
Admissibility as a Unifying Concept

Constraint, Reachability, and Relaxation
Across Biological, Cognitive, Industrial, and Cosmological Systems

A Critical and Theoretical Reading of the Monograph

Flyxion | June 2026

Abstract

The monograph *Constraint, Reachability, and Relaxation* advances a unified theoretical account of six empirical systems—cortical bistability, synovial joint homeostasis, genomic regulatory grammar, battery recycling logistics, reinforcement learning habit formation, and quasar accretion disk reverberation—under the organizing concept of the admissibility field $\mathcal{A}(x, t)$: the set-valued map over state space encoding which transitions a system can still make, given its history, boundary conditions, and spatial coupling. This essay engages that argument at the level of its formal structure, its primary empirical anchors, and its theoretical implications. Seven themes are developed in succession. First, the internal coherence of the unifying claim is assessed: whether the six domains genuinely share the same mathematical object or merely a family of analogous ones. Second, the distinction between relaxation and collapse—arguably the monograph’s sharpest conceptual contribution—is subjected to closer formal scrutiny. Third, the RSVP and CLIO frameworks that provide the theoretical vocabulary are examined as theoretical commitments, not merely as notation. Fourth, the research program of twelve predictions is evaluated with respect to what discriminating evidence would actually look like. Fifth, the Admissibility Display Principle—developed in derivative documents and absent from the monograph itself—is formalized as an application of CLIO projection logic to interface design. Sixth, a compact-field sensitivity principle is stated that unifies the quasar X-ray/infrared variability asymmetry with the cortical HR/non-HR unit pair decomposition and the stoichiometric collapse dynamics of the joint. Seventh, the convergent lining geometry of the synovial joint is analyzed as a trajectory-independent admissibility attractor, with implications for two formally distinct classes of therapeutic strategy. An eighth section develops the unified account and the framework’s appropriate epistemic status. A ninth section develops a concrete experimental program targeting the two most decisive tests of the framework: a parametric cortical crossing protocol establishing that $\chi(\gamma)$ rather than inhibition dose is the operative variable for relaxation, and a causal joint collapse protocol establishing that PI16+ density is rate-determining; this section also formalizes the strong-versus-weak versions of the maintainer inversion claim and describes the lineage-tracking experiment required to discriminate them.

Contents

1	The Unifying Claim and Its Logical Status	4
1.1	Structural Identity versus Analogy	4
1.2	The Form of the Argument	5
2	Relaxation, Collapse, and the Halorhodopsin Result	5
2.1	The Distinction as the Monograph's Sharpest Claim	5
2.2	The PI16+ Fibroblast as Threshold Maintainer	6
2.3	A Note on the Cortical Maintainer	7
3	The RSVP and CLIO Frameworks as Theoretical Commitments	7
3.1	What the Frameworks Provide	7
3.2	The Derivational Gap	8
3.3	The Sleep Pressure Functional	8
4	Residual Geometry and Historical Encoding	9
4.1	The Hierarchy of Memory	9
4.2	Non-Markovian Dynamics and the Epistemological Consequence .	10
4.3	The Battery Recycling Case and Dynamic Residual Geometry . .	10
5	Constraint Compression and the RL-Genomics Bridge	11
5.1	Habit as Admissibility Compression	11
5.2	SAGE-net and the Limits of Compression	11
6	Relaxation Across Scales and the Cosmological Connection	12
6.1	The Coupling Kernel as the Scale-Determining Object	12
6.2	The RSVP Cosmological Prediction	12
7	The Research Program and Its Twelve Predictions	13
7.1	The Predictions as Constraints on the Framework	13
7.2	The Open Mathematical Problem	14
8	The Admissibility Display Principle: Interface Design as Field Exposure	14
8.1	The Spinner as Collapsed Field	14
8.2	Productive Waiting and the CLIO Parallel	15
9	Compact Fields and Amplitude Sensitivity: A General Principle	16
9.1	The Variability Asymmetry as Admissibility Diagnostic	16
9.2	The Amplitude Ratio as a Quantitative Field Measure	17
9.3	Diagnostic Implications	17
10	Convergent Geometry and Therapeutic Attractors	18
10.1	The Convergent Lining Phenotype as an Attractor	18
10.2	Two Therapeutic Modes and Their Formal Distinction	19
10.3	The Stoichiometric Parameter as a Predictive Clinical Index	19
11	Toward a Unified Account	20

11.1	What the Framework Achieves	20
11.2	The Appropriate Epistemic Status	21
12	An Experimental Research Program: Testing the Two Central Claims	21
12.1	The Priority Ordering	21
12.2	The Cortical Crossing Protocol	22
12.3	The Joint Collapse Protocol	22
12.4	The Lineage-Tracking Experiment and the Strong/Weak Maintainer Inversion Claim	23
12.5	Open Methodological Questions	24

1. The Unifying Claim and Its Logical Status

1.1. Structural Identity versus Analogy

The monograph’s central claim is not that six systems are analogous but that they share an identical mathematical object. This is a much stronger claim and deserves to be assessed as such. An analogy between cortical bistability and joint homeostasis would say: both involve some notion of constraint and some notion of relaxation, and these notions bear a family resemblance that illuminates each domain when examined through the lens of the other. The monograph is committed to something different: that the same set-valued map

$$\mathcal{A} : \Omega \times \mathbb{R}_{\geq 0} \longrightarrow 2^{\mathcal{S}} \quad (1)$$

is the functionally determinative object in each domain, that its formal properties—bipartite decomposition, threshold surfaces, coupling kernels, relaxation and collapse operators—appear in each domain not by metaphor but by structure.

The case for structural identity is strongest in the three domains where the admissibility field has a computable or measurable correlate. In the cortical domain, the bipartite decomposition $\mathcal{A} = \mathcal{A}_H \cup \mathcal{A}_F$ is grounded by the 67%/33% split of homeostatically regulated unit pairs from non-regulated pairs in Driessen et al. (2026). This is not a conceptual partition but an empirical one: the fraction is consistent across mouse lines and across probe identity, which constitutes evidence that the decomposition reflects actual biological structure rather than a classification choice. In the reinforcement learning domain, the minimum planning depth $D \geq 2$ is proven mathematically by Okitsu and Sakai (2026) and is not a continuous parameter but a genuine discontinuity—the threshold surface is an exact boundary in parameter space. In the accretion disk domain, the thin-disk/slim-disk transition at the Eddington ratio is a sharp regime change with distinct spectral signatures below and above the threshold.

The case is somewhat weaker in the battery recycling and genomics domains. The 39.3% versus 65.68% utilization rates in Tian et al. (2026) constitute a threshold in a logistic and policy sense, but whether the system has a genuine dynamical threshold—a region of state space that must be traversed for the recycling structure to reorganize, not merely a bifurcation in steady-state outcomes depending on coordination level—is not established by the data. The genomics domain faces an analogous difficulty: the MAF threshold at approximately 0.05 is a modeling limitation as much as a structural feature of the regulatory grammar, since it reflects the training distribution of SAGE-net rather than a sharp boundary in the underlying sequence-to-function map.

These are not fatal objections. They suggest that the unifying claim is most rigorously established in four of the six domains and is more tentative in two, and the monograph is largely aware of this gradation. The honest assessment is that the claim succeeds as a structural claim in the core cases and functions as a productive organizing hypothesis in the others—which is not a minor achievement.

1.2. The Form of the Argument

The monograph argues by the following pattern: identify a set of formally defined properties (bipartite structure, threshold surface, relaxation operator, coupling kernel); then show that each domain exhibits an instance of each property. This is the correct form for a structural unification argument, and it is more rigorous than the alternative of showing only that certain qualitative features recur.

What the argument does not yet provide—and the monograph is explicit that it does not—is a derivation of the admissibility framework from a more fundamental theory. The RSVP and CLIO frameworks supply a vocabulary and a set of equations, but the admissibility field $\mathcal{A}(x, t)$ is introduced axiomatically (Definition 2.1 of the monograph) rather than derived from the RSVP field equations for $\Phi(x, t)$, $\mathbf{v}(x, t)$, and $S(x, t)$. Under the RSVP identification $|\mathcal{A}(x, t)| \propto e^{S(x, t)}$, one can read the admissibility size off the entropy density, but the structure of \mathcal{A} as a set-valued map—which transitions are actually available, not merely how many there are—is not directly recoverable from S alone. This is the correct location of the remaining theoretical gap.

2. Relaxation, Collapse, and the Halorhodopsin Result

2.1. The Distinction as the Monograph’s Sharpest Claim

Of all the conceptual distinctions introduced in the monograph, the separation between relaxation and collapse is the most important and the most precisely supported. It is worth reconstructing the argument carefully.

Relaxation is defined as an expansion of \mathcal{A}_H through threshold-surface traversal: the trajectory visits the minimum-admissibility state and returns, and the return trajectory carries expanded reachability. The critical point is that this expansion requires the crossing, not merely a reduction in activity. The halorhodopsin control in Driessen et al. (2026) establishes this experimentally by showing that matched firing-rate reduction—reduction that is statistically indistinguishable in magnitude from the SOM+ and ACR protocols—produces no downstream effects: no SWA reduction, no STTC reduction, no GluA1 or pGluA1 modulation, no memory rescue. The operative variable is the trajectory pattern, not the magnitude of suppression.

This is a clean experimental result, and its formal interpretation is exact. Define the crossing indicator $\chi(\gamma) = 1$ if and only if trajectory γ crosses the threshold surface Σ . The monograph’s Proposition 3.1 states:

$$\mathcal{R}_\tau[\mathcal{A}_H](x, t) \supsetneq \mathcal{A}_H(x, t) \iff \chi(\gamma) = 1. \quad (2)$$

The halorhodopsin result instantiates the contrapositive: $\chi(\gamma_{\text{tonic}}) = 0$ implies $\mathcal{R}_\tau[\mathcal{A}_H] = \mathcal{A}_H$, meaning no expansion occurs. The biconditional is not directly established by one experiment. The (\Leftarrow) direction would require showing that every trajectory crossing Σ produces expansion, not merely that the specific bistable protocols do. But the halorhodopsin control eliminates the principal confound—total

activity level—and provides the strongest possible evidence short of a complete traversal protocol.

Collapse is a qualitatively different phenomenon. It is not the failure to expand \mathcal{A}_H but the destruction of the mechanism that maintains \mathcal{A}_H . The arthritis data from Davidson et al. (2026) show that PI16+ fibroblasts, which under homeostatic conditions express a cassette maintaining structural integrity through WNT, BMP, and FGF signaling, invert their function under TNF and IL-1 β stimulation: they upregulate chemotactic attractants (CCL2, CCL3, CCL5, IL6, IL1B) and down-regulate the homeostatic signaling network. The maintaining mechanism does not merely fail; it becomes actively destabilizing. The monograph captures this with the formal definition of admissibility collapse as a strict reduction driven by functional inversion rather than constraint accumulation.

Definition 2.1 (Collapse versus Relaxation). Relaxation at (x, t) is the operator \mathcal{R}_τ that satisfies $\mathcal{R}_\tau[\mathcal{A}](x, t) \supseteq \mathcal{A}(x, t + \tau)$ with strict inclusion when $\chi(\gamma) = 1$. It is reversible on the timescale of the accumulation process. Collapse at (x, t) is the irreversible strict reduction

$$\mathcal{A}_{\text{collapse}}(x, t) \subsetneq \mathcal{A}(x, t) \quad (3)$$

driven by functional inversion of the mechanism that actively maintains \mathcal{A}_H . It is irreversible without external structural intervention.

The asymmetry between these two modes is not merely a matter of reversibility. Relaxation is trajectory-dependent: it requires the system to visit a specific region of state space. Collapse is mechanism-dependent: it requires the destruction or inversion of a specific cellular or structural element. The monograph is correct that these are qualitatively distinct and should not be conflated by a common term such as “loss of admissibility.”

2.2. The PI16+ Fibroblast as Threshold Maintainer

The arthritis paper’s most striking contribution is the identification of a specific cellular population—the PI16+ fibroblast—as the biological implementation of what the monograph calls a threshold maintainer: a structural element whose presence ensures that Σ exists and that the bipartite decomposition $\mathcal{A} = \mathcal{A}_H \cup \mathcal{A}_F$ is stable.

This is a strong claim because it assigns a precise formal role to a specific cell type. The empirical support is layered. First, the embryonic single-cell RNA sequencing data in Davidson et al. (2026) establish that PI16+ cells (cluster F2) occupy specific spatial niches in the developing joint: perivascular and adventitial positions, and interstitial tissue interfaces along complex tendon-ligament borders. Second, the stoichiometric enrichment of PI16+ cells at the PIP joint relative to the DIP joint—established in utero—provides a quantitative basis for the differential inflammatory vulnerability of these joints. Third, the cytokine stimulation experiments demonstrate the homeostatic-to-pathological inversion directly.

The formal consequence of this identification is that the admissibility field of the

joint is not passively shaped by its history but is actively maintained by a specific population. This makes collapse a different kind of event than relaxation: collapse requires perturbing the maintainer, not merely failing to perform threshold crossings. The monograph introduces the concept of a threshold maintainer (Definition 3.1) to name this class of element, and the PI16+ fibroblast is its clearest empirical instance across all six domains.

2.3. A Note on the Cortical Maintainer

In the cortical domain, the threshold maintainer is identified as the SOM+ interneuron network, through the established role of somatostatin interneurons as master regulators of cortical excitability documented in Funk et al. (2017) and through the chloride homeostasis mechanism characterized in Alfonsa et al. (2023). This identification is somewhat less precise than the PI16+ case, because the cortical system has no direct analog of the embryonic stoichiometry argument. The monograph treats the two on similar footing, but the structural evidence for SOM+ interneurons as threshold maintainers is of a different character: it is functional evidence (disrupting SOM+ function disrupts slow-wave generation) rather than stoichiometric and developmental evidence (more PI16+ cells at a site predicts faster collapse under fixed inflammatory challenge). This distinction matters for the research program: predictions about cortical collapse would need to specify what disruption of the maintaining mechanism looks like in the cortical case, which is less precisely defined than in the joint case.

3. The RSVP and CLIO Frameworks as Theoretical Commitments

3.1. What the Frameworks Provide

The RSVP framework enters the monograph as a source of formal objects—the scalar potential $\Phi(x, t)$, the lamphrodyne flow $\mathbf{v}(x, t)$, and the configurational entropy density $S(x, t)$ —that are identified with admissibility-field quantities. The central identification is $|\mathcal{A}(x, t)| \propto e^{S(x, t)}$, which allows constraint curvature to be written as

$$\kappa(x, t) = -\frac{\partial^2 \log |\mathcal{A}(x, t)|}{\partial t^2} = -\frac{\partial^2 S}{\partial t^2}. \quad (4)$$

This is a productive identification because it connects the size of the admissibility field to the entropy density in a way that is both mathematically natural and consistent with the thermodynamic interpretation of entropy as a measure of degrees of freedom.

The CLIO framework provides the projection operator $\pi : \mathcal{A} \rightarrow \hat{\mathcal{A}}$, which compresses the admissibility field, and the residual $\mathcal{A} \setminus \hat{\mathcal{A}}$, which is the information lost in compression. This is used to interpret habit formation as admissibility compression in the RL context: the model-based rollout is explicit computation of $\mathcal{A}(s, t)$ within the planning horizon, and the model-free value function $V(s) = \mathbf{v}^\top s$ is

the compressed approximation $\hat{\mathcal{A}}$ for states beyond the horizon. The unification claim of Okitsu and Sakai (2026)—that habits and Pavlovian responses arise from a single computational substrate—is interpreted as the claim that $\hat{\mathcal{A}}$ and $\mathcal{A} \setminus \hat{\mathcal{A}}$ are projection and residual of one field rather than outputs of two separate systems.

3.2. The Derivational Gap

The frameworks provide vocabulary and identification, but they do not yet derive the admissibility field from first principles. Specifically, the RSVP field equations govern the evolution of (Φ, \mathbf{v}, S) , and under the identification $|\mathcal{A}| \propto e^S$, they constrain the size of the admissibility field. But the admissibility field $\mathcal{A}(x, t)$ is a set-valued map, not a scalar, and recovering the set structure from S requires additional assumptions about which transitions are available at each level of entropy.

One natural approach would be to treat $\mathcal{A}(x, t)$ as a level set of the RSVP potential: transitions available to a system are those consistent with Φ remaining below some threshold. But this would require specifying the potential landscape $\Phi(x, t)$ explicitly for each domain, which the monograph does not do. The RSVP connection is therefore productive at the level of structural analogy—the potential landscape must have a barrier structure such that the off-period minimum must be visited for cortical relaxation to occur, as the monograph notes—but is not yet a derivation.

This is not a deficiency specific to this monograph; it is the characteristic position of a framework that is still being developed. The monograph is correct to use RSVP and CLIO as organizational vocabulary while acknowledging that the derivational program remains open. The open mathematical problem stated in Section 9.2—characterizing conditions under which threshold surfaces exist in terms of RSVP field equation properties—is precisely the problem of closing this gap.

3.3. The Sleep Pressure Functional

The RSVP sleep pressure functional introduced in Chapter 4 is worth examining in detail because it is the monograph’s most fully developed formal connection between the frameworks and empirical data. Sleep pressure at location x over duration T is defined as

$$P(x, T) = \int_0^T \max(0, S^*(x) - S(x, t)) dt, \quad (5)$$

where $S^*(x)$ is the homeostatic target entropy density. This functional is interpretable: it integrates the deficit between actual and target entropy density over time, measuring the accumulated cost of sustained constraint narrowing. The observable correlate—SWA in the first hour of recovery sleep—is well established in the synaptic homeostasis literature, and the connection $\mathcal{R}_\tau[P] \propto \Delta S \cdot (\tau_2 - \tau_1)$ for off-period induction is a concrete, testable relationship.

What makes this functional important is that it transforms the Threshold Hypothesis into a quantitative claim. If $\chi(\gamma_{\text{tonic}}) = 0$ implies $\Delta S_{\text{tonic}} \approx 0$, and if P is approximately the integral of these increments, then tonic inhibition—regardless of its depth—should produce no reduction in P unless it drives the trajectory across Σ . This is exactly what Driessen et al. (2026) observe, and the functional provides the formal link between the experimental outcome and the theoretical structure.

The RSVP constraint on the potential landscape—that $\Phi(x, t)$ must have a barrier structure such that visiting the off-period minimum is necessary for relaxation—is also a concrete claim. It predicts that graded approaches to Σ without crossing should produce no relaxation, which is the content of Prediction P1 in the research program.

4. Residual Geometry and Historical Encoding

4.1. The Hierarchy of Memory

Chapter 6 of the monograph introduces a hierarchy of three kinds of historical encoding: trajectory accumulation (reversible, hours), developmental encoding (irreversible, months), and stoichiometric encoding (irreversible without cell therapy). This is one of the most conceptually rich parts of the document and deserves elaboration.

The key formal contribution is the definition of residual geometry:

$$\mathcal{G}_{\text{res}}(x) = \mathcal{A}(x, t_{\text{adult}}) - \mathcal{A}_{\text{free}}(x, t_{\text{adult}}), \quad (6)$$

where $\mathcal{A}_{\text{free}}$ is the admissibility that would exist in the absence of historical encoding. This definition captures the idea that the present state of the system’s reachability field is not determined by its present state alone but carries the imprint of past trajectories that are no longer active.

The arthritis paper provides a particularly clean instance: the embryonic fibroblast signature is recoverable from adult tissue transcriptomics (Davidson et al., 2026; Zhang et al., 2019), which means the developmental trajectory is encoded in the adult admissibility structure as a recoverable residual. This is the empirical content of the matching between embryonic and adult AMP2 datasets.

Remark 4.1. The formal definition of \mathcal{G}_{res} requires $\mathcal{A}_{\text{free}}$ to be well-defined, which is not trivial. In the joint case, $\mathcal{A}_{\text{free}}$ can be approximated by the admissibility of a joint with the same gross anatomy but without the PI16+ stoichiometric bias—a counterfactual that is not directly observable but is operationally approximated by the DIP joint under the assumption that the DIP represents lower PI16+ density. This counterfactual interpretation is implicit in the monograph and should be made explicit in formal developments.

4.2. Non-Markovian Dynamics and the Epistemological Consequence

The monograph’s observation that systems with residual geometry are not Markovian at the observable level is important and deserves development beyond what the monograph provides. If $\mathcal{A}(x, t_{\text{adult}})$ depends on the developmental trajectory γ_{dev} rather than only on $s(x, t_{\text{adult}})$, then Bayesian inference on such systems requires a prior over developmental trajectories, not merely over current states.

This has a concrete consequence for the predictive framework. The monograph’s predictions about collapse dynamics in the joint (P4 through P6) are predictions about outcomes under perturbation of the current state, but the rate and character of collapse depend on the baseline PI16+ density, which is set developmentally. The prediction that collapse rate under TNF/IL-1 β is a monotone function of baseline PI16+ density (P4) is therefore implicitly a prediction about the consequences of a hidden variable—the developmental trajectory—not directly observable in adult tissue.

The Yarncrawler framework mentioned by the monograph is the correct epistemological response to this situation: world-state reconstruction as constraint closure. If the admissibility field encodes the developmental trajectory as a recoverable residual, and if that trajectory can be reconstructed from the adult field, then Bayesian inference can in principle integrate over the posterior on developmental history rather than treating the current state as sufficient. The matching of embryonic and adult signatures in Davidson et al. (2026) is the first step toward this kind of trajectory reconstruction.

4.3. The Battery Recycling Case and Dynamic Residual Geometry

The battery recycling application is the most policy-relevant of the six domains and the one where the residual geometry concept takes its most explicitly dynamic form. The spatial distribution of end-of-life batteries in 2030 is the residual of EV adoption diffusion from 2016 to 2030 (Tian et al., 2026), and the hotspot migration pattern across provinces—northeast to southwest from 2020 to 2026, southeast to northwest from 2026 to 2030—constitutes a moving admissibility field that recycling infrastructure must anticipate rather than serve.

This is a clean instance of the central claim in a domain where the policy implications are immediate: infrastructure sited to serve the current battery geography will be misaligned with the future battery geography by the time it comes online. The admissibility framework makes this precise by noting that the recycling system’s viability depends on spatial coupling between battery supply and processing capacity, and that the coupling kernel $K(x, x')$ in the spatial coupling equation

$$\frac{\partial \mathcal{A}(x, t)}{\partial t} = f[\mathcal{A}(x, t), s(x, t)] + \int_{\Omega} K(x, x') \mathbf{v}(x', t) dx' \quad (7)$$

is determined by logistics and geography rather than by physics or biology—and hence is more directly manipulable by policy intervention than the kernels in the cortical or joint systems.

5. Constraint Compression and the RL-Genomics Bridge

5.1. Habit as Admissibility Compression

The CLIO interpretation of the Okitsu-Sakai model is one of the monograph’s most elegant formal moves. The plan-until-habit framework unifies goal-directed and habitual control by showing that for planning depth $D \geq 2$, Pavlovian-instrumental transfer emerges from a single computational architecture. In CLIO terms, the model-based rollout explicitly represents $\mathcal{A}(s, t)$ for $t \in [0, D]$ and the model-free value function $V(s) = \mathbf{v}^\top s$ compresses the long-range admissibility $\hat{\mathcal{A}}(s, t > D)$.

The proposition that habits and Pavlovian responses are projection and residual of one field—rather than outputs of two separate systems—resolves a long-standing debate in computational neuroscience about whether goal-directed and habitual control are architecturally distinct or merely parameterically distinguished. The Okitsu-Sakai result says they are parameterically distinguished along the single axis of planning depth, with the threshold at $D = 2$ being a genuine discontinuity. This is the most mathematically explicit threshold surface in any of the six papers, and its proof by the authors rather than its identification post hoc makes it the cleanest instance of Hypothesis 3.1.

Proposition 5.1. Let $\mathcal{A}^{[D]}(s, t)$ denote the admissibility field computed by the plan-until-habit architecture with planning depth D . Then $\mathcal{A}^{[D]}(s, t)$ for $t \leq D$ is the model-based component, and $\hat{\mathcal{A}}^{[D]}(s, t)$ for $t > D$ is the model-free compression. The Pavlovian-instrumental transfer exists if and only if $D \geq 2$, corresponding to the threshold surface $\Sigma_D = \{D = 2\}$ in planning-depth parameter space.

The formal interest of this result extends beyond the RL domain. It establishes that the planning depth is not a continuous dial but a categorical threshold, which is exactly the structure the admissibility framework predicts for threshold surfaces: the relationship between \mathcal{A}_H -expansion and proximity to Σ is discontinuous (Proposition 3.1 of the monograph). In the RL case, the discontinuity is proven, not merely observed—it is a mathematical consequence of the matrix-multiplication structure of the planning rollout.

5.2. SAGE-net and the Limits of Compression

The genomics application centers on a negative result: personal genome training improves admissibility compression for seen genes but fails to improve it for unseen genes. This is interpreted as the failure of the SAGE-net architecture to learn the global constraint geometry of gene regulation—the locus-independent structure of how sequences determine function.

The seqlet distance analysis (p-SAGE-net mean 5,403 bp vs. r-SAGE-net 3,117 bp from TSS; $p = 2 \times 10^{-80}$) shows that personal genome training extends the effective range of the spatial coupling kernel $K(x, x')$, capturing distal regulatory interactions that the reference-genome model misses. But the failure to generalize across loci means that the extended kernel is locus-specific: the distal coupling

structure does not transfer.

This is interpretable in admissibility terms as a failure of $\hat{\mathcal{A}}$ to approximate \mathcal{A} globally. Personal genome training improves the local compression—it learns more of the locus-specific admissibility structure—but does not recover the global constraint geometry that would be needed for cross-locus generalization. DNA methylation partially generalizes at 100,000 training regions ($p = 4.96 \times 10^{-4}$), suggesting that the epigenomic admissibility field has lower curvature than the sequence-level field—it is smoother and therefore compresses better.

The connection between the SAGE-net result and the MAF threshold is the weakest formal link in the monograph. The claim that variants below $\text{MAF} \approx 0.05$ contribute negligibly to model predictions constitutes a threshold surface in genomic admissibility space only if the failure is structural—if no architecture could learn these contributions given current data—rather than merely a data limitation. The monograph acknowledges this ambiguity but does not resolve it.

6. Relaxation Across Scales and the Cosmological Connection

6.1. The Coupling Kernel as the Scale-Determining Object

The monograph’s Chapter 8 identifies three propagation regimes that differ primarily in the coupling kernel $K(x, x')$: local (cortical, $K \approx 0$ at inter-hemispheric distances), light-speed (quasar accretion, K propagating at c), and diffusive (battery adoption, propagating on year timescales). The coupling kernel is therefore the scale-determining object, and understanding why different systems have different kernels is a prerequisite for understanding their different dynamical regimes.

In the cortical system, the near-zero inter-hemispheric kernel is a consequence of the local architecture of SOM+ interneuron networks: they generate bistable patterns locally and do not directly entrain the contralateral hemisphere. In the quasar accretion system, the light-speed kernel is a consequence of the causal structure of spacetime: corona perturbations propagate at c through the disk, producing the wavelength-dependent lag $\tau \propto \lambda^{4/3}$ (Cackett et al., 2020; Fausnaugh et al., 2016). In the battery system, the diffusive kernel reflects the logistics and policy constraints of the recycling network.

What is interesting about this variety is that it suggests the coupling kernel is not an intrinsic feature of the admissibility framework but a domain-specific parameter that must be specified separately. This is consistent with the monograph’s approach—equation (2.5) presents the kernel as a given—but raises the question of whether there is a principled way to derive the coupling kernel from the system’s microscopic structure in each domain.

6.2. The RSVP Cosmological Prediction

The most ambitious claim in Chapter 8 is the RSVP cosmological prediction: that the thin-disk reverberation relation $\tau \propto \lambda^{4/3}$ should steepen to $\tau \propto \lambda^\alpha$ with

$\alpha > 4/3$ at high redshift ($z \gg 1$) if RSVP's reinterpretation of cosmological redshift as an entropic field response is correct.

This prediction is testable in principle with Rubin/Roman era data ($N > 100$ variable quasars at $z > 4$), and the current J0439+1634 result—consistent with $\alpha = 4/3$ within uncertainties at $z = 6.51$ (Leung et al., 2026)—does not discriminate between standard and RSVP interpretations. The prediction is therefore a genuine falsifiable claim about future observations.

Two aspects of this prediction deserve comment. First, the mechanism by which RSVP predicts steepening is not fully developed in the monograph. The claim is that cosmological redshift in RSVP arises from entropic field response rather than metric expansion, and that this produces a slight modification to the propagation-speed constraint. But the quantitative magnitude of the modification—how much steeper α should be relative to $4/3$, as a function of redshift—is not given. A prediction that does not specify the expected effect size is testable only at the level of the direction of the deviation, not its magnitude.

Second, the RSVP cosmological framework is not developed within the monograph; it is cited as an associated framework. The connection between the quasar reverberation relation and the cosmological claims of RSVP is therefore asserted rather than derived. This is appropriate for a monograph whose primary subject is the admissibility framework rather than RSVP cosmology, but it means that Prediction P11 rests on theoretical foundations that are not fully articulated in the present document.

7. The Research Program and Its Twelve Predictions

7.1. The Predictions as Constraints on the Framework

A theoretical framework's predictions are most valuable not as future confirmations but as present constraints. If the admissibility framework is correct, the twelve predictions must hold; if any of them fails, the framework requires revision. Examining what each prediction requires—both empirically and theoretically—clarifies what the framework is committed to.

Predictions P1 through P3 (threshold and relaxation in neuroscience) are the most cleanly derived from the Threshold Hypothesis. P1 requires that graded tonic inhibition approaching but not reaching global silence produces no SWA reduction regardless of depth. This is a strong claim: it says the threshold is sharp, not smeared, so that proximity to Σ produces no partial expansion of \mathcal{A}_H . P2 and P3 require that isolated threshold crossings produce discrete, linearly summable SWA reductions. Together, P1–P3 constrain the threshold surface to be a genuine discontinuity rather than a smooth transition, and they constrain the relaxation operator to be additive across crossings. Both are non-trivial commitments.

Predictions P4 through P6 (collapse dynamics in the joint) require that the PI16+ fibroblast stoichiometry is the rate-determining factor in collapse, not merely a correlated predictor. P4 is particularly demanding: it requires that collapse rate under fixed cytokine challenge is a monotone function of PI16+ baseline density

across joint sites. This would be violated if other factors—mechanical environment, vascular density, other fibroblast populations—contribute substantially to the collapse rate independently of PI16+ density. The monograph’s formal model (equation 5.2, in which collapse rate is proportional to $\rho(x)$) is a strong claim that should be tested against alternative models in which $\rho(x)$ is one of several predictors.

Predictions P9 and P10 (the RL threshold) are among the most directly testable because the Okitsu-Sakai model is computational and fully specified. P9 requires that probabilistic depth sampling between $D = 1$ and $D = 2$ produces a categorical rather than smooth transition in PIT magnitude. This is testable by simulating the model with non-integer depth using probabilistic rollout termination.

7.2. The Open Mathematical Problem

The monograph correctly identifies its central open problem: under what conditions does an admissibility field contain a threshold surface, and when does perturbation of the threshold maintainer produce collapse rather than relaxation? This is the right question because it is the one whose answer would transform the admissibility framework from a descriptive vocabulary into a predictive theory.

The empirical evidence suggests three conditions are associated with threshold surfaces: constraint accumulation is approximately irreversible on the accumulation timescale; the system has a natural minimum-admissibility state; and the relaxation mechanism requires visiting this minimum and returning. These conditions are descriptive rather than derivational, and the monograph acknowledges that a formal characterization in terms of RSVP field equation properties remains to be done.

Claim 7.1. A threshold surface $\Sigma \subset \mathcal{S}$ exists and produces \mathcal{A}_H -expansion upon crossing if and only if the potential landscape $\Phi(x, t)$ has a local minimum at the silence state s_{\min} such that the barrier height $\Delta\Phi = \Phi(s_{\min}) - \Phi(s_{\text{operating}})$ is strictly positive and the return trajectory from s_{\min} carries positive entropy increment $\Delta S > 0$.

This claim is offered as a candidate formalization of the conditions, not as a proven theorem. Its value is that it is falsifiable: a system with the structural features described (irreversible accumulation, natural minimum, required return) but without a positive barrier in Φ would fail to exhibit threshold surfaces. Deriving conditions on (Φ, \mathbf{v}, S) that guarantee the existence and stability of Σ is the primary mathematical problem for the next stage of the RSVP-admissibility research program.

8. The Admissibility Display Principle: Interface Design as Field Exposure

8.1. The Spinner as Collapsed Field

The derivative documents generated from the monograph introduce a concept absent from the original text but following naturally from its premises: the Admissi-

bility Display Principle. The core observation is that the standard computational interface exposes only state while concealing reachability. A spinning cursor or progress indicator reports $s = \text{busy}$ while collapsing the admissibility field $\mathcal{A}(x, t)$ to a single node. The user is told what the system is doing and told nothing about what it can still do, what it is choosing not to do, or what thresholds remain to be crossed.

This is a formally precise deficiency. In the admissibility framework, the informational value of a display is proportional to the extent to which it reveals $\mathcal{A}(x, t)$ rather than $s(x, t)$. A display that shows only s provides zero information about the system’s trajectory through state space, about which transitions have been explored and rejected, or about what the system would need from the user to narrow its field. The derivative materials call this “dead waiting”—waiting that is informationally inert from the perspective of the user’s ability to participate in the system’s constraint closure.

Definition 8.1 (Projection opacity). A display has projection opacity $\omega \in [0, 1]$ equal to the fraction of admissibility information it conceals. A display exposing only $s(x, t)$ has $\omega = 1$ (complete opacity). A display exposing $\mathcal{A}(x, t)$ fully has $\omega = 0$.

Current AI interfaces have $\omega \approx 1$ during generation: the user sees the current output state but not the distribution over candidate interpretations, not the rejected hypotheses, not the open questions whose resolution would accelerate convergence. This is not merely a usability problem; it is a structural mismatch between what the system’s internal state contains and what the interface makes available for human collaboration.

8.2. Productive Waiting and the CLIO Parallel

The alternative proposed in the derivative materials—“productive waiting” in which the interface exposes partial hypotheses, active conceptual branches, and rejected paths—is a direct application of the CLIO framework to interface design. A CLIO projection $\pi : \mathcal{A} \rightarrow \hat{\mathcal{A}}$ compresses the admissibility field; the residual $\mathcal{A} \setminus \hat{\mathcal{A}}$ is the information sacrificed. A dead-waiting interface performs maximal compression, retaining only the binary fact of current occupancy. A productive-waiting interface reduces compression, exposing the structure of $\hat{\mathcal{A}}$ and some representation of what has been eliminated.

What this means concretely is that an interface designed on admissibility principles should show three things: the interpretations currently under active consideration (the structure of $\hat{\mathcal{A}}$), the alternatives that have been closed and at what point they were eliminated (partial specification of $\mathcal{A} \setminus \hat{\mathcal{A}}$), and the open questions—the locations where user input could narrow the field and accelerate convergence to an answer. This third category is the most important: it transforms the waiting interval from a passive delay into a window for participatory constraint closure.

The parallel to CLIO inference dynamics is structural. In the CLIO framework, inference is the partial reconstruction of the residual $\mathcal{A} \setminus \hat{\mathcal{A}}$ from new observations. A user who sees the system’s partial hypotheses and contributes a disambiguation

is performing exactly this operation: they are supplying an observation that reduces the residual. The interface becomes a joint inference process rather than a unidirectional output channel.

Definition 8.2 (Admissibility Display Principle). An interface satisfies the Admissibility Display Principle if its display function $D : \mathcal{A}(x, t) \rightarrow \mathcal{I}$ (where \mathcal{I} is the space of displayable information) is not constant on the structure of \mathcal{A} —that is, if distinct admissibility geometries produce distinguishable displays. The principle fails whenever D factors through $s(x, t)$ alone.

The practical challenge is compression: $\mathcal{A}(x, t)$ is typically high-dimensional and not directly renderable. The interface designer must choose a projection $\pi : \mathcal{A} \rightarrow \hat{\mathcal{A}}_{\text{display}}$ that preserves decision-relevant structure. This is itself a CLIO problem, and the quality of the interface is determined by how much of the decision-relevant field structure survives the projection. The derivative materials propose that this projection should prioritize current trajectory, topology of rejected alternatives, and location of open thresholds—precisely the three components that most enable productive user participation.

Remark 8.1. The Admissibility Display Principle generates an empirical prediction that is independent of the framework’s biological claims: interfaces with lower projection opacity should produce better human-AI collaborative outcomes than interfaces with higher projection opacity, holding computational capability constant. This is testable through user studies measuring time-to-resolution and quality of outcomes under productive-waiting versus dead-waiting conditions. A confirmatory result would extend the admissibility framework’s empirical scope into human-computer interaction without requiring any commitment to the cortical or immunological claims.

9. Compact Fields and Amplitude Sensitivity: A General Principle

9.1. The Variability Asymmetry as Admissibility Diagnostic

The derivative presentation materials make explicit a structural observation that the monograph contains implicitly but does not state as a general principle. In the quasar data from Leung et al. (2026), the X-ray variability amplitude ($8.2 \pm 3.7\times$) vastly exceeds the infrared variability amplitude ($\sim 1.15\times$ in W1 band). The monograph notes that amplitude is inversely proportional to admissibility volume—amplitude $\propto |\mathcal{A}|^{-1}$ —and that the X-ray-emitting corona has a smaller admissibility volume than the outer disk, making it more sensitive to perturbations in Φ . The derivative slide deck connects this explicitly to the bipartite decomposition: the corona’s compactness is the astrophysical instance of the \mathcal{A}_H component, just as the 67% fraction of homeostatically regulated unit pairs in Driessen et al. (2026) is the cortical instance. In both cases, compactness of \mathcal{A}_H produces differential sensitivity: the high-constraint component responds more strongly to the same perturbation than the free component.

This connection is worth formalizing, because it provides a domain-independent

principle for using variability measurements as admissibility diagnostics.

Proposition 9.1 (Compact-field sensitivity). Let $\mathcal{A}(x, t) = \mathcal{A}_H(x, t) \cup \mathcal{A}_F(x, t)$ with $|\mathcal{A}_H| \ll |\mathcal{A}_F|$. For a perturbation $\delta\Phi$ of fixed magnitude applied uniformly to the potential landscape, the fractional change in transition probability is greater for transitions in \mathcal{A}_H than for transitions in \mathcal{A}_F . That is:

$$\frac{\delta P(\mathcal{A}_H)/P(\mathcal{A}_H)}{\delta P(\mathcal{A}_F)/P(\mathcal{A}_F)} > 1 \quad \text{for all } |\delta\Phi| > 0. \quad (8)$$

The argument is that if \mathcal{A}_H is narrow, the potential Φ must be steep near its boundaries, so a perturbation of fixed magnitude displaces a larger fraction of the probability mass within \mathcal{A}_H than within the shallower \mathcal{A}_F . The high-constraint component is therefore the primary amplifier of incoming perturbations, regardless of the domain.

9.2. The Amplitude Ratio as a Quantitative Field Measure

The compact-field sensitivity principle implies that the ratio of \mathcal{A}_H -component variability to \mathcal{A}_F -component variability is a quantitative indicator of the compactness of \mathcal{A}_H relative to \mathcal{A}_F . In the quasar case this ratio is approximately $8.2/1.15 \approx 7.1$, measured as the X-ray to infrared variability ratio. In the cortical case, the ratio is implicit in the 67%/33% bipartite decomposition—HR unit pairs show approximately twice the synchronization dynamics of non-HR pairs under equivalent conditions.

These ratios are not the same quantity measured in different units; they reflect genuinely different physical systems. But their structural role is identical: both quantify the degree to which the high-constraint component of the admissibility field amplifies perturbations relative to the free component. This suggests a diagnostic program: wherever a bipartite $\mathcal{A}_H/\mathcal{A}_F$ decomposition can be identified, measuring the variability ratio provides a quantitative index of the compactness of \mathcal{A}_H without requiring direct measurement of the admissibility field itself.

Applied to the joint system, this principle predicts that PI16+ fibroblasts—as the \mathcal{A}_H component of the synovial field—should show greater fractional response to inflammatory perturbation than other fibroblast populations. Davidson et al. (2026) report that PI16+ collapse under TNF/IL-1 β is indeed more dramatic than the shared proinflammatory response of other stromal populations, which is consistent with the prediction. But the quantitative version—measuring the PI16+ response amplitude relative to the non-PI16+ response amplitude and interpreting that ratio as an admissibility compactness index—has not been done and would provide a cross-domain validation of the compact-field sensitivity principle if it gave a ratio consistent with the stoichiometric enrichment argument.

9.3. Diagnostic Implications

If amplitude ratio is a general admissibility diagnostic, then measurement strategy across all six domains should be oriented toward the \mathcal{A}_H component rather than the full system. In the cortical domain, this means targeting HR unit pairs rather

than the full neural population when measuring sleep-pressure effects. In the astrophysical domain, it means using X-ray observations of the corona rather than infrared observations of the outer disk when characterizing accretion state. In the joint domain, it means tracking PI16+ gene expression rather than total synovial cytokine levels when monitoring the onset of collapse. In each case, the high-constraint component provides the most sensitive signal because it operates in the region where the potential landscape is steepest and perturbations produce the largest fractional effects.

This is a genuine unification result that does not depend on any specific feature of individual domains. It follows from the bipartite decomposition and the inverse relationship between admissibility volume and sensitivity, applied consistently across the six systems.

10. Convergent Geometry and Therapeutic Attractors

10.1. The Convergent Lining Phenotype as an Attractor

The clinical synthesis in the derivative materials develops an observation from Davidson et al. (2026) that the monograph treats briefly: the adult synovial lining layer emerges convergently from two developmentally distinct lineages—the interzone GDF5+ lineage and the non-interzone STF/PI16+ lineage—both converging on the same SOX5+ HBEGF+ phenotype under localized hypoxia and EGFR signaling. The derivative document describes this phenotype as a “functional attractor” in the admissibility landscape.

This description has formal content in the RSVP framework. An attractor corresponds to a basin in the potential landscape $\Phi(x, t)$ toward which multiple trajectories converge under the same dynamical conditions. If the hypoxic EGFR-signaling environment creates a sufficiently deep basin at the SOX5+ HBEGF+ phenotype, then developmental trajectories originating from different precursors will converge on that phenotype regardless of starting point, as long as the basin captures them both.

Definition 10.1 (Admissibility attractor). A state $s^* \in \mathcal{S}$ is an admissibility attractor for a system if there exists a neighborhood $U(s^*)$ such that all trajectories originating in $U(s^*)$ under the system’s natural dynamics converge to s^* , and the admissibility field at s^* is stable under small perturbations: $|\delta\mathcal{A}(s^*)| \rightarrow 0$ as $|\delta| \rightarrow 0$. An attractor is trajectory-independent if distinct trajectories from outside $U(s^*)$ converge to it under matching boundary conditions.

The convergent lining phenotype satisfies this definition empirically: two distinct lineages arrive at the same transcriptomic state under matched hypoxic EGFR conditions, and the resulting adult lining structure is stable under homeostasis. This trajectory-independence is important because it means the attractor is a property of the dynamical environment—the hypoxic gradient and EGFR signaling—rather than of any particular precursor population. The lining does not exist because a specific lineage produced it; it exists because the developmental field has a basin at that phenotype.

The matching of embryonic fibroblast signatures against adult AMP2 datasets confirms that this attractor is temporally stable: the same transcriptomic identity reached during embryogenesis persists into adulthood, constituting precisely the residual geometry $\mathcal{G}_{\text{res}}(x)$ defined in the monograph. The attractor is not transiently visited but stably occupied.

10.2. Two Therapeutic Modes and Their Formal Distinction

The convergent lining observation generates a distinction between two classes of therapeutic strategy that maps directly onto the relaxation-versus-collapse framework, and which the clinical synthesis documents make explicit.

The first class targets the threshold maintainer directly: preserve PI16+ homeostatic function under inflammatory challenge by stabilizing the WNT/BMP/FGF cassette. This is the strategy implied by Prediction P5, and it addresses collapse by protecting the mechanism that keeps Σ stable. If the maintaining mechanism is intact, the system can still perform threshold crossings and recovery is possible; if it inverts under cytokine pressure, the gateway to collapse opens regardless of other interventions. Therapeutic strategies of this class include WNT agonists, BMP pathway stabilizers, and agents that competitively block TNF/IL-1 β suppression of the homeostatic cassette.

The second class targets the attractor geometry: restore the convergent lining phenotype by reproducing the embryonic conditions—hypoxia and EGFR signaling—that drive both lineages to SOX5+ HBEGF+. This strategy addresses collapse not by preventing it but by re-establishing the attractor after it has been disrupted, driving endogenous precursors back toward the basin of the stable lining phenotype. Prediction P6 is the test of this strategy’s feasibility: if the attractor is genuinely trajectory-independent, then precursors from either lineage can be driven to it under matched conditions, and the lining’s homeostatic admissibility structure can be rebuilt without requiring the specific PI16+ population that originally produced it.

These two strategies are not alternatives but operate at different points in the collapse process. Maintainer protection is most relevant early, when the PI16+ cassette is still intact and the goal is to prevent the functional inversion threshold from being crossed. Attractor restoration is most relevant after collapse has occurred, when the maintainer has already inverted and the goal is to re-establish a stable admissibility geometry from whatever precursor populations remain. The formal distinction between the two—one prevents Σ from being crossed in the collapse direction, the other re-establishes the basin at s^* after crossing—clarifies why clinical strategies targeting the two objectives should be evaluated separately rather than treated as interchangeable.

10.3. The Stoichiometric Parameter as a Predictive Clinical Index

The clinical synthesis formalizes the collapse rate model as a quantitative prediction: under fixed inflammatory challenge, the rate at which $|\mathcal{A}_H(x, t)|$ decreases is proportional to the local PI16+ density $\rho(x)$:

$$\frac{d|\mathcal{A}_H(x, t)|}{dt} < 0 \quad \text{with rate} \propto \rho(x), \quad c(x, t) > c_{\text{threshold}}. \quad (9)$$

The clinical implication is that $\rho(x)$ measured during quiescent periods—when active inflammation does not confound the cell count—provides a predictive stoichiometric index of future collapse vulnerability. A joint site with high $\rho(x)$ is not currently pathological but carries more substrate for functional inversion; it is, in the language of the monograph, a site where the residual geometry $\mathcal{G}_{\text{res}}(x)$ encodes a larger admissibility maintenance capacity, and therefore a larger capacity for that maintenance to be inverted under sufficient inflammatory pressure.

This gives the monograph’s epigraph its clinical operational form. The PIP joint is more vulnerable because of what it was—because its embryonic trajectory allocated a higher $\rho(x)$ of PI16+ fibroblasts to its complex interstitial volume—and that historical allocation is measurable in the present tissue. The stoichiometric index $\rho(x)$ across joint sites is a residual geometry measurement: it reads the developmental trajectory from its contemporary trace, in precisely the sense that Bayesian inference on the Yarncrawler structure requires. A biopsy taken during remission is not a measure of current pathology; it is a measure of the developmental trajectory that determined the joint’s position in the admissibility landscape, and therefore of its future vulnerability under conditions that have not yet arrived.

11. Toward a Unified Account

11.1. What the Framework Achieves

The admissibility framework achieves three things that are genuinely valuable independent of whether all twelve predictions are confirmed. First, it provides a precise formal vocabulary for a class of system behavior—threshold-dependent relaxation and mechanism-dependent collapse—that was previously described only in domain-specific terms. This vocabulary is not merely relabeling; it makes cross-domain comparisons tractable and identifies structural features (bipartite decomposition, coupling kernels, maintainer populations) that would not be visible without a common formal language.

Second, it generates predictions that are non-trivial in the sense that they would not be predicted by simpler models. The prediction that tonic inhibition at any depth produces no SWA reduction (P1) would be false if the threshold were smeared or if the relationship between firing reduction and relaxation were monotone. The prediction that collapse rate is a monotone function of PI16+ density (P4) would be false if the joint’s inflammatory susceptibility were primarily determined by cytokine concentration rather than cellular composition. These predictions could fail, and their potential failure makes them scientifically valuable.

Third, it unifies the relaxation-versus-collapse distinction across domains in a way that has clinical and engineering implications. In the joint, the distinction implies that therapeutic strategies targeting relaxation (restoring the system’s ability to perform threshold crossings) are categorically different from strategies targeting collapse prevention (preserving the PI16+ homeostatic function). In the cortical

domain, it implies that sleep deprivation strategies that reduce activity without inducing bistability are categorically different from strategies that induce off-period patterns—a distinction with direct implications for cognitive performance interventions.

11.2. The Appropriate Epistemic Status

The appropriate epistemic status of the admissibility framework at the present stage is that of a well-structured theoretical hypothesis with significant empirical support in several domains, partial support in others, a clear identification of the remaining formal problems, and a specific set of predictions that could falsify it. It is not yet a derived theory—the admissibility field is introduced axiomatically rather than derived from the RSVP field equations—but it is more than a descriptive vocabulary because it generates non-trivial cross-domain predictions.

The monograph is intellectually honest about this position. It does not claim to have proven the framework; it claims to have established it as a productive research program with empirical content. The twelve predictions are the measure of that content. The open mathematical problem is the measure of what remains to be done. This is the correct posture for a theoretical framework at this stage of development, and it is the posture the monograph consistently maintains. The next section develops what a decisive experimental program targeting the framework’s two most vulnerable claims would look like.

12. An Experimental Research Program: Testing the Two Central Claims

12.1. The Priority Ordering

Among all twelve predictions in the research program, two are most important because they target the framework’s deepest structural claim—that relaxation and collapse are not merely different in degree but different in kind, governed by different operative variables. Prediction P1 (cortical threshold crossing as the operative variable for relaxation) and Prediction P4 (PI16+ density as the rate-determining parameter for collapse) are the two experiments that, if both confirmed, would establish the relaxation-collapse distinction as having genuine causal and predictive content rather than merely descriptive utility.

The priority ordering is: cortex first, joint second. The reason is that the cortical test is cleaner in its causal structure. The monograph already has a clean contrast—tonic firing-rate reduction does not produce relaxation, bistable traversal does—and the halorhodopsin experiment has isolated the critical variable. The joint test requires the additional step of manipulating $\rho(x)$ directly to establish causation rather than correlation. The cortical test can proceed with a well-specified parametric protocol; the joint test requires the cortical result as prior motivation before committing resources to the more complex manipulation arm.

12.2. The Cortical Crossing Protocol

The decisive cortical experiment varies three parameters independently: depth of inhibition, duration of inhibition, and number of discrete off-period crossings. The framework’s Proposition 3.1 predicts that SWA reduction, synaptic renormalization (GluA1 and pGluA1 levels), and memory rescue track crossing count $\chi(\gamma)$, not total inhibition dose measured as the product of depth and duration.

The formal prediction takes the following form. Let D denote inhibition depth, τ duration, and n the number of discrete threshold crossings. The admissibility framework predicts:

$$\Delta\text{SWA} = f(n), \quad \left. \frac{\partial(\Delta\text{SWA})}{\partial D} \right|_{n,\tau} \approx 0, \quad \left. \frac{\partial(\Delta\text{SWA})}{\partial \tau} \right|_{n,D} \approx 0. \quad (10)$$

That is, ΔSWA is a function of crossing count alone, and is approximately independent of depth and duration when crossing count is held fixed. A result showing that 100 shallow crossings produce relaxation while a single deep inhibition of equivalent total suppression-area does not would be decisive evidence for the discontinuity claim.

One refinement beyond the standard parametric design is important: the experiment should test whether crossings require completion—both entry into the minimum-admissibility silence region and return to the operating state—or whether mere entry suffices. The monograph’s relaxation operator definition specifies threshold-surface traversal, implying both crossing and return. A protocol that induces off-periods but prolongs or prevents the return trajectory (if physiologically achievable through targeted excitatory reinforcement at the silence minimum) would test whether the \mathcal{A}_H expansion is carried by the crossing itself or by the return trajectory. If expansion requires return, the relaxation operator is non-local in time in a specific sense: the system must visit the minimum and leave it before the field expands. If expansion occurs upon entry, the operator is simpler. This distinction would substantially constrain the RSVP potential landscape interpretation, since a return-dependent operator implies that the outgoing trajectory from s_{\min} carries positive entropy increment $\Delta S > 0$, while an entry-dependent operator does not.

12.3. The Joint Collapse Protocol

The joint test of Prediction P4 requires establishing that collapse rate under fixed TNF/IL-1 β challenge is a monotone function of baseline PI16+ cell density $\rho(x)$ across joint sites, not merely correlated with it. A correlational design—stimulating different joint sites with matched cytokine challenge and measuring collapse rate as a function of observed baseline $\rho(x)$ —would support the monotonicity claim but could not rule out confounders (mechanical environment, vascular density, other fibroblast populations) that covary with $\rho(x)$ across joint sites.

The causal design requires direct manipulation of $\rho(x)$. In animal models, this would be accomplished through embryonic stoichiometry perturbation: artificially increasing or decreasing PI16+ cell allocation to specific joint sites during the de-

velopmental window when the stoichiometric asymmetry is established (approximately 14–15 post-conception weeks in humans, with analogs in mouse fetal development). Measuring collapse vulnerability under adult inflammatory challenge in animals with experimentally altered $\rho(x)$ would establish the causal direction: if joints with artificially increased PI16+ density collapse faster under matched challenge, and joints with reduced density collapse more slowly, then $\rho(x)$ is the rate-determining parameter rather than a correlated predictor.

The table below summarizes the operative variable and test prediction for each experimental arm:

Mode	Operative variable	Test prediction
Relaxation	$\chi(\gamma)$ (crossing count)	$\Delta\text{SWA} \propto n$, not inhibition dose
Collapse	$\rho(x)$ (maintainer density)	Collapse rate \propto baseline PI16+ density

Table 1: Two experimental arms targeting the relaxation-collapse distinction. The framework survives if both predictions hold; it requires revision if either fails.

A framework that survived both tests would have established that the relaxation-collapse distinction is causally grounded, not merely descriptively useful. The cortical test proves that crossing is causally necessary for relaxation; the joint test proves that maintainer density is causally rate-determining for collapse.

12.4. The Lineage-Tracking Experiment and the Strong/Weak Maintainer Inversion Claim

The monograph’s characterization of PI16+ collapse as “functional inversion of the maintainer” rests on an implicit assumption that the same cells responsible for homeostatic maintenance under normal conditions are the ones that produce the chemotactic amplification under inflammatory challenge. This assumption—the strong version of the maintainer inversion claim—is empirically distinguishable from a weaker version in which PI16+ maintenance fails while a distinct subset expands to fill the inflammatory niche.

Definition 12.1 (Strong and weak maintainer inversion). The strong version of maintainer inversion holds if the cells expressing the homeostatic WNT/BMP/FGF cassette under normal conditions are the same cells that upregulate CCL2, CCL3, CCL5, IL6, and IL1B under inflammatory challenge. The weak version holds if PI16+ homeostatic function fails or is suppressed under challenge while a distinct cell population—not the same labeled cells—expands to produce the chemotactic response.

The strong version is required for the monograph’s formal claim that the threshold maintainer is “functionally inverted.” The weak version is consistent with collapse—the admissibility structure still fails—but the failure is better described as displacement or loss of the maintainer rather than its inversion. The therapeutic implications differ: strong inversion implies that preventing inversion requires protecting the same PI16+ cells from crossing Σ ; weak displacement implies that

the target is preserving PI16+ function while separately blocking the expansion of the inflammatory subset.

The decisive experiment is lineage-resolved state tracking. PI16+ fibroblasts are genetically labeled while in the homeostatic state—through a reporter construct tied to PI16 expression under homeostatic conditions—and then followed through TNF/IL-1 β challenge via single-cell RNA-sequencing combined with spatial transcriptomics, with trajectory reconstruction through lineage barcoding. The strong version predicts a continuous developmental trajectory:

$$\text{PI16}_{\text{homeostatic}}^+ \longrightarrow \text{PI16}_{\text{transitional}}^+ \longrightarrow \text{chemotactic inflammatory fibroblast,} \quad (11)$$

in which the labeled cells move through a transitional state and emerge with the inflammatory gene expression profile. The weak version predicts that labeled PI16+ cells disappear (through apoptosis, dedifferentiation, or transcriptional silencing of the PI16 reporter) while an unlabeled population expands with the inflammatory profile. The two trajectories are distinguishable in a lineage-barcoded single-cell dataset because barcode identity links the homeostatic and inflammatory states within each cell.

This experiment has implications beyond the immediate immunological question. If the strong version is confirmed—the same cells invert—it establishes that admissibility collapse is a genuine functional inversion of the maintaining mechanism, as the framework claims, rather than a replacement of one cell population by another. That distinction matters for the general theory because functional inversion implies that the threshold maintainer’s dual capacity (homeostatic protection under normal conditions, inflammatory amplification under perturbation) is an intrinsic feature of the cell type, not an accident of tissue composition. The cell is not simply removed from the system when the threshold is crossed; it becomes the instrument of collapse. This would be the strongest possible confirmation of the mechanism-dependent collapse model.

12.5. Open Methodological Questions

Three methodological questions remain open regardless of how the above experiments are designed. The first concerns the definition of “collapse rate” in the joint experiment. The monograph’s formal model gives $d|\mathcal{A}_H(x,t)|/dt \propto \rho(x)$, but $|\mathcal{A}_H|$ is not directly measurable. A proxy is needed—candidate proxies include the rate of CCL2 upregulation, the speed of leukocyte infiltration, or the time to measurable cartilage degradation—and the choice of proxy introduces assumptions about which downstream events faithfully track field collapse rather than reflecting independent pathological processes.

The second concerns the generalizability of the cortical result. Driessen et al. (2026) demonstrate threshold-dependent relaxation in mouse somatostatin-interneuron circuits. Whether the same discontinuity holds in human cortex, in other cortical regions, or in other species is not established. A framework-level claim that threshold surfaces are a general property of cortical admissibility fields requires

evidence across circuits and species. The predicted conservation of the bistability mechanism across mammalian cortex is itself an implication of the framework and is independently testable.

The third returns to the derivational gap identified earlier: even if both primary experiments succeed, the operational admissibility field—as measured by its downstream proxies in each domain—is not the same mathematical object as the set-valued map $\mathcal{A}(x, t)$ defined formally in the monograph. Confirming the predictions establishes that the predictions follow from the framework; it does not establish that the formal object is uniquely identified by the data. A framework that generates the correct predictions from an axiomatic field definition is not the same as a framework in which the field definition is derived from more fundamental principles. The experimental program can confirm the former; only theoretical development can establish the latter.

Reading of: *Constraint, Reachability, and Relaxation: Across Biological, Cognitive, Industrial, and Cosmological Systems* (Flyxion, June 2026, v3) · Primary sources: six papers, June 2026

Frameworks: RSVP (Relativistic Scalar-Vector Plenum) · CLIO (Constraint-Local Inference and Observation)

Sections 5–7 integrate the *Universal Admissibility Geometry* derivative documents (NotebookLM, June 2026)

Section 9 developed from critical dialogue on the experimental program, June 2026

References

- [1] M. A. Abramowicz, B. Czerny, J. P. Lasota, and E. Szuszkiewicz. Slim accretion disks. *Astrophysical Journal*, 332:646, 1988.
- [2] Hannah Alfonsa, Edward M. Merricks, Darpan L. Bhatt, et al. Intracellular chloride regulation mediates local sleep pressure in the cortex. *Nature Neuroscience*, 26:64–78, 2023.
- [3] Ziga Avsec, Vikram Agarwal, Daniel Visentin, et al. Effective gene expression prediction from sequence by integrating long-range interactions. *Nature Methods*, 18:1196–1203, 2021.
- [4] Bernard W. Balleine and John P. O’Doherty. Human and rodent homologs in action control. *Neuropsychopharmacology*, 35:48–69, 2010.
- [5] Christopher D. Buckley, Caroline Ospelt, Steffen Gay, and Kim S. Midwood. Location, location, location: how the tissue microenvironment affects inflammation in RA. *Nature Reviews Rheumatology*, 17:195–212, 2021.
- [6] Colin J. Burke et al. A characteristic optical variability time scale in astrophysical accretion disks. *Science*, 373:789–792, 2021.
- [7] Edward M. Cackett et al. Accretion disk reverberation mapping of Mrk 142. *Astrophysical Journal*, 896:1, 2020.
- [8] Adam P. Croft, João Campos, Kerstin Jansen, et al. Distinct fibroblast subsets drive inflammation and damage in arthritis. *Nature*, 570:246–251, 2019.
- [9] S. Davidson, D. Simone, K. Jansen, et al. The embryonic origins of site-specific arthritis. *Nature Immunology*, 2026.
- [10] Peter Dayan and Bernard W. Balleine. Reward, motivation, and reinforcement learning. *Neuron*, 36:285–298, 2002.
- [11] Carl G. de Boer and Jussi Taipale. Hold out the genome: a roadmap to solving the cis-regulatory code. *Nature*, 625:41–50, 2024.
- [12] Luisa De Vivo, Michele Bellesi, William Marshall, et al. Ultrastructural evidence for synaptic scaling across the wake/sleep cycle. *Science*, 355:507–510, 2017.
- [13] Amir Dezfouli and Bernard W. Balleine. Actions, action sequences and habits. *PLOS Computational Biology*, 10:e1003364, 2014.
- [14] Graham H. Diering, Raju S. Nirujogi, Richard H. Roth, Paul F. Worley, and Ananya Bhattacharyya. Homer1a drives homeostatic scaling-down of excitatory synapses during sleep. *Science*, 355:511–515, 2017.
- [15] Kort Driessen, Fabio Squarcio, Giulio Tononi, and Chiara Cirelli. Induction of cortical on/off periods in awake mice fulfills sleep functions. *Nature Neuroscience*, 2026.
- [16] Xiaohui Fan, Eduardo Bañados, and Robert A. Simcoe. Quasars and the

- intergalactic medium at cosmic dawn. *Annual Review of Astronomy and Astrophysics*, 61:373–426, 2023.
- [17] M. M. Fausnaugh et al. Optical continuum emission and broadband time delays in NGC 5548. *Astrophysical Journal*, 821:56, 2016.
- [18] Caitlin M. Funk, Kaspar Peelman, Michele Bellesi, et al. Role of somatostatin-positive cortical interneurons in the generation of sleep slow waves. *Journal of Neuroscience*, 37:9132–9148, 2017.
- [19] Samuel J. Gershman and Nathaniel D. Daw. Reinforcement learning and episodic memory in humans and animals. *Annual Review of Psychology*, 68:101–128, 2017.
- [20] Chengfei Huang et al. Personal transcriptome variation is poorly explained by current genomic deep learning models. *Nature Genetics*, 55:2056–2059, 2023.
- [21] Reto Huber, M. Felice Ghilardi, Marcello Massimini, and Giulio Tononi. Local sleep and learning. *Nature*, 430:78–81, 2004.
- [22] Alexander Karollus, Thomas Mauermeier, and Julien Gagneur. Current sequence-based models capture gene expression determinants in promoters but mostly ignore distal enhancers. *Genome Biology*, 24:56, 2023.
- [23] Mehdi Keramati, Amir Dezfouli, and Payam Piray. Speed/accuracy trade-off between the habitual and the goal-directed processes. *PLOS Computational Biology*, 7:e1002055, 2011.
- [24] Gene C. K. Leung, Anna-Christina Eilers, Christos Panagiotou, et al. Discovery of quasar variability and early accretion disk signatures at cosmic dawn. *Nature Astronomy*, 2026.
- [25] Johannes Linder, Divyanshi Srivastava, Han Yuan, Vikram Agarwal, and David R. Kelley. Predicting RNA-seq coverage from DNA sequence as a unifying model of gene regulation. *Nature Genetics*, 2025.
- [26] Marcello Massimini, Reto Huber, Fabio Ferrarelli, Sean Hill, and Giulio Tononi. The sleep slow oscillation as a traveling wave. *Journal of Neuroscience*, 24:6862–6870, 2004.
- [27] Daisuke Miyamoto et al. Top-down cortical input during NREM sleep consolidates perceptual memory. *Science*, 352:1315–1318, 2016.
- [28] Ida Momennejad, Evan M. Russek, Jin H. Cheong, et al. The successor representation in human reinforcement learning. *Nature Human Behaviour*, 1:680–692, 2017.
- [29] K. Okitsu and Y. Sakai. A single reinforcement learning model to unify habit formation and Pavlovian-instrumental interaction. *Scientific Reports*, 2026.
- [30] Bradley M. Peterson. Reverberation mapping of active galactic nuclei. *Publications of the Astronomical Society of the Pacific*, 105:247–268, 1993.
- [31] N. I. Shakura and R. A. Sunyaev. Black holes in binary systems: observa-

- tional appearance. *Astronomy and Astrophysics*, 24:337–355, 1973.
- [32] Anna E. Spiro, Xinming Tu, Yilun Sheng, et al. A scalable approach to investigating sequence-to-function predictions from personal genomes. *Nature Methods*, 2026.
- [33] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, 2nd edition, 2018.
- [34] Xi Tian, Fei Peng, Jon McKechnie, et al. Sustainable battery recycling through spatial and technological alignment. *Nature Sustainability*, 2026.
- [35] Giulio Tononi and Chiara Cirelli. Sleep and the price of plasticity. *Neuron*, 81:12–34, 2014.
- [36] Giulio Tononi and Chiara Cirelli. Sleep and synaptic down-selection. *European Journal of Neuroscience*, 51:413–421, 2020.
- [37] Vladyslav V. Vyazovskiy, Chiara Cirelli, Myriam Pfister-Genskow, Ugo Faraguna, and Giulio Tononi. Cortical firing and sleep homeostasis. *Neuron*, 63:865–878, 2009.
- [38] Vladyslav V. Vyazovskiy, Umberto Olcese, Erin C. Hanlon, et al. Local sleep in awake rats. *Nature*, 472:443–447, 2011.
- [39] Fan Zhang, Kevin Wei, Kamil Slowikowski, et al. Defining inflammatory cell states in rheumatoid arthritis joint synovial tissues. *Nature Immunology*, 20:928–942, 2019.
- [40] Jian Zhou and Olga G. Troyanskaya. Predicting effects of noncoding variants with deep learning-based sequence model. *Nature Methods*, 12:931–934, 2015.