

Nested Attractor Landscapes

A Multi-Scale Dynamical Framework for Conscious State Dynamics

Victoria Atarah Devika

Working Paper — v4.0

August 2026

Abstract

We present a dynamical framework in which conscious experience is modeled as a trajectory through a nested attractor landscape. Rather than treating consciousness as a discrete output of neural processes, we describe it as a continuous path through a state space shaped by stable and metastable attractors, with individual, social and environmental dynamics coupled across scales. The framework’s estimation machinery—how a landscape is defined as an empirical object, reconstructed from data, and reported—is developed and validated separately (Devika, 2026a); this paper takes that protocol as given and asks what follows for a theory of conscious dynamics. Three things do. First, features of subjective experience map onto measurable landscape geometry, and the mapping is stated at the level of robustness the estimator actually supports rather than uniformly. Second, the coupling between individual and collective scales has a computable synchronization threshold: the mean-field model undergoes a Desai–Zwanzig transition at a critical coupling $K_c(D)$ determined by the variance of the single-agent stationary density, so “collective attractor” is a claim with a number attached. Third, the framework yields three predictions that discriminate it from neighbouring accounts, and we withdraw two predictions of earlier drafts that do not. The sharpest test concerns the distribution of dwell times within a state, which no coordinate choice can adjust; published resting-alpha measurements contradict its unconditional form, and we state the conditional form that survives, together with the quantitative reason—smooth landscape drift saturates well short of the observed exponent—that the obvious rescue is insufficient.

Keywords: attractor dynamics, consciousness, metastability, nested systems, collective dynamics, phase transitions, dwell times, falsifiability

1 Introduction

1.1 The problem

Most existing frameworks for consciousness either reduce it to a single mechanism—neural correlates, information integration, quantum collapse—or remain abstract enough that no measurement bears on them. Physicalism maps brain states to conscious states but does not explain why those states feel like anything. Dualism preserves the distinctness of experience but offers no mechanism of interaction. Panpsychism distributes experience universally but does not explain how micro-level constituents combine.

What is missing is not another theory of what consciousness *is*, but a formal language for how it *behaves*: how it moves, stabilizes, transitions and reorganizes across time and across scales.

1.2 The proposal

Consciousness is modeled as a trajectory through a nested attractor landscape. The claim is three-fold.

- **Experience is a trajectory, not a state.** It corresponds to the continuous path of a system through a state space, not to any point in it.
- **Attractor structure constrains the space of possible experiences.** The state space is organized into basins that guide and stabilize trajectories into characteristic patterns.
- **These landscapes are nested.** An individual landscape is embedded in collective and environmental dynamics, each scale constraining and being influenced by its neighbours.

The framework is substrate-independent: it describes dynamics, not implementation.

1.3 What this paper contributes, and what it defers

The measurement apparatus on which everything here depends—the definition of the landscape as the effective potential $V_{\text{eff}} = -D \log P_{\text{ss}}$, the conditions under which that coincides with a modeled potential, the reconstruction pipeline, its blind validation, and the hierarchy of which reconstructed quantities survive which analytic choices—is developed in a companion methods paper (Devika, 2026a) and is not repeated. That separation is deliberate: the protocol is substrate-free and useful to fields with no interest in consciousness, and the theory below should be judged on what it adds to the protocol rather than on the protocol itself.

What this paper contributes is four things. A geometric vocabulary for experiential quality, stated at the robustness tier each item actually occupies (§3). A multi-scale architecture with an explicit coupling and a computable synchronization threshold (§4). A reading of therapeutic modalities as distinct operations on one landscape, with the timescale condition under which the static-landscape formalism does and does not apply (§2.3, §8). And three falsifiable predictions, together with the withdrawal of two that earlier drafts should not have made (§6).

1.4 Relation to existing work

The formalism of attractors, basins and bifurcations is inherited from dynamical systems theory (Strogatz, 1994; Abraham & Shaw, 1992); the role of neural attractors in memory and cognitive state dynamics from computational neuroscience (Hopfield, 1982; Amit, 1989; Freeman, 2000; Churchland et al., 2012), together with evidence for neural criticality (Beggs & Plenz, 2003; Cocchi et al., 2017); the reaction-rate and large-deviation formalism from the theory of rare events (Kramers, 1940; Freidlin & Wentzell, 2012); and energy-landscape analysis of brain states from neuroimaging (Watanabe et al., 2014; Ezaki et al., 2017). Among theories of consciousness, this framework is not a competitor but a complement: it describes the dynamics of state change and is silent on what makes a state conscious, which is the question Integrated Information Theory (Tononi, 2004) and Global Workspace Theory (Baars, 1988) address. It draws on embodied cognition (Thompson, 2007) and enactivism (Varela, Thompson & Rosch, 1991) for the treatment of somatic and relational coordinates as constitutive rather than peripheral.




Figure 1: A two-dimensional landscape with x read as valence and y as arousal. Contours show $V(x, y) = (x^2 - 1)^2 + y^2$; dots mark the two minima, the cross the saddle. The overlaid stochastic trajectory shows long confinement within a basin punctuated by noise-driven transitions—the geometry of a state that persists and then, without external cause, does not.

2 The Framework

2.1 State space and landscape

Let S be a smooth manifold of dimension n whose coordinates specify the state of the system. For a neural implementation these might include firing rates, neurotransmitter concentrations, connectivity patterns, and physiological variables such as heart-rate variability and autonomic state. A trajectory $\gamma(t)$ is a continuous path through S ; experience corresponds to the unfolding of $\gamma(t)$, not to any point in it.

Under noise, deterministic attractors are not invariant sets but *metastable states*: regions in which the stationary density concentrates and in which trajectories dwell for long periods before escaping. In this framework they represent stable or metastable patterns of experience—habitual emotional states, cognitive modes, identity configurations. The basin $B(\mathcal{A})$ is the set of initial conditions converging to $\mathcal{A}$; the landscape L is the complete topology of attractors and basins, encoding the repertoire of available states and the transitions between them.

The single-agent dynamics are the Itô equation

$$dx = [-\nabla V(x) + A(x)] dt + \sigma dW(t), \quad D = \sigma^2/2, \quad (1)$$

with V a confining potential and A a non-gradient field carrying rotational and oscillatory dynamics. The landscape as an *empirical* object is the effective potential $V_{\text{eff}} = -D \log P_{\text{ss}}$ obtained from the stationary density; this coincides with V when $A \perp \nabla V$, and in general is a finite-noise regularization of the Freidlin–Wentzell quasipotential, which governs escape times whether or not $A = 0$ (Devika, 2026a, §2).

2.2 The central model

The nested dynamics are a pair of relations:

$$dx_i = [-\nabla V_i(x_i) + A_i(x_i) + \alpha_i F(X, x_i)] dt + \sigma_i dW_i, \quad i = 1, \dots, N, \quad (2)$$

$$X(t) = G(\{x_j(t)\}_{j=1}^N). \quad (3)$$

Equation (2) governs individual dynamics, with $\alpha_i F(X, x_i)$ the top-down coupling from the collective state; (3) defines the collective state as an aggregate of individual states. We call the pair the Nested Attractor Dynamics equations, and this class of systems Nested Attractor Dynamical Systems (NADS).

Remark (On the collective equation). *Earlier drafts wrote the collective relation as a differential equation, $\dot{X} = G(\{x_j\}) + \sigma_X \eta_X$. With G an aggregation map such as the mean, that makes X the time integral of the aggregate rather than the aggregate itself: X grows without bound whenever the population mean is non-zero, and a coupling $F = \beta(X - x_i)$ then pulls every individual toward a runaway variable. Equation (3) is the intended relation. A collective variable with its own memory and exposure to environmental shocks is obtained by relaxing it to $dX = -\kappa[X - G(\{x_j\})] dt + \sigma_X dW_X$, which recovers (3) as $\kappa \rightarrow \infty$ and differs from it through the collective memory time $1/\kappa$ and the collective noise σ_X . Under (3) the fluctuations of X are inherited entirely from the individuals.*

Well-posedness of the joint system requires more than well-posedness of its parts: a coupling of linear growth contributes a term of order $|\alpha_i|c|x|^2$ that can overwhelm individual confinement. The attractive mean-field coupling of §4.3 is safe, since $\sum_i \langle x_i, \beta(X - x_i) \rangle = \beta(N|X|^2 - \sum_i |x_i|^2) \leq 0$ by Cauchy–Schwarz. Under joint dissipativity a unique stationary density exists on the product space; at fixed X each individual has a stationary distribution, which is what quasi-stationary reconstruction uses.

2.3 When the landscape moves

The landscape is not fixed. It is reshaped by experience, learning, perturbation and deliberate intervention, which we implement by allowing $V = V(x; \theta(t))$ with parameters θ evolving on a timescale τ_θ . Whether the static-landscape results apply is not a matter of taste; it is decided by comparing τ_θ against two internal timescales, and the two comparisons have different consequences.

The first is the intra-basin relaxation time $\tau_{\text{relax}} \sim 1/\lambda_{\min}$, where $\lambda_{\min}$ is the smallest eigenvalue of Hess V at the minimum. When $\tau_\theta \gg \tau_{\text{relax}}$ the within-basin distribution tracks θ quasi-statically and the geometry is well defined at each instant.

The second is the inter-basin escape time $\tau_{\text{esc}} \sim \exp(\Delta V/D)$. When $\tau_\theta \lesssim \tau_{\text{esc}}$ the landscape reshapes as fast as, or faster than, noise-driven transitions, and the system crosses basin boundaries *through landscape drift rather than stochastic escape*—a qualitatively different mechanism that Kramers’ formula does not describe. Since τ_{esc} grows exponentially with barrier height, this condition is most demanding precisely for the deep, stable attractors that matter most clinically.

Two consequences follow, and they are the reason this section is not an afterthought. First, therapeutic interventions almost certainly operate in the non-adiabatic inter-basin regime: they reshape the landscape faster than spontaneous escape would occur, which is the mechanism by which they work at all. Waiting for a deep depressive basin to be exited by noise is not a treatment. Second, τ_θ is measurable rather than postulated. Because a landscape drifting under a stationary analysis window produces a mixture of escape rates, and a mixture of exponentials is a stretched exponential, the exponent β estimated as a function of analysis-epoch length is an instrument for

Table 1: Landscape geometry, its experiential reading, cross-domain reading, and the robustness tier at which each quantity can be recovered (Devika, 2026a, §3).

Geometry	Experiential reading	Elsewhere	Estimated as	Tier
Barrier height ΔV_{eff}	persistence of a state; difficulty of leaving it	resilience of an ecological regime; stickiness of a market state	$D \log[P_{\text{ss}}(x^*)/P_{\text{ss}}(x_{\ddagger})]$	1
Basin topology	which states are adjacent in experience	reachable configurations	minima and index-one saddles of V_{eff}	1
Saddle location	critical junctures where small perturbations have large effects	tipping points	minimizer of V_{eff} on the basin boundary	1
Curvature at a minimum	how sharply a state resists small perturbation; recovery rate	relaxation to equilibrium	$\text{Hess } V_{\text{eff}}(x^*)$, with band-width deconvolution	2
Flatness	openness, flexibility, “beginner’s mind”	loss of regime structure	small eigenvalues of $\text{Hess } V_{\text{eff}}$ throughout	2
Dwell-time law	whether a state is left by accumulation or by chance	activated versus driven transitions	distribution of durations between state labels	0

τ_θ (Devika, 2026a, §6.3). Plasticity thus enters the framework as a quantity with a measurement procedure attached, not as a hedge.

3 Landscape Geometry and Experiential Quality

The central interpretive claim is that landscape geometry corresponds to qualitative features of experience. Table 1 states the correspondence, and adds two columns that earlier drafts lacked. The first gives the reading in other domains, which is how the substrate-independence claim is cashed out rather than merely asserted. The second gives the robustness tier of Devika (2026a): whether the quantity is recoverable from a state label alone (Tier 0), from a density estimate and invariant under any fixed linear change of coordinates (Tier 1), or dependent on the diffusion estimate and the choice of normalization (Tier 2). Claims should be made at the tier the data support, and the table makes visible that the framework’s most vivid interpretive claims sit at its least robust tier.

Of the entries in Table 1, barrier height carries the most interpretive weight, and it is the one that has been measured against a known answer. In blind reconstruction from simulated trajectories—the generating potential withheld from the estimation stage, the dwell times taken from held-out epochs—recovered barriers agreed with truth to within one to two percent, and the first-passage quadrature evaluated on the reconstructed density predicted *absolute* dwell times to a median error of 2.7% with no fitted parameters (Devika, 2026a, §5). The data were synthetic, and that qualification is not a small one: a synthetic system satisfies by construction the stationarity and noise assumptions that empirical recordings must be argued into. But the claim that basin depth predicts persistence is not, at least, an aspiration attached to a metaphor. It is a quantity that has been shown to predict, on data where the prediction was free to fail.

On accessible volume. Earlier drafts proposed $\varsigma = \sqrt{\det \Sigma_{\text{traj}}}$ as a measure of the effective range of available states, and predicted it would be larger in positive than in negative affect. We no longer make that prediction (§6). The quantity remains a reasonable descriptive statistic, with two

conditions: ς is an exact volume only for Gaussian trajectory distributions and otherwise an upper bound, so it is a valid monotone proxy only when non-Gaussianity is comparable across compared conditions; and by the law of total covariance $\Sigma_{\text{traj}} = \Sigma_W + \Sigma_B$, over which the determinant is not additive, so within- and between-basin components must be reported separately.

4 Multi-Scale Nesting

4.1 The individual scale

The individual landscape is shaped by neural architecture, developmental history, learned patterns and current physiological state. Four families of attractor are useful to distinguish. *Cognitive* attractors are stable patterns of thought, belief and reasoning style; at the network level the alternation between integrated and segregated configurations during task performance (Shine et al., 2016) is a candidate instance of metastable switching between them. *Emotional* attractors are stable or recurring affective states; depression is naturally modeled as a deep, wide basin with high barriers (Huys, Daw & Dayan, 2015). *Identity* attractors are the configurations constituting a sense of self. *Somatic* attractors are stable patterns of bodily state, interacting bidirectionally with the others—the reason state spaces built from neural variables alone are, on this account, underspecified.

4.2 The collective scale

Individual landscapes are coupled to collective ones. Three converging lines of evidence support the coupling as more than metaphor. Hyperscanning studies show inter-brain synchronization predicting empathy, cooperation and communicative success (Hasson et al., 2012; Dikker et al., 2017), and cardiac and respiratory rhythms synchronize between persons during joint activity such as choir singing (Müller & Lindenberger, 2011). Large-scale network studies show affective states propagating through social ties (Fowler & Christakis, 2008; Dodds et al., 2011; Kramer, Guillory & Hancock, 2014), and threshold models explain how individually small influences aggregate into collective shifts (Granovetter, 1978). Studies of collective behaviour reveal dynamical signatures consistent with approach to bifurcation (Scheffer et al., 2009).

4.3 Coupling and the synchronization threshold

The simplest concrete specification of the coupling is mean-field:

$$F(X, x_i) = \beta (X - x_i), \quad G(\{x_j\}) = \frac{1}{N} \sum_{j=1}^N x_j, \quad (4)$$

so each individual is pulled toward the collective mean with effective strength $K_i := \alpha_i \beta$. Note that α_i and β enter only through their product and are separately identifiable only when one is fixed by design.

Consider a homogeneous population ($V_i = V$, $K_i = K$, $\sigma_i = \sigma$, $A_i = 0$) as $N \rightarrow \infty$. The collective state converges to a deterministic m (the McKean–Vlasov limit), each individual evolves in the effective potential $V(x) + \frac{K}{2}|x - m|^2$, and the stationary single-individual density is $P_m(x) \propto \exp\{-[V(x) + \frac{K}{2}|x - m|^2]/D\}$. Self-consistency requires

$$m = \Phi(m) := \int x P_m(x) dx. \quad (5)$$




Figure 2: Nested dynamics under mean-field coupling (4), with $V(x) = (x^2 - 1)^2$, $D = 0.5$, $N = 500$. (a) The collective state $X(t)$ below threshold ($K = 0.4K_c$) and above it ($K = 2K_c$): below, X fluctuates about zero; above, the population settles into a shared attractor. (b) The stationary collective state m^* from (5) against the effective coupling $K = \alpha\beta$, with the pitchfork at $K_c \approx 0.74$ given by (6); dots mark time-averaged $|X|$ from the simulations in (a).

For symmetric V , $m = 0$ always solves (5). Differentiating under the integral, the Jacobian of Φ is $(K/D) \text{Cov}_{P_m}(x)$, and the disordered solution loses stability when its largest eigenvalue reaches one. In one dimension:

$$\frac{K}{D} \text{Var}_{P_0}(x) = 1, \quad P_0(x) \propto \exp\{-[V(x) + \frac{K}{2}x^2]/D\}. \quad (6)$$

Equation (6) defines a critical coupling $K_c(D)$. Below it the collective distribution is unimodal and individuals explore their personal landscapes independently; above it, (5) acquires a pair of ordered solutions $\pm m^*$ and the population synchronizes into a shared attractor. This is the Desai–Zwanzig transition (Desai & Zwanzig, 1978; Dawson, 1983), the diffusive counterpart of the Kuramoto transition for coupled phase oscillators (Strogatz, 2000). For the double well at $D = 0.5$, (6) gives $K_c \approx 0.74$ (Figure 2). At finite N , below threshold X fluctuates about zero with amplitude $O(N^{-1/2})$; above threshold it dwells near $\pm m^*$, and collective reversals occur on timescales exponentially long in N .

This is what distinguishes the nesting claim from a metaphor. “A group can have a shared attractor” is, in this framework, the statement that $K > K_c(D)$, with K_c computable from the individual landscape and the noise level, and with a predicted finite-size signature. It also supplies a formal response to the combination problem: the ordered state m^* exists only above a coupling threshold and is a property of the population rather than of any individual, so macro-level structure that is not the sum of micro-level dynamics is exhibited rather than postulated.

5 Transitions, Perturbation, and Coherence

Most moment-to-moment experience consists of trajectories within a single basin. Fixed-point attractors produce convergent, low-dimensional experience such as focused concentration; limit cycles produce periodic dynamics such as rumination or circadian variation; strange attractors produce complex, non-repeating but bounded dynamics.

Transitions between basins correspond to qualitative shifts: mood changes, perspective shifts, insight, identity transformation. Four mechanisms are distinguishable, and they differ in what would be observed. *Perturbation* pushes the trajectory across a boundary. *Basin erosion* gradually reduces the depth of the current attractor. *Landscape flattening* temporarily reduces overall structure, allowing freer exploration—the dynamics of neural annealing (Johnson, 2019). *Bifurcation* changes the number or arrangement of attractors outright. The first and third leave different signatures in the dwell-time statistics than the second and fourth, which is the basis of P7 below.

What is experientially described as *coherence*—alignment, flow, clarity—corresponds on this account to the alignment of attractors across subsystems and scales. When cognitive, emotional, somatic and relational attractors are mutually reinforcing, the joint landscape has a deep minimum at their intersection and the trajectory is stable. Incoherence is attractor conflict: the joint potential has no deep minimum where the subsystem basins meet, forcing the trajectory into regions of high potential or high curvature. Psychophysiological accounts of coherence in terms of cross-system synchrony (McCraty & Childre, 2010) are, on this reading, descriptions of the same alignment.

6 Predictions

Earlier drafts of this framework offered seven predictions. Three of them discriminate; two are consistency checks shared with neighbouring accounts; two do not work, and are withdrawn here. Stating this is not a concession. A framework that offers seven predictions of mixed quality is harder to test than one that offers three of known quality, and the withdrawal identifies exactly where the framework’s reach ends.

P1. Basin depth predicts state persistence. For a state whose dynamics are fixed-point dominated, as identified from the autocorrelation structure of the trajectory, the mean dwell time and the reconstructed landscape satisfy

$$\tau \propto \exp(\Delta V_{\text{eff}}/D) = \frac{P_{\text{ss}}(x^*)}{P_{\text{ss}}(x_{\ddagger})}, \quad (7)$$

where x^* is the minimum and $x_{\ddagger}$ the minimizer of V_{eff} on the basin boundary. The noise level cancels from the exponent, so the prediction relates two observables and requires no estimate of D . It tests whether the dynamics are diffusion in a landscape at all. For non-circularity the density ratio must come from a baseline epoch and τ from separate epochs.

Two qualifications are load-bearing. The regression form—that $\log \tau$ is linear in the log density ratio with unit slope—holds only up to the slowly varying logarithm of the Kramers prefactor, and in blind synthetic validation with near-exact reconstruction the fitted slope was 0.79, because curvature covaried with barrier height across the condition set (Devika, 2026a, §5). The test therefore carries a design requirement: hold curvatures approximately constant, or subtract the prefactor term computed from the reconstructed Hessians. And the prediction applies for $\Delta V_{\text{eff}} \gg D$ and not to limit-cycle attractors.

P4. Collective transitions show critical signatures. As K approaches K_c from below in (6), the collective state should show increased variance, slower recovery from perturbation, and flickering between configurations (Scheffer et al., 2009). This prediction is shared with the entire critical-transitions literature and does not by itself distinguish this framework; what is specific here is the *location* of the threshold, since (6) computes K_c from the individual landscape and noise level rather than fitting it. A measured K_c inconsistent with the individually reconstructed $\text{Var}_{P_0}(x)$ would be a failure of the nesting claim specifically.

P7. Dwell times approach exponentiality as the analysis epoch approaches stationarity.

Escape from a basin by noise activation is, in the high-barrier limit, a Poisson process, so dwell times are exponentially distributed. This is the framework’s sharpest test because it sits at Tier 0: it requires no landscape reconstruction, no estimate of D , no whitening, no bandwidth deconvolution and no defence of a projection, only an assignment of the trajectory to states. It is therefore untouched by the coordinate-dependence qualifications that can otherwise absorb a disconfirming result.

The prediction is stated conditionally because its unconditional form has been tested and fails. Freyer et al. (2009) found dwell times in resting alpha to be stretched-exponential with β averaging 0.49 for the low-power mode and 0.68 for the high, computed per participant so that population aggregation is not the explanation. The conditional form is:

The departure from exponentiality is an artefact of aggregating over a drifting landscape, and must therefore shrink as the analysis epoch shrinks toward the timescale on which θ is constant.

Two constraints travel with the test and cannot be omitted. Pooling across participants with different barrier heights gives $CV_{\text{pooled}}^2 = 1 + 2 CV_{\mu}^2$, which rejects a population of perfect Poisson escapers in 98.9% of samples at realistic heterogeneity unless each participant’s durations are rescaled by their own mean. And shortening the epoch manufactures the same recovery in a stationary control once epochs hold fewer than about ten mean dwell times, so $\beta(T_{\text{ep}})$ must be read against a matched control (Devika, 2026a, §6).

The rescue is quantitatively insufficient. It would be convenient to attribute $\beta \approx 0.5$ to landscape drift and declare the framework intact. Simulation forbids it: a smoothly drifting barrier saturates at $\beta \approx 0.74$ even when driven from 0.05 to 1.95, because the low-barrier phase contributes fast escapes that do not populate the tail (Devika, 2026a, §6.3). Three readings survive. The high-power mode may be a limit cycle, as Freyer et al. (2011) argue from a subcritical Hopf mechanism, in which case P7 excludes it by its own terms—but this does not cover the low mode. The state assignment is made through a wavelet power envelope whose time constant enters the dwell statistics directly, which is a violation of the one assumption Tier 0 does carry. Or the landscape drift is scale-free rather than bounded, in which case $\beta(T_{\text{ep}})$ rises without a plateau and the description “diffusion in a slowly varying landscape” fails qualitatively rather than by a correction. The three are distinguishable by the shape of the $\beta(T_{\text{ep}})$ curve and by varying the wavelet width, and no new data are required.

C1, C2: consistency checks. Two further expectations follow from the framework but are shared with other accounts, and we label them accordingly rather than advertising them as tests. *C1*: therapeutic interventions producing lasting change should be preceded by increased neural entropy or state-space exploration, followed by consolidation. *C2*: states of subjective coherence should correspond to alignment between measurable subsystems—increased zero-lag cross-correlation for fixed-point dynamics, increased phase locking for oscillatory dynamics.

Withdrawn. Two predictions of earlier drafts are withdrawn. The first asserted that computational implementations of nested attractor landscapes would exhibit behaviours paralleling biological consciousness dynamics; as stated it has no failure condition and is not a prediction. The second asserted that positive affect corresponds to expanded and negative affect to contracted accessible state-space volume; the framework’s own decomposition defeats it, since a negative state cycling

between two narrow attractors can yield a larger ς than a positive state confined to one broad basin, inverting the predicted direction. A prediction whose sign the framework cannot fix is not testable, and retaining it would have made the framework look more falsifiable than it is.

7 Falsifiability and Scope

7.1 What would falsify this framework

- Evidence that conscious state transitions are fundamentally discontinuous.
- Evidence that neural dynamics during stable states do not exhibit metastable behaviour.
- Evidence that dwell times are unrelated to reconstructed stationary-density ratios (P1): that persistence is not governed by a landscape.
- Evidence that the stretched-exponential exponent β does not approach unity as the analysis epoch is shortened toward stationarity, measured against a matched stationary control (P7). This is the most direct criterion and the least susceptible to rescue: it involves no reconstruction and no coordinate choice. A β independent of epoch length would count against the framework rather than against a particular implementation of it.
- Evidence that a measured collective threshold is inconsistent with the K_c computed from individually reconstructed landscapes (P4).
- Evidence that cross-scale coupling between individual and collective dynamics is negligible.

7.2 What the framework claims and does not claim

The positive claim is a formal, substrate-independent, multi-scale account of the *structure and dynamics* of conscious experience. It does not claim to solve the hard problem. It addresses the “landscape problem” rather than the “ground problem”: it maps the terrain without claiming to explain why there is terrain. It is agnostic about physical substrate and remains compatible with accounts that extend consciousness beyond conventional neural implementation. Just as thermodynamics was productive for a century before statistical mechanics supplied a micro-level explanation, a formal dynamics of experience can be productive regardless of whether the hard problem is solved—and if consciousness proves to be a fundamental feature of reality, attractor landscapes describe how that feature organizes at different scales rather than competing with the claim.

8 Applications

Mental health. The framework reframes conditions as landscape pathologies: depression as a deep, stable attractor with high barriers; anxiety as a limit-cycle attractor with rapid oscillation; trauma as a catastrophic deformation, in the technical sense of a sudden change in the number and arrangement of attractors (Thom, 1975). Interventions become landscape operations: medication may lower barriers, psychotherapy may reshape basin geometry, psychedelic-assisted therapy may temporarily flatten the landscape. The double-well parameters supply a vocabulary— ε -type interventions shift the balance between affective poles, λ -type interventions modify the coupling between valence and arousal. As §2.3 argues, therapeutic reshaping operates in the non-adiabatic regime relative to escape times; the vocabulary remains descriptively useful there, but the static-landscape rate formulas do not apply, and any quantitative claim about a treated landscape must be made on epochs short against τ_θ .

Collective dynamics. The multi-scale structure offers a formalism for social movements, cultural shifts and polarization (Castellano, Fortunato & Loreto, 2009; Sornette, 2003). Polarization is the deepening of two opposing basins with a growing barrier between them; sudden collective reversals of the kind analysed by Kuran (1991) are the finite- N collective flips of §4.3, occurring on timescales exponentially long in N and therefore appearing without proximate cause.

Neurotechnology and computational implementation. Rather than targeting regions or transmitter systems, interventions could be designed to reshape the landscape directly: lowering barriers between therapeutic target states, flattening maladaptive basins, or creating new basins through structured stimulation. Recurrent networks, reservoir architectures and agent-based models can instantiate nested landscapes computationally; the mean-field model of §4.3 is the simplest such implementation and already exhibits a collective phase transition.

9 Discussion

9.1 Open questions

What is the minimum dimensionality required to capture the essential features of conscious dynamics? How should the coupling functions F and G be estimated beyond the mean-field case, on a measured social network? Can the framework accommodate the binding problem, or does it require supplementary mechanisms? Can the double-well model reproduce specific clinical phenomena such as depressive relapse? And—sharpest of the open questions, because it has an experiment attached—does $\beta(T_{\text{ep}})$ in resting cortical rhythms plateau or not?

9.2 Limitations

- **Coordinate dependence.** The mapping from phenomenological variables to state-space coordinates is not uniquely specified, and a landscape reconstructed on a projection is a potential of mean force whose barriers depend on the coordinates chosen. Table 1 localizes the damage by tier, but does not remove it; every reported landscape must state its coordinates and normalization.
- **Static landscapes.** The rate results hold for a fixed landscape. Therapeutic reshaping is non-adiabatic and requires a treatment of the bifurcation structure of $V(x; \theta(t))$ that this paper does not provide.
- **Aggregation.** Distributional predictions tested on pooled populations are distorted by parameter heterogeneity in a way that mimics model failure. The correction is known for exponential dwell times; analogues for the framework’s other distributional claims are not worked out.
- **Mean-field idealization.** The collective model is homogeneous and all-to-all coupled. Heterogeneous α_i , network structure, and multiple collective scales are natural but unexecuted extensions.
- **Interpretive reach.** The correspondences in Table 1 are proposals, not results. Nothing in the formalism establishes that flatness is openness rather than merely low curvature.

10 Conclusion

This paper has presented a framework in which conscious experience is modeled as a trajectory through nested attractor landscapes: a formal, substrate-independent language for the structure, dynamics and multi-scale organization of experience. It defines the landscape as an empirical

object, states the assumptions under which it can be reconstructed, derives a collective threshold from individual dynamics, generates three predictions that can remove it, and offers applications in mental health, collective dynamics and neurotechnology.

More than a descriptive vocabulary, the framework offers a formal bridge between the mathematics of dynamical systems and the phenomenology of lived experience. Barrier height is not a metaphor for persistence; it is a ratio of stationary densities with predictive power. Landscape flattening is not a metaphor for openness; it is a dynamical operation with quantifiable effects on transition rates. Collective synchronization is not a metaphor for shared experience; it is a phase transition with a computable threshold.

And the framework can be wrong in a way it cannot argue with. The distribution of dwell times within a state is an observable that no reconstruction choice can adjust, and published measurements of that distribution in resting alpha do not match the simplest version of what this framework predicts. The convenient rescue—that the mismatch is an artefact of a drifting landscape—has been simulated rather than asserted, and it does not reach far enough. So P7 is stated not as a claim that the mismatch is absent but as a claim about how it must behave as the analysis epoch shortens. If it does not behave that way, this framework is describing alpha bistability with a formalism its own data contradict.

The nesting reflects a simple but consequential observation: no conscious system exists in isolation. Individual experience is shaped by collective dynamics, which are shaped by environmental and planetary-scale processes. Whether this nesting extends to a cosmic or ultimate scale is a question that lies beyond the scope of this paper—but the formalism is designed to accommodate it if the evidence warrants.

References

- Abraham, R.H. & Shaw, C.D. (1992). *Dynamics: The Geometry of Behavior*. Addison-Wesley.
- Amit, D.J. (1989). *Modeling Brain Function: The World of Attractor Neural Networks*. Cambridge University Press.
- Baars, B.J. (1988). *A Cognitive Theory of Consciousness*. Cambridge University Press.
- Beggs, J.M. & Plenz, D. (2003). Neuronal avalanches in neocortical circuits. *Journal of Neuroscience*, 23(35), 11167–11177.
- Castellano, C., Fortunato, S. & Loreto, V. (2009). Statistical physics of social dynamics. *Reviews of Modern Physics*, 81(2), 591–646.
- Chalmers, D.J. (1996). *The Conscious Mind*. Oxford University Press.
- Churchland, M.M. et al. (2012). Neural population dynamics during reaching. *Nature*, 487, 51–56.
- Cocchi, L., Gollo, L.L., Zalesky, A. & Breakspear, M. (2017). Criticality in the brain. *Neuroscience & Biobehavioral Reviews*, 75, 318–334.
- Dawson, D.A. (1983). Critical dynamics and fluctuations for a mean-field model of cooperative behavior. *Journal of Statistical Physics*, 31(1), 29–85.
- Desai, R.C. & Zwanzig, R. (1978). Statistical mechanics of a nonlinear stochastic model. *Journal of Statistical Physics*, 19(1), 1–24.
- Devika, V.A. (2026a). Effective landscapes from time series: a reconstruction protocol and a hierarchy of what it recovers. Companion methods paper.
- Dikker, S. et al. (2017). Brain-to-brain synchrony tracks real-world dynamic group interactions in the classroom. *Current Biology*, 27(9), 1375–1380.
- Dodds, P.S. et al. (2011). Temporal patterns of happiness and information in a global social network. *PLoS ONE*, 6(12), e26752.

- Ezaki, T., Watanabe, T., Ohzeki, M. & Masuda, N. (2017). Energy landscape analysis of neuroimaging data. *Philosophical Transactions of the Royal Society A*, 375, 20160287.
- Fowler, J.H. & Christakis, N.A. (2008). Dynamic spread of happiness in a large social network. *BMJ*, 337, a2338.
- Freeman, W.J. (2000). *Neurodynamics: An Exploration in Mesoscopic Brain Dynamics*. Springer.
- 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.
- Granovetter, M. (1978). Threshold models of collective behavior. *American Journal of Sociology*, 83(6), 1420–1443.
- Hasson, U., Ghazanfar, A.A., Galantucci, B., Garrod, S. & Keysers, C. (2012). Brain-to-brain coupling. *Trends in Cognitive Sciences*, 16(2), 114–121.
- Hopfield, J.J. (1982). Neural networks and physical systems with emergent collective computational abilities. *PNAS*, 79(8), 2554–2558.
- Huys, Q.J.M., Daw, N.D. & Dayan, P. (2015). Depression: a decision-theoretic analysis. *Annual Review of Neuroscience*, 38, 1–23.
- Johnson, M.E. (2019). Neural annealing: toward a neural theory of everything. Qualia Research Institute.
- Kramer, A.D.I., Guillory, J.E. & Hancock, J.T. (2014). Experimental evidence of massive-scale emotional contagion through social networks. *PNAS*, 111(24), 8788–8790.
- Kramers, H.A. (1940). Brownian motion in a field of force and the diffusion model of chemical reactions. *Physica*, 7(4), 284–304.
- Kuran, T. (1991). Now out of never: the element of surprise in the East European revolution of 1989. *World Politics*, 44(1), 7–48.
- McCraty, R. & Childre, D. (2010). Coherence: bridging personal, social, and global health. *Alternative Therapies in Health and Medicine*, 16(4), 10–24.
- Müller, V. & Lindenberger, U. (2011). Cardiac and respiratory patterns synchronize between persons during choir singing. *PLoS ONE*, 6(9), e24893.
- Scheffer, M., Bascompte, J., Brock, W.A. et al. (2009). Early-warning signals for critical transitions. *Nature*, 461, 53–59.
- Shine, J.M. et al. (2016). The dynamics of functional brain networks. *Neuron*, 92(2), 544–554.
- Sornette, D. (2003). *Why Stock Markets Crash*. Princeton University Press.
- Strogatz, S.H. (1994). *Nonlinear Dynamics and Chaos*. Addison-Wesley.
- Strogatz, S.H. (2000). From Kuramoto to Crawford. *Physica D*, 143, 1–20.
- Thom, R. (1975). *Structural Stability and Morphogenesis*. W.A. Benjamin.
- Thompson, E. (2007). *Mind in Life*. Harvard University Press.
- Tononi, G. (2004). An information integration theory of consciousness. *BMC Neuroscience*, 5, 42.
- Varela, F.J., Thompson, E. & Rosch, E. (1991). *The Embodied Mind*. MIT Press.
- Watanabe, T. et al. (2014). Energy landscapes of resting-state brain networks. *Frontiers in Neuroinformatics*, 8, 12.