different choices, when the correct option changes in time. Our analysis shows that ideal observers discount prior evidence at a rate determined by the volatility of the environment. Nondimensionalization shows that the dynamics of evidence accumulation is governed solely by the information gained over an average environmental epoch. A plausible neural implementation of an optimal observer in a changing environment shows that, in contrast with previous models, neural populations representing alternate choices are coupled through excitation.

Introduction. Mammals [1–4], insects [5, 6], and single cells [7] use incoming evidence to make choices. Perception is noisy, but it is possible to determine optimal policies for integrating uncertain information to make decisions [8–10]. Remarkably, such policies are consistent with the behavior of many animals faced with uncertain choices [3, 4, 11]. Stochastic accumulator models provide a plausible neural implementation of such decisions [12, 13]. These models are analytically tractable [2], and can implement optimal decision strategies [14]. However, a key assumption of most models is that the correct choice is fixed in time, *i.e.* decisions are made in a static environment. This assumption may hold in the laboratory, but natural environments are seldom static [15, 16].

Here we show that optimal stochastic accumulator models can be extended to a changing environment. These extensions reveal how to optimally discount old information, and provide explicit limits on the certainty that can be attained when the underlying truth changes. Moreover, they suggest a biophysical neural implementation for evidence integrators that differs considerably from those in a static environment.

We develop our model in a way that parallels the case of a static environment with two possible states, H_+ and H_- . The problem is to optimally integrate a stream of observations (measurements) to infer the present environmental state. In the static case, this can be done using sequential analysis [1, 9]: An observer makes a stream of independent, noisy measurements, $\xi_{1:n} = (\xi_1, \xi_2, \dots, \xi_n)$, at equally spaced times, $t_{1:n} = (t_1, t_2, \dots, t_n)$. Their distributions, $f_+(\xi_j) := \Pr(\xi_j|H_+)$, and $f_-(\xi_j) := \Pr(\xi_j|H_-)$, depend on the environmental state. Combined with the prior probability, $\Pr(H_\pm)$, of the states, this gives the *likelihood ratio*,

$$R_n = \frac{\Pr(H_+|\xi_{1:n})}{\Pr(H_-|\xi_{1:n})} = \frac{f_+(\xi_1)f_+(\xi_2)\cdots f_+(\xi_n)\Pr(H_+)}{f_-(\xi_1)f_-(\xi_2)\cdots f_-(\xi_n)\Pr(H_-)},$$

which can also be written recursively [9]:

$$R_n = \left(\frac{f_+(\xi_n)}{f_-(\xi_n)} \right) \cdot R_{n-1}, \quad \text{with} \quad R_0 = \frac{\Pr(H_+)}{\Pr(H_-)}. \quad (1)$$

With a fixed number of observations, the likelihood ratio can be used to make a choice that minimizes the total error rate [8], or maximizes reward [10]. Eq. (1) gives a recursive relation for the log likelihood ratio, $y_n = \ln R_n$,

$$y_n = y_{n-1} + \ln \frac{f_+(\xi_n)}{f_-(\xi_n)}. \quad (2)$$

When the time between observations, $\Delta t = t_j - t_{j-1}$, is small, we can approximate this stochastic process by the stochastic differential equation (SDE) [17],

$$dy = g_\pm dt + \rho_\pm dW_t, \quad (3)$$

where W_t is a Wiener process, and the constants $g_\pm = \frac{1}{\Delta t} E_\xi [\ln \frac{f_+(\xi)}{f_-(\xi)} | H_\pm]$ and $\rho_\pm^2 = \frac{1}{\Delta t} \text{Var}_\xi [\ln \frac{f_+(\xi)}{f_-(\xi)} | H_\pm]$ depend on the environmental state. Below we approximate other discrete time process, like Eq. (2), with SDEs. For details, see the Supplementary Material [18].

In state H_+ we have $g_+ \Delta t = \int_{-\infty}^{\infty} f_+(\xi) \ln \frac{f_+(\xi)}{f_-(\xi)} d\xi$. Thus the drift between two observations equals the Kullback–Leibler divergence between f_+ and f_- , *i.e.* the strength of the observed evidence from a measurement in favor of H_+ . Hence g_+ and g_- are the rates at which an optimal observer accumulates information.

Two alternatives in a changing environment. We assume that at time t the state, $H(t)$, of a changing environment is either H_+ or H_- . An observer infers the present state from a sequence of observations, $\xi_{1:n}$, made at equally spaced times, $t_{1:n}$, and characterized by distributions $f_\pm(\xi_n) := \Pr(\xi_n|H_\pm)$. The state of the environment changes between observations with probability $\epsilon_\pm \Delta t := \Pr(H(t_n) = H_\mp | H(t_{n-1}) = H_\pm)$ known to the observer. The likelihoods, $L_{n,\pm} = \Pr(H(t_n) = H_\pm | \xi_{1:n})$, then satisfy [18]

$$L_{n,\pm} \propto f_\pm(\xi_n) ((1 - \Delta t \epsilon_\pm) L_{n-1,\pm} + \Delta t \epsilon_\mp L_{n-1,\mp}), \quad (4)$$

* alanavc@math.uh.edu
[✉]equal contribution

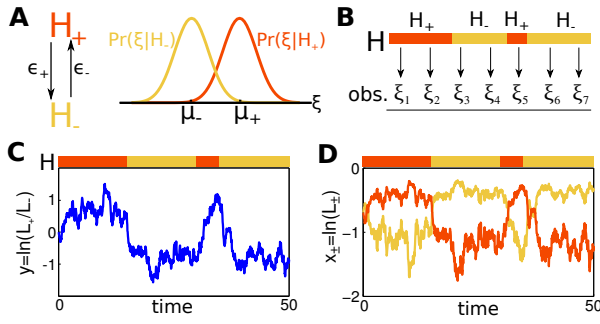


FIG. 1. Evidence accumulation in a dynamic environment. (A) The environmental state transitions from state H_+ to H_- and back with rates ϵ_+ and ϵ_- , respectively. Observations follow state dependent distributions $f_{\pm}(\xi) = \Pr(\xi|H_{\pm})$. (B) The distributions of the measurements, ξ_j , change with the environmental state. (C,D) The evolution of the log likelihood ratio, y , (panel C) and the log likelihoods x_{\pm} (panel D). At time t , evidence favors the environmental state H_+ if $y(t) > 0$, or, equivalently, $x_+(t) > x_-(t)$.

with proportionality constant $\Pr(\xi_{1:n-1})/\Pr(\xi_{1:n})$. The ratio of likelihoods of the two environmental states at time t_n , can be determined recursively as

$$R_n = \frac{L_{n,+}}{L_{n,-}} = \frac{f_+(\xi_n)}{f_-(\xi_n)} \frac{(1 - \Delta t \epsilon_+) R_{n-1} + \Delta t \epsilon_-}{\Delta t \epsilon_+ R_{n-1} + 1 - \Delta t \epsilon_-}. \quad (5)$$

This equation describes a variety of cases of evidence accumulation studied previously (Fig. S1 in [18]): If the environment is fixed ($\epsilon_{\pm} = 0$), we recover Eq. (1). If the environment starts in state H_- , changes to H_+ , but cannot change back ($\epsilon_- > 0, \epsilon_+ = 0$), we obtain

$$R_n = \frac{f_+(\xi_n)}{f_-(\xi_n)} \frac{R_{n-1} + \Delta t \epsilon_-}{1 - \Delta t \epsilon_-},$$

a model used in change point detection [19].

We can again approximate the stochastic process describing the evolution of the log likelihood ratio, $y_n = \ln R_n$, by an SDE [18]:

$$dy = [g(t) + \underbrace{\epsilon_-(e^{-y} + 1) - \epsilon_+(e^y + 1)}_{\text{nonlinearity}}]dt + \rho(t)dW_t, \quad (6)$$

where the drift $g(t) = \frac{1}{\Delta t} E_{\xi} \left[\ln \frac{f_+(\xi)}{f_-(\xi)} | H(t) \right]$, and variance $\rho^2(t) = \frac{1}{\Delta t} \text{Var}_{\xi} \left[\ln \frac{f_+(\xi)}{f_-(\xi)} | H(t) \right]$ are no longer constant, but depend on the state of the environment at time t .

The nonlinearity highlighted in Eq. (6) does not appear in Eq. (3), and serves to discount older evidence by a factor determined by environmental volatility. In previous work such discounting was modeled by a linear term [12, 20, 21], however our derivation shows that the resulting Ornstein-Uhlenbeck (OU) process does not model an optimal observer.

Equal switching rates. When $\epsilon := \epsilon_+ = \epsilon_-$, rates of switching between states are equal. Eq. (6) then becomes

$$dy = g(t)dt - 2\epsilon \sinh(y)dt + \rho(t)dW_t. \quad (7)$$

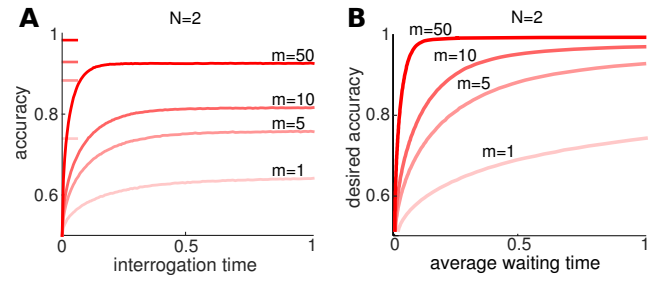


FIG. 2. Dependence of the probability of the correct response (accuracy) on normalized information gain, m , when $\epsilon_+ = \epsilon_-$. (A) Accuracy in an interrogation protocol increases with m and interrogation time, t , but saturates. Horizontal bars on left indicate the accuracy when the environment is in a single state for a long time, as in Eq. (10). (B) When the observer responds freely accuracy is similar, but saturates at 1. The increase in accuracy in time is exceedingly slow for low m . With multiple choices, $N > 2$, the accuracy behaves similarly (See Supplementary Material [18]).

Rescaling time using $\tau = \epsilon t$, we obtain

$$dy = [g(t)/\epsilon] d\tau - 2 \sinh(y) d\tau + [\rho(t)/\sqrt{\epsilon}] dW_t. \quad (8)$$

Since $g(t)$ is the rate of evidence accumulation, and ϵ^{-1} is the average time spent in each state, $[g(t)/\epsilon]$ can be interpreted as the *information gained over an average duration of an environmental state*.

When observations follow Gaussian distributions, $f_{\pm} \sim \mathcal{N}(\pm\mu, \sigma^2)$, then $g(t) = \pm 2\mu^2/\sigma^2$ and $\rho = 2\mu/\sigma$ and

$$dy = \text{sign}[g(t)] m d\tau - 2 \sinh(y) d\tau + \sqrt{2m} dW_t, \quad (9)$$

where $m = 2\mu^2/(\sigma^2\epsilon)$. The behavior of the system is completely determined by the single parameter m , the information gain over an average environmental epoch.

The probability of a correct response (accuracy) in both interrogation (Fig. 6A) and free response (Fig. 6B) protocols increase with m . When an optimal observer is interrogated about the state of the environment at time t , the answer is determined by the sign of y . The environment is changing, and observers discount old evidence at a rate increasing with $1/m$. Decisions are thus effectively based on a fixed amount of evidence, and accuracy saturates at a value smaller than 1 (Fig. 6A). If the environment remains in a single state for a long time, the log likelihood ratio, y , approaches a stationary distribution,

$$\mathcal{S}_{\pm}(y) = K e^{\pm y - \frac{2 \cosh(y)}{m}}, \quad H(t) = H_{\pm}, \quad (10)$$

where K is a normalization constant. In contrast, no stationary distribution exists when the environment is static ($\epsilon \equiv 0$). Since $\mathcal{S}_{\pm}(y)$ is obtained in the limit of many observations when the environment is trapped in a single state, $\int_0^{\infty} \mathcal{S}(y) dy$ provides an upper bound for accuracy (Fig. 6A). To achieve accuracy a in the free response protocol (Fig. 6B), we require $|y| \geq \ln \frac{a}{1-a}$ [2]. The waiting time for this accuracy steeply increases with a and decreases with m .

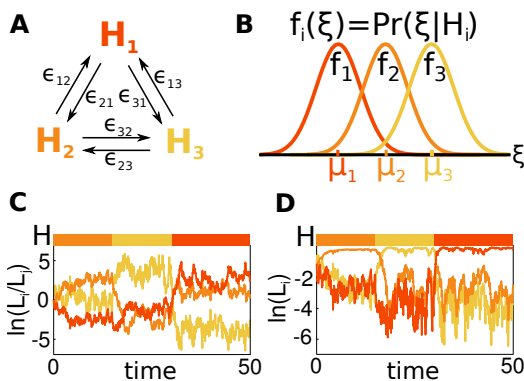


FIG. 3. Evidence accumulation with multiple choices in a changing environment. (A) The environment switches between N states (here $N = 3$). (B) Distributions $f_i(\xi) = \Pr(\xi|H_i)$ describing the probability of observation ξ in environmental state H_i . (C,D) Realization of the log likelihood ratios (panel C: $\ln(L_1/L_2)$, $\ln(L_2/L_3)$, $\ln(L_3/L_1)$) and log likelihoods (panel D).

We note that the hyperbolic sine term in Eq. (9) is superlinear. This means that, following an environmental switch, evidence is discounted more rapidly than in an OU model. We will see that similar observations hold in a changing environment with multiple states.

Likelihood update with multiple states. With multiple environmental states, H_i ($i = 1, \dots, N$), the optimal observer computes the present likelihood of each state from a sequence of measurements, $\xi_{1:n}$. Each measurement has distribution $f_i(\xi_n) := \Pr(\xi_n|H_i)$ dependent on the state [13, 22]. The environment switches from state j to i between two measurements with probability $\epsilon_{ij}\Delta t = \Pr(H(t_n) = H_i|H(t_{n-1}) = H_j)$ for $i \neq j$, and $\Pr(H(t_n) = H_i|H(t_{n-1}) = H_i) = 1 - \sum_{j \neq i} \Delta t \epsilon_{ji}$ (See Fig 3A).

We again use sequential analysis to obtain the likelihoods $L_{n,i} = \Pr(H(t) = H_i|\xi_{1:n})$ of the hypotheses H_i given n observations. The index that maximizes the likelihood, $\hat{i} = \text{argmax}_i L_{n,i}$, then gives the most likely state. Following the approach above, we obtain [18]:

$$L_{n,i} = \frac{\Pr(\xi_{1:n-1})}{\Pr(\xi_{1:n})} f_i(\xi_n) \times \left(\left(1 - \sum_{j \neq i} \Delta t \epsilon_{ji} \right) L_{n-1,i} + \sum_{j \neq i} \Delta t \epsilon_{ij} L_{n-1,j} \right). \quad (11)$$

Again after taking logarithms, $x_{n,i} = \ln L_{n,i}$, we can approximate the discrete stochastic process in Eq. (11), with an SDE [18],

$$dx = g(t)dt + \Lambda(t)d\mathbf{W}_t + K(x)dt, \quad (12)$$

where the drift has components $g_i(t) = \frac{1}{\Delta t} E_\xi [\ln f_i(\xi)|H(t)]$, $\Lambda(t)\Lambda(t)^T = \Sigma(t)$ with entries $\Sigma_{ij} = \frac{1}{\Delta t} \text{Cov}_\xi [\ln f_i(\xi), \ln f_j(\xi)|H(t)]$, components of \mathbf{W}_t are independent Wiener processes, and

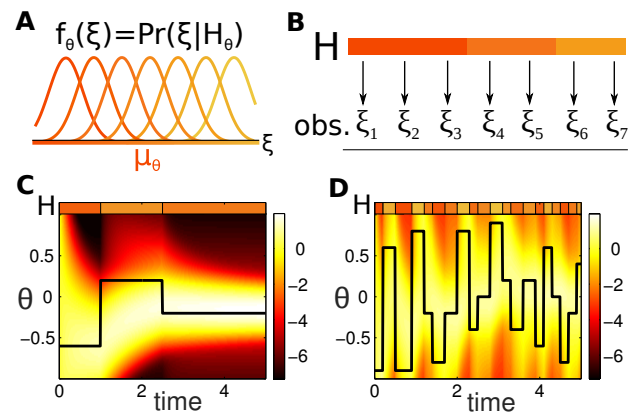


FIG. 4. Evidence accumulation with a continuum of choices. (A) The observer needs to infer the state of the environment, H_θ , where $\theta \in [-1, 1]$ from observations with distributions $f_\theta(\xi) = \Pr(\xi|H_\theta)$. (B) The environmental state changes stochastically at discrete times. (C) In slowly changing environments, the distribution of log likelihoods can nearly equilibrate prior to switches (solid line represents the true state of the environment at time t). (D) In quickly changing environments, the distribution remains flat, as limited evidence is accumulated between switches.

$K_i(x) = \sum_{j \neq i} (\epsilon_{ij} e^{x_j - x_i} - \epsilon_{ji})$. The drift g_i is maximized in environmental state H_i . To recover the $N = 2$ case, given in Eq. (6), we can exchange the numbers in Eq. (12) with \pm to obtain the log likelihood SDEs:

$$dx_\pm = [g_\pm(t) + (\epsilon_\mp e^{x_\mp - x_\pm} - \epsilon_\pm)]dt + dW_\pm, \quad (13)$$

$\langle W_i W_j \rangle = \Sigma_{ij} t$, and let $y = x_+ - x_-$. Analogous expressions for log likelihood ratios, $y_{ij} = \ln(L_i/L_j)$, can be derived [18]. The matrix of log likelihood ratios quantifies how much more likely one alternative is compared to others (e.g., Fig. 3C) [23].

A continuum of hypotheses. Lastly, we consider the case of a continuum of possible environmental states. This provides a tractable model for recent experiments with observers who infer the location of a hidden, intermittently moving target from noisy observations. Evidence suggests that humans update their beliefs quickly and near optimally when observations indicate that the target has moved [24].

Suppose the environmental state, $H(t)$, intermittently switches between a continuum of possible states, H_θ , where $\theta \in [a, b]$. An observer again computes the likelihood of each state from observations, $\xi_{1:n}$, with distributions $f_\theta(\xi_n) := \Pr(\xi_n|H_\theta)$. The environment switches from state θ' to state θ between observations with relative likelihood defined $\epsilon_{\theta\theta'} d\theta \Delta t := \Pr(H(t_n) = H_\theta|H(t_{n-1}) = H_{\theta'})$ for $\theta \neq \theta'$, and $\Pr(H(t_n) = H_\theta|H(t_{n-1}) = H_\theta) = 1 - \int_a^b \Delta t \epsilon_{\theta\theta'} d\theta'$ (See [18] for details). From Eq. (11) the expression for the likelihoods

$L_{n,\theta} = \Pr(H(t_n) = H_\theta | \xi_{1:n})$ [18] is:

$$L_{n,\theta} = \frac{\Pr(\xi_{1:n-1})}{\Pr(\xi_{1:n})} f_\theta(\xi_n) \times \left(\left(1 - \int_a^b \Delta t \epsilon_{\theta',\theta} d\theta' \right) L_{n-1,\theta} + \int_a^b \Delta t \epsilon_{\theta\theta'} L_{n-1,\theta'} d\theta' \right).$$

We again approximate the log likelihood, $\ln L_{n,\theta}$, by a temporally continuous process,

$$d\mathbf{x} = \mathbf{g}(t)dt + \Lambda(t)d\mathbf{W}_t + K(\mathbf{x})dt, \quad (14)$$

where $g_\theta(t) = \frac{1}{\Delta t} E_\xi [\ln f_\theta(\xi) | H(t)]$, the components of \mathbf{W}_t are independent Wiener processes, $K_\theta(\mathbf{x}) = \int_a^b (\epsilon_{\theta\theta'} e^{x_{\theta'} - x_\theta} - \epsilon_{\theta'\theta}) d\theta'$, and $\Lambda(t)\Lambda^*(t) = \Sigma(t)$ is the covariance function given by

$$\Sigma_{\theta\theta'}(t) = \frac{1}{\Delta t} \text{Cov}_\xi [\ln f_\theta(\xi), \ln f_{\theta'}(\xi) | H(t)]. \quad (15)$$

The drift $g_\theta(t)$ is maximal when θ agrees with the present environmental state. The most likely state, given observations up to time t , is $\hat{\theta} = \text{argmax}_\theta x_\theta(t)$.

Even when the environment is stationary for a long time, noise in the observations stochastically perturbs the likelihoods, $x_\theta(t)$, over the environmental states. In slowly changing environments, the likelihoods *nearly* equilibrate to a well peaked distribution between environmental switches (Fig. 4C). This does not occur in quickly changing environments (Fig. 4D). However, each log likelihood, $x_\theta(t)$, approaches a stationary distribution if the environmental state remains fixed for a long time. The term $K(\mathbf{x})$ in Eq. (14) provides for rapid departure from this quasi-stationary density when the environment changes, a mechanism proposed in [24].

Neural population model. Previous neural models of decision making typically relied on mutually inhibitory neural networks [12, 25, 26], with each population representing one alternative. In contrast, inference in dynamic environments with two states, H_+ and H_- , can be optimally performed by mutually excitatory neural populations with activities (firing rates) r_+ and r_- ,

$$dr_+ = [I_+(t) - \alpha r_+ + F_+(r_- - r_+)]dt + dW_+, \quad (16a)$$

$$dr_- = [I_-(t) - \alpha r_- + F_-(r_+ - r_-)]dt + dW_-, \quad (16b)$$

where the transfer functions are $F_\pm(x) = -\alpha x/2 + \epsilon_\pm e^x - \epsilon_\pm$, the external input $I_\pm(t) = I_\pm^0$ when $H(t) = H_\pm$ and vanishes otherwise, W_\pm are Wiener processes with covariance defined as in Eq. (13) [18]. When $\alpha > 0$ and sufficiently small, population activities are modulated by self-inhibition, and mutual excitation (Fig. 3A). Taking $y = r_+ - r_-$ reduces Eq. (49) to the SDE for the log likelihood ratio, Eq. (6). In the limit $\epsilon_\pm \rightarrow 0$, we obtain a linear integrator $dr_\pm = [I_\pm dt + dW_\pm] - \alpha(r_+ + r_-)dt/2$ [2, 26].

The model given by Eq. (49) is matched to the timescale of the environment determined by ϵ_\pm , solutions approach stationary distributions if input is constant. Due to mutual excitation, they are very sensitive

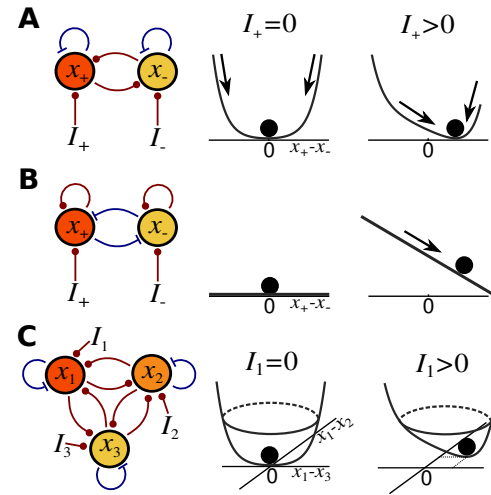


FIG. 5. Neural population models of evidence accumulation. (A) Two populations x_\pm receive a fluctuating stimulus with mean I_\pm ; they are mutually coupled by excitation (circles) and locally coupled by inhibition (flat ends) [18]. When $I_+ > 0$, the fixed point of the system has coordinates satisfying $x_+ > x_-$ as shown in the plots of the associated potentials. (B) Taking $\epsilon_\pm \rightarrow 0$ in Eq. (49) generates a mutually inhibitory network that perfectly integrates inputs I_\pm and has a flat potential function. (C) With $N = 3$ alternatives, three populations coupled by mutual excitation can still optimally integrate the inputs $I_{1,2,3}$, rapidly switching between the fixed point of the system in response to environmental changes.

to changes in inputs. These features are absent previous connectionist models [27]. Even when ϵ is small, Eq. (49) has a single attracting state determined by the mean inputs I_\pm^0 . We illustrate the response of the model to inputs using potentials (Fig. 5A). In contrast to the single attractor of Eq. (49), mutually inhibitory models can possess a neutrally stable line attractor that integrates inputs ($\epsilon_\pm \equiv 0$, Fig. 5B) [28]. In analogous systems for $N > 2$ dynamic states, coupling between populations is again excitatory (Fig. 5C) [18].

Discussion. We have derived a nonlinear stochastic model of optimal evidence accumulation in changing environments. Importantly, the resulting SDE is not an OU process, as suggested in previous heuristic models [12, 17, 21]. Rather, an exponential nonlinearity allows for optimal discounting of old evidence, and rapid adjustment of decision variables following environmental changes. Sequential sampling in dynamic environments with two states has been studied previously in special cases, such as adapting spiking model, capable of responding to environmental changes [29]. Likelihood update procedures have also been proposed for multiple alternative tasks in the limit $\epsilon_{ij} \rightarrow 0$ [23, 30]. Furthermore, Eq. (11) for the case $N = 2$ was derived in [31], but its dynamics were not analyzed. One important conclusion of our work is that $m = g/\epsilon$, the information gain over the characteristic environmental timescale, is the key parameter determining the model's dynamics and accuracy. It is easy to show that equivalent pa-

rameters govern the dynamics of likelihoods of multiple choices. We have thus presented tractable equations that illuminate how the belief of an optimal observer

evolves in a changing environment.

Acknowledgements. Funding was provided by NSF-DMS-1311755 (ZPK); NSF/NIGMS-R01GM104974 (AV-C and KJ); and NSF-DMS-1122094 (KJ).

-
- [1] J. R. Busemeyer and J. T. Townsend, Psychological review **100**, 432 (1993).
 - [2] R. Bogacz, E. Brown, J. Moehlis, P. Holmes, and J. D. Cohen, Psychol Rev **113**, 700 (2006).
 - [3] J. I. Gold and M. N. Shadlen, Annu Rev Neurosci **30**, 535 (2007).
 - [4] B. W. Brunton, M. M. Botvinick, and C. D. Brody, Science **340**, 95 (2013).
 - [5] T. D. Seeley, S. Camazine, and J. Sneyd, Behavioral Ecology and Sociobiology **28**, 277 (1991).
 - [6] I. D. Couzin, J. Krause, N. R. Franks, and S. A. Levin, Nature **433**, 513 (2005).
 - [7] G. Balázsi, A. van Oudenaarden, and J. J. Collins, Cell **144**, 910 (2011).
 - [8] J. Neyman and E. S. Pearson, *On the problem of the most efficient tests of statistical hypotheses* (Springer, 1992).
 - [9] A. Wald and J. Wolfowitz, The Annals of Mathematical Statistics **19**, 326 (1948).
 - [10] J. I. Gold and M. N. Shadlen, Neuron **36**, 299 (2002).
 - [11] R. Ratcliff and P. L. Smith, Psychological review **111**, 333 (2004).
 - [12] M. Usher and J. L. McClelland, Psychol Rev **108**, 550 (2001).
 - [13] J. M. Beck, W. J. Ma, R. Kiani, T. Hanks, A. K. Churchland, J. Roitman, M. N. Shadlen, P. E. Latham, and A. Pouget, Neuron **60**, 1142 (2008).
 - [14] R. Bogacz, M. Usher, J. Zhang, and J. L. McClelland, Philos Trans R Soc Lond B Biol Sci **362**, 1655 (2007).
 - [15] P. Domenici and R. W. Blake, J Exp Biol **200**, 1165 (1997).
 - [16] K. M. Passino, T. D. Seeley, and P. K. Visscher, Behavioral Ecology and Sociobiology **62**, 401 (2008).
 - [17] R. Ratcliff, Psychological review **85**, 59 (1978).
 - [18] See Supplementary Material for details of the continuum limit analysis and neural population models.
 - [19] J. Y. Angela, in *Advances in neural information processing systems* (2006), pp. 1545–1552.
 - [20] L. P. Sugrue, G. S. Corrado, and W. T. Newsome, science **304**, 1782 (2004).
 - [21] P. L. Smith and R. Ratcliff, Trends Neurosci **27**, 161 (2004).
 - [22] R. Bogacz and K. Gurney, Neural computation **19**, 442 (2007).
 - [23] V. Draglia, A. G. Tartakovsky, and V. V. Veeravalli, Information Theory, IEEE Transactions on **45**, 2448 (1999).
 - [24] J. T. McGuire, M. R. Nassar, J. I. Gold, and J. W. Kable, Neuron **84**, 870 (2014).
 - [25] X.-J. Wang, Neuron **36**, 955 (2002).
 - [26] T. McMillen and P. Holmes, Journal of Mathematical Psychology **50**, 30 (2006).
 - [27] K.-F. Wong, A. C. Huk, M. N. Shadlen, and X.-J. Wang, Front Comput Neurosci **1**, 6 (2007).
 - [28] C. K. Machens, R. Romo, and C. D. Brody, Science **307**, 1121 (2005).
 - [29] S. Deneve, Neural Comput **20**, 91 (2008).
 - [30] P. Armitage, Journal of the Royal Statistical Society. Series B (Methodological) **12**, 137 (1950).
 - [31] Q. Zhao and J. Ye, Signal Processing, IEEE Transactions on **58**, 5994 (2010).
 - [32] S. Bitzer, H. Park, F. Blankenburg, and S. J. Kiebel, Frontiers in human neuroscience **8** (2014).
 - [33] P. Billingsley, *Probability and Measure* (Wiley, 1995).
 - [34] M. Kac, *Probability and Related Topics in Physical Sciences* (American Mathematical Society, 1959).

Supplementary Material

Here we present the derivations for the likelihood update formulas and their approximations discussed in the main text. We begin by deriving the update expression for the likelihood ratio, R_n , in the case of two alternatives in a changing environment. The result is a nonlinear recursive equation. Subsequently, we show how to approximate the log likelihood ratio, $y_n = \ln R_n$, using a stochastic differential equation (SDE). To make the relation between the two precise, it is key to view the discrete equation for y_n as a family of equations parameterized by the time interval, Δt , over which each observation, ξ_n , is made [2]. Furthermore, we extend our derivations to multiple ($N > 2$) alternatives, and show that the log likelihood updates can be approximated by a nonlinear system of stochastic differential equations in the continuum limit. We emphasize that, with the appropriate scaling of the probabilities, $f_i(\xi) = \Pr(\xi|H_i)$, associated with each alternative H_i , there is a precise correspondence between the discretized version of the continuum limits and the log likelihood updates. Lastly, we present a derivation for the stochastic integrodifferential equation that represents the log likelihood for a continuum of possible environmental states, $\theta \in [a, b]$.

Note that throughout the supplementary material, we use notation involving a subscript Δt . This helps us define a family of stochastic processes indexed by the spacing between observations $\Delta t = t_n - t_{n-1}$. For instance, $f_{\Delta t, \pm}(\xi)$ represents the probability of an observation, ξ , in environmental state H_{\pm} (or, in the language of statistics, when hypothesis H_{\pm} holds). This probability changes with the timestep Δt . This approach allows us to properly take the continuum limit $\Delta t \rightarrow 0$. However, for simplicity we refrain from using this notation in the main text. Rather, we treat the limiting SDEs as approximations of discrete likelihood update processes. Also, we slightly abuse notation and write $f_i(\xi) = \Pr(\xi|H_i)$, even when ξ is a continuous random variable.

Likelihood ratio for two alternatives. We begin by deriving the recursive update equation for the likelihoods $L_{n, \pm} := \Pr(H(t_n) = H_{\pm} | \xi_{1:n})$ associated with each alternative H_{\pm} , where each observation (measurement), ξ_i , is made at time t_i . This is the probability that alternative H_{\pm} is true at time t_n , given that the series of observations $\xi_{1:n}$ has been made. Importantly, the underlying truth changes stochastically, and in a memoryless way, with transition probabilities given by $\epsilon_{\Delta t, \pm} := \Pr(H(t_n) = H_{\mp} | H(t_{n-1}) = H_{\pm})$, so that $\Pr(H(t_n) = H_{\pm} | H(t_{n-1}) = H_{\pm}) = 1 - \epsilon_{\Delta t, \pm}$. We begin by examining the likelihood $L_{n,+}$ associated with the alternative H_+ . Using Bayes' rule and the law of total probability we can relate the current likelihood $L_{n,+}$ to the conditional probabilities at the time of the previous observation, t_{n-1} :

$$\begin{aligned} L_{n,+} &= \frac{\Pr(H(t_n) = H_+)}{\Pr(\xi_{1:n})} \Pr(\xi_{1:n} | H(t_n) = H_+) \\ &= \frac{\Pr(H(t_n) = H_+)}{\Pr(\xi_{1:n})} \sum_{s \in \{+, -\}} \Pr(\xi_{1:n} | H(t_n) = H_+; H(t_{n-1}) = H_s) \Pr(H(t_{n-1}) = H_s | H(t_n) = H_+). \end{aligned}$$

To derive a recursive equation, with probabilities that are not conditioned on the state at t_n , we first use Bayes' rule again to write

$$\Pr(H(t_{n-1}) = H_+ | H(t_n) = H_+) = \frac{\Pr(H(t_n) = H_+ | H(t_{n-1}) = H_+) \Pr(H(t_n) = H_+)}{\Pr(H(t_{n-1}) = H_+)} = (1 - \epsilon_{\Delta t, +}) \frac{\Pr(H(t_{n-1}) = H_+)}{\Pr(H(t_n) = H_+)},$$

and

$$\Pr(H(t_{n-1}) = H_- | H(t_n) = H_+) = \frac{\Pr(H(t_n) = H_+ | H(t_{n-1}) = H_-) \Pr(H(t_{n-1}) = H_-)}{\Pr(H(t_n) = H_+)} = \epsilon_{\Delta t, -} \frac{\Pr(H(t_{n-1}) = H_-)}{\Pr(H(t_n) = H_+)}.$$

Plugging these formulas into our expression for $L_{n,+}$, we can then write

$$\begin{aligned} L_{n,+} &= \frac{1}{\Pr(\xi_{1:n})} \times \left((1 - \epsilon_{\Delta t, +}) \Pr(\xi_{1:n} | H(t_n) = H_+; H(t_{n-1}) = H_+) \Pr(H(t_{n-1}) = H_+) \right. \\ &\quad \left. + \epsilon_{\Delta t, -} \Pr(\xi_{1:n} | H(t_n) = H_+; H(t_{n-1}) = H_-) \Pr(H(t_{n-1}) = H_-) \right). \end{aligned}$$

The observation ξ_n is independent from the sequence of observations $\xi_{1:n-1}$ when conditioned on the states $H(t_n) =$

H_+ and $H(t_{n-1}) = H_{\pm}$, respectively. Thus, we obtain

$$\begin{aligned} L_{n,+} &= \frac{\Pr(\xi_n | H(t_n) = H_+)}{\Pr(\xi_{1:n})} \times \left((1 - \epsilon_{\Delta t,+}) \Pr(\xi_{1:n-1} | H(t_{n-1}) = H_+) \Pr(H(t_{n-1}) = H_+) \right. \\ &\quad \left. + \epsilon_{\Delta t,-} \Pr(\xi_{1:n-1} | H(t_{n-1}) = H_-) \Pr(H(t_{n-1}) = H_-) \right) \\ &= \frac{\Pr(\xi_n | H(t_n) = H_+)}{\Pr(\xi_{1:n})} \times \left((1 - \epsilon_{\Delta t,+}) \Pr(H(t_{n-1}) = H_+ | \xi_{1:n-1}) \Pr(\xi_{1:n-1}) \right. \\ &\quad \left. + \epsilon_{\Delta t,-} \Pr(H(t_{n-1}) = H_- | \xi_{1:n-1}) \Pr(\xi_{1:n-1}) \right) \\ &= \frac{\Pr(\xi_{1:n-1}) \Pr(\xi_n | H(t_n) = H_+)}{\Pr(\xi_{1:n})} \times \left((1 - \epsilon_{\Delta t,+}) \Pr(H(t_{n-1}) = H_+ | \xi_{1:n-1}) + \epsilon_{\Delta t,-} \Pr(H(t_{n-1}) = H_- | \xi_{1:n-1}) \right). \end{aligned}$$

Thus, by using our definition of the likelihoods $L_{n,\pm}$, we can write an update equation for $L_{n,+}$ in terms of the likelihoods $L_{n-1,\pm}$ at the previous time, t_{n-1} ,

$$L_{n,+} = \frac{\Pr(\xi_{1:n-1}) \Pr(\xi_n | H(t_n) = H_+)}{\Pr(\xi_{1:n})} ((1 - \epsilon_{\Delta t,+}) L_{n-1,+} + \epsilon_{\Delta t,-} L_{n-1,-}), \quad (17)$$

where $L_{0,+} = \Pr(H_+, t_0)$.

Similarly we obtain an update equation for the likelihood $L_{n,-}$ of the alternative H_- at time t_n :

$$L_{n,-} = \frac{\Pr(\xi_{1:n-1}) \Pr(\xi_n | H(t_n) = H_-)}{\Pr(\xi_{1:n})} (\epsilon_{\Delta t,+} L_{n-1,+} + (1 - \epsilon_{\Delta t,-}) L_{n-1,-}), \quad (18)$$

where $L_{0,-} = \Pr(H_-, t_0)$.

From Eqs. (17) and (18), the likelihood ratio $R_n = L_{n,+}/L_{n,-}$ is readily seen to satisfy the recursive equation

$$R_n = \frac{f_{\Delta t,+}(\xi_n)}{f_{\Delta t,-}(\xi_n)} \frac{(1 - \epsilon_{\Delta t,+}) R_{n-1} + \epsilon_{\Delta t,-}}{\epsilon_{\Delta t,+} R_{n-1} + 1 - \epsilon_{\Delta t,-}}, \quad (19)$$

where $f_{\Delta t,\pm}(\xi_n) = \Pr(\xi_n | H(t_n) = H_{\pm})$ is the distribution for each choice parameterized by the timestep $\Delta t = t_n - t_{n-1}$, and $R_0 = \frac{\Pr(H_+, t_0)}{\Pr(H_-, t_0)}$.

The continuum limit for the log likelihood ratio of two alternatives. In this section, we derive a continuum equation for the log likelihood ratio $y_n := \ln R_n$. We will proceed by first defining a family of stochastic difference equations for y_n , which are parameterized by the timestep $\Delta t = t_n - t_{n-1}$, between pairs of observations. By choosing an appropriate parameterization, we obtain a continuum limit that is a SDE. To begin, we divide both sides of Eq. (19) by R_{n-1} and take logarithms to yield

$$y_n - y_{n-1} = \ln \frac{f_{\Delta t,+}(\xi_n)}{f_{\Delta t,-}(\xi_n)} + \ln \frac{1 - \epsilon_{\Delta t,+} + \epsilon_{\Delta t,-} e^{-y_{n-1}}}{1 - \epsilon_{\Delta t,-} + \epsilon_{\Delta t,+} e^{y_{n-1}}}. \quad (20)$$

Following [2, 32], we assume that the time interval between individual observations, Δt , is small. Denote by $\Delta y_n = y_n - y_{n-1}$ the change in the log likelihood ratio due to the observation at time t_n . By assumption, the probability that the environment changes between two observations scales linearly with Δt up to higher order terms, so that $\epsilon_{\Delta t,\pm} := \Delta t \epsilon_{\pm} + o(\Delta t)$. Omitting higher order terms Δt , Eq. (20) can then be rewritten as

$$\Delta y_n = \ln \frac{f_{\Delta t,+}(\xi_n)}{f_{\Delta t,-}(\xi_n)} + \ln(1 + \Delta t(-\epsilon_+ + \epsilon_- e^{-y_{n-1}})) - \ln(1 + \Delta t(-\epsilon_- + \epsilon_+ e^{y_{n-1}})).$$

Since we assumed $\Delta t \ll 1$, we can use the approximation $\ln(1 + a) \approx a$ which is valid when $|a| \ll 1$. We also assume that the change in the log likelihood ratio, Δy_n , is small over the time interval Δt , so y_{n-1} can be replaced by y_n on the right-hand side of the equation. We obtain

$$\begin{aligned} \Delta y_n &\approx \ln \frac{f_{\Delta t,+}(\xi_n)}{f_{\Delta t,-}(\xi_n)} + \Delta t(\epsilon_- (e^{-y_n} + 1) - \epsilon_{\Delta t,+} (1 + e^{y_n})) \\ &= \mathbb{E}_{\xi} \left[\ln \frac{f_{\Delta t,+}(\xi_n)}{f_{\Delta t,-}(\xi_n)} \middle| H(t_n) \right] + \left(\ln \frac{f_{\Delta t,+}(\xi_n)}{f_{\Delta t,-}(\xi_n)} - \mathbb{E}_{\xi} \left[\ln \frac{f_{\Delta t,+}(\xi_n)}{f_{\Delta t,-}(\xi_n)} \middle| H(t_n) \right] \right) + \Delta t(\epsilon_- (e^{-y_n} + 1) - \epsilon_+ (1 + e^{y_n})). \end{aligned} \quad (21)$$

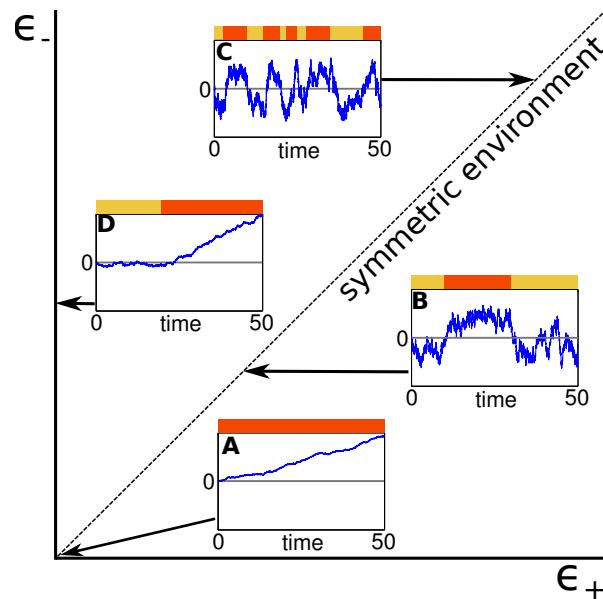


FIG. 6. In a dynamic environment, the dynamics of the log likelihood ratio, y , depends on the rates of switching between states. (A) When $\epsilon_{\pm} = 0$, the environment is static. (B) When the environment changes slowly, $|\epsilon_{\pm}| \ll 1$, the log likelihood ratio, y , can saturate. (C) In a rapidly changing environment, y tends not to equilibrate. (D) When $\epsilon_+ = 0$ and $\epsilon_- > 0$, the task becomes a change detection problem.

where we conditioned on the state of the environment, $H(t_n) = H_{\pm}$ at time t_n . Replacing the index n , with the time t , we can therefore write

$$\Delta y_t \approx \Delta t g_{\Delta t}(t) + \sqrt{\Delta t} \rho_{\Delta t}(t) \eta + \Delta t (\epsilon_- (e^{-y_t} + 1) - \epsilon_+ (1 + e^{y_t})), \quad (22)$$

where η is random variable with standard normal distribution, and

$$g_{\Delta t}(t) := \frac{1}{\Delta t} \mathbb{E}_{\xi} \left[\ln \frac{f_{\Delta t,+}(\xi)}{f_{\Delta t,-}(\xi)} \middle| H(t) \right] \quad \text{and} \quad \rho_{\Delta t}^2(t) := \frac{1}{\Delta t} \text{Var}_{\xi} \left[\ln \frac{f_{\Delta t,+}(\xi)}{f_{\Delta t,-}(\xi)} \middle| H(t) \right]. \quad (23)$$

Clearly, the drift $g_{\Delta t}$ and variance $\rho_{\Delta t}^2$ will diverge or vanish unless $f_{\Delta t,\pm}(\xi)$ are scaled appropriately in the $\Delta t \rightarrow 0$ limit. We discuss different ways of introducing such a scaling in the next section.

Assuming that we have well-defined limits $g(t) := \lim_{\Delta t \rightarrow 0} g_{\Delta t}(t)$ and $\rho^2(t) := \lim_{\Delta t \rightarrow 0} \rho_{\Delta t}^2(t)$, the discrete-time stochastic process, Eq. (22), approaches the stochastic differential equation

$$dy = g(t)dt + \rho(t)dW_t + (\epsilon_- (e^{-y} + 1) - \epsilon_+ (1 + e^y))dt, \quad (24)$$

where W_t is a standard Wiener process. This limit holds in the sense of distributions. Roughly, the smaller Δt is, the closer the distributions of the random variables y_n and $y(t_n)$ whose evolutions are described by Eq. (20), and Eq. (24), respectively. This correspondence can be made precise using the Donsker Invariance Principle [33].

In sum, Eq. (24), can be viewed as an approximation of the logarithm of the likelihood ratio whose evolution is given exactly by Eq. (19). For a fixed interval Δt , the parameters of the two equations are related via Eq. (23), and $\epsilon_{\Delta t,\pm}/\Delta t = \epsilon_{\pm}$. An illustration of the behaviors of the model, Eq. (24), is given in Fig. 6, showing how the volatility of the environment (given by ϵ_{\pm}) impacts the dynamics of evidence integration.

Precise correspondence. We now discuss two approaches in which the correspondence between Eqs. (20) and (24) can be made exact. We choose a specific scaling for the drift and variance arising from each observation, ξ_n . Suppose that over the time interval Δt , an observation, ξ_n , is a result of $r\Delta t$ separate observations – for example the measurement of the direction of $r\Delta t$ different dots [3]. In this case the estimate of the average of the individual measurements – *e.g.* the average of the velocities of dots in a display – will have both a mean and a variance that increase linearly with Δt .

As a concrete example we can compute $g(t)$ and $\rho(t)$ in SDE (24) when observations, ξ_n , follow normal distributions with mean and variance scaled by Δt ,

$$f_{\Delta t,\pm}(\xi) = \frac{1}{\sqrt{2\pi\Delta t\sigma^2}} e^{-(\xi - \Delta t\mu_{\pm})^2/(2\Delta t\sigma^2)}.$$

It is then straightforward to compute [2, 32],

$$g_{\Delta t} := \frac{1}{\Delta t} \mathbb{E}_{\xi} \left[\ln \frac{f_{\Delta t,+}(\xi)}{f_{\Delta t,-}(\xi)} \middle| H(t) \right] = \pm \frac{(\mu_+ - \mu_-)^2}{2\sigma^2} = g_{\pm}, \quad \rho_{\Delta t}^2(t) := \frac{1}{\Delta t} \text{Var}_{\xi} \left[\ln \frac{f_{\Delta t,+}(\xi)}{f_{\Delta t,-}(\xi)} \middle| H(t) \right] = \frac{(\mu_+ - \mu_-)^2}{\sigma^2} = \rho^2,$$

and note that $g(t) \in \{g_+, g_-\}$ is a telegraph process [34] with the probability masses $P(g_+, t)$ and $P(g_-, t)$ evolving according to the master equation $P_t(g_{\pm}, t) = \mp \epsilon_{\pm} P(g_{\pm}, t) \pm \epsilon_{\mp} P(g_{\mp}, t)$. In this case $\rho^2(t) = \rho^2$ remains constant.

More generally, we can obtain an identical result by considering that each observation made on a time interval consists of a number of sub-observations, each with statistics that scale with the length of the interval and the number of sub-observations. We define a family of stochastic processes parameterized by k , the number of sub-observations made in an interval of length Δt . As above, when $k = 1$, we assume that an observation ξ_n is the result of $r\Delta t$ separate observations. Assuming r is large, note that for $k > 1$ each of the k subobservations contain roughly $r_k = \lfloor r\Delta t/k \rfloor$ observations with mean and variance that scale linearly with $r_k \propto \Delta t/k$. We can achieve this by approximating $\ln \frac{f_{\Delta t,+}(\xi_n)}{f_{\Delta t,-}(\xi_n)}$ in Eq. (21) by the family of stochastic processes parameterized by k [2]

$$\sum_{l=1}^k \frac{\Delta t}{k} \ln \frac{f_+(\xi_l)}{f_-(\xi_l)} + \sum_{l=1}^k \frac{\sqrt{\Delta t}}{\sqrt{k}} \left(\ln \frac{f_+(\xi_l)}{f_-(\xi_l)} - \mathbb{E}_{\xi} \left[\ln \frac{f_+(\xi)}{f_-(\xi)} \middle| H(t) \right] \right).$$

The scaling in this approximation guarantees that $g(t) = \lim_{\Delta t \rightarrow 0} g_{\Delta t}(t) = \lim_{\Delta t \rightarrow 0} \frac{1}{\Delta t} \mathbb{E}_{\xi} \left[\ln \frac{f_{\Delta t,+}(\xi)}{f_{\Delta t,-}(\xi)} \middle| H(t) \right] = \mathbb{E}_{\xi} \left[\ln \frac{f_+(\xi)}{f_-(\xi)} \middle| H(t) \right]$ and $\rho^2(t) = \lim_{\Delta t \rightarrow 0} \rho_{\Delta t}^2 = \lim_{\Delta t \rightarrow 0} \frac{1}{\Delta t} \text{Var}_{\xi} \left[\ln \frac{f_{\Delta t,+}(\xi)}{f_{\Delta t,-}(\xi)} \middle| H(t) \right] = \text{Var}_{\xi} \left[\ln \frac{f_+(\xi)}{f_-(\xi)} \middle| H(t) \right]$. Furthermore, as $k \rightarrow \infty$, by the central limit theorem,

$$\Delta y_t \approx \sum_{l=1}^k \frac{\Delta t}{k} \ln \frac{f_+(\xi_l)}{f_-(\xi_l)} + \sum_{l=1}^k \frac{\sqrt{\Delta t}}{\sqrt{k}} \left(\ln \frac{f_+(\xi_l)}{f_-(\xi_l)} - \mathbb{E}_{\xi} \left[\ln \frac{f_+(\xi)}{f_-(\xi)} \middle| H(t) \right] \right) + \Delta t (\epsilon_- (e^{-y_t} + 1) - \epsilon_+ (1 + e^{y_t}))$$

converges in distribution to

$$\Delta y_t \approx \Delta t g(t) + \sqrt{\Delta t} \rho(t) \eta + \Delta t (\epsilon_- (e^{-y_t} + 1) - \epsilon_+ (1 + e^{y_t})),$$

where η is a standard normal random variable. Taking the limit $\Delta t \rightarrow 0$ yields Eq. (24).

Continuum limit for the log likelihood with multiple alternatives. We now describe the calculation of the continuum limit of the recursive system defining the evolution of the likelihoods $L_{n,i} = \Pr(H(t_n) = H_i | \xi_{1:n})$ of one among multiple alternatives (environmental states), H_i , $i = 1, \dots, N$. The state of the environment, and equivalently the correct choice at time t , again change stochastically. We assume that the transitions between the alternatives are memoryless, with transition rates $\epsilon_{\Delta t,ij} := \Pr(H(t_n) = H_i | H(t_{n-1}) = H_j)$. Using Bayes' rule and rearranging terms (analogous to the derivation of Eqs. (17) and (18)), we can express each likelihood $L_{n,i}$ in terms the likelihoods at the time of the previous observation, $L_{n-1,j}$,

$$L_{n,i} = \frac{\Pr(\xi_{1:n-1})}{\Pr(\xi_{1:n})} \Pr(\xi_n | H_i, t_n) \sum_{j=1}^N \epsilon_{\Delta t,ij} L_{n-1,j}.$$

Since we are only interested in comparing likelihoods, we can drop the common prefactor $\frac{\Pr(\xi_{1:n-1})}{\Pr(\xi_{1:n})}$, and use the fact that $\sum_{j=1}^N \epsilon_{\Delta t,ji} = 1$ (since $\epsilon_{\Delta t,ij}$ is a left stochastic matrix) to write $\epsilon_{\Delta t,ii} = 1 - \sum_{j \neq i} \epsilon_{\Delta t,ji}$ and obtain

$$L_{n,i} = f_{\Delta t,i}(\xi_n) \left(\left[1 - \sum_{j \neq i} \epsilon_{\Delta t,ji} \right] L_{n-1,i} + \sum_{j \neq i} \epsilon_{\Delta t,ij} L_{n-1,j} \right), \quad (25)$$

where $f_{\Delta t,i}(\xi_n) = \Pr(\xi_n | H_i, t_n)$. From Eq. (25), it follows that log of the rescaled likelihoods, $x_i := \ln L_i$, satisfies the recursive relation

$$x_{n,i} - x_{n-1,i} = \ln f_{\Delta t,i}(\xi_n) + \ln \left(1 - \sum_{j \neq i} \epsilon_{\Delta t,ji} + \sum_{j \neq i} \epsilon_{\Delta t,ij} e^{x_{n-1,j} - x_{n-1,i}} \right).$$

To derive an approximating SDE, we denote by $\Delta x_{n,i} = x_{n,i} - x_{n-1,i}$, the change in the log likelihood due to an observation at time t_n . As before, we assume $\epsilon_{\Delta t,ij} := \Delta t \epsilon_{ij} + o(\Delta t)$ for $i \neq j$, and drop the higher order terms, giving

$$\Delta x_{n,i} = \ln f_{\Delta t,i}(\xi_n) + \ln \left(1 - \sum_{j \neq i} \Delta t \epsilon_{ji} + \sum_{j \neq i} \Delta t \epsilon_{ij} e^{x_{n-1,j} - x_{n-1,i}} \right). \quad (26)$$

Assuming $\Delta t \ll 1$, we again use the approximation $\ln(1+a) \approx a$ which holds for $|a| \ll 1$. We also assume that the change in the log likelihood, $|\Delta x_{n,i}| \ll 1$, is small over the time interval Δt , so that

$$\begin{aligned} \Delta x_{n,i} &\approx \ln f_{\Delta t,i}(\xi_n) + \Delta t \sum_{j \neq i} (\epsilon_{ij} e^{x_{n,j} - x_{n,i}} - \epsilon_{ji}) \\ &= E_{\xi} [\ln f_{\Delta t,i}(\xi) | H(t_n)] + (\ln f_{\Delta t,i}(\xi_n) - E_{\xi} [\ln f_{\Delta t,i}(\xi) | H(t_n)]) + \Delta t \sum_{j \neq i} (\epsilon_{ij} e^{x_{n,j} - x_{n,i}} - \epsilon_{ji}), \end{aligned} \quad (27)$$

where we condition on the current state of the environment $H(t_n) \in \{H_1, \dots, H_N\}$.

Replacing the index n , by the time t , we can therefore write

$$\Delta x_{t,i} \approx \Delta t g_{\Delta t,i}(t) + \sqrt{\Delta t} \rho_{\Delta t,i}(t) \eta_i + \Delta t \sum_{j \neq i} (\epsilon_{ij} e^{x_{t,j} - x_{t,i}} - \epsilon_{ji}), \quad (28)$$

where η_i 's are correlated random variables with standard normal distribution

$$g_{\Delta t,i}(t) := \frac{1}{\Delta t} E_{\xi} [\ln f_{\Delta t,i}(\xi) | H(t)] \quad \text{and} \quad \rho_{\Delta t,i}^2(t) := \frac{1}{\Delta t} \text{Var}_{\xi} [\ln f_{\Delta t,i}(\xi) | H(t)]. \quad (29)$$

The correlation of η_i 's is given by

$$\text{Corr}_{\xi} [\eta_i, \eta_j] := \text{Corr}_{\xi} [\ln f_{\Delta t,i}(\xi), \ln f_{\Delta t,j}(\xi) | H(t)]. \quad (30)$$

Note that Eq. (28) is the multiple-alternative version of Eq. (22). Equivalently, we can write Eq. (28) as

$$\Delta x_{t,i} \approx \Delta t g_{\Delta t,i}(t) + \sqrt{\Delta t} \widehat{W}_{\Delta t,i} + \Delta t \sum_{j \neq i} (\epsilon_{ij} e^{x_{t,j} - x_{t,i}} - \epsilon_{ji}),$$

where $\widehat{W}_{\Delta t} := (\widehat{W}_{\Delta t,1}, \dots, \widehat{W}_{\Delta t,N})$ follows a multivariate Gaussian distribution with mean zero and covariance matrix $\Sigma_{\Delta t}$ given by

$$\Sigma_{\Delta t,ij} = \frac{1}{\Delta t} \text{Cov}_{\xi} [\ln f_{\Delta t,i}(\xi), \ln f_{\Delta t,j}(\xi) | H(t)]. \quad (31)$$

Finally, taking the limit $\Delta t \rightarrow 0$, and assuming that the limits

$$g_i(t) := \lim_{\Delta t \rightarrow 0} g_{\Delta t,i}(t), \quad \text{and} \quad \Sigma_{ij}(t) := \lim_{\Delta t \rightarrow 0} \Sigma_{\Delta t,ij}(t), \quad (32)$$

are well defined, we obtain the system of SDEs

$$dx_i = g_i(t)dt + d\widehat{W}_i(t) + \sum_{j \neq i} (\epsilon_{ij} e^{x_j - x_i} - \epsilon_{ji}) dt, \quad (33)$$

or equivalently as the vector system

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

where $\mathbf{g}(t) = (g_1(t), \dots, g_N(t))^T$ and $\Lambda(t)\Lambda(t)^T = \Sigma(t)$ are defined using the limits in Eq. (32), $K_i(\mathbf{x}) = \sum_{j \neq i} (\epsilon_{ij} e^{x_j - x_i} - \epsilon_{ji})$, and the components of \mathbf{W}_t are independent Wiener processes. We can recover Eq. (24) by taking $N = 2$, letting $y = x_1 - x_2$, and exchanging the indices 1 and 2 with + and -, respectively. The dependence of accuracy on time is shown in Fig. 7.

As in the case of two alternatives, Eq. (33) can be viewed as an approximation of the logarithm of the likelihood whose evolution is given exactly by Eq. (25). For a fixed interval Δt , the parameters of these equations are related via Eq. (33), and $\epsilon_{\Delta t,ij}/\Delta t = \epsilon_{ij}$.

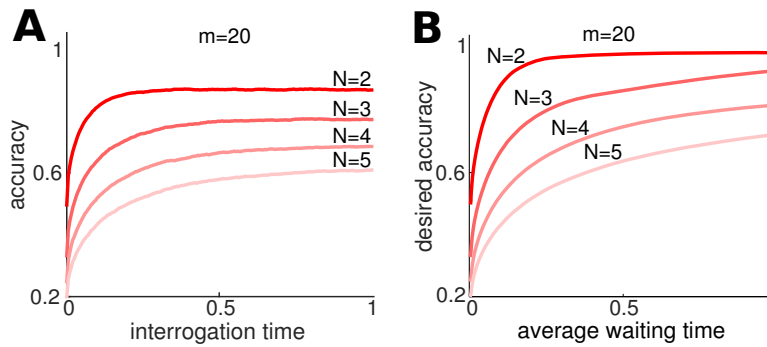


FIG. 7. Dependence of accuracy of responses on the number of alternatives N in Eq. (36). We fix $\epsilon_{ij} \equiv \epsilon$ for all $i \neq j$, $g_i \equiv g$, and set $m = g/\epsilon \equiv 20$. (A) Accuracy in an interrogation protocol decreases with the number of alternatives N , saturating at ever lower levels. (B) The free response protocol results in similar behavior, but the accuracy saturates at 1. The increase in accuracy in time is exceedingly slow for higher numbers of alternatives N .

The limits $g_i(t) := \lim_{\Delta t \rightarrow 0} g_{\Delta t, i}(t)$ and $\Sigma_{ij}(t) := \lim_{\Delta t \rightarrow 0} \Sigma_{\Delta t, ij}(t)$ are defined when the statistics of the observations scale with Δt . As we argued in section **Precise correspondence**, this can be obtained by considering observations drawn from a normal distribution with mean and variance scaled by Δt :

$$f_{\Delta t, i}(\xi) = \frac{1}{\sqrt{2\pi\Delta t\sigma^2}} e^{-(\xi - \Delta t\mu_i)^2 / (2\Delta t\sigma^2)}.$$

Alternatively, the required scaling can also be obtained when each observation made on a time interval consists of a number of sub-observations, (ξ_1, \dots, ξ_k) , with mean and variance scaled by $\frac{\Delta t}{k}$. To do so we approximate $\ln f_{\Delta t, i}(\xi_n)$ in Eq. (27) by

$$\sum_{l=1}^k \frac{\Delta t}{k} \ln f_i(\xi_l) + \sum_{l=1}^k \frac{\sqrt{\Delta t}}{\sqrt{k}} (\ln f_i(\xi_l) - \mathbb{E}_\xi [\ln f_i(\xi) | H(t)]).$$

Log likelihood ratio for multiple alternatives. We can also derive a continuum limit for the log likelihood ratio for any two choices $i, j \in \{1, 2, \dots, N\}$. From Eq. (25), the likelihood ratio $R_{n, ij} = L_{n, i} / L_{n, j}$. We note that this will provide us with a matrix of stochastic processes. We start with the recursive equation

$$R_{n, ij} = \frac{f_{\Delta t, i}(\xi_n)}{f_{\Delta t, j}(\xi_n)} \frac{(1 - \sum_{k \neq i} \epsilon_{\Delta t, ki}) R_{n-1, ij} + \sum_{k \neq i} \epsilon_{\Delta t, ik} R_{n-1, kj}}{1 - \sum_{k \neq j} \epsilon_{\Delta t, kj} + \sum_{k \neq j} \epsilon_{\Delta t, jk} R_{n-1, kj}}. \quad (34)$$

We can thus derive the continuum equation for the log likelihood ratio $y_{n, ij} := \ln R_{n, ij}$, as we did in the case of two alternatives. Since $y_{ij}(t)$ is the difference of log likelihoods $y_{ij}(t) = x_i(t) - x_j(t)$, from Eq. (33) we obtain

$$dy_{ij} = (g_i(t) - g_j(t))dt + d\widehat{W}_i(t) - d\widehat{W}_j(t) + \sum_{k \neq i} (\epsilon_{ik} e^{y_{ki}} - \epsilon_{ki}) dt - \sum_{k \neq j} (\epsilon_{jk} e^{y_{kj}} - \epsilon_{kj}) dt, \quad (35)$$

or

$$dy_{ij} = g_{ij}(t)dt + d\widehat{W}_{ij} + \left(\sum_{k \neq j} \epsilon_{kj} - \sum_{k \neq i} \epsilon_{ki} + \sum_{k \neq i} \epsilon_{ik} e^{y_{ki}} - \sum_{k \neq j} \epsilon_{jk} e^{y_{kj}} \right) dt, \quad (36)$$

where $g_{ij}(t) = \mathbb{E}_\xi \left[\ln \frac{f_i(\xi)}{f_j(\xi)} \middle| H(t) \right]$ and \widehat{W} is a Wiener process with covariance matrix given by $\text{Cov}_\xi [\widehat{W}_{ij}, \widehat{W}_{i'j'} | H(t)] = \text{Cov}_\xi \left[\ln \frac{f_i(\xi)}{f_j(\xi)}, \ln \frac{f_{i'}(\xi)}{f_{j'}(\xi)} \middle| H(t) \right]$. We can also write Eq. (36) in vector form

$$dy = gdt + \Lambda(t)d\mathbf{W}_t + \mathbf{K}(y)dt, \quad (37)$$

where $\mathbf{K}_{ij}(\mathbf{y}) = \sum_{k \neq j} \epsilon_{kj} - \sum_{k \neq i} \epsilon_{ki} + \sum_{k \neq i} \epsilon_{ik} e^{y_{ki}} - \sum_{k \neq j} \epsilon_{jk} e^{y_{kj}}$, $\Lambda(t)\Lambda(t)^T = \Sigma(t)$ is the covariance matrix, and the components of \mathbf{W}_t are independent Wiener processes.

Log likelihood for a continuum of alternatives. Finally, we examine the case where an observer must choose between a continuum of hypotheses H_θ where $\theta \in [a, b]$. Thus, we will first derive a discrete recursive equation for the evolution of the likelihoods $L_{n,\theta} = \Pr(H(t_n) = H_\theta | \xi_{1:n})$. The state of the environment, the correct choice at time t , again changes according to a continuous time Markov process. We define this stochastically switching process through its transition rate function $\epsilon_{\Delta t, \theta\theta'}$, which is given for $\theta' \neq \theta$ as

$$\int_{\theta_1}^{\theta_2} \epsilon_{\Delta t, \theta\theta'} d\theta' := \Pr(H(t_n) \in H_{[\theta_1, \theta_2]} | H(t_{n-1}) = H_{\theta'}),$$

where $H_{[\theta_1, \theta_2]}$ is the set of all states H_θ with θ in the interval $[\theta_1, \theta_2]$. Thus, $\epsilon_{\Delta t, \theta\theta'}$ describes the probability of a transition over a timestep, Δt , from state $H_{\theta'}$ to some state H_θ , with $\theta \in [\theta_1, \theta_2]$. This means that $\Pr(H(t_n) = H_\theta | H(t_{n-1}) = H_\theta) = 1 - \int_a^b \epsilon_{\Delta t, \theta\theta'} d\theta'$. As in the derivation of the multiple alternative $2 \leq N < \infty$ case, we can express each likelihood $L_{n,\theta}$ at time t_n in terms of the likelihoods $L_{n-1,\theta'}$ at time t_{n-1} , so

$$L_{n,\theta} = \frac{\Pr(\xi_{1:n-1})}{\Pr(\xi_{1:n})} \Pr(\xi_n | H(t_n) = H_\theta) \left(\Pr(H(t_n) = H_\theta | H(t_{n-1}) = H_\theta) L_{n-1,\theta} + \int_a^b \epsilon_{\Delta t, \theta\theta'} L_{n-1,\theta'} d\theta' \right).$$

Notice that the sum from the $N < \infty$ case, as in Eq. (25), has been replaced with an integral over all possible hypotheses $H_{\theta'}$, $\theta' \in [a, b]$ and a term corresponding to the probability of the environment not changing. Again we drop the common factor $\frac{\Pr(\xi_{1:n-1})}{\Pr(\xi_{1:n})}$, since we wish to compare likelihoods. We obtain

$$L_{n,\theta} = f_{\Delta t, \theta}(\xi_n) \left(\left[1 - \int_a^b \epsilon_{\Delta t, \theta\theta'} d\theta' \right] L_{n-1,\theta} + \int_a^b \epsilon_{\Delta t, \theta\theta'} L_{n-1,\theta'} d\theta' \right), \quad (38)$$

where $f_{\Delta t, \theta}(\xi_n) = \Pr(\xi_n | H(t_n) = H_\theta)$. From Eq. (38), we can thus derive a recursive relation for the log of the rescaled likelihoods $x_{n,\theta} := \ln L_{n,\theta}$ in terms of $x_{n-1,\theta}$ so

$$x_{n,\theta} - x_{n-1,\theta} = \ln f_{\Delta t, \theta}(\xi_n) + \ln \left(1 - \int_a^b \epsilon_{\Delta t, \theta\theta'} d\theta' + \int_a^b \epsilon_{\Delta t, \theta\theta'} e^{x_{n-1,\theta'} - x_{n-1,\theta}} d\theta' \right).$$

To approximate this discrete-time stochastic process with a SDE, we denote by $\Delta x_{n,\theta} = x_{n,\theta} - x_{n-1,\theta}$, the change in log likelihood due to the observation at time t_n . Furthermore, we assume $\epsilon_{\Delta t, \theta\theta'} := \Delta t \epsilon_{\theta\theta'} + o(\Delta t)$ and drop higher order terms,

$$\Delta x_{n,\theta} = \ln f_{\Delta t, \theta}(\xi_n) + \ln \left(1 - \int_a^b \Delta t \epsilon_{\theta\theta'} d\theta' + \int_a^b \Delta t \epsilon_{\theta\theta'} e^{x_{n-1,\theta'} - x_{n-1,\theta}} d\theta' \right). \quad (39)$$

Assuming $\Delta t \ll 1$, we can utilize the approximation $\ln(1+a) \approx a$ which holds when $|a| \ll 1$. Assuming $|\Delta x_{n,\theta}| \ll 1$,

$$\Delta x_{n,\theta} \approx \ln f_{\Delta t, \theta}(\xi_n) + \Delta t \int_a^b (\epsilon_{\theta\theta'} e^{x_{n-1,\theta'} - x_{n-1,\theta}} - \epsilon_{\theta\theta'}) d\theta' \quad (40)$$

$$= \mathbb{E}_\xi [f_{\Delta t, \theta}(\xi) | H(t_n)] + (\ln f_{\Delta t, \theta}(\xi_n) - \mathbb{E}_\xi [\ln f_{\Delta t, \theta}(\xi) | H(t_n)]) + \Delta t \int_a^b (\epsilon_{\theta\theta'} e^{x_{n-1,\theta'} - x_{n-1,\theta}} - \epsilon_{\theta\theta'}) d\theta' \quad (41)$$

conditioned on the current state of the environment $H(t_n) = H_\varphi$ where $\varphi \in [a, b]$.

Exchanging the index n with the time, t , we can therefore write

$$\Delta x_{t,\theta} \approx \Delta t g_{\Delta t, \theta}(t) + \sqrt{\Delta t} \rho_{\Delta t, \theta}(t) \eta_\theta + \Delta t \int_a^b (\epsilon_{\theta\theta'} e^{x_{t,\theta'} - x_{t,\theta}} - \epsilon_{\theta\theta'}) d\theta', \quad (42)$$

where η_θ 's are correlated random variables which marginally follow a standard normal distribution, and

$$g_{\Delta t, \theta}(t) := \frac{1}{\Delta t} \mathbb{E}_\xi [\ln f_{\Delta t, \theta}(\xi) | H(t)], \quad \text{and} \quad \rho_{\Delta t, \theta}^2(t) := \frac{1}{\Delta t} \text{Var}_\xi [\ln f_{\Delta t, \theta}(\xi) | H(t)]. \quad (43)$$

The correlation of η_i 's is given by

$$\text{Corr}_\xi[\eta_\theta, \eta_{\theta'}] := \text{Corr}_\xi[\ln f_{\Delta t, \theta}(\xi), \ln f_{\Delta t, \theta'}(\xi) | H(t)]. \quad (44)$$

Equivalently, we can write Eq. (42) as

$$\Delta x_{t, \theta} \approx \Delta t g_{\Delta t, \theta}(t) + \sqrt{\Delta t} \widehat{W}_{\Delta t, \theta} + \Delta t \int_a^b (\epsilon_{\theta \theta'} e^{x_{t, \theta'} - x_{t, \theta}} - \epsilon_{\theta' \theta}) d\theta',$$

where $\widehat{W}_{\Delta t} := (\widehat{W}_{\Delta t, \theta})_{\theta \in [a, b]}$ follows a multivariate Gaussian distribution with mean zero and covariance function $\Sigma_{\Delta t, \theta \theta'}$ given by

$$\Sigma_{\Delta t, \theta \theta'} = \frac{1}{\Delta t} \text{Cov}_\xi[\ln f_{\Delta t, \theta}(\xi), \ln f_{\Delta t, \theta'}(\xi) | H(t)]. \quad (45)$$

Finally, taking the limit $\Delta t \rightarrow 0$, and assuming that the limits

$$g_\theta(t) := \lim_{\Delta t \rightarrow 0} g_{\Delta t, \theta}(t), \quad \text{and} \quad \Sigma_{\theta \theta'}(t) := \lim_{\Delta t \rightarrow 0} \Sigma_{\Delta t, \theta \theta'}(t), \quad (46)$$

are well defined, we obtain the system of SDEs

$$dx_\theta = g_\theta(t)dt + d\widehat{W}_\theta(t) + \int_a^b (\epsilon_{\theta \theta'} e^{x_{\theta'} - x_\theta} - \epsilon_{\theta' \theta}) d\theta' dt, \quad (47)$$

or equivalently as the system of SDEs

$$d\mathbf{x} = \mathbf{g}(t)dt + \Lambda(t)d\mathbf{W}_t + K(\mathbf{x})dt, \quad (48)$$

where $\mathbf{g}(t) = (g_\theta(t))_{\theta \in [a, b]}$ and $\Lambda(t)\Lambda(t)^T = \Sigma(t)$ are defined using the limits in Eq. (46), $K_\theta(\mathbf{x}) = \int_a^b (\epsilon_{\theta \theta'} e^{x_{\theta'} - x_\theta} - \epsilon_{\theta' \theta}) d\theta'$, and the components of \mathbf{W}_t are independent Wiener processes.

Neural population model for $N = 2$. Inference in dynamic environments with two states, H_+ and H_- , can be optimally performed by mutually excitatory neural populations r_+ and r_- ,

$$dr_+ = [I_+(t) - \alpha r_+ + F_+(r_- - r_+)]dt + dW_+, \quad (49a)$$

$$dr_- = [I_-(t) - \alpha r_- + F_-(r_+ - r_-)]dt + dW_-, \quad (49b)$$

where the external input $I_\pm(t) = I_\pm^0$ when $H(t) = H_\pm$ and vanishes otherwise, $F_\pm(r) = -\alpha r/2 + \epsilon_\mp e^r - \epsilon_\pm$, and $\alpha > 0$ so that the system is stable. Note, the parameter α provides the leak in the activity of each individual population, which depends on both the time constants and recurrent architecture of the local network [25].

When the environment has not changed for a long time, the noise-free system approaches a fixed point given by

$$(\bar{r}_+, \bar{r}_-) = \left(\frac{I_+ + \epsilon_- e^{-\bar{y}} - \epsilon_+}{\alpha} + \frac{\bar{y}}{2}, \frac{I_- + \epsilon_+ e^{\bar{y}} - \epsilon_-}{\alpha} - \frac{\bar{y}}{2} \right),$$

where $\bar{y} = \ln \left[\frac{I_+ - I_- + \epsilon_- - \epsilon_+}{2\epsilon_+} + \sqrt{\frac{(I_+ - I_- + \epsilon_- - \epsilon_+)^2}{4\epsilon_+^2} + \frac{\epsilon_-}{\epsilon_+}} \right].$

Note that by increasing (decreasing) α , the fixed points (\bar{r}_+, \bar{r}_-) move closer (farther) from the origin $(0, 0)$.

We now demonstrate that coupling between populations described by Eq. (49) can be excitatory and coupling within populations inhibitory, as stated in the main text. To do so, note that the Jacobian matrix of (F_+, F_-) has the form:

$$J(r_+, r_-) = \begin{bmatrix} \alpha/2 - \epsilon_- e^{r_- - r_+} & -\alpha/2 + \epsilon_- e^{r_- - r_+} \\ -\alpha/2 + \epsilon_+ e^{r_+ - r_-} & \alpha/2 - \epsilon_+ e^{r_+ - r_-} \end{bmatrix}.$$

For $\epsilon_\pm > 0$, taking $\alpha < 2 \min\{\epsilon_- e^{-\bar{y}}, \epsilon_+ e^{\bar{y}}\}$ will guarantee that the sign of the Jacobian matrix is $\begin{bmatrix} - & + \\ + & - \end{bmatrix}$ on a region that contains the fixed point. This corresponds to a neural network with self-inhibition and mutual excitation. We illustrate this point with the network wiring diagrams we have drawn in Fig. 5 of the main text.

Neural population model for multiple alternatives. We can extend our results for the $N = 2$ case by deriving a neural population model of decision making in changing environments. In [13], the reliability of motion information

was assumed to vary during a trial, and the optimal model encoded the posterior probability distribution over the possible stimulus space. Here, we assume the true hypothesis, $H(t)$, changes in time. For an arbitrary number of possible states, $\{H_1, \dots, H_N\}$, decisions can be performed optimally by neural populations x_1, \dots, x_N coupled by mutual excitation

$$dr_i = \left[I_i(t) - \alpha r_i + \sum_{j \neq i} F_{ij}(r_j - r_i) \right] dt + d\widehat{W}_i(t), \quad (50)$$

where the external input $I_i(t) = I_i^0$ when $H(t) = H_i$ and 0 otherwise and $(d\widehat{W}_1(t), \dots, d\widehat{W}_N(t))^T = \Lambda(t)d\mathbf{W}_t$ with $\Lambda(t)$ defined as in Eq. (48). Population firing rates are again determined by inhibition within each population and excitation between populations as described by the arguments of the firing rate function

$$F_{ij}(r) = -\alpha r/N + \epsilon_{ij}e^r - \epsilon_{ji}.$$

Note that by taking $y_{ij} = r_i - r_j$, we can convert the neural population model, Eq. (50), to the SDE for the log likelihood ratio, Eq. (36). Also notice that, as in the case of $N = 2$ alternatives, in the limit $\epsilon_{ij} \rightarrow 0$, we obtain linear integrators [26]

$$dr_i = \left[I_i(t) - \alpha \sum_{j=1}^N r_j/N \right] dt + d\widehat{W}_i(t).$$