

Effective Landscapes from Time Series

A Reconstruction Protocol and a Hierarchy of What It Recovers

Victoria Atarah Devika

Working Paper — v1.0
August 2026

Abstract

Many fields fit “landscapes” to metastable time series: ecology, neuroscience, cell biology, finance, collective behaviour. The fitted object is rarely defined precisely, and the question of which of its features survive the analyst’s choices is rarely asked. We define the landscape as an empirical object—the effective potential $V_{\text{eff}} = -D \log P_{\text{ss}}$ obtained from the stationary density—state the conditions under which it coincides with a modeled potential and why it governs transition times even when it does not, and specify a six-step reconstruction pipeline. We then validate the assembled pipeline blind on simulated trajectories, recovering barrier heights to within two percent and predicting absolute dwell times to a median error of three percent with no fitted parameters. The central organizing result is a *hierarchy of recoverable quantities*, established by running the same data through deliberately different normalizations: the shape of the dwell-time distribution needs only a state label; barrier heights are invariant under any fixed linear change of coordinates; absolute rates are invariant under nothing, and the sign of their error flips between whitening (+20%) and z -scoring (−26%). Geometry is a property of the system; kinetics is a property of the system and the coordinates through which it is observed. Two aggregation artefacts are quantified and corrected. Pooling dwell times across units with different barrier heights gives $\text{CV}_{\text{pooled}}^2 = 1 + 2 \text{CV}_{\mu}^2$, so that a heterogeneous population of perfect Poisson escapers is rejected as non-exponential in 98.9% of samples; rescaling by each unit’s own mean restores the nominal rate. Shortening the analysis epoch to resolve a drifting landscape recovers exponentiality as intended, but the same operation manufactures it in a stationary control once epochs hold fewer than about ten mean dwell times. Applied to published resting-EEG measurements, the protocol yields a quantitative constraint: smooth barrier drift saturates at $\beta \approx 0.74$ and cannot account for the observed $\beta \approx 0.5$, which points to scale-free rather than bounded landscape motion.

Keywords: potential landscape, quasipotential, Boltzmann inversion, Kramers escape, metastability, dwell-time distributions, reconstruction, coordinate invariance, critical transitions

1 Introduction

A recurring move across the quantitative sciences is to look at a time series that lingers in a few regions of its state space, and to describe it as motion in a landscape: basins, barriers, ridges, escape. The move is made in ecology for regime shifts in lakes and rangelands, in cell biology for differentiation, in neuroimaging for resting-state dynamics (Watanabe et al., 2014; Ezaki et al., 2017), in finance for market states, and in collective behaviour for polarization and social tipping (Scheffer et al., 2009). The vocabulary transfers effortlessly. What does not transfer is a shared answer to three questions.

- **Which potential is being estimated?** The decomposition of a drift into gradient and non-gradient parts is not unique, so “the potential of the system” is not a well-defined object until a normalization is fixed.
- **What survives a change of coordinates?** Landscapes are almost always reconstructed on projections, embeddings, or standardized variables. Some reported quantities are invariant under those choices and some are not, and the literature does not distinguish them.
- **What is an artefact of pooling?** Distributional predictions are routinely tested on data aggregated across participants, cells, or lakes, whose parameters differ. Aggregation distorts distributional shape in a direction that mimics model failure.

This paper answers all three, and treats the third as the more consequential. Sections 2–4 define the estimand and the pipeline. Section 3 states the hierarchy that organizes the rest of the paper. Section 5 validates the pipeline blind against known ground truth, including a deliberate comparison of two normalizations. Section 6 treats dwell-time distributions, derives the pooling identity, and reports a new result on the use of analysis-epoch length as an instrument for landscape drift.

Nothing here is committed to a substrate. The worked system is a double well; the reader should substitute their own bistable observable.

2 The Landscape as an Empirical Object

2.1 Setup

Let S be a smooth manifold of dimension n whose coordinates are the measured variables. In most of what follows $S = \mathbb{R}^n$; when some coordinates are phases, $S = \mathbb{R}^{n-m} \times \mathbb{T}^m$ and the results below hold unless stated otherwise. The evolution on S is governed by the Itô stochastic differential equation

$$dx = [-\nabla V(x) + A(x)] dt + \sigma dW(t), \quad (1)$$

where W is a standard n -dimensional Brownian motion, $\sigma > 0$ is a constant noise amplitude, and $D = \sigma^2/2$. The potential V drives the system toward local minima; A is a non-gradient field representing rotational flows and oscillatory behaviour. The dynamics are overdamped: there is no inertial term. Two assumptions about the noise are made throughout.

- (N1) *Additive and state-independent.* Because σ does not depend on x , the Itô and Stratonovich interpretations of (1) coincide.
- (N2) *Isotropic in the coordinates of S .* The noise amplitude is a scalar. Isotropy is a property of a coordinate system, not of a system; §2.4 explains how coordinates are chosen so that (N2) holds and what happens when it fails.

The density $P(x, t)$ obeys the Fokker–Planck equation

$$\partial_t P = -\nabla \cdot [(-\nabla V + A) P] + D \nabla^2 P, \quad (2)$$

whose stationary solution P_{ss} , when it exists, carries the current $J_{ss} = (-\nabla V + A)P_{ss} - D\nabla P_{ss}$ with $\nabla \cdot J_{ss} = 0$. Conditions for existence, uniqueness and geometric ergodicity are collected in Appendix A; they matter for correctness but not for use.

2.2 Boltzmann inversion

For $A = 0$ the stationary solution of (2) is $P_{ss} \propto \exp(-V/D)$ (Gardiner, 2009): detailed balance holds and $J_{ss} = 0$, so the potential can be read off from the stationary distribution,

$$V(x) = -D \log P_{ss}(x) + \text{const.} \quad (3)$$

When $A \neq 0$, substituting $P = e^{-V/D}$ into the current gives $J_{ss} = A P_{ss}$ and

$$\nabla \cdot J_{ss} = P_{ss} \left(\nabla \cdot A - \frac{A \cdot \nabla V}{D} \right),$$

so $e^{-V/D}$ remains stationary if and only if $\nabla \cdot A = A \cdot \nabla V/D$; for a divergence-free A this requires $A \cdot \nabla V = 0$ everywhere, i.e. the non-gradient flow must be tangent to the level sets of V . Solenoidality alone is not sufficient. In general (3) applied to data returns the *effective potential*

$$V_{\text{eff}}(x) := -D \log P_{ss}(x) + \text{const.}, \quad (4)$$

which equals V exactly when the modeled non-gradient field is everywhere perpendicular to ∇V . Equation (4), not (3), is the estimand of this paper.

2.3 Which potential? Existence and uniqueness of the decomposition

Given any smooth drift f , a decomposition $f = -\nabla V + A$ with $\nabla \cdot A = 0$ *exists*: take V to solve $\nabla^2 V = -\nabla \cdot f$, which is solvable on $\mathbb{R}^n$ for every smooth right-hand side (Malgrange–Ehrenpreis; Hörmander, 1983, Ch. 10). But it is never *unique*: for any harmonic h , the pair $(V+h, A+\nabla h)$ is also admissible, since ∇h is simultaneously a gradient and divergence-free. “The gradient component of the drift” is not a well-defined object until a normalization is fixed. The difficulty is non-uniqueness, not existence.

We fix the normalization by the stationary density—the potential–flux decomposition of nonequilibrium dynamics (Wang, Xu & Wang, 2008):

$$V := V_{\text{eff}} = -D \log P_{ss}, \quad A := f + \nabla V_{\text{eff}} = J_{ss}/P_{ss}. \quad (5)$$

This decomposition is unique given the dynamics and the noise level, reduces to the modeled potential under detailed balance, and identifies the residual with the normalized stationary current. The residual satisfies $\nabla \cdot (P_{ss}A) = 0$ rather than $\nabla \cdot A = 0$; by §2.2 the two coincide exactly when $A \perp \nabla V_{\text{eff}}$. Ao’s (2004) construction $f = -(\mathbf{D} + \mathbf{Q})\nabla\phi$ yields the same potential ($\phi = V_{\text{eff}}/D$, $\mathbf{D} = D\mathbf{I}$ under (N2)) whenever the stationary current is orthogonal to ∇P_{ss} —as for linear dynamics with isotropic noise (Kwon, Ao & Thouless, 2005)—but not in general.¹

Two consequences follow. First, a modeler who writes down a potential and a non-gradient field is choosing a decomposition; the landscape recovered from data is V_{eff} , and the two agree if and only if the chosen A is perpendicular to ∇V . Second, the distinction is usually immaterial for what the landscape is used for. In the weak-noise limit $\lim_{D \rightarrow 0} D \log P_{ss}(x) = -W(x) + \text{const}$ locally uniformly, where W is the Freidlin–Wentzell quasipotential (Freidlin & Wentzell, 2012, Ch. 4 and 6), and W controls escape in the sense of logarithmic equivalence,

$$\tau \asymp \exp(\Delta W/D), \quad \Delta W = \min_{\partial B} W - W(x^*), \quad (6)$$

¹A tilted periodic potential on a circle carries a constant stationary current parallel to ∇P_{ss} wherever P_{ss} is non-uniform, so no antisymmetric $\mathbf{Q}$ with $\mathbf{D} = D\mathbf{I}$ reproduces it.

whether or not $A = 0$. V_{eff} is thus a finite-noise regularization of precisely the quantity that governs transitions, which is the reason for making it primary. Two caveats: the point on the basin boundary that controls escape is the minimizer of W (or V_{eff}) there, which need not be a saddle of a modeled V ; and the Eyring–Kramers prefactor takes a different form for irreversible dynamics (Bouchet & Reygner, 2016). Detailed balance is directly testable: if the trajectory statistics are time-reversal symmetric then $J_{\text{ss}} = 0$ and (3) is exact.

2.4 Coordinates, metric, and noise

(a) Metric correction. If (1) with $A = 0$ is the covariant Langevin equation on a Riemannian manifold (S, g) and P is the density with respect to Lebesgue measure in local coordinates—what a kernel density estimator returns—then $P = e^{-V/D} \sqrt{g}$ with $g = \det(g_{ij})$, and

$$V = -D \log P + \frac{D}{2} \log g + \text{const.} \quad (7)$$

Check: for an isotropic Gaussian in the plane, $V = r^2/2$, the density in polar coordinates is $P(r, \theta) \propto r e^{-r^2/2D}$, so $-D \log P = r^2/2 - D \log r$; since $g = r^2$, the correction $+\frac{D}{2} \log g = +D \log r$ restores V . Under a global linear change of coordinates g is constant and (7) reduces to (3) up to a constant. If P is defined with respect to Riemannian volume, no correction is needed.

(b) Anisotropic noise. Suppose that in raw coordinates $dx = f dt + \sqrt{2\mathbf{D}} dW$ with constant symmetric positive-definite $\mathbf{D}$. In whitened coordinates $y = \mathbf{D}^{-1/2}x$ the diffusion matrix is the identity; if $f = -\mathbf{D}\nabla_x U + \dots$, the drift in y is $-\nabla_y U$, the stationary density is $P_{\text{ss}} \propto e^{-U}$, and (3)–(5) apply with $D = 1$. The procedure is: estimate $\hat{\mathbf{D}} = \frac{1}{2\Delta t} \langle \Delta x \Delta x^\top \rangle$ from short-lag increments, whiten, then reconstruct. Standardizing by total variance is *not* equivalent: total variance mixes the landscape with the noise, so z -scoring yields a different, non-canonical effective potential. Section 5.2 measures the difference.

(c) State-dependent noise. A diffusion matrix $\mathbf{D}(x)$ that varies with state cannot be whitened globally; P_{ss} then contains a noise-induced drift term and the Itô–Stratonovich distinction becomes material. This case lies outside the present formulation.

(d) Projections. When the landscape is reconstructed on a reduced coordinate $z = \pi(x)$, $-D \log P_{\text{ss}}(z)$ is a potential of mean force: it contains an entropic contribution from the integrated-out dimensions, and barriers along a poorly chosen coordinate typically underestimate the barriers that govern dwell times (Best & Hummer, 2005). The reconstructed landscape is a landscape *in the chosen coordinates*.

(e) Stationarity. The inversion assumes an approximately stationary, ergodic epoch. Where the landscape parameters θ drift on a timescale τ_θ , a quasi-stationary approximation holds only for epochs short against τ_θ ; §6.3 turns this requirement into a measurement.

3 A Hierarchy of Recoverable Quantities

The qualifications of §2.4 are usually presented as a list of caveats. They are better read as a structure. Different quantities computed from the same reconstruction depend on the analyst’s choices to sharply different degrees, and the dependence can be measured (§5.2). Table 1 states the




Figure 1: A landscape in a two-dimensional state space. Contours show $V(x, y) = (x^2 - 1)^2 + y^2$, with two minima (dots) at $(\pm 1, 0)$ and a saddle (cross) at the origin. The overlaid trajectory is a numerical solution of (1) with $A = 0$ and $D = 0.3$ over $t \in [0, 90]$: long confinement within basins punctuated by noise-driven transitions through the saddle region.

resulting hierarchy; it organizes the remainder of this paper and, we suggest, should organize the reporting of any landscape analysis.

The practical value is that a reader can locate their own data. A long ecological record with forty annual observations supports Tier 0 and nothing above it. A twenty-minute continuous physiological recording supports Tier 1 and, with care, Tier 2. A high-frequency financial series supports all of Tier 2. Claims should be made at the tier the data supports, and a disconfirming result at Tier 0 cannot be rescued by an appeal to coordinate choice, because no coordinate choice was made.

The one assumption at Tier 0. It is tempting to describe Tier 0 as assumption-free. It is not. Tier 0 depends on the state assignment, and a state assignment made through a filtered observable inherits the filter’s memory. If states are labelled from wavelet power in a narrow band, the wavelet’s time constant enters the dwell-time statistics directly, and a dwell time comparable to that constant is not a measurement of the system. The assumption should be stated with the tier: *the state label must be recoverable without a filter whose time constant is comparable to the dwell times being measured*. This bears on §6.3, where the published comparison data are wavelet power envelopes.

4 The Reconstruction Pipeline

Whichever coordinates are used, the landscape is estimated in six steps.

1. Fix the coordinates and their dimensionality across all conditions to be compared.
2. Estimate the diffusion matrix from short-lag increments and whiten (§2.4(b)). Note that after whitening the leading principal component is not in general the reaction coordinate; select the projection axis by a bimodality criterion, not by variance (§5.2).

Table 1: What a landscape reconstruction recovers, by robustness. Each tier requires everything below it. The empirical basis for the Tier 1/Tier 2 boundary is §5.2.

Tier	Quantity	Requires	Invariant under
0	Shape of the dwell-time distribution	a state label and time-stamps	any coordinate change preserving the labelling
1	Barrier heights, $\Delta V_{\text{eff}} = D \log[P_{\text{ss}}(x^*)/P_{\text{ss}}(x_{\ddagger})]$; basin topology	a stationary density estimate	any fixed linear change of coordinates
2	Absolute rates and dwell-time predictions; curvatures	$\hat{\mathbf{D}}$, whitening, a defended reaction coordinate	nothing; the error changes sign with the normalization
3	The meaning of the coordinates	domain theory	—

3. Estimate P_{ss} on a stationary epoch by kernel density or Gaussian-mixture estimation, form $V_{\text{eff}} = -\log P_{\text{ss}}$, and correct curvatures and extremal densities for the estimator bandwidth h . To first order $\kappa \rightarrow \kappa(1+h^2\kappa)$ for an eigenvalue κ of $\text{Hess } V_{\text{eff}}$, with the matching factor $\prod_i (1+h^2\kappa_i)^{1/2}$ on the density at the extremum: kernel smoothing flattens peaks and fills saddles, and §5 finds the uncorrected bias in density ratios to be of order ten percent.
4. Locate the minima of V_{eff} and the index-one saddles between them, by a string or minimum-energy-path method on V_{eff} , or a watershed segmentation of the density.
5. Report the normalization and bandwidth used, since the reconstructed geometry depends on both, and report which tier of Table 1 the data support.
6. Before empirical deployment, run the assembled pipeline blind on synthetic surrogates with known ground truth.

Step 3 is the continuous-state analogue of the maximum-entropy energy-landscape analysis applied to binarized brain states by Watanabe et al. (2014) and Ezaki et al. (2017). Step 6 is the subject of the next section, and is the step most often omitted.

5 Blind Validation

5.1 Isotropic protocol and findings

The pipeline was applied, blind, to simulated trajectories of the double-well reference system of Appendix B: the reconstruction stage received sample paths and nothing else, and its output was compared with the potential that generated them.

Protocol. Seven conditions varied the barrier scale c in $V(x, y) = c(x^2 - 1)^2 + y^2 + \varepsilon x$ ($c \in \{0.7, 1.0, 1.3\}$), the noise level ($D \in \{0.30, 0.45\}$), and the tilt (one condition with $\varepsilon = 0.35$, whose two inequivalent wells contribute separate data points), giving eight (condition, well) pairs spanning $\Delta V/D$ from 1.6 to 4.3. Each condition comprised 64 replicas of duration $T = 1500$ (Euler–Maruyama, $dt = 0.005$). The first half of each recording was used for landscape estimation and the second half, only, for dwell times (hysteric milestone at $x = \pm 0.6$), so that the test is non-circular. Reconstruction followed §4: histogram kernel estimate of P_{ss} (bandwidth $h = 0.06$), $V_{\text{eff}} = -\log P_{\text{ss}}$, minima and connecting saddle located on the reconstructed surface, Hessians by

local quadratic fits with first-order bandwidth deconvolution, and $\hat{D}$ by Richardson extrapolation of short-lag increments.

Findings. Three results, in increasing order of strength (Table 2, Figure 2).

- *Reconstruction accuracy (Tier 1).* Recovered barriers agree with the generating potential to within 0.02 in units of ΔV , one to two percent across all conditions—but only after bandwidth deconvolution. The raw kernel estimate compresses density ratios by roughly ten percent at the deepest barrier; the first-order correction of step 3 removes almost all of it.
- *The regression form and its qualifier.* Regressing $\log \tau$ on $\log[\hat{P}_{ss}(x^*)/\hat{P}_{ss}(x_{\ddagger})]$ gives slope 0.79 (bootstrap 95% interval $[0.66, 0.92]$), not unity, *with reconstruction near-exact*. The discrepancy is the prefactor: across these conditions curvature covaries with barrier height, because the wells deepen by stiffening, so the slowly varying logarithm is anticorrelated with the log-ratio and depresses the fitted slope. Any exponent-only regression therefore carries a design requirement: hold curvatures approximately constant across the comparison set, or compute the slowly varying term from the reconstructed Hessians and subtract it before fitting.
- *Absolute prediction (Tier 2).* With curvatures and $\hat{D}$ estimated from data, the asymptotic Langer rate (10) predicts absolute dwell times to a median error of 13%—the accuracy floor of the asymptotic formula itself, as Table 3 documents against the exact quadrature. Replacing it by the exact mean-first-passage quadrature evaluated on the *reconstructed* marginal closes the gap: median error 2.7%, maximum 10%, with zero fitted parameters. On the reconstructed density the quadrature reads

$$\hat{\tau} = \frac{1}{\hat{D}} \int_{x^*}^{x_b} \frac{du}{\hat{p}(u)} \int_{-\infty}^u \hat{p}(v) dv,$$

in which $\hat{D}$ enters only as an overall time scale: writing $V/D = -\log p$ makes the exponential structure of the integrand a measured quantity rather than a modeled one. Here the x -marginal is an exact reaction coordinate because y decouples; on a general projection the quadrature inherits the potential-of-mean-force caveat of §2.4(d) and should be evaluated along the identified minimum-energy path.

The residual errors are consistent with the known small biases of the estimators— $\hat{D}$ high by two to five percent at this integration step, and the milestoning convention differing from minimum-to-minimum passage by a boundary-layer term—not with failure of the landscape description.

5.2 Anisotropic noise: the Tier 1 / Tier 2 boundary

The isotropic validation exercises every stage of the pipeline except whitening, since its noise is isotropic by construction. We therefore repeated it with a constant but anisotropic and correlated diffusion matrix (four-to-one anisotropy, off-diagonal coupling), on dynamics $dx = -\mathbf{D}\nabla U dt + \sqrt{2}\mathbf{D}^{1/2}dW$ for which $P_{ss} \propto e^{-U}$ exactly, so barriers remain known while dwell times become non-trivial functions of the diffusion. Two branches were run blind from identical data: whitening by the estimated $\hat{\mathbf{D}}$, as §2.4(b) prescribes, and z -scoring by per-coordinate standard deviation, the common alternative.

Three findings (Figure 3). First, a technical trap with practical consequences: after whitening, the leading principal component is not in general the reaction coordinate. Whitening equalizes noise, not geometry, and in these systems it stretched the transverse mode until it dominated the variance; the first attempt at reconstruction failed completely for this reason. Selecting the projection axis

Table 2: Blind reconstruction versus ground truth across the eight (condition, well) pairs. n is the number of dwell events in the held-out epoch; $\Delta\hat{V}_{\text{eff}}$ is the reconstructed barrier (true value beside it); $\hat{\tau}$ is the absolute dwell-time prediction from the first-passage quadrature on the reconstructed marginal. The error column is $\tau/\hat{\tau} - 1$.

c	ε	D	n	$\Delta\hat{V}_{\text{eff}}/\Delta V$	τ (obs)	$\hat{\tau}$ (quad)	error
0.7	0	0.45	5715	0.692 / 0.700	8.19	7.79	+5.1%
1.0	0	0.45	4171	1.007 / 1.000	11.18	10.88	+2.8%
0.7	0	0.30	2562	0.687 / 0.700	17.89	17.42	+2.7%
1.0	0	0.30	1356	0.981 / 1.000	32.21	31.70	+1.6%
1.3	0	0.30	641	1.289 / 1.300	61.51	64.89	−5.2%
1.3	0	0.45	2653	1.292 / 1.300	17.29	15.68	+10.2%
1.0	0.35	0.35 (L)	640	1.353 / 1.373	54.86	53.52	+2.5%
1.0	0.35	0.35 (R)	684	0.662 / 0.673	8.93	9.04	−1.2%

by a bimodality criterion repairs it. Any implementation that reaches for PCA immediately after whitening should expect this.

Second, **barrier recovery was insensitive to the normalization**: both branches recovered ΔV_{eff} to within 0.04 in natural units. Third, **rate prediction was not**, and the two errors differ in sign and in kind: +20% for whitening, an approximation error of the projection that can in principle be reduced; −26% for z -scoring, an error of a different kind, since the discarded off-diagonal diffusion information cannot be recovered from the standardized coordinates at all.

This is the empirical content of Table 1, and it is worth stating as a claim rather than as a caveat:

The topology of the landscape is a property of the system. Its kinetics is a property of the system and of the coordinates through which it is observed.

A barrier is a difference of V_{eff} between critical points, and any fixed linear change of coordinates leaves it invariant. A rate is not, and the direction of the resulting bias depends on which normalization was chosen. Reports that give barriers without stating a normalization are incomplete; reports that give rates without stating one are uninterpretable.

5.3 A design constraint on block-structured recordings

A consequence of the validation constrains study design rather than analysis. The absolute-rate prediction requires enough basin transitions per recording to estimate a mean dwell time, of order fifty. Published recordings of alpha mode switching that yield enough events to fit a dwell-time distribution run to fourteen minutes or more of continuous acquisition per participant (Freyer et al., 2009), which places mean dwell times in the range of seconds and event counts per participant in the tens. A protocol segmented into one-minute blocks, with a further split between estimation and test epochs, will therefore deliver transition counts of order ten per participant per condition: enough to detect a state, marginal for estimating a rate, insufficient for a per-unit Tier 2 test. On such data the estimand must be moved to the population level, or the recordings must be continuous and long. Distributional tests at Tier 0 are the appropriate instrument in this regime, since they pool across units by design—which is exactly why the next section is necessary.




Figure 2: Blind validation. (a) Effective potential reconstructed from trajectory data alone for the tilted condition ($c = 1$, $\varepsilon = 0.35$, $D = 0.35$); dots and cross mark the minima and saddle located on the reconstructed surface. (b) Observed dwell times (held-out epochs) against absolute predictions from the blind reconstruction. Squares: asymptotic Langer rate with reconstructed curvatures (median error 13%). Circles: exact first-passage quadrature on the reconstructed marginal (median error 2.7%). No fitted parameters in either case.

6 Tier 0: Dwell-Time Distributions and Two Aggregation Artefacts

6.1 Exponentiality

Escape from a basin by noise activation is, in the high-barrier limit, a Poisson process: the probability of leaving in the next instant does not depend on how long the system has already been there, and dwell times are exponentially distributed. This is the sharpest available test of the landscape description, because it requires no reconstruction, no $\hat{D}$, no whitening, no bandwidth deconvolution and no defence of a projection. A heavy-tailed dwell-time distribution—power-law, log-normal, or stretched-exponential—is evidence against landscape diffusion that cannot be absorbed by a change of coordinates.

Two qualifications belong to the test itself. Exponentiality is asymptotic in $\Delta V_{\text{eff}}/D$, so it should be applied to the tail above a stated cutoff, discarding short dwells that reflect recrossing of the state boundary rather than genuine escape. And it does not apply to limit-cycle attractors, whose exit times are not activation-limited. Both must be stated in advance; invoked afterwards they are unfalsifiable rescues.

6.2 The aggregation identity

Dwell times are almost always pooled across units, and a mixture of exponentials with different rates is not exponential. If $T \mid \lambda \sim \text{Exp}(\lambda)$ and the unit mean dwell $\mu = 1/\lambda$ varies across a population, then

$$\text{CV}_{\text{pooled}}^2 = 1 + 2\text{CV}_{\mu}^2, \quad (8)$$

so a population whose mean dwell times vary with coefficient of variation 0.5 yields a pooled coefficient of variation of 1.22 even when every unit is a perfect Poisson escaper. Simulation confirms (8) to three significant figures and shows the consequence to be severe rather than marginal: at




Figure 3: Whitening under anisotropic, correlated noise. (a) Effective potential reconstructed blind in whitened, bimodality-aligned coordinates. (b) Observed against predicted dwell times for the whitened branch (circles) and the z -scored branch (crosses). Barrier heights are recovered accurately by both; rates are not, and the errors have opposite signs.

that level of heterogeneity, an Anderson–Darling test of exponentiality applied to naively pooled durations rejects the null in 98.9% of samples drawn from a population that satisfies the prediction exactly (Figure 4a). Barrier heights differ between units; heterogeneity of this magnitude is the expected case, not a pathological one.

The correction is to rescale each unit’s dwell times by that unit’s own mean before pooling, which removes the rate heterogeneity identically and restores the nominal error rate across the full range examined (3.9–6.5% against a nominal 5%). The cost is a loss of sensitivity to alternatives distinguished chiefly by overall dispersion when the per-unit means are themselves poorly estimated. At a scale typical of an open cohort—two hundred units contributing about five dwell events each—the rescaled test has power 98% against a stretched-exponential alternative (Weibull, $k = 0.7$) and 100% against a light-tailed one (gamma, $k = 2$), but only 5% against a log-normal alternative of comparable dispersion; ten events per unit restores the latter to 98% (Figure 4b). A null result must therefore be reported against a stated power, since at small per-unit counts the test is genuinely blind to one family of alternatives.

Equation (8) is not specific to any substrate. Any framework whose predictions concern distributional shape, and whose tests are conducted on populations aggregated across units with differing parameters, faces the same distortion. Stating the aggregation correction alongside the prediction is part of what it means for a method to be usable by another field.

6.3 Epoch length as an instrument, and a second artefact

The motivating observation. The test has been performed on real data. Freyer et al. (2009) recorded eyes-closed resting EEG from sixteen participants for fourteen to thirty minutes, estimated wavelet power at each participant’s peak alpha frequency at an occipital channel, and found clearly bimodal power distributions in thirteen of them. They then examined the dwell-time distributions of each mode, rescaled to their means in the manner prescribed above and computed separately for each participant, so that (8) does not apply. The distributions were not power-law, but neither were




Figure 4: The distributional test and its aggregation correction. (a) Rejection rate of an Anderson–Darling test of exponentiality applied to data in which *every* unit is a perfect Poisson escaper, as a function of across-unit heterogeneity in mean dwell time. Naive pooling (red) manufactures apparent falsification; rescaling each unit’s dwell times by their own mean (blue) holds the nominal rate (dotted line). (b) Power of the rescaled test against heavy- and light-tailed alternatives, as a function of cohort size and dwell events per unit.

they exponential: they were well described by a stretched exponential $\exp(-at^\beta)$, in which $\beta = 1$ recovers the simple exponential. Across the thirteen participants β averaged 0.49 for the low-power mode (range 0.40–0.60) and 0.68 for the high-power mode (range 0.50–0.90).

A stretched exponential is the canonical signature of a superposition of exponential processes with a distribution of rates, and (8) is the same phenomenon in its second moment. A barrier that drifts over a twenty-minute recording generates precisely this mixture, which suggests a measurement rather than a caveat: *estimate β as a function of the length of the analysis epoch*. If the deviation is a mixture over a drifting landscape, $\beta \rightarrow 1$ as the epoch shortens toward the timescale on which θ is effectively constant, and the rate of that approach measures τ_θ . If the deviation is a property of the escape process itself, β is flat in epoch length.

Probe. We tested this directly. One-dimensional double well, $D = 0.35$, barrier $c(t) = c_0[1 + a \sin(2\pi t/\tau_\theta)]$ with $c_0 = 1$ and $\tau_\theta = 800$; 250 replicas, $T = 5000$, hysteretic milestoneing at ± 0.6 ; dwell times assigned to windows of length T_{ep} , rescaled by the window mean, pooled, and β fitted by maximum likelihood. A control run with $a = 0$ passed through the identical windowing pipeline.

Result 1: the instrument works. With drift, β climbs from 0.861 at $T_{\text{ep}} = 4000$ through 0.894, 0.981 and 1.045 as the epoch shortens to 120. Without drift, β is flat at 1.016 from $T_{\text{ep}} = 4000$ down to 250 (Figure 5a). The measurement discriminates, and it is calibratable: β crosses 0.95 near $T_{\text{ep}} \approx 250$ for $\tau_\theta = 800$, i.e. at roughly $\tau_\theta/3$, so the epoch length at which β recovers is a usable estimator of the drift timescale.

Result 2: a censoring floor. In the *no-drift* control, β also rises once epochs fall below roughly five mean dwell times: $1.016 \rightarrow 1.035$ at $T_{\text{ep}} = 120$ and 1.143 at $T_{\text{ep}} = 60$. Requiring a dwell to fit inside a window truncates the tail and manufactures exactly the signal being sought. This

is a second aggregation artefact, structurally identical to (8) in that the analyst’s pooling choice generates the appearance of the effect under test, and it must travel with the measurement in the same way. **Rule: epochs must hold at least ten mean dwell times, and $\beta(T_{\text{ep}})$ must be reported against a matched stationary control.** A $\beta(T_{\text{ep}})$ curve alone is uninterpretable.

Result 3: smooth drift saturates. Scanning the drift amplitude at whole-recording scale gives $\beta = 1.014, 0.858, 0.793, 0.759, 0.742$ for $a = 0, 0.5, 0.7, 0.85, 0.95$ (Figure 5b). Even driving the barrier from 0.05 to 1.95—opening and closing it almost completely—does not bring β below about 0.74. The mechanism has a ceiling: mixing over a bounded, smoothly varying rate range produces bounded overdispersion, because the low-barrier phase contributes fast escapes that do not populate the tail.

Consequence. The observed $\beta \approx 0.49$ in the low-power alpha mode lies outside what smooth barrier drift can produce. The explanation “the landscape drifted and stationarity was violated” is therefore quantitatively insufficient on its own, and three readings remain.

1. *The high-power mode is a limit cycle.* Freyer et al. (2011) attribute the switching to a subcritical Hopf bifurcation in a corticothalamic model. If the high mode is a suprathreshold oscillation rather than a second fixed point, activated escape does not describe it, and the exclusion stated in advance applies. This reading does not cover the low mode.
2. *The state assignment is filtered.* The observable is a wavelet power envelope, whose time constant enters the dwell statistics directly (§3). This is a Tier 0 assumption violation, not a Tier 1 or Tier 2 one, and it is testable by varying the wavelet width.
3. *The landscape drift is scale-free.* If $\theta(t)$ itself has $1/f$ structure—entirely plausible for arousal or vigilance over twenty minutes—the mixture is over a scale-free rate distribution and β can fall well below the bounded-drift ceiling.

The third reading is the most consequential and makes the sharpest prediction, because it differs from bounded drift in the *shape* of the $\beta(T_{\text{ep}})$ curve rather than in its endpoint:

Bounded drift $\Rightarrow \beta(T_{\text{ep}})$ **plateaus** at unity below τ_θ .

Scale-free drift $\Rightarrow \beta(T_{\text{ep}})$ rises **without a plateau**, because there is no timescale to separate.

Figure 5a supplies the bounded-drift template and the stationary null against which an empirical curve is to be read. If the plateau is absent, the description “diffusion in a landscape that happens to be slowly varying” fails qualitatively rather than by a correction term, and the object being measured is not a landscape plus drift but something with no characteristic timescale at all.

Limitations of the probe. One drift law (sinusoidal, single τ_θ), one potential, one noise level, one milestone threshold. The saturation at $\beta \approx 0.74$ is suggestive, not proof; a different drift law, and in particular a heavy-tailed one, is exactly the alternative it is offered against. Threshold choice affects short dwells and hence β , and should be swept in any application.

6.4 Detecting drift directly: a surrogate null for the barrier

The $\beta(T_{\text{ep}})$ instrument diagnoses drift through the dwell-time distribution. The barrier can also be interrogated directly, and doing so exposes a third artefact of the same family as (8) and the censoring floor.




Figure 5: Analysis-epoch length as an instrument for landscape drift (§6.3). (a) β against epoch length in units of the mean dwell time, for a drifting barrier ($a = 0.5$, $\tau_\theta = 800$, dashed line) and a stationary control. The control is flat until epochs fall below about ten mean dwells, where window censoring manufactures the same signal (shaded). (b) β at whole-recording scale against drift amplitude. Smooth drift saturates near 0.74 and cannot reach the range observed in resting alpha (shaded band, Freyer et al., 2009).

The naive test compares ΔV_{eff} reconstructed on the first and second halves of a recording. It is blind to oscillating drift: a barrier that completes several cycles within the recording averages to the same value in both halves. Scanning window lengths repairs that blindness and introduces a second problem, because short windows carry few samples, the density estimate becomes noisy, and estimator variance inflates the between-window spread exactly as drift would. In our stationary reference condition a scan down to sixty-four windows reports a spread of nineteen percent with no drift present at all.

The correction is a surrogate null. For each window length, the observed spread is compared with the spread obtained from the same data, the same window length and the same estimator after permuting contiguous blocks an order of magnitude shorter than the window. Block permutation destroys variation slower than a window while preserving local dynamics, so the surrogate carries the estimator’s own noise and nothing else. Drift is claimed only from the excess. On the stationary condition the excess vanishes; on a barrier driven at $a = 0.7$, $\tau_\theta = 600$ the excess reaches thirty-eight percent, and the window length at which it peaks estimates τ_θ . The general lesson is the one this section keeps meeting: every diagnostic of non-stationarity in this framework must be read against a matched stationary control computed through the identical pipeline, because each stage of the pipeline generates, on stationary data, a weak version of the signal the diagnostic is looking for.

7 Discussion

What travels. The pipeline, the hierarchy, and the two aggregation corrections are substrate-free. Any field with a metastable observable and a state label can run the Tier 0 test today; the entry cost is a list of dwell durations. Fields with continuous, long recordings can add Tier 1, and the barrier estimates they obtain are comparable across laboratories provided the normalization is reported. Tier 2 is worth the additional assumptions only where the absolute timescale is itself the object of interest.

What does not. Tier 3 is where the disciplines differ and where this paper is silent. The mapping from measured coordinates to whatever the coordinates are supposed to mean is a domain question, and no amount of estimator care substitutes for it. Two reconstructions of the same system on different observables are two landscapes, both real, related by the projection algebra of §2.4(d).

Limitations. State-dependent noise is outside the formulation. The rate results assume a fixed landscape; where the landscape drifts fast relative to escape, the mechanism of basin change is drift rather than activation and Kramers’ formula does not describe it—§6.3 measures this regime but does not model it. The correction (8) is specific to exponential dwell times; analogous corrections for other distributional predictions are not worked out here. Finally, the validations are on a single family of double-well systems, and higher-dimensional landscapes with many competing basins will exercise the saddle-location step of the pipeline far harder than anything reported above.

Software. The pipeline is distributed as a small package rather than described only in prose, because a protocol that must be reimplemented from a methods section is not portable in any useful sense. The entry point takes a time series and a sampling interval and returns the reconstructed barriers in both conventions, the absolute dwell-time prediction, the dwell-time exponent with a bootstrap interval, and—the part that matters—the tier the data actually support, with every refusal logged and its reason stated. It declines to report a barrier on a unimodal density, declines to report an exponent when no two-state structure was found, declines Tier 2 under z -scoring or when the surrogate test detects drift, and declines a heavy-tail verdict when the bootstrap interval on β includes unity. Nothing is dropped silently. On the reference condition of §5 it recovers $\Delta V = 1.004$ and 0.953 against a true value of unity.

Reproducibility. The validation and probe scripts accompany this paper, with fixed seeds and a single output root. All numbers quoted in the text and tables are produced by them.

A Well-posedness, Ergodicity, and the Deterministic Limit

Assumptions. (A1) $V \in C^2(S)$ is a Morse function (all critical points non-degenerate; Milnor, 1963) and is confining: $V(x) \rightarrow +\infty$ as $|x| \rightarrow \infty$ along the Euclidean factor of S . (A2) A is locally Lipschitz; when a modeled decomposition is used, A is taken divergence-free. (A3) The full drift is dissipative: there exist $a, b > 0$ with $\langle x, -\nabla V(x) + A(x) \rangle \leq a - b|x|^2$ for all x .

Under (A1)–(A3) and (N1), solutions of (1) exist, are unique, and do not explode in finite time; a stationary density P_{ss} exists, is unique, strictly positive, and as regular as the drift permits; and the law of $x(t)$ converges to P_{ss} geometrically fast in total variation (Khasminskii, 2012, Ch. 3–4; Mattingly, Stuart & Higham, 2002; Meyn & Tweedie, 1993).

Remark (On Lipschitz assumptions). *A globally Lipschitz drift is incompatible with (A1): a confining potential of super-quadratic growth—such as the quartic double well of Appendix B, whose gradient $4x(x^2 - 1)$ is only locally Lipschitz—cannot have a globally Lipschitz gradient. Condition (A3) is the standard replacement, and must be verified for the full drift, since a large non-gradient field can defeat the confinement provided by V .*

Deterministic limit. Set $\sigma = 0$. For a pure gradient system, $\dot{V} = -|\nabla V|^2 \leq 0$ with equality only at critical points: V is a strict Lyapunov function, so periodic orbits are impossible in any dimension and gradient systems admit only fixed-point attractors. By LaSalle’s invariance principle

(LaSalle, 1960) every trajectory converges to the critical set of V ; under (A1) the critical points are isolated and hyperbolic, so each trajectory converges to a single critical point.² A critical point of Morse index $k \geq 1$ has a stable manifold of dimension $n - k$, of Lebesgue measure zero; hence for almost every initial condition the trajectory converges to a local minimum. Basins are open; their boundaries are closures of unions of stable manifolds of critical points of index ≥ 1 , whose codimension-one part consists generically of the stable manifolds of index-one saddles (Palis & de Melo, 1982). When $A \neq 0$ the system can exhibit limit cycles and strange attractors, to which the escape-rate results of this paper do not apply.

B The Double-Well Reference System

The system used throughout for validation is

$$V(x, y) = c(x^2 - 1)^2 + \mu y^2 + \lambda xy + \varepsilon x, \quad (9)$$

with dynamics (1). For $c = \mu = 1$ and $\lambda = \varepsilon = 0$ it has minima at $(\pm 1, 0)$, a saddle at the origin, and barrier $\Delta V = 1$. The Hessian at the minima is diagonal with eigenvalues $12x^2 - 4 = 8$ and $2\mu = 2$, so $\det \text{Hess } V|_{\min} = 16$; at the saddle the eigenvalues are -4 and 2 , so $|\det \text{Hess } V|_{\ddagger}| = 8$.

Escape rate. In the high-barrier limit $\Delta V \gg D$ the transition rate is given by the multidimensional overdamped Kramers formula (Kramers, 1940; Langer, 1969; Hänggi, Talkner & Borkovec, 1990),

$$k = \tau^{-1} = \frac{|\lambda_{\ddagger}^{\dagger}|}{2\pi} \sqrt{\frac{\det \text{Hess } V|_{\min}}{|\det \text{Hess } V|_{\ddagger}|}} e^{-\Delta V/D}, \quad (10)$$

where $\lambda_{\ddagger}^{\dagger}$ is the negative eigenvalue at the saddle. For the symmetric well with $\mu = 1$ this gives $k = (2\sqrt{2}/\pi)e^{-\Delta V/D} \approx 0.900 e^{-1/D}$: the y -curvatures at minimum and saddle are equal and cancel, so μ affects neither barrier nor rate. Because (9) is separable for $\lambda = 0$, the mean first-passage time can be computed exactly by the quadrature $T = \frac{1}{D} \int_{-1}^1 e^{V(u)/D} \int_{-\infty}^u e^{-V(v)/D} dv du$ (Gardiner, 2009). Table 3 compares the asymptotic formula, the exact quadrature, and direct simulation; the asymptotic rate is accurate to about 10% at $\Delta V/D = 5$ and degrades gracefully, with the expected $O(D/\Delta V)$ correction. This is the accuracy floor referred to in §5.

Table 3: Mean escape time from one well of (9) ($c = \mu = 1$, $\lambda = \varepsilon = 0$, $\Delta V = 1$), in natural time units. “Langer” is $1/k$ from (10); “exact” is the one-dimensional first-passage quadrature; “simulation” is the mean first-passage time from $(-1, 0)$ to $x = +1$ over 2000 Euler–Maruyama trajectories with step 0.004 (mean $\pm$ standard error).

D	$\Delta V/D$	Langer	exact	simulation
0.20	5.00	164.8	182.4	189.6 ± 4.2
0.25	4.00	60.6	69.3	69.2 ± 1.5
0.30	3.33	31.1	36.6	37.2 ± 0.8
0.40	2.50	13.5	16.6	16.8 ± 0.4

²For non-Morse V , convergence to the critical *set* is all LaSalle’s principle provides; convergence to a single point can fail for smooth V (Palis & de Melo, 1982), although it holds for real-analytic V by the Łojasiewicz gradient inequality (Łojasiewicz, 1984).

Parameter dependence. For $\lambda \neq 0$ the minima move to $x^2 = 1 + \lambda^2/8\mu$, $y = -\lambda x/2\mu$; the wells remain symmetric, the saddle stays at the origin, and the barrier becomes $\Delta V = (1 + \lambda^2/8\mu)^2$. For $\varepsilon \neq 0$ the tilt deepens one well relative to the other by approximately 2ε ; at $|\varepsilon| = 8/(3\sqrt{3}) \approx 1.54$ the shallower well merges with the saddle in a saddle–node bifurcation. Approaching that point, the smallest Hessian eigenvalue at the disappearing minimum tends to zero and the intra-basin relaxation time $1/\lambda_{\min}$ diverges, the signature of critical slowing down. This intra-basin slowing is distinct from the inter-basin escape time, which *shortens* as $\Delta V \rightarrow 0$; near the bifurcation the system becomes sluggish within its basin while transitions out of it become more likely. Equation (10) is not valid there: its derivation requires $\Delta V \gg D$, its prefactor vanishes as $\lambda_{\min} \rightarrow 0$, and (A1) fails at the degenerate critical point. At the saddle–node point the mean exit time scales as $D^{-1/3}$, by the scaling $x \propto D^{1/3}$, $t \propto D^{-1/3}$ of the normal form $dx = x^2 dt + \sqrt{2D} dW$.




Figure 6: (a) The one-dimensional double well $V(x) = (x^2 - 1)^2$: two minima at $x = \pm 1$ separated by a barrier of height $\Delta V = 1$. (b) The two-dimensional extension $V(x, y) = (x^2 - 1)^2 + \mu y^2$ with $\mu = 1$, $\lambda = \varepsilon = 0$.

References

- Anderson, T.W. & Darling, D.A. (1954). A test of goodness of fit. *Journal of the American Statistical Association*, 49(268), 765–769.
- Ao, P. (2004). Potential in stochastic differential equations: novel construction. *Journal of Physics A*, 37(3), L25–L30.
- Best, R.B. & Hummer, G. (2005). Reaction coordinates and rates from transition paths. *PNAS*, 102(19), 6732–6737.
- Bouchet, F. & Reygner, J. (2016). Generalisation of the Eyring–Kramers transition rate formula to irreversible diffusion processes. *Annales Henri Poincaré*, 17(12), 3499–3532.
- Ezaki, T., Watanabe, T., Ohzeki, M. & Masuda, N. (2017). Energy landscape analysis of neuroimaging data. *Philosophical Transactions of the Royal Society A*, 375, 20160287.
- Freidlin, M.I. & Wentzell, A.D. (2012). *Random Perturbations of Dynamical Systems* (3rd ed.). Springer.
- Freyer, F., Aquino, K., Robinson, P.A., Ritter, P. & Breakspear, M. (2009). Bistability and non-Gaussian fluctuations in spontaneous cortical activity. *Journal of Neuroscience*, 29(26), 8512–8524.

- Freyer, F., Roberts, J.A., Becker, R., Robinson, P.A., Ritter, P. & Breakspear, M. (2011). Biophysical mechanisms of multistability in resting-state cortical rhythms. *Journal of Neuroscience*, 31(17), 6353–6361.
- Gardiner, C.W. (2009). *Stochastic Methods: A Handbook for the Natural and Social Sciences* (4th ed.). Springer.
- Hänggi, P., Talkner, P. & Borkovec, M. (1990). Reaction-rate theory: fifty years after Kramers. *Reviews of Modern Physics*, 62(2), 251–341.
- Hörmander, L. (1983). *The Analysis of Linear Partial Differential Operators II*. Springer.
- Kantz, H. & Schreiber, T. (2004). *Nonlinear Time Series Analysis*. Cambridge University Press.
- Khasminskii, R. (2012). *Stochastic Stability of Differential Equations* (2nd ed.). Springer.
- Kramers, H.A. (1940). Brownian motion in a field of force and the diffusion model of chemical reactions. *Physica*, 7(4), 284–304.
- Kwon, C., Ao, P. & Thouless, D.J. (2005). Structure of stochastic dynamics near fixed points. *PNAS*, 102(37), 13029–13033.
- Langer, J.S. (1969). Statistical theory of the decay of metastable states. *Annals of Physics*, 54(2), 258–275.
- LaSalle, J.P. (1960). Some extensions of Liapunov’s second method. *IRE Transactions on Circuit Theory*, 7(4), 520–527.
- Łojasiewicz, S. (1984). Sur les trajectoires du gradient d’une fonction analytique. *Seminari di Geometria 1982–1983*, Università di Bologna, 115–117.
- Mattingly, J.C., Stuart, A.M. & Higham, D.J. (2002). Ergodicity for SDEs and approximations: locally Lipschitz vector fields and degenerate noise. *Stochastic Processes and their Applications*, 101(2), 185–232.
- Meyn, S.P. & Tweedie, R.L. (1993). *Markov Chains and Stochastic Stability*. Springer.
- Milnor, J. (1963). *Morse Theory*. Princeton University Press.
- Palis, J. & de Melo, W. (1982). *Geometric Theory of Dynamical Systems: An Introduction*. Springer.
- Scheffer, M., Bascompte, J., Brock, W.A. et al. (2009). Early-warning signals for critical transitions. *Nature*, 461, 53–59.
- Strogatz, S.H. (1994). *Nonlinear Dynamics and Chaos*. Addison-Wesley.
- Wang, J., Xu, L. & Wang, E. (2008). Potential landscape and flux framework of nonequilibrium networks. *PNAS*, 105(34), 12271–12276.
- Watanabe, T. et al. (2014). Energy landscapes of resting-state brain networks. *Frontiers in Neuroinformatics*, 8, 12.