

**From Neural Fluctuations to Perceptual Choice:
A Fluctuation–Response Test of Evidence Accumulation**

PIATRA . INSTITUTE

October 2026

Abstract

In models of choice, temperature is usually a parameter fitted to the choices it is meant to explain. For a decision variable that accumulates evidence to a bound, a quantity with the role of a temperature can be measured instead. The probability of a choice is a logistic function of the evidence h divided by $D/(MB)$, where D is the diffusion of the decision variable, M its mobility (the change of drift per unit of evidence) and B the bound, so the psychometric slope is MB/D . The ratio is unchanged by any affine change of the decoded coordinate and by the time unit, so a decoder in arbitrary units can be compared with behaviour. We give the exact factors by which leak, instability, trial-to-trial drift variability and start-point bias change the slope; leak can at most double it. For any stationary linear system, the fluctuation-response ratio along the Fisher discriminant of the stationary response is constant at every lag whether or not the system is at equilibrium, so its constancy along a trained decoder tests nothing. We estimate D and M from all trials through optional stopping and split-half cross moments, and test the estimators on a simulated population of 200 Poisson neurons carrying a bounded accumulator. One-bin increment variances predict a slope of 1.40 against a fitted 10.02, and variograms of undecided trials predict 12.29, while the stopped split-half estimators predict 9.89. Under leak, instability and drift variability the mean log ratio of fitted to predicted slope stays between -0.044 and 0.051; urgency and a slow signal along the coding axis that does not drive the choice raise it to 0.108 and 0.097. With 8,000 trials the run-to-run SD of the log ratio is 0.080; an equivalence test with a $\pm 25\%$ margin reaches a power of 0.8 at about 8,700 trials, and a $\pm 10\%$ margin needs about 48,000.

1. Introduction

Temperature appears in models of perception and choice as a parameter. A softmax rule with an inverse temperature is fitted to choices in reinforcement-learning tasks (Daw et al. 2006), and precision, an inverse variance that scales prediction errors, is inferred in models of attention (Feldman and Friston 2010; Friston 2010). In each case the value comes from the behaviour that the model then explains. In physics temperature is a measured state variable. Einstein showed that the diffusion coefficient of a suspended particle and its mobility under a known force fix the temperature, $D = \mu k_B T$ (Einstein 1905), and the fluctuation-dissipation theorem generalizes the relation to any observable conjugate to a weak perturbation (Callen and Welton 1951; Kubo 1966). Out of equilibrium the ratio of fluctuation to response still defines an effective temperature, which can depend on time scale (Cugliandolo et al. 1997; Cugliandolo 2011); hair bundles of the inner ear violate the equilibrium relation, and their effective temperature diverges near the frequency of spontaneous oscillation (Martin et al. 2001). The size of the violation measures dissipation (Harada and Sasa 2005; Baiesi et al. 2009; Seifert 2012), and the same tests detect active processes in cytoskeletal networks and cells (Mizuno et al. 2007; Gnesotto et al. 2018) and broken detailed balance in human brain recordings (Lynn et al. 2021). Responses of cortex to stimuli have been linked to its ongoing fluctuations as well (Arieli et al. 1996; Churchland et al. 2010; Sarracino et al. 2020).

A perceptual decision offers a setting where the analogue can be stated with all three quantities observable. In random-dot motion discrimination, neurons of the lateral intraparietal area (LIP)

ramp at a rate set by motion strength and reach a common level before the saccade that reports the choice (Shadlen and Newsome 2001; Roitman and Shadlen 2002). Brief motion pulses have persistent effects on both choices and LIP activity (Huk and Shadlen 2005), accumulation stops at a bound even when the environment controls viewing time (Kiani et al. 2008), and the accumulated quantity behaves like a log likelihood ratio (Gold and Shadlen 2001, 2007; Yang and Shadlen 2007; Kira et al. 2015). The account is a bounded diffusion with roots in sequential analysis (Wald 1945; Ratcliff 1978; Ratcliff and McKoon 2008; Bogacz et al. 2006; Palmer et al. 2005; Shadlen and Kiani 2013). The variance and correlation of LIP activity across trials match those of a diffusion (Churchland et al. 2011), single-trial decision variables can be decoded from populations (Kiani et al. 2014; Peixoto et al. 2021), and simultaneous recordings of up to 203 LIP neurons with Neuropixels probes show a drift-diffusion signal on single decisions (Steinemann et al. 2024b; Stine et al. 2023). Single-neuron spike trains have also been read as discrete steps (Latimer et al. 2015), LIP representations depend on task (Okazawa et al. 2021), and in a pulse-based task in rats the noise lay in sensory processing rather than in the accumulator (Brunton et al. 2013).

For a bounded accumulator the choice probability depends on diffusion, mobility and bound only through $D/(MB)$, a standard first-passage result (Palmer et al. 2005; Redner 2001; Risken 1989). Each of the three can be estimated from a decoded neural decision variable on trials whose choices are not used, and the psychometric slope can be estimated from choices on other trials. The comparison is a test of the bounded-accumulation account with no parameter fitted to the outcome it predicts.

The relation is invariant under every affine change of the decoded coordinate and the time unit, and it transforms with the evidence unit as the fitted slope does. Leak, instability, drift variability and start-point bias change the slope by factors we derive in closed form; leak can at most double it. The equilibrium form of the test, a fluctuation-response ratio that is constant in lag, is uninformative along a trained decoder: for any stationary linear system it holds along the Fisher discriminant whether or not the system is at equilibrium. Optional stopping yields estimators of D and M that use all trials without bias from the absorbing bounds, and split-half cross moments remove observation noise. In a simulated spiking population the estimators predict the fitted slope without detectable bias, naive estimators fail by large factors, and the prediction stays close to behaviour under several departures from the integrator, which limits what agreement can show. We give the number of trials needed for an equivalence test and an analysis plan for a public dataset.

2. The first-passage relation

Let X be a scalar decision variable that starts at $z = 0$ and follows

$$dX_t = (Mh(t) - \lambda X_t) dt + \sqrt{2D} dW_t$$

until it reaches $+B$ or $-B$; the bound reached gives the choice. Here h is the signed evidence (motion coherence in the examples), M the mobility, D the diffusion and λ a leak, zero for the

pure integrator. The decision time τ is the first exit time; the reaction time adds a non-decision time.

Proposition 1. For the pure integrator with constant evidence h , the probability of the choice $+$ is

$$P(+ | h) = \frac{1}{1 + e^{-h/\Theta}}, \quad \Theta = \frac{D}{MB},$$

so the psychometric function is logistic with slope $\beta = MB/D$ per unit of evidence.

Proof. The exit probability $u(x)$ through $+B$ from a start x solves $Du'' + Mh u' = 0$ with $u(-B) = 0$ and $u(B) = 1$, so $u(x) = (1 - e^{-Mh(x+B)/D}) / (1 - e^{-2MhB/D})$. At $x = 0$ this is $1/(1 + e^{-MhB/D})$. $\square$

Proposition 2. Let the decoded variable be $A = aX + b$ with $a \neq 0$, let time be measured in a unit κ times larger, and let evidence be measured in a unit γ times larger. Then D , M and B measured in the new units give the same $D/(MB)$ when the coordinate or the time unit changes. When the evidence unit changes, $D/(MB)$ is divided by γ and the fitted slope is multiplied by γ , so their product is unchanged. If $A = g(X)$ with g nonlinear, the diffusion and mobility measured in A are $g'(X)^2 D$ and $g'(X)M$, and their ratio varies with the state.

Proof. Per new unit of time the drift of A is $a\kappa Mh$ and its diffusion $a^2\kappa D$, and the bound sits at distance $|a|B$ from the start, so $D/(MB)$ is multiplied by $a^2\kappa/(a\kappa|a|) = \pm 1$, the sign recording a relabelling of the choices. With evidence $h' = h/\gamma$ in the new unit the mobility per unit of h' is γM , so $D/(MB)$ is divided by γ , and $P(+)$ has slope $\gamma\beta$ per unit of h' . For a nonlinear g , Itô's formula gives a local diffusion $g'(X)^2 D$ and a local drift coefficient $g'(X)M$ plus a curvature term. $\square$

The proposition is what makes a neural test possible. A linear decoder of a population is known only up to scale and offset, and its output is in arbitrary units; D , M and B measured in those units still give $D/(MB)$ in units of evidence. A nonlinear readout destroys the invariance, so the decoder has to be linear in the decision variable over its range.

The relation is the Einstein relation in the coordinates of the decision. With evidence in the role of a force and M in the role of mobility, D/M has the role of $k_B T$, and the choice probability has the two-state Boltzmann form with the energy difference h and the temperature Θ , measured in units of evidence once divided by the bound. The analogy is exact for the pure integrator and fails for the departures below.

3. Departures from the integrator

The relation $\beta = MB/D$ holds for a particular stochastic process. We give the slope at zero evidence, $\beta_0 = 4 dP/dh$ at $h = 0$, which equals the logistic slope when the function is logistic, for the standard departures.

Proposition 3. Relative to MB/D , the slope at zero evidence is multiplied by:

(a) for a leak λ , with the leak number $\Lambda = \lambda B^2/(2D)$,

$$r(\Lambda) = \frac{e^\Lambda - 1}{\Lambda \int_0^1 e^{\Lambda u^2} du},$$

which increases from 0 as $\Lambda \rightarrow -\infty$ (instability) through 1 at $\Lambda = 0$ to 2 as $\Lambda \rightarrow \infty$;

(b) for a drift that varies across trials with SD η , with $s = \eta B/D$, the factor $4 E[\sigma'(sZ)]$, where σ is the logistic function and Z a standard normal variable;

(c) for a start point z , the factor $1 - z^2/B^2$;

(d) for a start point uniform on $(-s_z, s_z)$, the factor $1 - s_z^2/(3B^2)$.

Proof. (a) With the scale density $\exp(-(vy - \lambda y^2/2)/D)$, the exit probability from 0 is the integral of the density over $(-B, 0)$ divided by its integral over $(-B, B)$. Differentiating at $v = 0$ gives $dP/dv = J/(DG)$ with $J = \int_0^B y e^{\lambda y^2/(2D)} dy = (B^2/(2\Lambda))(e^\Lambda - 1)$ and $G = 2B \int_0^1 e^{\Lambda u^2} du$; the integrator has $J/G = B/4$. For $\Lambda \rightarrow \infty$ the integral tends to $e^\Lambda/(2\Lambda)$, giving 2; for $\Lambda \rightarrow -\infty$ it tends to $\sqrt{\pi/(4|\Lambda|)}$, giving $2/\sqrt{\pi|\Lambda|} \rightarrow 0$. (b) $P = E[\sigma((Mh + \eta Z)B/D)]$; differentiate at $h = 0$ and use $\sigma'(0) = 1/4$. (c) Expanding the exit probability of Proposition 1 from $x = z$ to first order in Mh gives $dP/dh = M(B^2 - z^2)/(4BD)$. (d) Average (c) over the start point. $\square$

The leak limit has a physical reading. With strong leak the decision variable equilibrates in the well around zero and leaves by escape over the barrier at either bound (Kramers 1940); the odds of the two exits then follow the Boltzmann ratio between the two bounds, whose energies differ by $2MBh$ in units of D , and the slope is $2MB/D$. The pure integrator has half that slope because it starts at the midpoint and exits before it equilibrates.

The closed forms agree with quadrature of the scale density to within 5×10^{-7} for Λ from -2 to 5 and with Monte Carlo simulation of 40,000 trials per condition. Table 1 gives values. A leak number of 1, which means the leak acts on the same time scale as diffusion across the bound, multiplies the slope by 1.175, the same instability multiplies it by 0.846, and drift variability with $s = 1$ by 0.826.

Table 1: Slope at zero evidence relative to the pure integrator, Proposition 3.

departure	size	slope relative to MB/D
leak	$\Lambda = 0.5, 1, 2, 5$	1.086, 1.175, 1.351, 1.717
instability	$\Lambda = -0.5, -1, -2$	0.920, 0.846, 0.723
drift variability	$s = 0.5, 1, 2$	0.944, 0.826, 0.606
start-point offset	$z/B = 0.3$	0.910
start-point spread	$s_z/B = 0.6$	0.880

Bounds that collapse with time have no closed form. With $B(t) = B_0(1 - t/(t + t_{50}))$ and the reference parameters ($D = 0.8$ per s and $M = 8$ per unit coherence per s with $B_0 = 1$, a slope of 10 per unit coherence), the logistic slope fitted to 66,000 simulated choices over eleven coherences is 8.54, 7.82 and 6.94 for t_{50} of 2, 1 and 0.5 s, against 9.90 for the fixed bound. Replacing B_0 by the mean bound $\bar{B}$ at the time of decision on trials below 10% coherence gives $M\bar{B}/D$ of 8.32, 7.52 and 6.61, slightly below the fitted slopes. A bound measured as the average termination level of the decoded variable therefore accounts for most of the effect of a collapse.




Figure 1: The first-passage relation. (a) Choice probability against evidence for four parameter sets; the first three share $D/(MB) = 0.1$ and fall on one curve (lines exact, symbols Monte Carlo with 6,000 trials per point). (b) Slope at zero evidence relative to MB/D against the leak number; negative values are instability, and the slope tends to 2 with strong leak. (c) Drift variability across trials. (d) Start-point offset and spread.

4. What the equilibrium form of the test measures

An equilibrium version of the test would record a decoded variable A without evidence, compute its autocovariance $C(\tau)$, perturb with weak evidence to obtain the response function $R(\tau)$, and ask whether $-C'(\tau)/R(\tau)$ is constant in lag, as the fluctuation-dissipation theorem requires at equilibrium for the observable conjugate to the perturbation. For a one-dimensional stationary process $dX = (Mh - \lambda X) dt + \sqrt{2D} dW$ the ratio is D/M at every lag, since $C(\tau) = (D/\lambda)e^{-\lambda\tau}$ and $R(\tau) = Me^{-\lambda\tau}$. For populations the result depends on the readout.

Proposition 4. Let x follow $dx = (-Kx + fh) dt + \text{noise}$ with stationary covariance S , so that $C_w(\tau) = w^\top e^{-K\tau} S w$ and $R_w(\tau) = w^\top e^{-K\tau} f$ for a readout w and $\tau \geq 0$. Along $w^* = S^{-1} K^{-1} f$, the Fisher discriminant of the stationary responses to a constant h , $-C'_{w^*}(\tau) = R_{w^*}(\tau)$ for every τ , with or without detailed balance. Under detailed balance $KS = D_x$, where D_x is the diffusion matrix, and w^* is proportional to $D_x^{-1} f$, the readout conjugate to the perturbation.

Proof. $Sw^* = K^{-1} f$, so $-C'_{w^*}(\tau) = w^{*\top} K e^{-K\tau} K^{-1} f = w^{*\top} e^{-K\tau} f$ because K commutes with $e^{-K\tau}$. Under detailed balance $S^{-1} K^{-1} = (KS)^{-1} = D_x^{-1}$. $\square$

The proposition removes the test from consideration. A decoder trained to discriminate responses to opposite evidence approximates w^* , and along w^* the ratio is constant whatever the dynamics, so constancy says nothing about equilibrium or about a state variable. Along other readouts the ratio varies with lag even at equilibrium. In a two-dimensional example with a common station-

ary covariance (Figure 2), the conjugate readout and the Fisher discriminant coincide at detailed balance and the ratio is constant to rounding error along both, while along the first coordinate it varies by 29% and along the mean-response direction by 95% over the lags shown. With a rotational component that breaks detailed balance, the Fisher discriminant departs from the conjugate readout by 70° , the ratio along it remains constant, and the response along the conjugate readout changes sign. The first-passage relation avoids the difficulty, because the bound converts D/M into evidence units and the prediction concerns choices, an outcome the decoder was not trained on.




Figure 2: The fluctuation-response ratio of stationary linear systems, relative to its value at zero lag, along four readouts. (a) Detailed balance: the conjugate readout and the Fisher discriminant coincide and give a constant ratio; other readouts do not. (b) Broken detailed balance with the same stationary covariance: the ratio is still constant along the Fisher discriminant; the responses along the conjugate readout and the first coordinate change sign and are not shown.

5. Estimating diffusion, mobility and bound from a population

The decoded variable of a recorded population differs from X in four ways: it is $A = cX + o + e$ with unknown scale c and offset o ; the observation noise e from spiking is large in short bins; the trajectory stops at the bound, so trials still undecided at a late time are a selected subset; and the time of the bound crossing is known only through the reaction time.

Proposition 5. For the pure integrator started at zero, with τ the exit time and any fixed t ,

$$E[X_{t \wedge \tau}] = E\left[\int_0^{t \wedge \tau} Mh(s) ds\right], \quad E[X_{t \wedge \tau}^2] = 2D E[t \wedge \tau] \text{ when } h = 0.$$

If two disjoint parts of the population give $A_k = c_k X + o_k + e_k$ with observation noise independent between the parts and over time, then $E[(A_1(t \wedge \tau) - A_1(0))(A_2(t \wedge \tau) - A_2(0))] = c_1 c_2 E[X_{t \wedge \tau}^2]$.

Proof. $X_t - \int_0^t Mh$ and, at $h = 0$, $X_t^2 - 2Dt$ are martingales; $t \wedge \tau$ is a bounded stopping time, so optional stopping applies. The cross moment follows from the independence of e_1 , e_2 and X . $\square$

The proposition turns the absorbing bound from a source of bias into part of the estimator. A state-space model with point-process observations would also separate the two noises (Kalman 1960; Smith and Brown 2003), at the cost of a model of the dynamics whose errors enter D and M . Each trial contributes its decoded value frozen at the decision, and the time measure is the decision time when the decision came first. In data the frozen value is taken from the post-decision period, before the saccade, and the decision time is the reaction time minus a non-decision time. Three estimators follow. D comes from zero-evidence trials, as half the slope of the stopped split-half cross moment against the mean stopped time. M comes either from sustained low evidence, as the slope of the mean of $h(A(t \wedge \tau) - A(0))$ against the mean of $h^2(t \wedge \tau)$, or from pulses, as the slope of the onset-aligned stopped displacement against the stopped exposure to the pulse. B is half the difference between the mean decoded levels before the saccade for the two choices. The scale factors cancel in $\hat{D}_{12}/(\hat{M}_1 \hat{B}_2)$, which we average with the estimate in which the parts are swapped. With binned data, a bin average of a Brownian path has second moment $2D(t_c - \Delta/6)$ for a bin of width Δ centred at t_c , and we use that time measure for bins before the decision.

We tested the estimators on a simulated population. The decision variable followed the integrator with $D = 0.8$ per s, $M = 8$ per unit coherence per s and $B = 1$, a slope of 10 per unit coherence and a mean decision time of 0.625 s at zero evidence, simulated in 2 ms steps with a Brownian-bridge test for crossings between steps. A population of 200 neurons had baseline rates between 15 and 35 Hz and gains on the decision variable proportional to their baselines, bounded so that rates stayed positive, and emitted Poisson counts in 20 ms bins. After the crossing the represented variable stayed at the bound until the saccade, 0.3 s later. Linear discriminants with strong shrinkage were trained on 400 trials at $\pm 51.2\%$ coherence, for the full population and for two random halves, and frozen. Separate trial sets served each estimate: 1,000 zero-coherence trials for D , 2,000 zero-coherence trials with a 100 ms pulse of $\pm 20\%$ coherence at a random onset between 0.1 and 0.5 s for the pulse estimate of M , 2,000 trials at ± 6.4 , 12.8 and 25.6% for the sustained estimate of M , and 3,000 trials over eleven coherences from -51.2 to 51.2% whose choices gave the fitted slope. No choice entered any neural estimate. Training trials always used Poisson counts; for the test trials most runs replaced the counts of each decoder by a Gaussian variable with the same mean and variance, and one comparison below checks the approximation.

Over 48 simulated datasets the slope fitted to choices was 10.02 (SD 0.44) and the stopped split-half estimators predicted 9.89 (SD 0.66). The mean log ratio of fitted to predicted slope was 0.014 (SE 0.011), with an SD of 0.079 across datasets. The estimates themselves were $D/B^2 = 0.800$ against 0.8 and $M/B = 7.90$ against 8; the pulse estimate of M/B was 8.28, with twice the spread of the sustained estimate. The naive estimator, the variance of one-bin increments of the full decoder on undecided trials, predicted a slope of 1.40, because spiking noise in a 20 ms bin dominates the diffusion of the decision variable. A variogram on undecided trials, which removes white observation noise through its intercept, predicted 12.29, too steep because trajectories still undecided at a given time have been selected away from the bounds and diffuse less (Figure 3).

Three checks bound the remaining sources of error. Misstating the non-decision time by 50 ms shifts the mean log ratio from -0.018 to 0.029 (for 0.25 s) or -0.056 (for 0.35 s), so the non-decision




Figure 3: Estimators on the simulated population, pure integrator. (a) The stopped split-half cross moment from one dataset, in units of the bound, against the stopped time; the dashed line has the true slope $2D/B^2$. (b) Slopes predicted by four estimators over 48 datasets (mean $\pm$ SD, logarithmic axis) against the slope fitted to choices (band).

time must be known to a few tens of milliseconds, for example from the neural signature of termination. With 50 neurons instead of 200 the SD of the log ratio rises from 0.079 to 0.133, and 400 neurons reduce it only to 0.074, because the latent diffusion dominates once observation noise is averaged out. Exact Poisson counts and the Gaussian approximation used for most runs give mean log ratios of -0.002 and -0.013 over 16 datasets each.

6. Prediction under six generating models

To ask what agreement shows, the same pipeline was applied to data generated by five departures from the integrator: leak ($\Lambda = 1$), instability ($\Lambda = -1$), drift variability across trials ($s = 1$), an urgency signal that multiplies the represented variable by $1 + t/(1 \text{ s})$ against a fixed neural bound, and a slow signal along the coding axis (SD 0.3 of the bound, time constant 1 s) that enters the neural readout but not the decision (Table 2, Figure 4).

The instantaneous parameters predict a slope of 10 in every case, against fitted slopes of 11.48, 8.66 and 8.51 for the leak, instability and drift-variability models. The stopped estimators do better because they measure an effective diffusion and mobility over the decision window: with leak the stopped second moment grows more slowly than $2Dt$, with instability faster, and the prediction moves in the direction of behaviour. Under these three models the mean log ratio stays between -0.044 and 0.051, and the prediction falls within $\pm 25\%$ of the fitted slope in at least 98% of datasets. Agreement between the neural prediction and behaviour is therefore evidence for bounded accumulation with the measured effective ratio; it does not single out the pure integrator.

Table 2: The neural prediction and the slope fitted to choices, mean over 48 simulated datasets per model (leak $\Lambda = 1$, instability $\Lambda = -1$, drift variability $s = 1$); the last two columns are diagnostics computed without choices.

generating model	fitted slope	predicted slope	mean log ratio (SE)	late/early $D/(MB)$	pulse/sustained M
integrator	10.02	9.89	0.014 (0.011)	1.03	1.05
leak	11.48	12.01	-0.044 (0.012)	0.97	1.12
instability	8.66	8.23	0.051 (0.010)	1.04	0.94
drift	8.51	8.74	-0.026 (0.012)	1.12	1.03
variability					
urgency	7.87	7.07	0.108 (0.010)	1.23	1.00
aligned	10.05	9.15	0.097 (0.014)	0.98	1.03
nuisance					

Two departures leave larger gaps. Urgency (Churchland et al. 2008; Ditterich 2006; Drugowitsch et al. 2012; Hanks et al. 2014; Cisek et al. 2009) makes the neural slope and diffusion grow during the trial, and the prediction is 7.07 against a fitted 7.87; the ratio of $D/(MB)$ in a late window (0.3 to 0.6 s) to an early one (0.1 to 0.3 s) rises to 1.23 against 1.03 for the integrator, which flags it; it also rises, to 1.12, under drift variability. A signal along the coding axis that does not drive the choice inflates the measured diffusion and lowers the prediction to 9.15 against 10.05, and no diagnostic computed without choices separates it from a noisier accumulator; it shows only as the gap itself. The ratio of pulse to sustained mobility is 1.05 under the integrator, 1.12 under leak and 0.94 under instability, so evidence delivered briefly is integrated differently from evidence delivered throughout when the accumulator leaks or is unstable.

7. Sample size

The comparison is an equivalence test: two one-sided tests at level 0.05 on the log ratio of fitted to predicted slope, with a margin chosen in advance (Schuirmann 1987; Lakens 2017). Over 96 simulated datasets at each of 4,000, 8,000, 16,000 and 32,000 trials, with every trial set scaled in proportion, the SD of the log ratio was 0.106, 0.067, 0.069 and 0.041 (Figure 5). One curve $\kappa/\sqrt{n}$, with n the number of trials and $\kappa = 7.12$, describes the four values, which scatter around it by up to 23%, and we use the curve for the design statements. It gives an SD of 0.080 at 8,000 trials. Two one-sided tests with a margin of $\pm 25\%$ then reach a power of 0.8 at about 8,700 trials, and a margin of $\pm 10\%$ needs about 48,000 trials. Most of the uncertainty is on the neural side: at 8,000 trials the SD of the log fitted slope is 0.037 and that of the log predicted slope 0.066.




Figure 4: Six generating models. (a) Slope fitted to choices against the slope predicted from the neural estimates (filled, mean $\pm$ SE over 48 datasets) and against the instantaneous $MB/D = 10$ (open); the band is $\pm 10\%$. (b) The ratio of late to early $D/(MB)$ and of pulse to sustained mobility.




Figure 5: Sample size for the comparison. (a) SD across 96 simulated datasets of the log fitted slope, the log predicted slope and their difference, against the number of trials, with the pooled curve $\kappa/\sqrt{n}$. (b) Power of the equivalence test at margins of $\pm 10\%$ and $\pm 25\%$ when the true ratio is 1.

8. An analysis plan for public data

The data of Steinemann and colleagues fit the test closely (Steinemann et al. 2024b, 2024a). Two monkeys made reaction-time decisions about random-dot motion while up to 203 LIP neurons were recorded simultaneously, in eight sessions with 9,684 trials for one monkey and 8,142 for the other; on half of the trials a 100 ms pulse of weak motion appeared at a random time and biased choices and reaction times by the equivalent of 1.4% coherence. The authors fitted a race between two negatively correlated accumulators (correlation -0.71) with a linear urgency signal. The trial counts per monkey are close to the reference design of section 5, but the allocation of trials across coherences and pulses differs, so power has to be recomputed by simulation for the actual design. We propose the following plan, fixed on one monkey and run once on the other, with the code published before the confirmation run.

1. Decoder. A linear readout trained on strong-motion trials of the development monkey, frozen, with two disjoint halves of the population; the readout is checked for linearity by comparing termination levels across coherences.
2. Diffusion. The stopped split-half cross moment on 0% coherence trials without pulses, against the stopped time, with the non-decision time taken from the neural termination signature; the slope is reported for an early and a late window.
3. Mobility. The sustained estimate from low-coherence trials and the pulse estimate from pulse trials, reported separately; their ratio is a diagnostic.
4. Bound. Half the difference between choice-conditioned levels before the saccade, overall and by quintile of decision time, since the fitted urgency implies a level that changes with time.
5. Prediction. The slope MB/D with a bootstrap interval over trials (Efron and Tibshirani 1993), compared with the slope fitted to held-out choices with a lapse rate (Wichmann and Hill 2001) by two one-sided tests with a margin of $\pm 25\%$, together with the late/early and pulse/sustained diagnostics.

The outcomes have distinct readings. Agreement within the margin with flat diagnostics supports bounded accumulation with the measured effective ratio. A prediction that is too low with a rising late/early ratio points to urgency, which the race fit already contains, and calls for a time-resolved prediction with the bound measured by decision time. A prediction that is too low with flat diagnostics points to a component of the readout that does not drive the choice, or to choice noise outside LIP. A prediction that is too high points to drift variability across trials or to readout noise that the split halves share. A manipulation of stimulus noise at fixed mean motion (Zylberberg et al. 2016) raises D without changing M and, if the bound stays fixed, should lower the slope in proportion, which tests the dependence on diffusion directly. Because the fitted model is a race, the relation of Proposition 1 holds for the difference of the accumulators only when their correlation is -1, so the plan has to reduce to a single variable by decoding that difference or replace Proposition 1 by the exit probability of the fitted race, computed numerically.

9. Limits

The propositions concern a one-dimensional decision variable; races of partially correlated accumulators and multidimensional representations need the exit probability of the fitted model, and the invariance of Proposition 2 holds only for linear readouts. The simulations use one reference parameter set, one population model with Poisson spiking and gains linear in the decision variable, and a decision variable that stays at the bound until the saccade; real LIP activity bursts before saccades and may not hold a constant level. The departures were applied one at a time, and combinations can cancel in the comparison. The estimators assume that observation noise is independent between the two halves of the population, which shared fluctuations of arousal or of other latent variables would violate; the aligned-nuisance model is one such case. The evidence unit, coherence, is a property of the stimulus, so the measured ratio is not comparable across tasks or modalities until evidence is expressed in a common unit such as a log likelihood ratio (Kira et al. 2015). Precision in active-inference models (Feldman and Friston 2010) could be compared with $1/\Theta$ in the same way, but whether those models make that identification is a question about the models (Aguilera et al. 2022; Isomura et al. 2023). No data were analyzed. The analysis plan applies to a public dataset (Steinemann et al. 2024a), and the same pipeline could run on recordings from other species and tasks, such as the mouse visual task of the International Brain Laboratory (The International Brain Laboratory 2021), once a stimulus variable with a defined evidence scale is available.

References

- Aguilera, Miguel, Beren Millidge, Alexander Tschantz, and Christopher L. Buckley. 2022. "How Particular Is the Physics of the Free Energy Principle?" *Physics of Life Reviews* 40: 24–50. <https://doi.org/10.1016/j.plrev.2021.11.001>.
- Arieli, Amos, Alexander Sterkin, Amiram Grinvald, and Ad Aertsen. 1996. "Dynamics of Ongoing Activity: Explanation of the Large Variability in Evoked Cortical Responses." *Science* 273 (5283): 1868–71. <https://doi.org/10.1126/science.273.5283.1868>.
- Baiesi, Marco, Christian Maes, and Bram Wynants. 2009. "Fluctuations and Response of Nonequilibrium States." *Physical Review Letters* 103 (1). <https://doi.org/10.1103/physrevlett.103.010602>.
- Bogacz, Rafal, Eric Brown, Jeff Moehlis, Philip Holmes, and Jonathan D. Cohen. 2006. "The Physics of Optimal Decision Making: A Formal Analysis of Models of Performance in Two-Alternative Forced-Choice Tasks." *Psychological Review* 113 (4): 700–765. <https://doi.org/10.1037/0033-295x.113.4.700>.
- Brunton, Bingni W., Matthew M. Botvinick, and Carlos D. Brody. 2013. "Rats and Humans Can Optimally Accumulate Evidence for Decision-Making." *Science* 340 (6128): 95–98. <https://doi.org/10.1126/science.1233912>.

- Callen, Herbert B., and Theodore A. Welton. 1951. "Irreversibility and Generalized Noise." *Physical Review* 83 (1): 34–40. <https://doi.org/10.1103/physrev.83.34>.
- Churchland, Anne K, Roozbeh Kiani, and Michael N Shadlen. 2008. "Decision-Making with Multiple Alternatives." *Nature Neuroscience* 11 (6): 693–702. <https://doi.org/10.1038/nn.2123>.
- Churchland, Anne. K., R. Kiani, R. Chaudhuri, Xiao-Jing Wang, Alexandre Pouget, and M.N. Shadlen. 2011. "Variance as a Signature of Neural Computations During Decision Making." *Neuron* 69 (4): 818–31. <https://doi.org/10.1016/j.neuron.2010.12.037>.
- Churchland, Mark M, Byron M Yu, John P Cunningham, et al. 2010. "Stimulus Onset Quenches Neural Variability: A Widespread Cortical Phenomenon." *Nature Neuroscience* 13 (3): 369–78. <https://doi.org/10.1038/nn.2501>.
- Cisek, Paul, Geneviève Aude Puskas, and Stephany El-Murr. 2009. "Decisions in Changing Conditions: The Urgency-Gating Model." *The Journal of Neuroscience* 29 (37): 11560–71. <https://doi.org/10.1523/jneurosci.1844-09.2009>.
- Cugliandolo, Leticia F. 2011. "The Effective Temperature." *Journal of Physics A: Mathematical and Theoretical* 44 (48): 483001. <https://doi.org/10.1088/1751-8113/44/48/483001>.
- Cugliandolo, Leticia F., Jorge Kurchan, and Luca Peliti. 1997. "Energy Flow, Partial Equilibration, and Effective Temperatures in Systems with Slow Dynamics." *Physical Review E* 55 (4): 3898–914. <https://doi.org/10.1103/physreve.55.3898>.
- Daw, Nathaniel D., John P. O'Doherty, Peter Dayan, Ben Seymour, and Raymond J. Dolan. 2006. "Cortical Substrates for Exploratory Decisions in Humans." *Nature* 441 (7095): 876–79. <https://doi.org/10.1038/nature04766>.
- Ditterich, Jochen. 2006. "Stochastic Models of Decisions about Motion Direction: Behavior and Physiology." *Neural Networks* 19 (8): 981–1012. <https://doi.org/10.1016/j.neunet.2006.05.042>.
- Drugowitsch, Jan, Rubén Moreno-Bote, Anne K. Churchland, Michael N. Shadlen, and Alexandre Pouget. 2012. "The Cost of Accumulating Evidence in Perceptual Decision Making." *The Journal of Neuroscience* 32 (11): 3612–28. <https://doi.org/10.1523/jneurosci.4010-11.2012>.
- Efron, Bradley, and R.J. Tibshirani. 1993. *An Introduction to the Bootstrap*. Chapman and Hall. <https://doi.org/10.1201/9780429246593>.
- Einstein, A. 1905. "Über Die von Der Molekularkinetischen Theorie Der Wärme Geforderte Bewegung von in Ruhenden Flüssigkeiten Suspendierten Teilchen." *Annalen Der Physik* 322 (8): 549–60. <https://doi.org/10.1002/andp.19053220806>.
- Feldman, Harriet, and Karl J. Friston. 2010. "Attention, Uncertainty, and Free-Energy." *Frontiers in Human Neuroscience* 4. <https://doi.org/10.3389/fnhum.2010.00215>.
- Friston, Karl. 2010. "The Free-Energy Principle: A Unified Brain Theory?" *Nature Reviews Neuroscience* 11 (2): 127–38. <https://doi.org/10.1038/nrn2787>.

- Gnesotto, F S, F Mura, J Gladrow, and C P Broedersz. 2018. "Broken Detailed Balance and Non-Equilibrium Dynamics in Living Systems: A Review." *Reports on Progress in Physics* 81 (6): 066601. <https://doi.org/10.1088/1361-6633/aab3ed>.
- Gold, Joshua I., and Michael N. Shadlen. 2001. "Neural Computations That Underlie Decisions about Sensory Stimuli." *Trends in Cognitive Sciences* 5 (1): 10–16. [https://doi.org/10.1016/s1364-6613\(00\)01567-9](https://doi.org/10.1016/s1364-6613(00)01567-9).
- Gold, Joshua I., and Michael N. Shadlen. 2007. "The Neural Basis of Decision Making." *Annual Review of Neuroscience* 30 (1): 535–74. <https://doi.org/10.1146/annurev.neuro.29.051605.113038>.
- Hanks, Timothy, Roozbeh Kiani, and Michael N Shadlen. 2014. "A Neural Mechanism of Speed-Accuracy Tradeoff in Macaque Area LIP." *eLife* 3. <https://doi.org/10.7554/elife.02260>.
- Harada, Takahiro, and Shin-ichi Sasa. 2005. "Equality Connecting Energy Dissipation with a Violation of the Fluctuation-Response Relation." *Physical Review Letters* 95 (13). <https://doi.org/10.1103/physrevlett.95.130602>.
- Huk, Alexander C., and Michael N. Shadlen. 2005. "Neural Activity in Macaque Parietal Cortex Reflects Temporal Integration of Visual Motion Signals During Perceptual Decision Making." *The Journal of Neuroscience* 25 (45): 10420–36. <https://doi.org/10.1523/jneurosci.4684-04.2005>.
- Isomura, Takuya, Kiyoshi Kotani, Yasuhiko Jimbo, and Karl J. Friston. 2023. "Experimental Validation of the Free-Energy Principle with in Vitro Neural Networks." *Nature Communications* 14 (1). <https://doi.org/10.1038/s41467-023-40141-z>.
- Kalman, R. E. 1960. "A New Approach to Linear Filtering and Prediction Problems." *Journal of Basic Engineering* 82 (1): 35–45. <https://doi.org/10.1115/1.3662552>.
- Kiani, Roozbeh, Christopher J. Cueva, John B. Reppas, and William T. Newsome. 2014. "Dynamics of Neural Population Responses in Prefrontal Cortex Indicate Changes of Mind on Single Trials." *Current Biology* 24 (13): 1542–47. <https://doi.org/10.1016/j.cub.2014.05.049>.
- Kiani, Roozbeh, Timothy D. Hanks, and Michael N. Shadlen. 2008. "Bounded Integration in Parietal Cortex Underlies Decisions Even When Viewing Duration Is Dictated by the Environment." *The Journal of Neuroscience* 28 (12): 3017–29. <https://doi.org/10.1523/jneurosci.4761-07.2008>.
- Kira, Shinichiro, Tianming Yang, and Michael N. Shadlen. 2015. "A Neural Implementation of Wald's Sequential Probability Ratio Test." *Neuron* 85 (4): 861–73. <https://doi.org/10.1016/j.neuron.2015.01.007>.
- Kramers, H.A. 1940. "Brownian Motion in a Field of Force and the Diffusion Model of Chemical Reactions." *Physica* 7 (4): 284–304. [https://doi.org/10.1016/s0031-8914\(40\)90098-2](https://doi.org/10.1016/s0031-8914(40)90098-2).
- Kubo, R. 1966. "The Fluctuation-Dissipation Theorem." *Reports on Progress in Physics* 29 (1): 255–84. <https://doi.org/10.1088/0034-4885/29/1/306>.

- Lakens, Daniël. 2017. "Equivalence Tests." *Social Psychological and Personality Science* 8 (4): 355–62. <https://doi.org/10.1177/1948550617697177>.
- Latimer, Kenneth W., Jacob L. Yates, Miriam L. R. Meister, Alexander C. Huk, and Jonathan W. Pillow. 2015. "Single-Trial Spike Trains in Parietal Cortex Reveal Discrete Steps During Decision-Making." *Science* 349 (6244): 184–87. <https://doi.org/10.1126/science.aaa4056>.
- Lynn, Christopher W., Eli J. Cornblath, Lia Papadopoulos, Maxwell A. Bertolero, and Danielle S. Bassett. 2021. "Broken Detailed Balance and Entropy Production in the Human Brain." *Proceedings of the National Academy of Sciences* 118 (47). <https://doi.org/10.1073/pnas.2109889118>.
- Martin, P., A. J. Hudspeth, and F. Jülicher. 2001. "Comparison of a Hair Bundle's Spontaneous Oscillations with Its Response to Mechanical Stimulation Reveals the Underlying Active Process." *Proceedings of the National Academy of Sciences* 98 (25): 14380–85. <https://doi.org/10.1073/pnas.251530598>.
- Mizuno, Daisuke, Catherine Tardin, C. F. Schmidt, and F. C. MacKintosh. 2007. "Nonequilibrium Mechanics of Active Cytoskeletal Networks." *Science* 315 (5810): 370–73. <https://doi.org/10.1126/science.1134404>.
- Okazawa, Gouki, Christina E. Hatch, Allan Mancoo, Christian K. Machens, and Roozbeh Kiani. 2021. "Representational Geometry of Perceptual Decisions in the Monkey Parietal Cortex." *Cell* 184 (14): 3748–3761.e18. <https://doi.org/10.1016/j.cell.2021.05.022>.
- Palmer, John, Alexander C. Huk, and Michael N. Shadlen. 2005. "The Effect of Stimulus Strength on the Speed and Accuracy of a Perceptual Decision." *Journal of Vision* 5 (5): 1. <https://doi.org/10.1167/5.5.1>.
- Peixoto, Diogo, Jessica R. Verhein, Roozbeh Kiani, et al. 2021. "Decoding and Perturbing Decision States in Real Time." *Nature* 591 (7851): 604–9. <https://doi.org/10.1038/s41586-020-03181-9>.
- Ratcliff, Roger. 1978. "A Theory of Memory Retrieval." *Psychological Review* 85 (2): 59–108. <https://doi.org/10.1037/0033-295x.85.2.59>.
- Ratcliff, Roger, and Gail McKoon. 2008. "The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks." *Neural Computation* 20 (4): 873–922. <https://doi.org/10.1162/neco.2008.12-06-420>.
- Redner, Sidney. 2001. *A Guide to First-Passage Processes*. Cambridge University Press. <https://doi.org/10.1017/cbo9780511606014>.
- Risken, Hannes. 1989. *The Fokker-Planck Equation*. 2nd ed. Springer. <https://doi.org/10.1007/978-3-642-61544-3>.
- Roitman, Jamie D., and Michael N. Shadlen. 2002. "Response of Neurons in the Lateral Intraparietal Area During a Combined Visual Discrimination Reaction Time Task." *The Journal of Neuroscience* 22 (21): 9475–89. <https://doi.org/10.1523/jneurosci.22-21-09475.2002>.

- Sarracino, A., O. Arviv, O. Shriki, and L. de Arcangelis. 2020. "Predicting Brain Evoked Response to External Stimuli from Temporal Correlations of Spontaneous Activity." *Physical Review Research* 2 (3). <https://doi.org/10.1103/physrevresearch.2.033355>.
- Schuirmann, Donald J. 1987. "A Comparison of the Two One-Sided Tests Procedure and the Power Approach for Assessing the Equivalence of Average Bioavailability." *Journal of Pharmacokinetics and Biopharmaceutics* 15 (6): 657–80. <https://doi.org/10.1007/bf01068419>.
- Seifert, Udo. 2012. "Stochastic Thermodynamics, Fluctuation Theorems and Molecular Machines." *Reports on Progress in Physics* 75 (12): 126001. <https://doi.org/10.1088/0034-4885/75/12/126001>.
- Shadlen, Michael N., and Roozbeh Kiani. 2013. "Decision Making as a Window on Cognition." *Neuron* 80 (3): 791–806. <https://doi.org/10.1016/j.neuron.2013.10.047>.
- Shadlen, Michael N., and William T. Newsome. 2001. "Neural Basis of a Perceptual Decision in the Parietal Cortex (Area LIP) of the Rhesus Monkey." *Journal of Neurophysiology* 86 (4): 1916–36. <https://doi.org/10.1152/jn.2001.86.4.1916>.
- Smith, Anne C., and Emery N. Brown. 2003. "Estimating a State-Space Model from Point Process Observations." *Neural Computation* 15 (5): 965–91. <https://doi.org/10.1162/089976603765202622>.
- Steinemann, Natalie, Gabriel M Stine, Eric Trautmann, Ariel Zylberberg, Daniel M Wolpert, and Michael N Shadlen. 2024a. "Data for 'Direct Observation of the Neural Computations Underlying a Single Decision'." Zenodo. <https://doi.org/10.5281/zenodo.13207505>.
- Steinemann, Natalie, Gabriel M Stine, Eric Trautmann, Ariel Zylberberg, Daniel M Wolpert, and Michael N Shadlen. 2024b. "Direct Observation of the Neural Computations Underlying a Single Decision." *eLife* 12: RP90859. <https://doi.org/10.7554/elife.90859>.
- Stine, Gabriel M., Eric M. Trautmann, Danique Jeurissen, and Michael N. Shadlen. 2023. "A Neural Mechanism for Terminating Decisions." *Neuron* 111 (16): 2601–2613.e5. <https://doi.org/10.1016/j.neuron.2023.05.028>.
- The International Brain Laboratory. 2021. "Standardized and Reproducible Measurement of Decision-Making in Mice." *eLife* 10. <https://doi.org/10.7554/elife.63711>.
- Wald, A. 1945. "Sequential Tests of Statistical Hypotheses." *The Annals of Mathematical Statistics* 16 (2): 117–86. <https://doi.org/10.1214/aoms/1177731118>.
- Wichmann, Felix A., and N. Jeremy Hill. 2001. "The Psychometric Function: I. Fitting, Sampling, and Goodness of Fit." *Perception & Psychophysics* 63 (8): 1293–313. <https://doi.org/10.3758/bf03194544>.
- Yang, Tianming, and Michael N. Shadlen. 2007. "Probabilistic Reasoning by Neurons." *Nature* 447 (7148): 1075–80. <https://doi.org/10.1038/nature05852>.

Zylberberg, Ariel, Christopher R Fetsch, and Michael N Shadlen. 2016. "The Influence of Evidence Volatility on Choice, Reaction Time and Confidence in a Perceptual Decision." *eLife* 5. <https://doi.org/10.7554/elife.17688>.