

The Copy And Paste Model: Refined A Unified Path-Integral Framework for Informational Propagation, Substrate Adaptation, and Biomedical Application to Immunological Memory

Claudio Serdar-Petkovski

Independent Researcher, Systems Geometry Group

Theoretical Extension & Biomedical Formalization

Developed in collaborative dialogue, June 2026

† Primary Architect and Originator of the Dynamic Replication Framework.
Biomedical extension and statistical mechanics refinements developed through
iterative theoretical dialogue.

Signature of Primary Author



Date: 10.08.2026

Abstract

We present the refined Copy And Paste Model (CAPM), extending the modularity of Claudio Serdar-P (CSP) Probability Engine into the domain of biomedical systems, with particular application to immunological substrate memory and allergic sensitization dynamics. The core framework governing instantaneous transmission, validation, and macroscopic lifecycle dynamics of replicated informational configurations across adaptive network topologies is extended in three principal directions: (1) a continuous field reformulation of the discrete path-integral, replacing the discrete sum with a Feynman-style functional integral over tissue geometry; (2) a thermodynamic diffusion operator augmenting the macroscopic population wave, bridging CAPM with reaction-diffusion systems in biochemistry; and (3) a statistical mechanics correction to the compatibility tensor $c(s)$, expressing receptor-allergen binding affinity as a thermodynamic free-energy landscape measurable via calorimetry. A critical biological clarification separates the allergen (static blueprint) from the immune network (active propagation engine), grounding the framework in cellular reality. A fourth parameter, structural resolution ρ_s grounded in T-cell receptor diversity and HLA binding groove specificity extends the phase space to three dimensions $(\alpha, \gamma_{\text{mem}}, \rho_s)$ and generates a polysensitization risk prediction testable against existing immune repertoire sequencing datasets. Applied to allergic disease, the refined engine formalizes a phase space over the immunological learning rate α and forgetting rate γ_{mem} , within which a hard boundary separates tolerance-recovery trajectories from persistent attack-memory states. A resensitization term κ identified empirically from the GRASS immunotherapy trial (Shamji et al., 2021) is added to the Memory-Dissipation Law. Numerical evaluation of the resulting system motivates three further developments reported here. First, the ambient environmental encounter rate is not observable, whereas untreated baseline severity P^* is; we therefore invert the equilibrium condition to derive the former from the latter, which also yields a reachability constraint on fitted parameter sets. Second, with γ_{mem} held constant no protocol can move the post-treatment equilibrium, which invalidates the prescription earlier versions of this work derived; we therefore apply the Memory-Dissipation Law a second time to the regulatory apparatus itself, making γ_{mem} a state variable and yielding a closed-form durable endpoint that reduces exactly to the constant- γ case. Third, this reinterpretation reconciles the two arms of the GRASS trial, whose constant- γ fits are shown to be mutually incompatible with randomisation, as differing in route-dependent induction efficiency rather than in patient properties. The framework generates a further mathematically grounded prediction: an individually optimized exposure protocol connected to the clinical concept of metronomic dosing that drives substrate memory $P_{\text{substrate}}$ toward the healthy baseline P_0 , together with

a schedule-invariance test in which outcome depends on cycle-averaged dose rate alone. Two structural extensions follow. The framework’s parallel memory systems are formalised as a subsystem tensor with cross-population coupling χ_{ij} , which is shown to be a prerequisite rather than a refinement: a single substrate memory makes novelty common to every pathway, where it cancels identically from the partition function. The coupled formulation also generates a sensitization threshold not present in any scalar version, and absorbs tolerance plasticity as one of its entries. The geometric and thermodynamic definitions of receptor compatibility, previously offered as alternatives, are shown to be one quantity in two coordinate systems, which supplies both a consistency condition linking them and a bound on immune escape. The coupling magnitudes of that tensor, unavailable in the literature in any directly usable form, are shown to reduce to two scalars recoverable from standard cytokine-blockade and regulatory-depletion experiments. Resolving the novelty divergence over epitope feature bins rather than over scalars removes a degeneracy in which allergens of differing composition but equal mean were indistinguishable, and establishes the independence of the structural resolution parameter, which the compatibility channel alone leaves confounded. Finally, pattern mutation is formalised as diffusion on the unit sphere of configuration space the projection required by the norm constraint on the identity vector yielding a closed-form decay of receptor compatibility at rate $\eta^2 d$, an escape time, and a dimensionless mutation-to-adaptation ratio. That ratio supplies the framework’s account of why allergic memory persists: an allergen is a static blueprint, the ratio vanishes, and an encoding written against a target that never moves is never invalidated. Parameters that are not separately identifiable from reactivity data are identified as such throughout. The model further predicts a pre-symptomatic acoustic signature resolved here into a substrate-displacement core surrounded by a coherence-loss annulus, with the requirement that three channels agree recast as rejection of independent instrument noise rather than coincidence of independent phenomena detectable by quantitative ultrasound prior to clinical immune activation, connecting the theoretical framework to emerging measurement capabilities in biomedical acoustics.

Keywords: path-integral formalism, informational propagation, substrate memory, immunological dynamics, allergy, quantitative ultrasound, statistical mechanics, free-energy landscape, reaction-diffusion, phase space, resonance window, metronomic dosing, resensitization dynamics, structural resolution, polysensitization, TCR diversity, cognitive geometry, thought as substrate propagation, tolerance plasticity, durable endpoint, parameter identifiability, schedule invariance, numerical falsification, subsystem tensor, coupling matrix, sensitization threshold, pattern mutation, spherical diffusion, escape time, mutation-adaptation ratio.

Contents

1	Introduction and Foundational Philosophy	6
1.1	Thought as Translated Geometry: A Cognitive Template	6
1.2	Dimensional Occupancy Propagation: Native Substrate Detection . .	9
1.3	Vision Origin and Development Methodology	17
1.4	Biomedical Lens Extension: Blueprint and Execution Engine	19
2	The Original Microscopic Transmission Engine	22
2.1	The CSP Probability Engine	22
2.2	The Information Action Functional	22
2.3	Normalisation Discards What the Population Wave Needs	22
2.4	Path-Length Normalisation: The Hybrid Action	24
2.5	The Noise Scale as Temperature	25
2.6	Coupled Micro-Coefficients	26
3	Substrate Memory Dynamics and Noise Coupling	29
3.1	The Memory-Dissipation Adaptation Law	29
3.2	Bounded Systemic Congestion Noise	30
3.3	The Macroscopic Population Wave	30
3.4	Stochastic Seeding	30
4	Refinements for Biomedical Application	32
4.1	Refinement I: Continuous Field Reformulation	32
4.2	Refinement II: Thermodynamic Diffusion Augmentation	34
4.3	Refinement III: Statistical Mechanics Correction to $c(s)$ with Structural Resolution	35
4.4	The Two Forms of $c(s)$ Are One Quantity	37
5	Allergy as a CAPM Substrate Rewriting Problem	39
5.1	The Core Mapping	39
5.2	The Observable Baseline and the Ambient Exposure Inversion	40

5.3	Why α and γ_{mem} Are Pathological in Allergy	43
5.4	The Phase Space and the Tolerance Boundary	43
5.5	Structural Resolution and the Three-Dimensional Phase Space	45
6	The Resonance Window: A Novel Prediction	48
6.1	Derivation	48
6.2	The Counterintuitive CAPM Warning	49
6.3	The Optimal Waveform	50
6.4	Tolerance Plasticity and the Durable Endpoint	50
6.5	Empirical Reassessment: The GRASS Trial	53
7	The Subsystem Tensor	55
7.1	Formulation	55
7.2	The Regulatory Subsystem Absorbs Tolerance Plasticity	55
7.3	The Rest State Is Algebraic	56
7.4	Pathways Belong to Compartments	56
7.5	Novelty Recovered	56
7.6	What Constrains χ	57
7.7	An Unplanned Consequence: The Ignition Threshold	59
8	Pattern Mutation and Escape from Substrate Memory	61
8.1	The Constraint Forces a Projection	61
8.2	The Decay Law for Compatibility	61
8.3	Numerical Verification	62
8.4	Escape Time and the Mutation–Adaptation Ratio	63
8.5	Why Allergy Is the $\eta \rightarrow 0$ Limit	63
9	Pre-Symptomatic Detection via Quantitative Ultrasound	65
9.1	The Pre-Symptomatic Window: Detecting the Execution Engine	65
9.2	What Quantitative Ultrasound Would Detect	65
9.3	The Pre-Symptomatic Time Window	68

10 Cross-Substrate Parameter Calibration	70
10.1 The Invariance Claim, Tested	70
10.2 The Multi-Agent Column: An Independent Confirmation	71
11 Epistemological Boundaries and Limitations	75
11.1 Mutations and Evolving Pattern Identity	75
11.2 The Novelty Channel Is Structurally Weak	75
11.3 Parallel Immune Subsystems	77
11.4 Sensitization Cannot Occur	78
11.5 A Recurring Failure Mode	79
11.6 Assumptions Adopted Without Derivation	80
11.7 Parameter Identifiability	81
11.8 Self-Referential Predictability	82
12 Conclusion and Future Directions	83
A Summary of Key Equations	87

1 Introduction and Foundational Philosophy

A profound structural regularity manifests across macroscopic distribution systems throughout the observable universe. The branching layouts of river deltas, the vascular networks of biological organisms, the layout of human metropolitan road systems, and the dense latent spaces of artificial neural networks all display striking geometric commonalities. These parallel structures are not coincidental; they point toward a fundamental governing principle: the continuous optimization of pattern replication across physical media.

Traditional scientific frameworks frequently analyze informational systems in isolation. Computer science models data routing via abstract network protocols; sociology monitors the spread of ideas via qualitative cultural indicators; and ecology tracks population growth via classical resource models. The Copy And Paste Model (CAPM) bridges these divided silos. The foundational philosophy rests upon a universal evolutionary pipeline:

$$[\text{COPY}] \longrightarrow [\text{PASTE}] \longrightarrow [\text{ADAPT}] \longrightarrow [\text{REPEAT}] \quad (1)$$

Every idea, piece of software, or biological blueprint behaves as an informational configuration seeking survival. To survive, it must absorb environmental energy, recognize the structural layout of its current domain, replicate its pattern onto a new substrate, and continuously optimize its internal code against ambient resistance.

1.1 Thought as Translated Geometry: A Cognitive Template

Before descending into the formal machinery, it is worth establishing a cognitive template that makes the entire framework easier to see. The insight is this: *thought itself is geometry, translated through successive substrates.*

Consider the act of recognizing a fan. The object’s physical geometry enters the visual system as structured light. The visual cortex does not receive a label it receives *shape*: edges, rotational symmetry, spatial periodicity. This geometric signal propagates inward, is matched against stored geometric templates, crosses the boundary into linguistic substrate, and produces the word “Fan.” The word is not the primary event. It is the *compressed label* the final substrate assigns to the geometry that arrived before it. Thought, on this account, is what geometry looks like after passing through enough substrate crossings to become conscious.

This can be formalized as a propagation pipeline:

$$\begin{aligned}
 &[\text{GEOMETRY}] \xrightarrow{\chi_1} [\text{VISION}] \xrightarrow{\chi_2} [\text{NOVEL PATTERN}] \xrightarrow{\chi_3} [\text{CONTEXT}] \\
 &\xrightarrow{\chi_4} [\text{LINGUISTIC FORM}] \xrightarrow{\chi_5} [\text{LANGUAGE}] \xrightarrow{\chi_6} [\text{THOUGHT (GEOMETRY)}]
 \end{aligned} \tag{2}$$

Each arrow is a substrate crossing carrying a Bridge Operator Cost χ_j from Equation (12) the energy paid to re-encode the pattern in the new medium. The remarkable observation is that the pipeline is *closed*: Thought at the end of the chain is itself geometric in structure. It feeds back into perception, primes future pattern-matching, and alters the familiarity profile $P_{\text{substrate}}$ of the cognitive network for all subsequent encounters with that geometry.

Definition 1 (Cognitive Substrate Propagation). *A thought is an informational configuration whose pattern $\vec{p}$ originates as physical geometry, propagates across n sequential substrate crossings $\{\chi_j\}_{j=1}^n$, and arrives in conscious representation as a re-encoded geometric structure whose compatibility profile $C(\gamma) = \prod_{s \in \gamma} c(s)$ reflects the cumulative alignment between the original geometry and every intermediate processing medium.*

Three consequences follow directly from this framing and from the CAPM machinery introduced in subsequent sections:

- (i) **Mathematical intuition is geometry recognizing itself.** When a mathematician “sees” why a proof is true before articulating it, the geometric pattern has propagated far enough to activate high-compatibility pathways in the cognitive substrate but has not yet completed the full translation into linguistic form. The feeling of insight is the system catching the pattern mid-propagation, before the final χ_j is paid.
- (ii) **Thought is a self-modifying substrate.** The Memory-Dissipation Law (Equation (24), introduced in Section 3) governs not only immunological and network substrates but cognitive ones. Each time a geometric pattern completes the pipeline and becomes a named thought, the substrate updates its familiarity profile $P_{\text{substrate}}$ reducing the novelty $V(s, t)$ of that pattern on all future encounters and lowering the action cost $S[\gamma, t]$ of the path that recognized it. This is the CAPM formalization of *learning*.
- (iii) **Language is not the origin of thought but its final substrate.** The pipeline in Equation (2) clarifies a long-standing debate in cognitive science. Language does not generate thought; it is the last medium through which geometric thought passes before becoming communicable. Thoughts that resist

full linguistic encoding musical intuition, spatial reasoning, aesthetic perception are patterns whose χ_j cost at the geometry-to-language boundary is high: the Bridge Operator resists the translation, and the pattern remains partially encoded in an earlier substrate.

This template is not a metaphor layered onto the mathematics. It is a direct application of the same CAPM engine that governs information propagation across silicon networks and immune tissue. The substrate changes; the governing equations do not. That invariance is the model's deepest claim.

The Coherence Law Is Not an Additional Postulate

The propagation of coherence across substrate crossings does not require an axiom of its own. Equation (11) weights each pathway by $e^{-S[\gamma,t]/\sigma}$, and Equation (12) carries the bridge costs inside S as a bare sum $\sum_{j \in B_t} \chi_j$, unmodulated by novelty or compatibility. The factor contributed by the boundaries alone is therefore

$$C_{\text{total}} = \exp\left(-\frac{1}{\sigma} \sum_{j=1}^n \chi_j\right) \quad (3)$$

This is the boundary term of the action functional pushed through the path-integral weight, not a separate law, and the attenuation constant that earlier versions of this work fitted is thereby fixed: $k \equiv 1/\sigma$, the inverse of the same environmental noise scale that governs Sections 2–7.

The substitution converts a free parameter into a testable consistency requirement, and the requirement is met. Reading each reported coherence value backwards for the σ it implies gives 1.6548, 1.6650, 1.6672 and 1.6755 for the four patterns treated below a spread of 0.4% about $\sigma = 1.666$, hence $k = 0.600$. A single noise scale reproduces every published figure. That value also lies inside the range the biomedical sections operate in, so the cognitive and immunological substrates are consistent with one temperature. The invariance claim of Section 1 could have failed this check and does not.

Two Pipelines, One Notation

Equation (2) and the dimensional cycle of the following subsection are different objects written in the same symbols, and the difference matters once Equation (3) fixes what those symbols do. Equation (2) carries six crossings, χ_1 through χ_6 ; the cycle carries four. Since coherence attenuates as $\exp(-\sum_j \chi_j/\sigma)$, counting the same

pattern through the two gives different answers, and not by a negligible margin.

At the per-crossing cost implied by the published figures, $\chi \approx 0.115$, four crossings give $C_{\text{total}} = 0.759$ and six give 0.661. Under the regime boundaries of the following subsection that is the difference between *Resonant* and *Coherent* a change of category for the same pattern, produced entirely by how many boundaries were counted. Every published coherence in this paper is a four-crossing figure, and should be read as applying to the cycle rather than to Equation (2).

The resolution is not to reconcile the counts, which describe genuinely different traversals, but to state the crossing set alongside any coherence value. A coherence without its pipeline is not a well-defined number.

What Equation (2) Does and Does Not Deliver

The closure claim of Equation (2) that thought at the end of the chain is itself geometric, feeds back into perception, and lowers the novelty of that geometry on subsequent encounters is not decoration. It is Equation (24) applied to a cognitive substrate, and it can be checked. It holds: iterating the traversal, novelty falls monotonically from 0.210 to 0.050 over seven encounters, with diminishing decrements throughout. The forward pass does what the section claims.

The *selectivity* it claims does not hold. Equation (2) asserts that a completed traversal alters the familiarity profile for *that* geometry. In the formalism $P_{\text{substrate}}$ is a single scalar per node, so learning one pattern raises familiarity with every pattern equally: two profiles diverging by 0.490 evaluate to novelties of 0.0335 and 0.0348 at the same substrate state, and the states “has learned *A*” and “has learned *B*” are indistinguishable. A substrate cannot be familiar with one geometry and naive to another at the same node.

The repair is available and already in the paper. Equation (19) resolves the substrate over feature bins precisely so that familiarity can be selective, and is applied only to the allergen case. Extending it to the cognitive substrate would make selectivity structural rather than asserted; we note it as available rather than claiming it, since the feature bins of a geometric pattern are not obviously the epitope bins of a protein.

1.2 Dimensional Occupancy Propagation: Native Substrate Detection

A critical question emerges immediately from the framework: *where does an informational configuration naturally enter the dimensional cycle?* Not all patterns enter through the same gateway. A musical melody enters through 3D acoustic

space; written text enters through 2D visual space; a mathematical principle enters through 4D abstract space; a gesture enters through 3D kinesthetic space. Yet the CAPM equations remain invariant across all substrates. This section establishes the principle by which the model auto-detects the native substrate of any pattern.

The Native Substrate is Measured, Not Chosen

Consider two words of similar phonological structure: “onomatopoeia” and “brum-bartianoal.” Both contain 11–16 letters. Both are polysyllabic sequences. Yet one is remarkably sticky (you remember it forever) while the other is genuinely forgettable. The CAPM predicts this difference entirely through dimensional occupancy and substrate coherence.

The Native Substrate is Measured, Not Chosen (But Observer-Relative)

Onomatopoeia is fundamentally a word about *sound*. For a hearing observer, its native substrate is 3D acoustic space the actual phonemic sequence that the word itself mimics. When the pattern cycles through all dimensional boundaries:

$$[3D \text{ acoustic}] \rightarrow [2D \text{ visual}] \rightarrow [4D \text{ abstract}] \rightarrow [1D \text{ linguistic}] \rightarrow [3D \text{ acoustic}] \quad (4)$$

the bridge costs χ_j at each crossing remain minimal. Why? Because each transition respects the pattern’s native structure: acoustic-to-visual is natural (you write down what you hear), visual-to-abstract is efficient (the written form clearly encodes sound-meaning isomorphism), abstract-to-linguistic is direct (the concept compresses to the word itself), and linguistic-to-acoustic recovery is nearly perfect (you hear the word and the sound returns spontaneously). The final coherence C_{total} is high, and by Equation (7) the forgetting rate γ_{mem} is correspondingly low. You never forget onomatopoeia.

But for a **deaf observer**, the 3D acoustic substrate is inaccessible. The optimal native substrate within their available set S_{obs} must be different. The cycle might instead begin at 2D visual:

$$[2D \text{ visual}] \rightarrow [3D \text{ spatial}] \rightarrow [4D \text{ abstract}] \rightarrow [1D \text{ linguistic}] \rightarrow [2D \text{ visual}] \quad (5)$$

This achieves lower coherence than the acoustic route, because the pattern “onomatopoeia” (literally “sound-making”) has no direct visual referent. Yet this is still the *optimal native substrate for the deaf observer* given their constraints. The CAPM predicts they will remember onomatopoeia less effortlessly than a hearing

observer not because the model is wrong, but because their substrate availability is fundamentally different.

In contrast, *brumbartianoal* is a phonetically complex nonsense word with zero semantic grounding across all dimensional substrates. Regardless of which substrate is the entry point, every crossing incurs massive bridge costs because there is no meaningful content to carry through the cycle. The cycle beginning at 3D acoustic (the “best” entry point for this word):

$$[3D \text{ acoustic}] \rightarrow [2D \text{ visual}] \rightarrow [4D \text{ abstract}] \rightarrow [1D \text{ linguistic}] \rightarrow [3D \text{ acoustic}] \quad (6)$$

incurs catastrophically high bridge costs at every transition: acoustic-to-visual costs energy because there is no semantic anchor ($\chi_1 = 0.58$), visual-to-abstract is even worse because there is no meaning to extract ($\chi_2 = 0.65$), abstract-to-linguistic fails because there is no concept ($\chi_3 = 0.62$), and linguistic-to-acoustic recovery returns only meaningless noise ($\chi_4 = 0.59$). The pattern hemorrhages coherence at **every boundary**. C_{total} collapses. γ_{mem} drops to near zero. You forget *brumbartianoal* immediately and completely.

Critically: the native substrate is *not chosen arbitrarily by the observer*. It is *measured by the model* given the observer’s specific substrate availabilities S_{obs} . The procedure tests all possible entry substrates within that set, runs the complete 4-crossing dimensional cycle from each, and identifies which yields the highest final coherence for *that observer*. The prediction follows directly from the coherence value and the observer’s geometry.

This is neither absolute nor arbitrary. It is constrained by biology, neurology, and sensory capacity. An observer without auditory access cannot optimize through acoustic substrates. A child without abstract reasoning cannot use 4D. A visual learner will spontaneously encode through visual substrates. The model makes precise predictions about how substrate availability shapes cognition and memory.

Substrate Detection Algorithm

Remark 1 (The algorithm is not runnable as specified). *The procedure below ranks five candidate entry substrates by the coherence each yields, and requires as input the four bridge costs χ_j along each of those five routes twenty per pattern. The paper supplies four: the costs along the route that wins.*

For the one worked example with several routes, Control-Alt-Delete, the reported quantities are the final coherences of the four candidates, 0.26, 0.27, 0.30 and 0.79. Inverting them through Equation (3) gives per-crossing costs of 0.561, 0.545, 0.501

and 0.098: the kinesthetic route is asserted to be some five times cheaper per boundary than the others, and nothing in the framework derives that ratio. **The coherences the algorithm is meant to compute are supplied as its inputs**, so the ranking is a restatement of an assumption rather than a derivation from one.

This does not make the conclusion wrong Control-Alt-Delete plainly is remembered kinesthetically but it does mean the framework predicts nothing here that was not put in. Making the procedure runnable requires χ_j to be derived rather than assigned, and the natural candidate is the one the framework already uses for the same job elsewhere: a divergence between the pattern as represented in the source substrate and as represented in the target, in the manner of Equation (17). That is a research task rather than an adjustment, and until it is done the substrate-detection algorithm should be read as an illustration of the coherence law rather than as an application of it.

For any informational pattern (a word, a gesture, a concept, an immune signature, a symptom), the native substrate is determined through the following procedure:

1. Generate the full dimensional cycle (all four substrate crossings) from each of the five possible starting substrates: 2D visual, 3D spatial, 3D acoustic, 3D kinesthetic, 1D linguistic, 4D abstract.
2. For each starting substrate, compute the final coherence C_{total} after traversing all four bridge-operator costs.
3. Rank the substrates by their final coherence values.
4. The substrate yielding the highest C_{total} within the observer's accessible substrate set S_{obs} **is the optimal native substrate** for that specific observer and pattern.
5. Predict the memory persistence rate:

$$\gamma_{\text{mem}} = \omega_r (1 - C_{\text{total}}) \quad (7)$$

where $\omega_r > 0$ is the rate at which the substrate attempts recall. The form follows from Equation (11): a pattern is retrieved with Boltzmann weight C_{total} per attempt and is therefore lost at a rate proportional to $1 - C_{\text{total}}$.

Note the direction. γ_{mem} is the **forgetting** rate throughout this paper (Section 3.1), so a highly coherent pattern must carry a *low* γ_{mem} . Earlier versions of this section wrote $\gamma_{\text{mem}} \approx 0.8 C_{\text{total}}$ and described γ_{mem} as a persistence parameter, inverting the convention used everywhere else: under that relation

onomatopoeia, at $C = 0.7596$, would be forgotten more than three times faster than *brumbartianoal* at $C = 0.2331$, which is the opposite of the claim the surrounding text makes. Equation (7) restores the intended direction, giving retention times $1/\gamma_{\text{mem}}$ of 4.16 and 1.30 in units of $1/\omega_r$ respectively.

Critical clarification: Native substrate is observer-relative. The optimal encoding depends entirely on which dimensional substrates are available to the specific observer. A deaf learner cannot access 3D acoustic substrates; their optimal encoding of “onomatopoeia” will necessarily be visual or kinesthetic, with correspondingly lower coherence than for a hearing observer. A prelinguistic infant lacks access to 4D abstract substrates; memory encoding must occur in earlier substrates. A blind individual cannot use 2D visual substrates; their optimal encoding will be acoustic or kinesthetic.

This observer-relativity is not a weakness of the framework. Rather, it becomes a strength: the CAPM now predicts precisely how available substrate set determines optimal encoding. Different observers with different substrate availabilities will achieve different coherence values for the same pattern not because the pattern changed, but because their accessible substrate geometry changed. The model predicts this variability.

This algorithm produces a striking empirical correspondence: patterns with high final coherence are exactly those that persist in memory. Those with low coherence are forgotten readily. The “stickiness” of a pattern is not determined by its emotional salience, its frequency of exposure, or its semantic importance. Rather, stickiness is determined entirely by how cleanly and efficiently the pattern propagates through the dimensional cycle beginning from its native substrate.

Empirically, three coherence regimes emerge:

- **Resonant** ($C_{\text{total}} > 0.70$): Patterns that achieve near-perfect coherence across all dimensional boundaries. These patterns are essentially unforgettable. The pattern returns intact through every crossing. γ_{mem} is high. Example: *onomatopoeia* achieves $C = 0.7596$ in its native 3D acoustic substrate.
- **Coherent** ($0.60 < C_{\text{total}} \leq 0.70$): Patterns that maintain moderate coherence but with some information loss. These patterns are memorable but can fade with disuse. γ_{mem} is moderate.
- **Incoherent** ($C_{\text{total}} < 0.23$): Patterns with no semantic grounding across dimensional boundaries. These patterns lose coherence catastrophically and are forgotten immediately. γ_{mem} is near zero. Example: *brumbartianoal* achieves only $C = 0.2331$ even in its optimal 3D acoustic substrate.

The three-word spectrum illustrates the prediction power:

- Onomatopoeia (3D acoustic): $C = 0.7596$ (RESONANT)
- Snowman (2D visual): $C = 0.6970$ (COHERENT)
- Brumbartianoal (3D acoustic): $C = 0.2331$ (INCOHERENT)

Onomatopoeia persists in memory forever because it achieves resonance. Brumbartianoal vanishes from memory immediately because it is incoherent across all substrates. The model makes precise, quantitative predictions about which patterns stick and which are forgotten.

Observer-Relativity: The Power and the Constraint

The insight that native substrate is observer-relative transforms the model from a purely geometric claim into a framework that predicts neurodevelopmental, sensory, and individual cognitive differences. Consider the implications:

Sensory Loss: A person who loses hearing midlife has an existing optimal substrate set that included 3D acoustic. Post-deafness, that substrate becomes inaccessible. The CAPM predicts they must re-encode patterns through remaining substrates (visual, kinesthetic, abstract). Memory coherence for acoustic-primary patterns will drop. This is not failure of the system but adaptation to new substrate geometry.

Learning Modality: A student who is strongly visual will spontaneously choose 2D visual as their preferred native substrate for abstract mathematical concepts. A kinesthetic learner will prefer 3D spatial/kinesthetic routes. Both achieve memory encoding, but through different substrates. The model explains why “learning styles” matter without reducing to pseudoscience: different substrates achieve different coherence levels for the same content.

Neurodevelopmental Progression: A young child (age 3–5) lacks fully developed abstract reasoning (4D). Their substrate set S_{obs} is limited to 2D visual, 3D spatial, 3D kinesthetic, and 1D linguistic. They cannot achieve the coherence available to older children who access 4D. Yet they *can* achieve optimal encoding within their available set. The CAPM predicts they will remember visual and kinesthetic patterns better than abstract ones which is empirically true.

Clinical Intervention Design: For an allergic patient with low immune substrate resolution ρ_s (poor discrimination), the accessible substrate set for encoding the allergen pattern may be limited. Treatment should be designed to disrupt the

pattern specifically within the patient’s native substrate to minimize coherence within their available encoding space. One-size-fits-all protocols fail because they ignore observer-relative substrate geometry.

The observer-relative framework is therefore **not a limitation but a prediction engine**. It predicts how cognitive differences, sensory differences, developmental stage, and individual learning geometry will affect memory and encoding. This connects the abstract mathematics to observable human diversity.

Kinesthetic Occupancy: The Control-Alt-Delete Case Study

A particularly illuminating case is the computer command sequence “Control-Alt-Delete.” On the surface, this is a three-word linguistic phrase. Yet human memory does *not* treat it linguistically. Instead, it is remembered as *kinesthetic substrate occupancy* as a muscle memory, as the simultaneous pressing of three physical keys, as an embodied gesture that your hands execute without conscious mediation.

When the substrate detection algorithm is applied, testing all five possible starting substrates:

- **2D visual** (reading the text “Control-Alt-Delete” on a screen): Final coherence ≈ 0.26 (weak)
- **1D linguistic** (speaking the words aloud): Final coherence ≈ 0.27 (weak)
- **4D abstract** (understanding the concept “system interrupt command”): Final coherence ≈ 0.30 (weak)
- **3D kinesthetic** (pressing the three keys simultaneously): Final coherence ≈ 0.79 (very high)

The native substrate is unambiguously 3D kinesthetic. The prediction: you will remember Control-Alt-Delete through embodied muscle memory, not through reading, speaking, or understanding. You learned it *by doing*. Your hands remember it when your conscious mind forgets it. The CAPM model predicts this empirical fact from first principles: the substrate with highest coherence is the substrate where the pattern naturally lives.

This case demonstrates that the dimensional occupancy framework does not assume that patterns entered verbally or rationally. Patterns enter through whatever substrate minimizes initial bridge cost χ_j . For computer users, the “Control-Alt-Delete” pattern entered through the kinesthetic learning pathway, yielding the highest coherence and therefore the strongest encoding in γ_{mem} .

Biological and Clinical Implications for Immunological Substrate

The principle of native substrate detection extends directly to immunological systems. An allergen pattern may be encoded across multiple substrates simultaneously: (i) its chemical molecular structure, (ii) the immune cell activation signature it triggers, (iii) the inflammatory cascade it initiates in tissue. Yet the CAPM predicts that the immune system's encoding will achieve maximal coherence in exactly one of these substrates the one representing the initial route of sensitization and the one matching the immune network's native resolution ρ_s .

Individual variability in allergy severity, sensitization trajectory, and treatment response becomes directly explicable: two patients with apparently identical allergen exposure histories may have radically different disease courses if their immune substrates initially occupied different dimensional spaces. Patient A may have acquired sensitization primarily through respiratory (3D acoustic-analogue) exposure, achieving high coherence in a particular immune substrate dimension. Patient B may have acquired it through cutaneous (2D visual-analogue) exposure, encoding it differently.

The CAPM predicts that treatment efficacy should be optimized by targeting the encoded pattern specifically at its native substrate dimension where χ_j is lowest and coherence is highest. This offers a mathematical framework for understanding why individual immunotherapy protocols, tailored to the patient's likely native sensitization route, may outperform generic population-average protocols.

The Detection Procedure Has an Unspecified Input

The procedure above ranks candidate substrates by the coherence each yields and selects the highest. With Equation (3) in hand the ranking step is now exactly computable. The step before it is not.

Nothing in this framework assigns χ_j to a given crossing. The bridge costs are the sole input to the coherence, so the ranking is determined entirely by a quantity the framework does not derive, measure, or constrain. Reading the published Ctrl-Alt-Delete figures backwards for the costs they imply makes the position plain:

candidate substrate	C_{total}	implied $\sum_j \chi_j$	per crossing
2D visual	0.26	2.244	0.561
1D linguistic	0.27	2.181	0.545
4D abstract	0.30	2.005	0.501
3D kinesthetic	0.79	0.393	0.098

Table 1: Bridge costs implied by the published coherences at $\sigma = 1.666$.

The kinesthetic route is assigned a per-crossing cost roughly five times lower than its competitors. **That factor of five is the prediction, and it is an input rather than an output.** Told that pressing three keys is five times cheaper per boundary than reading the words, the conclusion that the pattern lives in kinesthetic substrate follows immediately; but the framework supplies the ranking machinery and not the thing being ranked.

This does not make the procedure wrong. It makes it a *conditional* prediction of the form already familiar from Section 11.5: given the bridge costs, the native substrate follows. Closing it requires a rule assigning χ_j from properties of the two substrates being crossed their dimensionalities, their representational capacities, or the information discarded in translation between them which this paper does not provide. Until then the worked examples should be read as illustrations of the ranking step, not as derivations of their conclusions.

1.3 Vision Origin and Development Methodology

The CSP Probability Engine did not emerge through conventional axiomatic derivation or systematic literature review. Its origin and development methodology are documented here not as biography, but as epistemology: the nature of a framework’s origin constrains the kind of claims it can legitimately make, and transparency on this point is a prerequisite for honest scientific evaluation.

The Core Geometric Observation

The framework originates in a structural observation made across macroscopic distribution systems of radically different physical composition. Terrestrial road networks, biological vascular architectures, neural pathway formations, colony foraging routes, and large-scale human migration patterns all exhibit a striking geometric regularity: each system evolves pathways that minimize cumulative resistance along routes of propagation. Street networks converge toward layouts that minimize travel cost.

Neural circuits form to minimize signal latency. Biological colonies establish routes requiring minimum energetic expenditure. The observation is domain-independent.

This regularity was recognized as an instance of a deeper organizing principle: *least-action optimization governs the geometry of information propagation across physical media*. The CSP Probability Engine is, at its core, a formal statement of this principle. The Boltzmann-weighted path sum in Equation (11) is the mathematical encoding of least-action selection; the action functional $S[\gamma, t]$ in Equation (12) quantifies the cumulative resistance a propagating configuration encounters along each candidate route.

Intuition-First Formalization

The development methodology inverts the conventional sequence of mathematical inquiry. Rather than deriving equations from established axioms and subsequently seeking empirical domains of application, the geometric intuition was observed first and the mathematical framework was constructed afterward to encode it with precision. The formalism is therefore best understood as a *language for expressing an underlying geometric regularity*, not as a proof structure descended from prior theory.

This approach has notable historical precedent. Feynman’s path-integral formulation of quantum mechanics originated in a geometric intuition about least-action trajectories, with rigorous measure-theoretic foundations constructed subsequently. The present framework makes no claim to equivalent mathematical depth; it does claim an equivalent epistemological structure: geometry first, axioms second.

Development Process

The framework crystallized through an extended period of iterative cross-domain examination spanning many months of structured dialogue across multiple analytical systems and research contexts during which the geometric regularity described above was stress-tested against progressively more distant domains. The central question driving each iteration was whether the least-action propagation principle retained predictive structure when the substrate changed entirely: from silicon network routing, to sociological information diffusion, to immunological memory dynamics. The mathematical formalism was refined iteratively to preserve structural invariance across these crossings, with each substrate translation serving as a progressively stronger theoretical stress test of the underlying geometric claim.

This development process is itself an instance of the CAPM framework: the

geometric pattern propagated across successive analytical substrates, paying a Bridge Operator Cost χ_j at each domain crossing, with the invariant core least-action optimization surviving each translation as the framework’s compatibility tensor $c(s)$ with each new domain remained non-zero.

Epistemological Boundaries

The strength of the framework rests not on mathematical elegance alone but on whether the geometric principle it encodes has predictive validity across novel domains not used in its construction. This is the appropriate standard for an intuition-first theoretical framework: not internal consistency, which is necessary but insufficient, but *external predictive reach*.

The remainder of this paper documents how the geometric intuition scales from discrete network topologies (Sections 2–3) through continuous field reformulations (Section 4) to biomedical application (Sections 5–7), and advances testable predictions (Sections 6–7) that would validate or refute the underlying geometric principle on empirical grounds. The framework invites falsification. That openness is not a concession but a commitment.

1.4 Biomedical Lens Extension: Blueprint and Execution Engine

The original CAPM formulation implicitly treated the propagating pattern and the mechanism responsible for its propagation as closely coupled. The biomedical substrate exposes a limitation in this assumption: an allergen does not need to replicate physically in order for information about its molecular structure to propagate through an immune system.

In the biomedical interpretation, the allergen is therefore treated as a **pattern blueprint**, while the immune network functions as the **execution and propagation engine**. The distinction is important. The allergen provides the molecular configuration encountered by the organism; immune cells, signaling pathways, antibody-producing cells, and memory populations determine whether and how that configuration is encoded into the biological substrate.

The relevant CAPM propagation event therefore occurs at the substrate crossing:

$$P_{\text{pattern}} \xrightarrow{\chi_{\text{immune}}} P_{\text{substrate}} \quad (8)$$

where P_{pattern} represents the encountered allergen configuration and $P_{\text{substrate}}$ represents the corresponding state of the immune system.

This formulation does not imply that the allergen itself becomes physically reproduced inside the immune system. Rather, the information necessary to recognize and respond to the molecular pattern is propagated into the biological substrate. In an IgE-mediated response, allergen exposure can initiate antigen presentation, T-helper-cell-associated signaling, allergen-specific IgE production, and subsequent sensitization; subsequent encounters can then reactivate the established response. CAPM interprets this sequence as a substrate-level propagation process rather than as physical replication of the external allergen.

The propagation event should therefore be distinguished from the initial physical encounter. Exposure initiates the substrate crossing through which an external molecular configuration can become encoded into an internal biological state. Propagation subsequently occurs through the immune network as the response is amplified, maintained, modified, or dissipated.

CAPM consequently interprets sensitization as a form of **substrate memory formation**. The external allergen acts as the blueprint, while the immune network performs the copying, amplification, and persistence operations. The resulting state can subsequently influence how the substrate responds to future encounters with the same or related molecular patterns.

The blueprint should not necessarily be interpreted as a single static molecular object. A biologically relevant encounter may instead be represented as a context-dependent pattern:

$$P_{\text{pattern}} = P(s, d, r, c, t, H) \quad (9)$$

where s denotes molecular structure, d exposure dose, r route of exposure, c local biological context, t temporal state, and H host-specific state variables. The same nominal allergen may therefore constitute different effective pattern instances under different substrate conditions.

This distinction resolves the apparent contradiction between CAPM and the conventional description of allergens as non-replicating molecular structures. The object being propagated is not the allergen itself, but the biological information encoded from its encounter.

The resulting interpretation provides a direct bridge to the Memory-Dissipation Law developed later in the model:

$$\partial_t P_{\text{substrate}} = \alpha N(t) P_{\text{pattern}} - \gamma_{\text{mem}} (P_{\text{substrate}} - P_0). \quad (10)$$

Here, $N(t)$ represents the active immune execution population, α the effective learning or encoding coefficient, and γ_{mem} the rate at which the induced substrate state relaxes toward its baseline configuration. Equation (10) is not presented as an established biological law; it is the CAPM hypothesis to be tested against immunological data.

This reformulation constitutes an important boundary test of CAPM: a propagating pattern and the mechanism that propagates it need not be physically identical. If this distinction permits the same mathematical structure to generate empirically testable predictions in biological memory without requiring physical replication of the allergen, the biomedical substrate provides a meaningful extension of the framework. If the proposed mapping fails against measured immune dynamics, the corresponding CAPM component must be revised or rejected.

2 The Original Microscopic Transmission Engine

2.1 The CSP Probability Engine

Let $\Gamma_t = (V_t, E_t)$ define the discrete topological network at time t , where V_t represents system nodes and E_t the active connectivity edges. The replication probability from origin node $s_0 \in V_t$ to destination node $s_f \in V_t$ is:

$$R(t) = \frac{\sum_{\gamma \in \mathcal{P}(s_0 \rightarrow s_f; \Gamma_t)} \gamma_{\text{len}}^{|\gamma|} \cdot C(\gamma) \cdot e^{-S[\gamma, t]/\sigma(t)}}{\sum_{s_{\text{alt}} \in V_t} \sum_{\gamma' \in \mathcal{P}(s_0 \rightarrow s_{\text{alt}}; \Gamma_t)} \gamma_{\text{len}}^{|\gamma'|} \cdot C(\gamma') \cdot e^{-S[\gamma', t]/\sigma(t)}} \quad (11)$$

where $\gamma_{\text{len}} \in (0, 1)$ is the path-length discount factor, $C(\gamma)$ the global structural compatibility of path γ , $\sigma(t)$ the environmental congestion noise, and $S[\gamma, t]$ the accumulated informational action. The denominator is the systemic partition function $\mathcal{Z}$, ensuring $0 \leq R(t) \leq 1$.

2.2 The Information Action Functional

$$S[\gamma, t] = \sum_{s \in \gamma} \frac{E(s)}{1 + \tanh(V(s, t) \cdot L(s) \cdot c(s))} + \sum_{j \in B_t} \chi_j \quad (12)$$

where $B_t \subset \gamma$ is the set of boundary interfaces, $E(s) > 0$ the baseline energy at node s , and $\chi_j \geq 0$ the Substrate-Bridge Operator Cost for cross-substrate translation.

2.3 Normalisation Discards What the Population Wave Needs

Equation (11) defines $R(t)$ as a share: the weight reaching the nominated destination divided by the partition function over all destinations. Any factor common to every pathway therefore cancels exactly, and this has a consequence that is easy to miss and hard to live with.

The symptom. Degrading compatibility *raises* $R(t)$. Sweeping $\Delta G_0/k_B T$ from -4 to $+6$, so that $c(s)$ falls from 0.982 to 0.0025, R rises from 0.592 to 0.9987. Worse matching, higher replication probability. Reversing which pathway is nominated reverses the direction ($0.492 \rightarrow 0.000$), and equalising all path lengths flattens it entirely ($0.5447 \rightarrow 0.5445$). The sign is set by whether the target is shorter or

longer than its alternatives, not by compatibility at all: $c(s)$ does not raise or lower propagation, it amplifies whatever length advantage already exists.

The obstruction is structural. No normalised share can express “compatibility is poor, so less propagates.” That is a statement about the absolute rate, and normalisation is precisely the operation that discards it. Moving $\gamma_{\text{len}}^{|\gamma|}$ and $C(\gamma)$ inside the exponent does not help: the ratio $c^{|\gamma_1|}/c^{|\gamma_2|}$ still diverges as $c \rightarrow 0$, merely more slowly.

The information is not missing, only unused. The partition function itself retains it. Across the same sweep $\mathcal{Z}$ falls from 8.5×10^{-1} to 2.0×10^{-11} ten orders of magnitude, in exactly the direction required. Equation (26) reads only R , and so never sees it.

We therefore let the population wave read both outputs of the path integral: R for *where* propagation goes, and a propagation efficiency for *whether* it goes at all. Since $\mathcal{Z}$ scales as $c^{|\gamma|}$, the per-node efficiency is the quantity dimensionally suited to multiplying a per-node rate:

$$\Theta = \left(\frac{\mathcal{Z}}{\mathcal{Z}_{\text{ideal}}} \right)^{1/\bar{n}}, \quad \frac{dN}{dt} = \omega_0 \Theta R N \left(1 - \frac{N}{K} \right) + \dots \quad (13)$$

with $\mathcal{Z}_{\text{ideal}}$ the partition function at $c \equiv 1$ and $\bar{n}$ the mean path length. $\Theta \rightarrow 1$ as compatibility approaches ideal, so Equation (13) reduces to Equation (26) exactly and every result in this paper is unaffected at $\Theta = 1$.

Remark 2 (The inversion must be re-derived, not merely re-run). *Introducing Θ changes the growth rate the ambient exposure inversion is solving against. Proposition 1 derives ν from $dN/dt = 0$, and with Equation (13) that condition reads $\omega_0 \Theta R(1 - \bar{n})$ rather than $\omega_0 R(1 - \bar{n})$. Solving the unmodified expression while running the modified dynamics leaves the substrate short of its stated baseline by up to 0.083 a silent error, since the simulation converges perfectly well to the wrong place. With the substitution made, the inversion is exact again to 10^{-14} .*

The general point is worth more than the particular correction. Every extension in this paper has required the machinery downstream of it to be re-derived rather than merely re-executed: γ_{mem} becoming a state variable changed what Equation (46) must be evaluated at, and Θ changes what Proposition 1 must be solved against. A framework whose central thesis is copy, paste, adapt, repeat should perhaps expect this of itself.

Remark 3 (Systematic verification across model variants). *The extensions in this*

paper are independently switchable, which makes the model a family rather than a single system, and errors of the kind above are invisible when variants are checked one or two at a time. We therefore verified the invariants across the family systematically: every single-switch variant, every pair, and a random sample of higher-order combinations, 632 in all, each required to reproduce its stated baseline under the inversion, to keep R , Θ and $P_{\text{substrate}}$ within $[0, 1]$, and to remain finite throughout.

The exercise found two defects that single-variant testing had not. First, the inversion of Proposition 1 subtracts a growth term that the model permits to be disabled; with logistic expansion switched off the subtraction is spurious and the substrate settles at 0.522 against a stated baseline of 0.700. Second, and not a coding error, the regularity form of induction (Equation (58)) makes the inversion implicit: at $\zeta = 0$ the inducing exposure is still $a_\nu \nu > 0$, so ambient encounter alone raises P_{tol} and with it γ_{mem} , which is the quantity the closed form assumes fixed at γ_0 . Measured γ_{mem} at rest was 0.560 against $\gamma_0 = 0.05$, and the baseline landed at 0.223. That branch now solves numerically, as the seasonal branch already did.

With both corrected, all 632 variants pass, as do 1750 preset-by-variant combinations, and schedule invariance holds to 8.3×10^{-4} across every combination tested. We record the procedure as much as the result: a model assembled from switchable extensions should be verified over the family it defines, not over the configuration its author happens to be running.

With Θ in place the sweep behaves as both substrate readings require. The equilibrium density falls monotonically with compatibility 0.435, 0.407, 0.214, 0, 0 where without it the same sweep produces 0.483, 0.606, 0.697, rising as the match degrades. It also couples the two halves of the framework in a way they were not previously coupled: the supercriticality condition becomes $\omega_0 \Theta R > \lambda$, so poor compatibility can move a substrate from self-sustaining to driven. That is the distinction Section 10.1 found separating silicon from tissue, now reachable by degrading the match rather than only by changing rates.

2.4 Path-Length Normalisation: The Hybrid Action

Equation (12) sums $E(s)$ over the nodes of γ , so the action grows with path length even when every node along the route is highly compatible. Route *quality* and route *length* are therefore confounded in a single quantity. The confound is compounded by the compatibility product $C(\gamma) = \prod_{s \in \gamma} c(s)$, which penalises length a second time, and by the discount $\gamma_{\text{len}}^{|\gamma|}$ in Equation (11), which penalises it a third.

Numerical evaluation of the partition function shows the practical consequence.

At $\sigma = 0.2$ the shortest admissible route receives essentially the entire weight of $\mathcal{Z}$, and $R(t)$ pins at unity regardless of every other parameter in the model. The path integral is then formally present but carries no information.

Where routes of differing length must be compared on equal terms, we therefore adopt the **hybrid action**:

$$S_{\text{hybrid}}[\gamma, t] = \frac{S[\gamma, t]}{|\gamma|} + \lambda|\gamma| \quad (14)$$

The first term is the mean action density per node, which measures route quality independently of length. The second restores length as an explicit preference with its own coefficient $\lambda \geq 0$. Setting $\lambda = 0$ leaves the length preference entirely to $\gamma_{\text{len}}^{|\gamma|}$, which is temperature-independent; a nonzero λ places part of it inside the exponent, where it softens as σ rises. At $\lambda = 0$ and uniform $|\gamma|$ the hybrid action is a monotone reparameterisation of Equation (12) and selects the same route ordering.

Under the hybrid form the same partition function that collapsed to 100/0/0 at $\sigma = 0.2$ retains a spread of 88.6/10.8/0.6 across three candidate routes, and the sensitivity of $R(t)$ to the binding free energy ΔG_0 widens from 0.247 to 0.384 over the range $\Delta G_0/k_B T \in [-4, 3]$.

2.5 The Noise Scale as Temperature

Equation (28) places $\sigma(t)$ exactly where $\hbar$ sits in the Feynman path integral, which makes σ the temperature of the path sum rather than an incidental nuisance parameter. Nothing in the original formulation fixes its scale, and the behaviour of $R(t)$ depends on that scale qualitatively rather than quantitatively: below $\sigma \approx 1$ the partition function is degenerate in the sense above.

We therefore impose a floor

$$\sigma(t) \geq \sigma_{\min} \quad (15)$$

with $\sigma_{\min}$ expressed in per-node action units, so that it carries the same meaning under Equation (12) and under Equation (14). Under the raw action the floor must be scaled by the typical path length $|\overline{\gamma}|$ to preserve that meaning. The floor is a guard against degeneracy, not a physical claim; determining what actually fixes σ in a biological substrate remains open.

2.6 Coupled Micro-Coefficients

Definition 2 (Localized Node Compatibility).

$$c(s) = \max(0, \vec{p}^T M_s \vec{p}), \quad \|\vec{p}\|_2 = 1, \quad \|M_s\|_2 \leq 1 \quad (16)$$

with global pathway compatibility $C(\gamma) = \prod_{s \in \gamma} c(s)$.

Definition 3 (Time-Adaptive Novelty with Structural Resolution).

$$V(s, t, \rho_s) = \rho_s \cdot \text{JSD}(P_{\text{pattern}} \parallel P_{\text{substrate}}(s, t)) \quad (17)$$

where JSD denotes the symmetric Jensen-Shannon Divergence and $\rho_s > 0$ is the structural resolution of the substrate at node s . When $\rho_s = 1$ the original formulation is recovered exactly.

What P_{pattern} and $P_{\text{substrate}}$ Range Over

Equation (17) is a divergence between distributions, and the choice of the space they are distributions over is not incidental. Treating both as scalars the natural first reading reduces JSD to its Bernoulli case, and that case is degenerate in a way worth stating plainly: **two allergens with different epitope compositions but the same scalar mean have a novelty of exactly zero relative to one another**. Not small; zero. The formulation is not merely imprecise about such a pair, it cannot see the difference at all.

We therefore resolve both distributions over B epitope feature bins. Let $\vec{w}$ be the allergen’s epitope profile, normalised over the bins. A substrate of resolution ρ_s cannot register arbitrarily fine structure in $\vec{w}$; what it can register is the profile convolved with a kernel whose width narrows as resolution improves,

$$\vec{b}^\rho = \mathcal{N}[\vec{w} * G_{w_0/\rho_s}], \quad G_\sigma(d) = e^{-d^2/2\sigma^2} \quad (18)$$

with $\mathcal{N}$ denoting normalisation and w_0 the kernel width at unit resolution. The encoded state interpolates from the uniform distribution, when nothing has been learned, to $\vec{b}^\rho$, when everything learnable has been:

$$q_i = P_{\text{substrate}} b_i^\rho + \frac{1 - P_{\text{substrate}}}{B}, \quad V = \rho_s \cdot \text{JSD}(\vec{w} \parallel \vec{q}) \quad (19)$$

The convex combination in Equation (19) is automatically normalised, since both components sum to unity.

This is what makes ρ_s an independent axis. Under the scalar reading, $V = \rho_s f(P_{\text{substrate}})$ is separable: ρ_s multiplies a function of the substrate state and does nothing else. The ratio $V(0.9)/V(0.1)$ is then 0.0107 at *every* value of ρ_s , which is precisely the confound Remark 9 identifies. Under Equation (19), ρ_s also sets the kernel width, so the shape of V as a function of $P_{\text{substrate}}$ depends on it. The same ratio now runs from 0.85 at $\rho_s = 0.25$ to 0.013 at $\rho_s = 4$. V is no longer separable, and ρ_s has acquired a channel that ΔG_0 cannot imitate.

The kernel width is recoverable from cross-reactivity. Before discussing what w_0 trades off, it is worth establishing that it need not be assumed. The kernel governs how much of a profile a substrate can resolve, so it governs how strongly a substrate trained on one allergen recognises another which is cross-reactivity, and cross-reactivity panels are among the most abundant data in the field.

Let $\vec{b}^\rho$ be the resolvable profile of Equation (18) for a substrate trained to saturation on $\vec{w}$, and define

$$X = \left\langle \frac{1 - \text{JSD}(\vec{w}' \| \vec{b}^\rho)}{1 - \text{JSD}(\vec{w} \| \vec{b}^\rho)} \right\rangle_{\vec{w}'} \quad (20)$$

the recognition of unrelated profiles $\vec{w}'$ relative to the trained one. Since JSD is bounded, $1 - \text{JSD}$ reads directly as an overlap. X is strictly increasing in w_0 a coarsely resolving substrate confuses more and therefore invertible: bisection recovers w_0 from a measured X to machine precision. The value $X = 0.708$ returns $w_0 = 0.5006$, which is the figure used below.

This closes the last hand-set constant on which a result of this paper depends. Remark 9 requires w_0 to lie in an intermediate range, and without Equation (20) that requirement would amount to choosing a parameter so that a conclusion follows from it.

The kernel width sets an exchange rate. w_0 is not free of consequence. A wide kernel leaves a large irreducible mismatch between $\vec{w}$ and the best the substrate could encode; that mismatch is common to every pathway, so it cancels in the normalisation of Equation (11) and contributes nothing except to drown the learning signal. Narrowing w_0 from 1.6 to 0.5 lowers the residual at perfect learning from 0.167 to 0.022 and raises the span of V across the learning range from 0.032 to 0.173, while leaving the resolution dependence above intact. Narrowing it much further, to $w_0 \approx 0.2$, eliminates the mismatch entirely and returns ρ_s to a pure multiplier reinstating the very confound the construction exists to remove. The two properties

are obtained together only for intermediate w_0 , and $w_0 \approx 0.5$ is used throughout.

Definition 4 (Structural Resolution ρ_s). *The structural resolution ρ_s quantifies how finely the substrate at node s can distinguish meaningful structural differences between informational configurations. In the biomedical context it is grounded in three measurable cellular properties of the immune execution engine:*

$$\rho_s = f(D_{TCR}, H_{MHC}, F_{DC}) \quad (21)$$

where D_{TCR} is the T-cell receptor repertoire diversity (measurable via immune repertoire sequencing), H_{MHC} is the HLA allele binding groove specificity (genetically fixed, measurable via HLA typing), and F_{DC} is the dendritic cell antigen presentation fidelity (the number of distinct epitopes presented per allergen protein). In the simplest empirically accessible form:

$$\rho_s \approx \frac{D_{TCR}}{D_{TCR}^{ref}} \quad (22)$$

normalized to a population reference mean, so $\rho_s = 1$ represents average immune resolution, $\rho_s > 1$ high resolution, and $\rho_s < 1$ low resolution. The parameter is bounded below at $\rho_s > 0$ and has no theoretical upper bound, though practical immune repertoire diversity imposes a physiological ceiling.

Definition 5 (Logic Coherence Engine).

$$L(s) = e^{-\varepsilon(s)} \quad (23)$$

where $\varepsilon(s)$ counts semantic exceptions or compilation errors at node s .

3 Substrate Memory Dynamics and Noise Coupling

Remark 4 (On the role of numerical evaluation in what follows). *Several statements in Sections 3–6 are attributed to numerical integration of the model rather than to closed-form analysis. These come from a reference implementation of the complete system path integral, population wave, substrate memory, and the extensions introduced below built to make each equation independently switchable, so that the consequence of any single term can be isolated by disabling it. Where that implementation contradicted a claim in an earlier version of this work, the claim is corrected here and the correction is marked. Quantities reported from simulation are stated in the model’s own dimensionless units and are not calibrated against clinical measurement; they describe the behaviour of the equations, not of a patient.*

Two integration details matter enough to record, because both were initially wrong in ways that produced plausible output. First, the controlled dose $\zeta(t)$ is a rectangular pulse train, and no Runge–Kutta scheme exceeds first order across a discontinuity in the forcing; clipping every step to the next pulse edge is what allows a fourth-order method to deliver its order. With that in place the measured convergence ratio is 15.9 against a theoretical 16, and RK4 at $\Delta t = 0.4$ is five orders of magnitude more accurate than explicit Euler at $\Delta t = 0.0125$. Second, quantities averaged over a dose cycle must be integrated with the window boundary treated as an event and with a trapezoidal rule; a left-endpoint rectangle rule caps the reported quantity at first order however accurate the underlying trajectory is.

Both apply to the spatial system of Equation (31) as well as to the well-mixed reduction, and the gain is the same order. Integrating the tissue strip, the regulatory substrate and the subsystem vector together under a fourth-order scheme with edge-clipped steps gives observed convergence ratios of 24–41 against explicit Euler’s 2.0, and at $\Delta t = 0.2$ the result is a hundred times more accurate than Euler at $\Delta t = 0.0125$. The additional cost is 27% per step rather than the fourfold one might expect, because the path integral is evaluated once per stage rather than once per spatial cell. One consequence is worth noting for anyone reproducing this: with steps clipped to pulse edges, a fixed number of steps no longer corresponds to a fixed interval of model time, and convergence must be assessed at fixed t rather than at fixed step count.

3.1 The Memory-Dissipation Adaptation Law

The substrate is not a static background. When a wave of copies floods a network node, the substrate updates its internal cache to recognize the configuration. This is governed by:

$$\frac{\partial P_{\text{substrate}}(s, t)}{\partial t} = \alpha \cdot N(t) \cdot P_{\text{pattern}}(s) - \gamma_{\text{mem}}(P_{\text{substrate}}(s, t) - P_0) \quad (24)$$

where $\alpha \in (0, 1)$ is the learning rate, $N(t)$ the macro population size, $\gamma_{\text{mem}} \in [0, 1]$ the forgetting rate, and P_0 the pristine initial state of the medium. As $N(t) \rightarrow 0$, the substrate resets toward P_0 , allowing novelty to regenerate over time.

Remark 5 (Normalisation of the population term). *In the biomedical formulation of Section 5 the population enters as the normalised density $\bar{n} = N/K \in [0, 1]$ rather than as N . With N unnormalised, $P_{\text{substrate}}$ scales with the carrying capacity and leaves the unit interval on which the divergence in Equation (17) is defined. The substitution is dimensional bookkeeping and does not alter the structure of the law.*

3.2 Bounded Systemic Congestion Noise

$$\sigma(t) = \sigma_0 + (\sigma_{\text{max}} - \sigma_0) \cdot \tanh\left(\beta \frac{N(t)}{K}\right) \quad (25)$$

As $N(t) \rightarrow K$, noise saturates at σ_{max} , driving $R(t) \rightarrow 0$ and creating a stabilizing self-regulatory loop.

3.3 The Macroscopic Population Wave

$$\frac{dN}{dt} = \max\left(0, \omega_0 \cdot R(t) \cdot N(t) \left(1 - \frac{N(t)}{K}\right)\right) - \lambda N(t) \quad (26)$$

All parameters are bounded and dimensionally consistent: $\omega_0 > 0$ (inverse time), $K > 0$ (count), $\lambda > 0$ (inverse time).

3.4 Stochastic Seeding

For out-of-distribution initializations where $\lim_{P_{\text{substrate}} \rightarrow \emptyset} C(\gamma_{\text{anomaly}}) \rightarrow 0$, the system is augmented with a stochastic injection term:

$$\frac{dN}{dt} = \max\left(0, \omega_0 R(t) N(t) \left(1 - \frac{N(t)}{K}\right)\right) + \zeta \cdot \xi(t) \cdot \delta(t - t_{\text{injection}}) - \lambda N(t) \quad (27)$$

where δ is the Dirac delta function, $\xi(t) \in [0, 1]$ the structural coherence of the anomalous pattern, and $\zeta > 0$ the injection seed size.

Remark 6 (Equation (27) seeds; it does not switch). *Earlier versions titled this subsection Stochastic Phase Transitions. That overstates what Equation (27) can do, and the correction is worth making precisely because the term is otherwise useful.*

A phase transition induced by a perturbation requires bistability: two stable states at one parameter set, with a separatrix between them that a large enough kick can cross. The model has none. Injecting seeds spanning a 256-fold range, from $\zeta = 0.5$ to 128, every trajectory returns to the same equilibrium 14.25926 in engine units, identical to five decimal places in every case. The peak excursion scales linearly with ζ , with the ratio of excess to seed size constant at 0.0391 across the range. This is the signature of a linear transient in a globally stable system, not of a threshold.

The reason is structural rather than parametric. Linearising Equation (27), the growth term is $N[\omega_0 R(1 - N/K) - \lambda]$, which has a single interior root; and both routes by which R responds to N through the congestion noise $\sigma(N)$ of Equation (25), and through $P_{\text{substrate}}$ into the novelty V are negative feedbacks. A unique equilibrium follows, and a unique equilibrium admits no switching. The coupled subsystem model of Section 7 does not rescue it: starting the same parameter set from quiescent and from activated initial conditions gives the same aggregate to five decimals, so it is not bistable either.

The framework does contain a sharp transition, but it is a different object. The ignition threshold of Section 7.7 is a fold in the control parameter $\bar{\nu}$ the state jumps from 0.150 to 0.541 across a change of 7×10^{-4} and it lives in the coupling matrix χ , not in the injection term. Crossing it requires changing the environment, not delivering a seed. Producing a genuine perturbation-induced transition would require bistability to be built in, most naturally through an Allee term of the form $\omega_0 R N(N - a)(1 - N/K)$, which no equation in this paper contains.

Equation (27) remains correct and useful for what it is: a way to introduce a configuration that the substrate has not encountered, so that its subsequent propagation can be followed. It is a seeding term. Nothing in this paper's conclusions depends on it being more.

4 Refinements for Biomedical Application

The original CAPM operates on discrete network topologies with passive environmental background. Biomedical application requires three fundamental extensions, developed in Sections 4.1–4.3. A fourth subsection then reconciles the compatibility tensor introduced by the third with the geometric definition it appears to replace; that reconciliation is not itself a refinement but a consistency requirement, and Section 8 depends on it.

4.1 Refinement I: Continuous Field Reformulation

The discrete path sum in Equation (11) must become a genuine functional integral over continuous tissue geometry:

$$R(t) = \frac{1}{\mathcal{Z}} \int \mathcal{D}\gamma \exp\left(-\frac{S[\gamma, t]}{\sigma(t)}\right) \quad (28)$$

The correspondence with the Feynman path integral is close, with $\sigma(t)$ playing the role of $\hbar$, but it is not an identity and the difference is the part Feynman had to construct by hand. Equation (28) writes $\mathcal{D}\gamma$ without specifying it, and without a measure on path space the discrete sum of Equation (11) has no unique limit to approach.

The point is demonstrable rather than pedantic. Enumerating many paths within each destination class and computing the class share, three equally natural enumerations lengths drawn uniformly, lengths drawn geometrically, and lengths held fixed each converge, but to 0.676, 0.718 and 0.713 respectively. The answer depends on how paths are counted, not merely on how many are counted. Until $\mathcal{D}\gamma$ is fixed, Equation (28) names an object rather than defining one.

The Framework Already Contains the Normalisation

The missing ingredient is largely present, unrecognised, in Equation (11). Feynman’s measure is defined by time-slicing with a normalisation per slice; the analogue here is a normalisation per node, and $\gamma_{\text{len}}^{|\gamma|}$ is exactly that. Its role has been described in this paper as a path-length discount, which understates it.

On a substrate graph of out-degree b , the number of distinct walks of length n grows as b^n , while each carries $\gamma_{\text{len}}^n C(\gamma) e^{-S/\sigma}$ with $C(\gamma) \leq c_{\text{max}}^n$. The class sum is

therefore a geometric series, and converges if and only if

$$b \cdot \gamma_{\text{len}} \cdot c_{\text{max}} \cdot e^{-E_{\text{min}}/\sigma} < 1 \quad (29)$$

Equation (29) is the condition for the path sum to exist at all, and hence for Equation (28) to have anything to converge to. It also gives γ_{len} a determinate role: it is not free, but bounded above by the topology it is applied to.

What Failing the Bound Actually Means

The bound is not comfortable at the values used here. With $\gamma_{\text{len}} = 0.92$, $c_{\text{max}} \approx 0.62$ and $\sigma \approx 1.67$, the ratio in Equation (29) is 0.944 at $b = 3$ and 1.259 at $b = 4$, while vascular, lymphatic and neural networks all exceed a branching factor of three.

It would be wrong, however, to read this as the framework’s central expression being undefined. No physical substrate admits walks of unbounded length: signals degrade, cells have finite lifetimes, networks have finite diameter. With any cutoff $n \leq N_{\text{max}}$ the sum in Equation (11) is finite whatever the ratio. What Equation (29) actually decides is whether the answer depends on where that cutoff is placed.

b	ratio	share of the sum from the longest decade
2	0.627	0.9%
3	0.940	35.1%
3.5	1.097	71.6%
4	1.254	90.6%
8	2.507	99.99%

Table 2: Contribution of the longest paths to the class sum. Below ratio one the tail is negligible and the cutoff is irrelevant; above it the cutoff supplies essentially the whole answer.

The correct statement is therefore sharper than a convergence condition. **When Equation (29) holds, $R(t)$ is a property of the substrate. When it fails, $R(t)$ is a property of how long paths were allowed to be.** A model in the second regime is not meaningless, but it is reporting a modelling choice as though it were a result, and any such model must state N_{max} and justify it.

Restoring the bound at realistic branching is possible but not free. At $b = 8$ convergence requires $\gamma_{\text{len}} \lesssim 0.3$, which concentrates 90% of the partition function on the shortest route pressing back toward the degeneracy that motivated the hybrid action of Section 2.4. The two requirements pull in opposite directions, and the

honest summary is that the framework has a narrow window in which the path sum is both substrate-determined and non-degenerate. Locating that window on a real topology is unfinished work, and is recorded here as such.

The action functional $S[\gamma, t]$ in tissue becomes:

$$S[\gamma, t] = \int_{\gamma} \frac{E(\mathbf{x})}{1 + \tanh(V(\mathbf{x}, t) \cdot L(\mathbf{x}) \cdot c(\mathbf{x}))} d\ell(\mathbf{x}) + \sum_{j \in B_t} \chi_j \quad (30)$$

where $\mathbf{x} \in \mathbb{R}^3$ is a continuous tissue coordinate and $d\ell$ the arc-length element along path γ . The energy $E(\mathbf{x}) > 0$ now represents binding energy the thermodynamic cost of allergen interfacing with each receptor node directly measurable via calorimetry. The novelty $V(\mathbf{x}, t)$ becomes the structural divergence between allergen geometry and local receptor memory, computable from protein structure databases.

4.2 Refinement II: Thermodynamic Diffusion Augmentation

Allergens lack an explicit replication drive but are subject to thermodynamic diffusion. The macroscopic population wave is augmented with a diffusion operator:

$$\frac{dN}{dt} = \max\left(0, \omega_0 R(t) N(t) \left(1 - \frac{N(t)}{K}\right)\right) + \mu \nabla^2 N(t) - \lambda N(t) \quad (31)$$

where $\mu > 0$ is the molecular diffusion coefficient, and $\nabla^2 N(t)$ is the Laplacian over tissue geometry. The new term bridges CAPM with standard reaction-diffusion equations used in biochemistry, producing a unified model in which both active informational propagation and passive thermodynamic spreading are formally represented. When $\mu \rightarrow 0$ the original CAPM population wave is recovered exactly.

Remark 7 (Front velocity, corrected and verified). *Equation (31) is of Fisher–KPP type when $R(t)$ is held constant, and therefore admits a travelling front. Earlier versions of this remark gave its velocity as $v = 2\sqrt{\mu\omega_0 R}$. That is incorrect: the Fisher–KPP velocity is set by the linear growth rate of the state being invaded, and linearising Equation (31) about $N = 0$ gives $r = \omega_0 R - \lambda$, not $\omega_0 R$. The clearance term was omitted.*

$$v = 2\sqrt{\mu(\omega_0 R - \lambda)} \quad (32)$$

*Two conditions must hold before Equation (32) applies at all, and neither was previously stated. First, the ambient encounter rate must vanish: with $\nu > 0$, $N = 0$ is not an equilibrium, the strip fills uniformly, and there is no empty state to invade. Second, the empty state must be unstable, $\omega_0 R > \lambda$. **At the reference parameters used throughout this paper the model is subcritical $\omega_0 R - \lambda = -0.236$ so no***

front exists there and the population is sustained entirely by ambient exposure. The prediction is not merely untested at those settings but inapplicable to them.

Where the conditions do hold, Equation (32) is accurate. Seeding an invasion on a 3000-cell strip and tracking the half-maximum contour with sub-cell interpolation, the speed converges from below as expected for a pulled front to 1.018 times the predicted value, agreement within 2%. The superseded expression overstates it by 29% at the same parameters. Front speed remains directly measurable via acoustic microscopy, and is now a quantitative test rather than a scaling claim.

4.3 Refinement III: Statistical Mechanics Correction to $c(s)$ with Structural Resolution

At the molecular scale of allergen-receptor interaction, the classical compatibility tensor requires a physically grounded refinement. Rather than invoking quantum mechanical wavefunctions which, while theoretically interesting, introduce dependencies that are difficult to constrain experimentally we ground $c(s)$ in the well-established framework of **statistical mechanics and thermodynamic free-energy landscapes**, augmented by the structural resolution parameter ρ_s .

Receptor-allergen binding is fundamentally a thermodynamic process whose sensitivity depends on how finely the receptor population can discriminate between molecular configurations. We therefore write:

$$c(s) = \frac{1}{1 + \exp(\Delta G(s, \rho_s) / k_B T)} \quad (33)$$

where $\Delta G(s, \rho_s)$ is the resolution-dependent binding free energy:

$$\Delta G(s, \rho_s) = \Delta G_0(s) - k_B T \ln(\rho_s) \quad (34)$$

Here $\Delta G_0(s)$ is the baseline binding free energy at average resolution $\rho_s = 1$, and the $-k_B T \ln(\rho_s)$ term encodes the thermodynamic effect of resolution: a substrate with higher ρ_s resolves finer free-energy differences between molecular configurations, effectively lowering the apparent binding free energy for structurally similar but distinct patterns. Substituting:

$$c(s) = \frac{\rho_s}{\rho_s + \exp(\Delta G_0(s) / k_B T)} \quad (35)$$

Remark 8. At $\rho_s = 1$, Equation (35) reduces to the standard Fermi-Dirac occupancy form. As $\rho_s \rightarrow \infty$, $c(s) \rightarrow 1$ regardless of ΔG_0 a perfectly resolving substrate finds

everything compatible, which is the mathematical expression of the polysensitization risk discussed in Section 5.5. As $\rho_s \rightarrow 0$, $c(s) \rightarrow 0$ the substrate cannot distinguish the pattern at all.

This formulation is directly experimentally accessible: $\Delta G_0(s)$ is measurable via isothermal titration calorimetry and surface plasmon resonance; ρ_s is estimable from TCR repertoire sequencing. Furthermore, the substrate memory $P_{\text{substrate}}(s, t)$ under this formulation becomes a *resolution-weighted stochastic state-space* analogous to a neural network where both the learning rate α and the effective input dimensionality ρ_s jointly determine the richness of representable patterns.

4.4 The Two Forms of $c(s)$ Are One Quantity

Definition 2 gives compatibility as a geometric quantity, $c(s) = \max(0, \vec{p}^T M_s \vec{p})$, and Equation (35) replaces it with a thermodynamic occupancy in ΔG_0 . These have been presented as alternatives the second refining the first for biomedical use. They are not alternatives. They are the same quantity in different coordinates, and the change of variables between them is explicit.

Inverting Equation (35) for the free energy implied by a given compatibility:

$$\frac{\Delta G_{\text{eff}}(s)}{k_B T} = \ln \left(\frac{\rho_s (1 - c(s))}{c(s)} \right) \quad (36)$$

Every geometric configuration therefore *has* a binding free energy, and every binding free energy *has* an implied alignment. The content of Equation (36) is that structural alignment and thermodynamic affinity are not two mechanisms acting on $c(s)$ but two readings of one number.

The consistency condition. Using both forms in one model requires them to agree where they overlap. Under the parameterisation $M_s = a \vec{u}_s \vec{u}_s^T + b \mathbb{I}$, which satisfies $\|M_s\|_2 \leq 1$ for $a + b \leq 1$ and gives $c(s) = a (\vec{u}_s \cdot \vec{p})^2 + b$, the requirement that a perfectly aligned pattern reproduce the thermodynamic value fixes

$$a + b = \frac{1}{1 + \exp(\Delta G_0(s)/k_B T - \ln \rho_s)} \quad (37)$$

Equation (37) is not cosmetic. Imposed, the geometric formulation reproduces the thermodynamic one to machine precision and the mutation dynamics of Section 8 become a strict extension of the static model. Omitted, the two forms disagree at $t = 0$ by an amount depending on nothing in particular, and every downstream quantity shifts.

Mutation is a climbing affinity. Combining Equation (36) with the decay law of Proposition 9, a pattern drifting under mutation presents the receptor population with a monotonically rising effective binding free energy. The geometric and thermodynamic pictures of immune escape loss of shape complementarity, and loss of binding affinity are the same process.

Escape is bounded by promiscuity. Because $c(s)$ relaxes to $\bar{c} = a/d + b$ rather than to zero, the effective free energy saturates:

$$\frac{\Delta G_{\text{eff}}^{\text{max}}}{k_B T} = \ln \left(\frac{\rho_s(1 - \bar{c})}{\bar{c}} \right), \quad \bar{c} \geq b \quad (38)$$

A substrate with nonzero promiscuity b therefore retains residual recognition of an arbitrarily mutated pattern. Complete escape requires $b = 0$, that is a receptor population that binds nothing it was not shaped for. This is the framework’s account of persistent cross-reactivity, and it also bounds the polysensitization mechanism of Section 5.5: the same b that prevents escape is what admits structurally unrelated allergens.

An experimental consequence. Equations (37) and (38) connect two measurement modalities that would otherwise be separate. Calorimetric determination of ΔG_0 constrains $a + b$; the residual binding of a mutated or unrelated ligand constrains b alone; and the two together fix a and hence the configuration dimension d through $\bar{c}$. The geometric parameters of Definition 2, which have no direct experimental interpretation as written, become accessible through binding assays.

5 Allergy as a CAPM Substrate Rewriting Problem

5.1 The Core Mapping

In classical immunology, allergy is characterized as immune system hypersensitivity. In the CAPM biomedical lens, allergy is precisely characterized as follows but with a critical biological clarification that makes the mapping both more accurate and more defensible.

Biological Role Separation: Blueprint vs. Execution Engine

A key refinement over naive interpretations of the framework is the explicit separation of biological roles. The allergen itself is not the propagating configuration. It is the *static informational template* the blueprint. The **immune network** dendritic cells presenting antigen, T-cells coordinating response, B-cells differentiating and secreting antibodies, mast cells degranulating constitutes the *active execution engine*. It is the immune network that copies, expands, and propagates the response pattern across tissue, writing and rewriting $P_{\text{substrate}}$ with each exposure event.

This framing resolves the apparent tension with classical biochemistry. The allergen does not need an active replication drive; it provides the pattern. The immune cells replicate the *response* to that pattern and it is this response propagation that the CAPM engine governs. The substrate is altered not by the allergen directly, but by the immune network's expanding population $N(t)$ as it encodes the allergen's signature into immunological memory.

Definition 6 (Allergic State). *An allergy is a state in which the immune network has encoded $P_{\text{pattern}}(\text{allergen})$ via clonal expansion of allergen-specific lymphocytes and IgE class-switching as highly compatible with the attack pathway destination node s_{attack} in the partition function $\mathcal{Z}$, such that:*

$$R(t) \big|_{s_f=s_{\text{attack}}} \approx 1 \quad (39)$$

despite the allergen posing no systemic threat. The immune execution engine has written a false compatibility profile into the substrate.

The biological parameter correspondences, with role separation explicit, are:

CAPM Parameter	Biological Role	Cellular Equivalent	Notes
P_{pattern}	Blueprint	Allergen protein structure	Static; does not propagate
$N(t)$	Execution engine	Dendritic, T-, B-cells	Active propagator of response
$P_{\text{substrate}}(s, t)$	Substrate memory	IgE / T-cell encoding	Written by $N(t)$, not allergen
α	Learning rate	Clonal expansion efficiency	How fast $N(t)$ writes memory
γ_{mem}	Forgetting rate	IgE antibody half-life	Natural memory decay
κ	Resensitization	Re-exposure activation	$N(t)$ re-encodes on contact
ρ_s	Structural resolution	TCR diversity / HLA specificity	How finely immune system discriminates
P_0	Healthy baseline	Naive immune state	Pre-sensitization reference
K	Carrying capacity	Max lymphocyte count	Resource ceiling
ζ	Dose magnitude	Allergen load per exposure	Controls $N(t)$ activation
χ_j	Crossing cost	Cell membrane / tissue boundary	Energy per substrate transition

Table 3: Parameter correspondences with explicit biological role separation. The allergen functions as a static blueprint; the immune network is the active CAPM propagation engine.

5.2 The Observable Baseline and the Ambient Exposure Inversion

Neither Equation (24) nor its extended form Equation (50) below contains a term representing uncontrolled environmental encounter. Without one the execution population $N(t)$ decays to zero whenever controlled dosing stops, the resensitization term $\kappa\bar{n}(1 - P_{\text{substrate}})$ vanishes with it, and the model cannot represent an untreated allergic patient at all every substrate relaxes to P_0 and the disease disappears by

construction. We therefore add an ambient encounter rate $\nu \geq 0$ to the population wave:

$$\frac{dN}{dt} = \max\left(0, \omega_0 R(t) N \left(1 - \frac{N}{K}\right)\right) + \mu \nabla^2 N + \nu \left(1 - \frac{N}{K}\right) - \lambda N \quad (40)$$

ν is not clinically observable. What is observable is the reactivity a patient exhibits when nobody is treating them. We therefore take that **baseline severity** P^* as the primitive parameter and derive ν from it, rather than the reverse.

Proposition 1 (Ambient Exposure Inversion). *Let P^* denote the untreated equilibrium of the substrate memory law (Equation (50)) with $\zeta = 0$. Setting $\partial_t P_{\text{substrate}} = 0$ gives the engine density the baseline requires,*

$$\bar{n}^* = \left[\frac{\gamma_{\text{mem}}(P^* - P_0)}{\kappa(1 - P^*)} \right]^{1/q}, \quad (41)$$

and setting $dN/dt = 0$ in Equation (40) at $\bar{N}^* = K\bar{n}^*$ gives

$$\nu = \frac{\bar{N}^* [\lambda - \omega_0 R^*(1 - \bar{n}^*)]}{1 - \bar{n}^*}. \quad (42)$$

Since $\bar{n}^*$ does not depend on R , the pair closes in one pass without iteration.

The inversion has three consequences that are not optional readings of the model but direct algebraic facts about it.

The environmental pressure is a single constant. Equation (41) states that $\kappa \bar{n}^{*q}$ equals

$$A = \frac{\gamma_{\text{mem}}(P^* - P_0)}{1 - P^*} \quad (43)$$

identically. A combines the patient's constitutive forgetting rate with their environment and is the quantity that actually governs the dynamics.

The relaxation rate loses κ . Substituting Equation (43) into Equation (50) gives a relaxation rate toward P^* of $\gamma_{\text{mem}} + A$, in which κ does not appear. Numerical integration confirms the predicted time constant to three significant figures. The consequence for Section 6.5 is developed there.

Not every baseline is dynamically reachable. Equation (41) requires $\bar{n}^* < 1$. A patient with high γ_{mem} and low κ cannot sustain a high P^* at any ambient exposure rate: the combination $\kappa = 0.24$, $\gamma_{\text{mem}} = 0.50$ admits no baseline above $P^* \approx 0.375$,

whereas $\kappa = 9.79$, $\gamma_{\text{mem}} = 0.01$ admits baselines up to $P^* \approx 0.99$. This is a falsifiable constraint on any set of fitted parameters, and Section 6.5 applies it.

The constraint also bears on how phase-space diagrams of this model should be read, and the effect is larger than it appears. Sweeping $(\gamma_{\text{mem}}, \alpha)$ over the unit square at a fixed baseline $P^* = 0.85$, Equation (41) has no solution over roughly 63% of that square: a patient who forgets quickly cannot sustain a high untreated baseline at any ambient exposure. Those parameter combinations are not regions of the phase space in which the model behaves differently they are combinations the model does not possess. A tolerance boundary $\partial\mathcal{T}$ plotted across the full square is therefore drawn mostly over states that do not exist, and any conclusion about the shape of $\partial\mathcal{T}$ away from the feasible region is an artefact. Diagrams of this kind should mask the unreachable set rather than colour it, and the fraction masked should be reported alongside.

The Activation Gain Is Not Separately Identifiable

Equation (40) and its subsystem form Equation (59) carry two source terms, the ambient rate ν and the controlled contribution $g_N\zeta(t)$. The gain g_N never appears except multiplied by ζ , and it is added to a quantity in the same units. It has therefore no meaning on its own, and fitting it is fitting a scale that the data cannot see.

What is meaningful is the factor by which a protocol raises total antigen encounter above background:

$$\Xi = \frac{g_N \bar{\zeta}}{\nu}, \quad \bar{\zeta} = \frac{\zeta \tau}{T} \quad (44)$$

Ξ is dimensionless, invariant under the schedule reparameterisations of Proposition 6, and is a quantity a clinical protocol could in principle state. We therefore treat Ξ as the primitive and derive $g_N = \Xi\nu/\bar{\zeta}$, exactly as P^* is treated as primitive and ν derived from it in Proposition 1. In both cases the substitution replaces an unobservable parameter with an observable one and leaves the dynamics unchanged.

Seasonal Environments

Equation (40) treats ν as constant. For the aeroallergens to which the framework is most directly applied this is false: grass pollen is strongly periodic, and the GRASS trial is a grass trial. We therefore allow

$$\nu(t) = \bar{\nu} \left[1 + a \sin \left(2\pi \left(\frac{t}{T_{\text{yr}}} + \varphi \right) \right) \right]_+ \quad (45)$$

with $a \in [0, 1]$ the relative swing, T_{yr} the environmental period and φ the phase at which observation begins.

Under Equation (45) the untreated state is a limit cycle rather than a fixed point, and the closed form of Proposition 1 no longer applies: because Equation (50) is nonlinear in $P_{\text{substrate}}$, the cycle mean is not the mean of the constant- ν solutions. The inversion becomes a one-dimensional root find for the $\bar{\nu}$ whose cycle-averaged substrate memory equals P^* , seeded from the closed-form answer. It converges to within 10^{-5} for amplitudes up to $a = 0.95$ and baselines from 0.3 to 0.9, and reduces to the closed form exactly at $a = 0$.

The nonlinearity has a consequence worth stating: *a seasonal environment requires a higher mean exposure than a constant one to sustain the same baseline.* The saturating factor $(1 - N/K)$ discounts the high-exposure months more than it credits the low ones, and the required $\bar{\nu}$ rises by approximately 60% between $a = 0$ and $a = 0.95$. Cross-sectional studies that infer environmental exposure from measured baseline reactivity, without accounting for the seasonality of the source, will therefore underestimate it.

5.3 Why α and γ_{mem} Are Pathological in Allergy

In allergic individuals, Equation (24) reveals the core problem immediately. Inserting the allergic parameter profile:

- $\alpha \approx 1$: The immune system learns the allergen aggressively.
- $\gamma_{\text{mem}} \approx 0$: Immunological memory almost never decays.

Every exposure drives $\alpha \cdot N(t) \cdot P_{\text{pattern}}$ upward, deepening the attack encoding, while the restoring term $-\gamma_{\text{mem}}(P_{\text{substrate}} - P_0)$ is too small to decay the false memory between exposures. This is the unified mathematical explanation for why accidental re-exposure worsens sensitization.

5.4 The Phase Space and the Tolerance Boundary

Definition 7 (Tolerance Boundary $\partial\mathcal{T}$). *The tolerance boundary in $(\gamma_{\text{mem}}, \alpha)$ phase space is the locus of parameter pairs where the rewriting signal exactly balances the resensitization pressure:*

$$(\alpha\zeta + \gamma_{\text{mem}})(P_{\text{substrate}} - P_0) = \kappa \cdot \bar{n}(t) \cdot (1 - P_{\text{substrate}}) \quad (46)$$

where $\kappa > 0$ is the resensitization coefficient and $\bar{n} = N/K \in [0, 1]$ the normalised execution-engine density. Above $\partial\mathcal{T}$ lies the tolerance zone; below it the persistent attack zone.

Orientation of the difference. Both rewriting terms drive $P_{\text{substrate}}$ toward P_0 and are therefore negative contributions to $\partial_t P_{\text{substrate}}$ whenever $P_{\text{substrate}} > P_0$, which is the entire clinically relevant regime. Writing the balance with $(P_0 - P_{\text{substrate}})$ on the left would set a negative quantity equal to a positive one and the boundary would be empty. The difference is oriented as $(P_{\text{substrate}} - P_0)$ throughout this paper, and Equations (51), (47) and (54) follow the same convention.

γ_{mem} here is the current value, not the constitutive one. Equation (46) is a balance among instantaneous rates, so every quantity in it is evaluated at the state the substrate currently occupies. Once Section 6.4 makes γ_{mem} a state variable this ceases to be a pedantic distinction: evaluating the boundary at the constitutive γ_0 instead of the current $\gamma_{\text{mem}}(t)$ misplaces the equilibrium substrate memory by 0.31 in a reference run, while evaluating it at the value the run actually reaches agrees to 3×10^{-4} .

This has a consequence for the phase space of Section 5.4, which is drawn over $(\gamma_{\text{mem}}, \alpha)$. Under plasticity a patient does not occupy a point in that plane treatment translates them through it. The coordinate that is fixed is the constitutive rate γ_0 ; γ_{mem} is where they have been moved to. Measured displacements over a 300-unit protocol run from 0 at $\alpha_{\text{tol}} = 0$ to 0.39 at $\alpha_{\text{tol}} = 0.2$, which is a large fraction of the axis. Read correctly this strengthens rather than weakens the construction: immunotherapy is literally motion through the phase space, and the tolerance boundary is a contour a patient can be driven across rather than a wall they sit behind.

A local balance is not a finite-horizon contour. Equation (46) defines $\partial\mathcal{T}$ as the locus where the instantaneous rates cancel. A phase diagram produced by integrating the model and shading the sign of the net change over a finite horizon defines a related but different object, because the state and with it $\bar{n}$ and γ_{mem} moves during the integration. The two agree to 1% when the horizon is short enough that the state does not travel, and differ by about 7% at long horizons. Both are legitimate; which is meant should be stated whenever such a diagram is shown.

Normalisation of the engine term. $\bar{n}$ rather than N appears here so that the resensitization pressure is dimensionless and $P_{\text{substrate}}$ remains confined to $[0, 1]$ independently of the carrying capacity K . With N unnormalised the substrate state scales with K and leaves its own domain.

The phase space contains three biologically distinct populations:

Population I — High α , high γ_{mem} (top-right). Aggressive immune learners who also forget quickly. Paradoxically the easiest to treat with immunotherapy. Most likely to naturally outgrow childhood allergies. The substrate rewrites in both directions rapidly.

Population II — Low α , low γ_{mem} (bottom-left). Rarely develop allergies. If sensitized, difficult to treat because γ_{mem} is too low attack memory persists with minimal decay. Immunotherapy must supply the forgetting the substrate cannot generate endogenously.

Population III — High α , low γ_{mem} (top-left). The severe allergy population. Deep in the attack zone. Each accidental exposure at dose $N(t)$ re-deepens the false memory. Farthest from $\partial\mathcal{T}$ and most resistant to standard fixed-interval immunotherapy.

5.5 Structural Resolution and the Three-Dimensional Phase Space

The two-dimensional phase space $(\gamma_{\text{mem}}, \alpha)$ characterizes *how the immune system learns and forgets*. The structural resolution parameter ρ_s adds a third, independent axis characterizing *how finely the immune system discriminates* between molecular configurations. Together they define a three-dimensional phase space $(\gamma_{\text{mem}}, \alpha, \rho_s)$ within which every patient's immune substrate occupies a unique position.

The tolerance boundary $\partial\mathcal{T}$ becomes a surface in this space. Its ρ_s -dependence follows from the resolution-weighted novelty $V(s, t, \rho_s)$ in Equation (17): as ρ_s increases, the effective novelty of any allergen encounter increases, which increases the action $S[\gamma, t]$ along attack pathways, pushing the tolerance boundary outward. Formally, the resonance condition (Equation (51)) acquires a resolution correction:

$$(\alpha\bar{\zeta} + \gamma_{\text{mem}})(\bar{P}_{\text{substrate}} - P_0) > \kappa\bar{n}(1 - \bar{P}_{\text{substrate}}) \cdot \rho_s \quad (47)$$

Equation (47) is Equation (51) with a single factor ρ_s applied to the resensitization side; the bracketing of $\alpha\bar{\zeta}$ and the orientation of the difference are identical in both, as they must be for the two to be the same inequality under $\rho_s \rightarrow 1$.

The ρ_s factor on the right-hand side encodes the key insight: a high-resolution immune system requires a proportionally larger rewriting drive to achieve the same tolerance outcome, because its resensitization pressure is amplified by its own discriminatory precision.

The Polysensitization Prediction

The most clinically significant prediction generated by the ρ_s extension concerns **polysensitization** the clinical phenomenon in which patients are simultaneously allergic to multiple structurally unrelated allergens. Prior to the introduction of ρ_s the framework had no parameter capable of explaining why two patients with identical α and γ_{mem} present with entirely different allergen profiles.

The structural resolution parameter provides the explanation directly. A patient with high ρ_s has an immune execution engine that resolves fine structural differences between molecular configurations. The same resolution that makes this system effective against structurally diverse pathogens also makes it more likely to misclassify harmless molecules as threatening because it recognizes *more* structural detail in each molecular encounter, generating higher novelty signals $V(s, t, \rho_s)$ for a broader range of environmental molecules.

Formally, the compatibility $c(s)$ under Equation (35) increases monotonically with ρ_s for any fixed $\Delta G_0(s)$. This means high- ρ_s substrates assign higher compatibility to a wider range of molecular patterns not just the primary allergen but structurally adjacent molecules generating broader sensitization across the allergen landscape.

Proposition 2 (Polysensitization Risk). *A patient at position $(\alpha, \gamma_{\text{mem}}, \rho_s)$ in the three-dimensional phase space exhibits polysensitization risk proportional to:*

$$\Pi \propto \rho_s \cdot \frac{\alpha}{\gamma_{\text{mem}} + \epsilon} \quad (48)$$

where $\epsilon > 0$ prevents divergence at $\gamma_{\text{mem}} = 0$. Polysensitization is therefore not a separate pathological mechanism but a predictable consequence of high structural resolution combined with high learning rate and low forgetting rate.

Remark 9 (ρ_s and ΔG_0 are not separately identifiable). Equation (35) can be rewritten

$$c(s) = \frac{1}{1 + \exp(\Delta G_0(s)/k_B T - \ln \rho_s)}, \quad (49)$$

which is Equation (34) substituted into Equation (33) and read as an identity: $c(s)$ depends on ρ_s and ΔG_0 only through the single combination $g_{\text{eff}} = \Delta G_0/k_B T - \ln \rho_s$. Holding g_{eff} fixed while sweeping ρ_s over a fourteenfold range changes the simulated trajectory by a root-mean-square of 1.2×10^{-3} , or about 0.1%.

In the compatibility channel, therefore, structural resolution is not a third axis but a reparameterisation of binding affinity. Raising ρ_s is indistinguishable from lowering ΔG_0 . The phase space $(\gamma_{\text{mem}}, \alpha, \rho_s)$ is three-dimensional only if ρ_s possesses a channel that ΔG_0 does not share.

Equation (17) supplies exactly that channel, since ρ_s multiplies the novelty V as well as entering $c(s)$ but only if the distributions it compares are resolved finely enough for resolution to matter. Under a scalar reading of P_{pattern} and $P_{\text{substrate}}$ it does not: $V = \rho_s f(P_{\text{substrate}})$ is separable, ρ_s acts as a pure multiplier, and the confound is complete in both channels at once.

Section 2.6 resolves this. Once the distributions range over epitope feature bins and ρ_s sets the width of the kernel through which the substrate registers them, V ceases to be separable: the ratio $V(0.9)/V(0.1)$, constant at 0.0107 for every ρ_s under the scalar reading, runs from 0.85 at $\rho_s = 0.25$ to 0.013 at $\rho_s = 4$. **The independence of ρ_s is therefore established rather than merely unrefuted**, and the three-dimensional phase space $(\gamma_{\text{mem}}, \alpha, \rho_s)$ stands. Proposition 2 may be read as written.

Two qualifications remain. The independence lives entirely in the novelty channel; in the compatibility channel ρ_s and ΔG_0 are still exactly confounded through g_{eff} , so an experimental determination of ρ_s must proceed through a measurement sensitive to V a discrimination assay rather than a binding assay. And the strength of the effect depends on the kernel width w_0 , which is itself unmeasured.

Remark 10. This prediction is testable against existing clinical datasets. Polysensitized patients are routinely identified in allergy practice via skin prick test panels. If TCR diversity measurements (the primary correlate of ρ_s) are systematically higher in polysensitized patients than monosensitized patients with matched α and γ_{mem} profiles, this would constitute direct empirical validation of the ρ_s extension. Immune repertoire sequencing datasets paired with allergen sensitization profiles already exist in the literature and could be used to test this prediction without new data collection.

6 The Resonance Window: A Novel Prediction

The most significant novel prediction generated by the refined CAPM is the existence of an *individually optimized exposure protocol* a specific sequence of doses $\{\zeta_i\}$ and intervals $\{\Delta t_i\}$ that drives $P_{\text{substrate}} \rightarrow P_0$ via **systemic feedback disruption** of the encoded false-attack pattern. We term this the **resonance window**. This concept connects directly to the clinical framework of *metronomic dosing* in mathematical oncology the finding that precisely timed low-dose continuous exposure can exhaust memory cell populations or trigger receptor desensitization before a full clonal expansion wave $N(t)$ can establish itself. The CAPM provides the unified mathematical substrate for this intuition.

6.1 Derivation

From the full Memory-Dissipation Law incorporating tolerance rewriting, natural forgetting, and the resensitization pressure identified empirically from the GRASS trial (Shamji et al., 2021):

$$\frac{\partial P_{\text{substrate}}}{\partial t} = \underbrace{\alpha \cdot \zeta(t) \cdot (P_0 - P_{\text{substrate}})}_{\text{tolerance rewriting}} + \underbrace{\gamma_{\text{mem}}(P_0 - P_{\text{substrate}})}_{\text{natural forgetting}} + \underbrace{\kappa \cdot N(t) \cdot (1 - P_{\text{substrate}})}_{\text{resensitization}} \quad (50)$$

where $\kappa > 0$ is the resensitization coefficient quantifying how strongly immune cell activation $N(t)$ re-encodes the attack pathway into substrate memory between exposure doses. This term was absent from earlier formulations and was identified as necessary through empirical fitting to the GRASS immunotherapy trial data, where it accounts for the clinically observed post-treatment rebound in allergen reactivity.

For a periodic pulsed exposure protocol with dose ζ and interval T :

Proposition 3 (Resonance Condition). *The exposure protocol achieves net rewriting of $P_{\text{substrate}}$ toward P_0 if and only if, averaged over one exposure cycle $[t_k, t_k + T]$:*

$$(\alpha \bar{\zeta} + \gamma_{\text{mem}})(\bar{P}_{\text{substrate}} - P_0) > \kappa \bar{n}(1 - \bar{P}_{\text{substrate}}) \quad (51)$$

where $\bar{P}_{\text{substrate}}$ and $\bar{n}$ are the cycle-averaged values and $\bar{\zeta} = \zeta \tau / T$ is the cycle-averaged dose rate for a pulse of width τ .

Sketch. Integrate Equation (50) over $[t_k, t_k + T]$ and require $\Delta P_{\text{substrate}} < 0$. The two rewriting terms contribute $-(\alpha \bar{\zeta} + \gamma_{\text{mem}})(\bar{P}_{\text{substrate}} - P_0)$, which is negative in

the regime $P_{\text{substrate}} > P_0$; the resensitization term contributes $+\kappa\bar{n}(1 - \bar{P}_{\text{substrate}})$, which is positive. Requiring the sum to be negative gives the stated inequality. $\square$

Remark 11 (Cycle-phase averaging). *The cycle averages in Equation (51) are averages over a full period T , not instantaneous values. Under a pulsed protocol the instantaneous state depends on where in the dose cycle it is sampled, and reading a single instant produces an artefactual dependence of every reported quantity on the sampling phase. Numerical evaluation of this model confirmed the effect: sampling at the final instant made the durable outcome non-monotone in treatment duration, while averaging over the final complete cycle restored monotonicity.*

6.2 The Counterintuitive CAPM Warning

Equation (51) constrains which protocols rewrite the substrate and which fail to. The direction of that constraint requires care, and an earlier formulation of this section stated it incorrectly.

What the equation gives as written. Both the rewriting drive and the dose-driven engine activity scale as $\bar{\zeta} = \zeta\tau/T$. Lengthening the interval at fixed pulse dose therefore weakens *both* sides of Equation (51) together, and the protocol converges on doing nothing rather than on doing harm. Numerical integration of Equation (50) confirms this: with the interval swept from $T = 1$ to $T = 32$ at fixed cumulative dose, the outcome varies monotonically and no interior optimum appears. Harm, where it appears at all, appears at *short* intervals and high κ , where engine activity accumulates faster than it clears.

What would be required for the opposite direction. For long intervals to worsen sensitization the clinically reported phenomenon of paradoxical worsening resensitization must be *superlinear* in engine activity, so that brief high spikes of $\bar{n}$ cost more than the same time-integrated activity spread out. Replacing $\kappa\bar{n}$ with $\kappa\bar{n}^q$ in Equation (50) achieves this, and an interior optimum in T emerges for $q \gtrsim 1.5$. We record q as an explicit parameter of the model rather than an assumption folded into it:

$$\left. \frac{\partial P_{\text{substrate}}}{\partial t} \right|_{\text{resens}} = \kappa \cdot \bar{n}(t)^q \cdot (1 - P_{\text{substrate}}), \quad q \geq 1 \quad (52)$$

At $q = 1$ this is Equation (50) unchanged. Whether $q > 1$ is biologically warranted whether, mechanistically, mast-cell priming or IgE crosslinking responds superlinearly to activation density is an empirical question this model does not settle.

The resonance window as a *sweet spot* rather than a monotone threshold depends on the answer, and that dependence should be stated rather than assumed.

6.3 The Optimal Waveform

The optimal protocol minimizes total treatment duration subject to the resonance condition being satisfied at every cycle. For a patient at position (α_p, γ_p) in phase space, the optimal interval is:

$$T^* = \frac{\ln(\kappa/(\alpha_p \zeta^* + \gamma_p))}{\lambda} \quad (53)$$

and the optimal dose:

$$\zeta^* = \frac{\kappa \bar{n}(1 - \bar{P}_{\text{substrate}})/(\bar{P}_{\text{substrate}} - P_0) - \gamma_p}{\alpha_p} \quad (54)$$

The difference is oriented as $(\bar{P}_{\text{substrate}} - P_0)$ in accordance with the convention fixed at Equation (46); with the opposite orientation ζ^* returns a negative dose throughout the clinically relevant regime.

These are individually patient-specific. Current immunotherapy protocols use population-averaged fixed intervals and doses, and the CAPM predicts that individualizing the waveform would improve outcomes.

Remark 12 (The optimal dose is not identifiable as written). *Equation (54) requires α_p as an input. But α and ζ enter Equation (50) only through the product $\alpha\zeta(t)$, and no trajectory of $P_{\text{substrate}}$ can separate them: numerical integration confirms that $(\alpha, \zeta) = (0.85, 1.2)$ and $(0.2125, 4.8)$ produce identical trajectories to six decimal places. Any procedure that fits α to reactivity data and then computes ζ^* from it recovers whatever value it started from.*

This does not defeat the prediction, but it changes its form. The identifiable quantity is the product $\alpha\bar{\zeta}$, and the clinically actionable statement is the drive a protocol must supply, not the dose it must deliver. Converting drive to dose requires an independent determination of α for instance a direct measurement of clonal expansion efficiency and cannot be obtained from the reactivity time course being fitted.

6.4 Tolerance Plasticity and the Durable Endpoint

Equations (24) and (50) treat γ_{mem} as a constant. That choice has a consequence which was not visible until the model was integrated numerically, and which invali-

dates the clinical prescription earlier versions of this work derived in the empirical-validation section.

Proposition 4 (Treatment cannot move a fixed- γ_{mem} endpoint). *With γ_{mem} constant and ν determined by Equation (42), the equilibrium of Equation (50) at $\zeta = 0$ is P^* for every protocol. A protocol holds $P_{\text{substrate}}$ away from P^* while it runs and the substrate returns to P^* when it stops, regardless of dose, duration, or schedule.*

Numerical integration confirms this across a tenfold range of dose and a tenfold range of treatment duration: the post-treatment plateau is identical to six significant figures in every case. The equilibrium is a property of the patient and their environment and never of the treatment.

This is fatal to the prescription *extend treatment until γ_{mem} exceeds κ* , because no equation in the framework permits treatment to change γ_{mem} . Either the prescription is unfounded or γ_{mem} must become a state variable. We take the second route, and the structural-invariance claim of Section 1 dictates its form.

The Memory-Dissipation Law Applied Twice

The regulatory apparatus that immunotherapy induces IgG₄ blocking antibodies, allergen-specific regulatory T cells, sustained mucosal IgA is itself an informational configuration, written onto its own substrate by the same execution engine, by the same controlled exposure. If the governing law is genuinely substrate-invariant, it applies there without modification. Let $P_{\text{tol}} \in [0, 1]$ denote that regulatory substrate:

$$\frac{dP_{\text{tol}}}{dt} = \alpha_{\text{tol}} \cdot \zeta(t) \cdot (1 - P_{\text{tol}}) - \gamma_{\text{tol}} \cdot P_{\text{tol}} \quad (55)$$

$$\gamma_{\text{mem}}(t) = \gamma_0 + (\gamma_{\text{max}} - \gamma_0) \cdot P_{\text{tol}} \quad (56)$$

Here γ_0 is the patient's constitutive forgetting rate the quantity ν was inverted from in Equation (42), and the one that does not move while α_{tol} measures how efficiently a given exposure route writes the regulatory substrate and $1/\gamma_{\text{tol}}$ sets how long the induced apparatus persists. Equation (55) is Equation (24) with P_{tol} in place of $P_{\text{substrate}}$; no new dynamical form is introduced.

Proposition 5 (Durable Endpoint). *With ν fixed by Equation (42) and hence $\kappa\bar{n}^q = A$ constant, the equilibrium of Equation (50) at a raised forgetting rate γ_f is*

$$P_{\text{end}} = \frac{\gamma_f P_0 + A}{\gamma_f + A} \quad (57)$$

with A given by Equation (43). The expression returns P^* exactly when $\gamma_f = \gamma_0$ and tends to P_0 as $\gamma_f \rightarrow \infty$.

The limit $\gamma_f = \gamma_0$ is the point of the construction: setting $\alpha_{\text{tol}} = 0$ recovers the constant- γ_{mem} model exactly, so Equations (55)–(57) extend the framework rather than replacing it.

Proposition 6 (Schedule Invariance). *P_{tol} integrates only the cycle-averaged dose rate $\bar{\zeta} = \zeta\tau/T$. Two protocols with equal $\bar{\zeta}$ therefore reach the same P_{end} whatever their pulse dose and interval.*

Numerical integration confirms this to within 2×10^{-4} across (ζ, T) from $(0.6, 3)$ to $(4.8, 24)$. Proposition 6 is the cheapest falsification test the framework offers: any immunotherapy trial that varied dose and interval together at matched cumulative rate and found a schedule effect refutes Equation (55) as written and requires an induction term with its own nonlinearity.

Proposition 7 (Phase Invariance of the Endpoint). *Under Equations (55) and (45), the durable endpoint is independent of the phase φ at which a protocol begins relative to the environmental cycle.*

Sketch. P_{tol} integrates $\zeta(t)$ alone; $\nu(t)$ enters only the population equation and never Equation (55). The trajectory of P_{tol} , and hence γ_f and P_{end} , is therefore independent of φ . $\square$

Numerical integration confirms the split this implies: sweeping φ through a full cycle at fixed protocol and fixed cumulative dose, on-treatment reactivity spans 0.41 to 0.58 while the durable endpoint is identical to every reported digit. **The model therefore predicts that pre-seasonal and co-seasonal immunotherapy differ in symptom burden during treatment and not at all in durable outcome.**

This is a sharper test than Proposition 6, because pre-seasonal immunotherapy is established clinical practice and the comparison has been made.

The Test Selects Between Two Formulations

Proposition 7 follows from the assumption that only the controlled dose induces tolerance. That assumption has a natural rival, and stating it makes the experiment decisive rather than merely falsifying.

The rival discriminates on the *regularity* of exposure rather than its origin: a sustained encounter induces tolerance, a spike does not, whatever its source. Writing

$E(t) = \zeta(t) + a_\nu \nu(t)$ for total encounter and $\bar{E}$ for its low-pass at timescale τ_{reg} ,

$$\frac{dP_{\text{tol}}}{dt} = \alpha_{\text{tol}} \min(E, \bar{E})(1 - P_{\text{tol}}) - \gamma_{\text{tol}} P_{\text{tol}} \quad (58)$$

Steady input gives $E = \bar{E}$ and full induction; a brief spike is clipped to the running level and contributes almost nothing. Setting $a_\nu = 0$ recovers Equation (55) exactly, so the two are nested.

They differ precisely where the experiment looks. Sweeping the seasonal phase φ at fixed protocol and fixed cumulative dose, the durable endpoint spread is *exactly zero* under Equation (55) the content of Proposition 7 and 6.2×10^{-3} under Equation (58), because ambient seasonal exposure now contributes to induction and its phase relative to the protocol matters.

The pre-seasonal comparison therefore does not merely test the framework; it selects between two versions of it. A demonstrated durable advantage to pre-seasonal dosing selects Equation (58); its absence selects Equation (55). Either outcome removes an assumption that Section 11.6 would otherwise have to carry indefinitely.

Finally, *durable* is not *permanent*. After cessation P_{tol} decays at γ_{tol} and $P_{\text{substrate}}$ drifts back toward P^* on the timescale $1/\gamma_{\text{tol}}$. Equation (57) gives the level held while the apparatus persists. Whether γ_{tol} is small enough for that to constitute clinical cure is an empirical question, and $1/\gamma_{\text{tol}}$ is the quantity to measure.

6.5 Empirical Reassessment: The GRASS Trial

The GRASS trial (Shamji et al., 2021; $n = 84$; baseline, Year 1, Year 2 on-treatment, Year 3 off-treatment) reports a post-treatment rebound in both arms. Fitting Equation (50) to those data with γ_{mem} held constant yields $\kappa = 9.79$, $\gamma_{\text{mem}} \approx 0$ for SCIT and $\kappa = 0.24$, $\gamma_{\text{mem}} = 0.50$ for SLIT.

The constant- γ_{mem} reading is inconsistent

Applying the reachability constraint of Equation (41) to those two fits shows that they cannot describe one trial population. The SCIT parameters sustain baselines up to $P^* \approx 0.99$; the SLIT parameters admit none above $P^* \approx 0.375$. Randomised arms of a single trial cannot have untreated baseline severities separated that widely. Read as fixed patient properties, the fitted values are mutually incompatible.

A second objection is independent of the first. Once ν is determined by Equation (42), $\kappa \bar{n}^q$ equals A identically and the relaxation rate is $\gamma_{\text{mem}} + A$, in which κ does not appear. κ survives only inside $\bar{n}^*$ an *activated-cell density*, not a reactivity.

Fitting κ to an IgE or symptom-score time course therefore estimates a parameter to which that time course is structurally insensitive. The fitted values should be regarded as unidentified rather than measured.

The plastic reading is consistent

Under Equations (55)–(56) the two arms share a constitutive $\gamma_0 = 0.01$ and a common baseline $P^* = 0.85$, and differ only in the route-dependent induction efficiency α_{tol} . Both parameter sets are then reachable. The difference the original fit attributed to γ_{mem} reappears as $\alpha_{\text{tol}} \approx 0.004$ for SCIT against $\alpha_{\text{tol}} \approx 0.163$ for SLIT, and Equation (57) separates the durable endpoints to $P_{\text{end}} \approx 0.59$ and 0.17 respectively.

The reinterpretation is the substantive claim. Under the constant- γ reading, the arms differ because the patients differ, which a randomised trial forbids. Under the plastic reading they differ because the *routes* differ in how efficiently they write the regulatory substrate consistent with the sustained mucosal IgA response characteristic of sublingual delivery, and consistent with randomisation. The rebound is then not a failure of the resonance condition but the decay of an induced apparatus at rate γ_{tol} .

Remark 13. *The prescription this yields differs from the one in earlier versions of this work. Extend treatment until γ_{mem} exceeds κ is not implementable, since κ is not identifiable from reactivity data and γ_{mem} is not a treatment variable under a constant- γ model. The implementable statement is: extend treatment until P_{tol} approaches its plateau $\alpha_{\text{tol}}\bar{\zeta}/(\alpha_{\text{tol}}\bar{\zeta} + \gamma_{\text{tol}})$, which takes a time of order $(\alpha_{\text{tol}}\bar{\zeta} + \gamma_{\text{tol}})^{-1}$, and select the route that maximises α_{tol} . Both quantities are estimable from serial IgG₄ or regulatory T-cell measurement rather than from symptom scores.*

7 The Subsystem Tensor

Section 11.3 lists the absence of parallel memory systems as a limitation. It is better understood as the framework's central structural gap, because a second and independent argument demands the same repair: Section 11.2 shows that a single substrate memory makes V common to every route, and a common shift cancels in the normalisation of Equation (11). Novelty cannot function without heterogeneity, and heterogeneity is what parallel subsystems supply.

7.1 Formulation

Replace the scalar pair $(P_{\text{substrate}}, N)$ with vectors over M subsystems indexed by i , each carrying its own memory P_i , normalised engine density $n_i = N_i/K_i$, structural resolution ρ_i , learning rate α_i , constitutive forgetting rate $\gamma_{0,i}$ and resensitization coefficient κ_i . Let χ_{ij} be the cross-subsystem bridge operator: the rate at which activity in subsystem j recruits ($\chi_{ij} > 0$) or suppresses ($\chi_{ij} < 0$) subsystem i . Then

$$\frac{dn_i}{dt} = \omega_i R_i n_i (1 - n_i) - \lambda_i n_i + \underbrace{\left(\sigma_i^\nu \bar{\nu} + \sigma_i^\zeta g_N \zeta(t) \right)}_{\text{exogenous}} + \underbrace{\sum_{j \neq i} \chi_{ij} n_j}_{\text{coupling}} (1 - n_i) \quad (59)$$

$$\frac{dP_i}{dt} = \alpha_i \zeta(t) (P_0 - P_i) + \gamma_i (P_0 - P_i) + \kappa_i n_i^q (1 - P_i) \quad (60)$$

where σ_i^ν and σ_i^ζ are the shares of ambient and controlled exposure reaching subsystem i . Crucially each subsystem sees the path integral through its *own* resolution and *own* memory:

$$V_i = \rho_i \cdot \text{JSD}(P_{\text{pattern}} \| P_i), \quad R_i = R[V_i, \rho_i] \quad (61)$$

which is Equation (17) with the node index s realised as a subsystem index rather than discarded. Clinically observed reactivity is the weighted aggregate $\bar{P} = \sum_i w_i P_i / \sum_i w_i$, with w_i the symptomatic weight of each population.

7.2 The Regulatory Subsystem Absorbs Tolerance Plasticity

Designate one subsystem regulatory. Its memory obeys Equation (55) rather than Equation (60), and it sets the forgetting rate of every effector:

$$\gamma_i = \gamma_{0,i} + (\gamma_{\max,i} - \gamma_{0,i}) P_{\text{reg}} \quad (62)$$

Equations (55)–(56) are therefore not a separate extension but one entry of the tensor. This was not imposed; it is what results from writing the subsystems down and asking which one immunotherapy acts on. Numerically, zeroing the regulatory column of χ raises the treated aggregate from 0.13 to 0.76: in this formulation immunotherapy operates through suppression of the effector engines rather than by rewriting effector memory directly.

7.3 The Rest State Is Algebraic

With $\zeta = 0$ the regulatory memory decays to zero, each effector runs at its constitutive $\gamma_{0,i}$, and Equation (60) solves in closed form for a given engine density:

$$P_i^{\text{rest}} = \frac{\gamma_{0,i}P_0 + \kappa_i n_i^q}{\gamma_{0,i} + \kappa_i n_i^q} \quad (63)$$

Only the engine densities require solving, and they relax on timescales of order $1/\lambda_i$ rather than $1/\gamma_{0,i}$. The distinction is practical: with $\gamma_0 \approx 0.02$ for the IgE compartment, a time-marched equilibrium search converges two orders of magnitude more slowly than Equation (63) and will report equilibria it has not reached.

7.4 Pathways Belong to Compartments

Equations (59)–(61) give each subsystem its own novelty, but the path integral only benefits if the pathways of Equation (11) are attached to them. Each candidate route is therefore assigned a compartment index, and reads that compartment’s memory and resolution:

$$V_\gamma = \rho_{\text{sub}(\gamma)} \cdot \text{JSD}(\vec{w} \parallel \vec{q}^{\text{sub}(\gamma)}) \quad (64)$$

The assignment is not cosmetic. It is what converts compartmental heterogeneity into pathway heterogeneity, and pathway heterogeneity is the only kind the normalisation in Equation (11) does not cancel. An attack route reading the mast-cell compartment and a tolerance route reading the regulatory compartment see genuinely different novelty because those compartments hold genuinely different memories; two routes reading one scalar cannot.

7.5 Novelty Recovered

The purpose of the construction is served. Because the compartments have genuinely different constitutive forgetting rates IgE at $\gamma_0 = 0.02$ settles near $P = 0.98$, mast populations at $\gamma_0 = 0.35$ near $P = 0.72$ their novelties differ. The measured spread

in V across effector subsystems rises from effectively zero in the scalar model to 3.9×10^{-2} , and the induced spread in R from 2×10^{-4} to 9.2×10^{-2} , roughly two orders of magnitude. Only with the tensor in place does Equation (17) carry the weight the framework assigns it.

The consequence for $R(t)$ follows from Equation (64): across a sweep of the substrate state, R spans 9.1×10^{-3} with the compartments separated against 3.8×10^{-4} without them. The path integral responds to substrate memory roughly twenty-four times more strongly, which is the difference between a term that participates in the dynamics and one that is present but inert.

7.6 What Constrains χ

The coupling matrix is the framework's largest unconstrained input, and the number of compartments M is a modelling choice rather than a measurement. Both admit a check that does not require any of χ 's entries to be known.

Sign structure is fixed by immunology; magnitudes are not. The topology of χ which entries are positive, which negative, and which vanish follows from established cellular immunology: Th2 cytokines drive B-cell class switching to IgE, IgE arms mast cells through FcεRI, and regulatory populations suppress both arms. The *magnitudes* are a different matter. Published cytokine and receptor data do not exist in a form that maps onto the rate constants χ_{ij} requires, and we do not assign values to them here. The parameters used in this paper are plausible in sign and ordering and are not fitted to anything.

The magnitudes are identifiable from blockade data. Unpublished is not the same as unmeasurable. Although no source reports χ_{ij} in these units, two interventions that are routine in experimental immunology constrain them directly: blocking upstream cytokine signalling reduces effector activation, and depleting the regulatory compartment raises it. Both responses are ratios of activation levels, so both are dimensionless and comparable across preparations.

Hold the topology fixed it is determined by the sign structure above and write the matrix as two scalar magnitudes acting on it,

$$\chi = \psi^+ A^+ - \psi^- A^- \tag{65}$$

where $A^\pm$ are the adjacency patterns of the recruiting and suppressing entries. Let

$\bar{n}_{\text{eff}}$ denote mean effector engine density at the untreated equilibrium, and define

$$\Phi^+ = 1 - \bar{n}_{\text{eff}}|_{\psi^+=0}/\bar{n}_{\text{eff}} \quad (\text{upstream blockade}) \quad (66)$$

$$\Phi^- = \bar{n}_{\text{eff}}|_{\psi^-=0}/\bar{n}_{\text{eff}} - 1 \quad (\text{regulatory depletion}) \quad (67)$$

Two observables, two unknowns. The map $(\psi^+, \psi^-) \mapsto (\Phi^+, \Phi^-)$ is not monotone everywhere strong suppression can extinguish the recruitment response entirely but its Jacobian is non-singular in the physiological region, and a coarse grid followed by a pattern search inverts it reliably. Round-trip testing confirms this: choosing (ψ^+, ψ^-) , computing the two responses, and solving back recovers the originals to three decimal places across the cases tested.

This completes a pattern the framework has followed throughout. The untreated baseline P^* replaces the ambient rate ν (Proposition 1); the dose-to-ambient ratio Ξ replaces the activation gain g_N (Section 5.2); the noise scale σ replaces the coherence attenuation constant (Section 1.1); and now two blockade responses replace the coupling magnitudes. In each case an unobservable parameter is expressed through a quantity an experiment yields. What remains genuinely unconstrained in χ is not its scale but its *pattern* and Proposition 8 shows that the pattern is constrained from a different direction entirely.

Convergence in M is checkable without knowing the magnitudes. A model whose predictions move when the compartment count is refined is reporting a property of its own discretisation. We therefore rebuild the system at $M \in \{3, 4, 6, 8, 12\}$, interpolating effector properties across the range the four-compartment version spans, assigning equal weights, and preserving the total drive reaching any compartment. The result separates two topologies sharply.

	star (parallel)		chain (relay)	
M	aggregate	Δ	aggregate	Δ
3	0.80145	—	0.80145	—
4	0.79248	0.00896	0.75656	0.04489
6	0.78831	0.00417	0.65058	0.10597
8	0.78702	0.00129	0.59636	0.05422
12	0.78604	0.00098	0.53602	0.06035

Table 4: Aggregate reactivity above the ignition threshold as the compartment count is refined, for two coupling topologies. Δ is the change from the preceding row.

Under a star topology, in which the effectors are activated in parallel from a common source, successive refinements change the answer by less than 10^{-3} and the model has a well-defined limit. Under a chain, in which each compartment relays activation to the next, the answer does not settle: lengthening the relay attenuates terminal activation, so every refinement lowers the aggregate again.

Proposition 8 (Topology, not magnitude, determines well-posedness). *The existence of a compartment-count limit for the coupled system is a property of the sign-and-adjacency structure of χ , not of the size of its entries. No rescaling of a chain topology produces a limit; a star topology possesses one for any magnitudes tested.*

The practical consequence is that χ is constrained from an unexpected direction. Its magnitudes remain unmeasured, but its *admissible topologies* are decided by a test requiring no measurements at all. The parameters used in Section 7 form a hybrid a two-step relay $\text{Th2} \rightarrow \text{IgE} \rightarrow \text{mast}$ alongside a direct $\text{Th2} \rightarrow \text{mast}$ term and therefore sit between the two behaviours, which is a reason for caution rather than confidence.

7.7 An Unplanned Consequence: The Ignition Threshold

The coupling matrix generates behaviour that no scalar version of the model exhibits. Sweeping ambient exposure, aggregate reactivity remains at 0.150 up to $\bar{\nu} = 0.3230$ and stands at 0.530 by $\bar{\nu} = 0.3237$: a jump of 0.38 across a change of 7×10^{-4} . Between those values there is no equilibrium at all, and baselines within the gap cannot be represented by the model.

The mechanism is the positive loop $\text{Th2} \rightarrow \text{IgE} \rightarrow \text{mast}$ opposed by regulatory suppression. Damping the positive entries of χ to a third of their value removes the discontinuity and the response becomes gradual. What the tensor produces, in other words, is a **sensitization threshold**: below a critical environmental load the network stays quiescent, above it the loop runs away.

Three remarks are in order. First, a threshold is a biologically reasonable property, and its appearance is a point in favour of the coupled formulation. Second, *nothing in CAPM predicted it*. It is a property of χ , and it establishes that the coupling matrix is not bookkeeping but carries dynamical content of its own. Any inversion of clinical data through this model must first check that the target baseline lies outside the ignition gap, since parameter fits inside it are fitting to a state the model does not possess.

Third, and following Section 7.6, the threshold requires a qualification that the convergence test supplies. The *location* of the threshold is robust: $\bar{\nu}_{\text{crit}}$ varies by

under 2% across $M \in \{3, \dots, 12\}$ under either topology. The *state the network ignites into* is not: under a star topology the post-ignition aggregate converges to within 10^{-3} , while under a chain it continues to fall with every refinement. The claim that a sensitization threshold exists is therefore better supported than any claim about what lies above it, and the two should be reported separately.

8 Pattern Mutation and Escape from Substrate Memory

The framework as developed so far holds the identity vector $\vec{p}$ fixed, while biological patterns evolve; Section 11.1 records this as a limitation and proposes a mutation term $d\vec{p}/dt = \eta \xi(t)$ (Equation (78)). That proposal requires one correction before it can be used, after which it yields more than was asked of it.

8.1 The Constraint Forces a Projection

Definition 2 requires $\|\vec{p}\|_2 = 1$: the identity vector is a direction in configuration space, not a magnitude. An unconstrained additive perturbation does not respect this. Integrating $d\vec{p} = \eta \xi dt$ with $\eta = 0.4$ for 2000 steps drives $\|\vec{p}\|$ from 1 to 8.15, at which point $c(s) = \max(0, \vec{p}^T M_s \vec{p})$ has lost its meaning as a bounded compatibility.

Mutation must therefore be a diffusion *on the sphere* S^{d-1} , obtained by projecting the noise onto the tangent space at $\vec{p}$:

$$d\vec{p} = \eta (\mathbb{I} - \vec{p} \vec{p}^T) \circ d\vec{W} \quad (68)$$

in the Stratonovich convention, with $\vec{W}$ a standard d -dimensional Wiener process. The projector $(\mathbb{I} - \vec{p} \vec{p}^T)$ annihilates the radial component, so $\|\vec{p}\| = 1$ is preserved exactly; numerically the norm holds to twelve decimal places over 2×10^4 steps. Equation (68) is the corrected form of the proposal, and η retains its intended meaning as the mutation rate.

8.2 The Decay Law for Compatibility

Equation (68) generates Brownian motion on S^{d-1} with generator $\frac{\eta^2}{2} \Delta_{S^{d-1}}$. The compatibility $c(s) = \vec{p}^T M_s \vec{p}$ is a quadratic form, which decomposes into a constant part and a degree-two spherical harmonic; the Laplace–Beltrami eigenvalue on degree- ℓ harmonics is $\ell(\ell + d - 2)$, giving $2d$ at $\ell = 2$. Hence:

Proposition 9 (Compatibility Decay Under Mutation). *Under Equation (68),*

$$\mathbb{E}[c(s)](t) = \bar{c} + (c(s)|_{t=0} - \bar{c}) e^{-\eta^2 d t}, \quad \bar{c} = \frac{\text{tr } M_s}{d} \quad (69)$$

Compatibility relaxes exponentially, at rate $\eta^2 d$, from its initial value to the isotropic average $\text{tr } M_s / d$ the compatibility a randomly oriented pattern would have.

Monte Carlo integration of Equation (68) over 4000 replicas reproduces the predicted rate to within 2% for $d \in \{4, 8, 16\}$ and $\eta \in \{0.15, 0.3\}$. Note the

dimensional dependence: escape accelerates in richer configuration spaces, because there are more directions in which to differ.

8.3 Numerical Verification

Proposition 9 is a closed-form claim and is therefore checkable. It was implemented in the reference simulator described in Section 3 and tested directly, using the parameterisation $M_s = a \vec{u}_s \vec{u}_s^T + b \mathbb{I}$ of Section 4.4, whose exact isotropic mean $\bar{c} = a/d + b$ is what makes the prediction falsifiable rather than merely plottable.

d	η	predicted $\eta^2 d$	measured	ratio
4	0.15	0.0900	0.0896	0.996
6	0.20	0.2400	0.2440	1.017
8	0.30	0.7200	0.6960	0.967
12	0.20	0.4800	0.4780	0.996
20	0.25	1.2500	1.2376	0.990
8	0.60	2.8800	2.7041	0.939

Table 5: Decay rate of $\mathbb{E}[c(s)]$ measured over ensembles of 1500 independent realisations of Equation (68), against the rate predicted by Proposition 9.

Agreement is within 1–6% across a fivefold range of dimension and a fourfold range of mutation rate, the residual being ensemble noise and the finite integration step.

Remark 14 (Proposition 9 predicts an expectation, not a path). *Equation (69) governs $\mathbb{E}[c(s)]$, and a single realisation of Equation (68) is a sample of size one from a distribution whose spread is not small. At $\eta = 0.25$ and $t = 20$, one realisation gives $c = 0.216$ while the ensemble mean over sixty-four gives 0.167 with a standard deviation of 0.077: the single path sits two-thirds of a standard deviation from the expectation, which is unremarkable for one draw and useless as a measurement. Any quantity in this framework that depends on $c(s)$ under mutation $R(t)$ in particular must therefore be evaluated over an ensemble, and reported with its spread. This is a requirement on the numerics rather than a limitation of the theory, but it is easy to violate without noticing, because a single trajectory looks like a result.*

Two further checks bear on the correction argued in Section 8.1. The projected form preserves $\|\vec{p}\| = 1$ to twelve decimal places over 5×10^4 steps at $\eta = 0.5$; the

unprojected form $d\vec{p} = \eta \boldsymbol{\xi} dt$ reaches $\|\vec{p}\| = 8.15$ within 2×10^3 steps at $\eta = 0.4$. The projection is a requirement, not a refinement.

Finally, with the anchoring condition Equation (37) imposed, the mutating model at $\eta = 0$ reproduces the static model to 1.1×10^{-16} over a full protocol simulation. Section 8 is therefore an extension of the preceding sections and not a replacement for them, in the same sense as tolerance plasticity and the subsystem tensor.

8.4 Escape Time and the Mutation–Adaptation Ratio

Equation (69) gives the time for a pattern to escape an established recognition:

$$\tau_{\text{escape}} = \frac{1}{\eta^2 d} \ln \left(\frac{c_0 - \bar{c}}{c_{\text{thr}} - \bar{c}} \right) \quad (70)$$

for a recognition threshold c_{thr} . The substrate, meanwhile, relaxes toward its own equilibrium at rate $\gamma_{\text{mem}} + A$ (Section 5). The competition between the two is governed by a single dimensionless group:

$$\text{Mu} = \frac{\eta^2 d}{\gamma_{\text{mem}} + A} \quad (71)$$

- $\text{Mu} \ll 1$: the substrate tracks the pattern. Memory, once written, remains valid, and Sections 5–6 apply as written.
- $\text{Mu} \gg 1$: the pattern escapes faster than the substrate can encode it. $c(s)$ decays toward $\bar{c}$ before the memory consolidates, novelty V never falls, and no stable encoding is reached.

At the reference parameters of the simulator ($d = 8$, $\gamma_{\text{mem}} = 0.05$, $P^* = 0.85$) the three regimes are separated by roughly one order of magnitude in η : $\eta = 0.05$ gives $\text{Mu} = 0.065$ with $\tau_{\text{escape}} = 45.3$, so the substrate tracks; $\eta = 0.15$ gives $\text{Mu} = 0.59$ with $\tau_{\text{escape}} = 5.0$, a contested regime in which encoding and escape proceed on comparable timescales; and $\eta = 0.6$ gives $\text{Mu} = 9.4$ with $\tau_{\text{escape}} = 0.31$, in which no encoding survives. The transition is sharp because Mu depends on η^2 .

8.5 Why Allergy Is the $\eta \rightarrow 0$ Limit

This resolves something the framework asserted but did not explain. Section 1.4 separates the allergen, a static blueprint, from the immune network that propagates the response. Equation (71) shows why that separation has consequences: an allergen is a protein of fixed sequence, so $\eta \approx 0$ and $\text{Mu} \approx 0$.

The persistence of allergic memory is therefore not a peculiarity of γ_{mem} but a direct consequence of the target not moving. A pathogen with $\text{Mu} \gtrsim 1$ escapes recognition and forces re-encoding; an allergen never escapes, so the false-attack pattern encoded against it is never invalidated and persists indefinitely. The same engine produces chronic sensitization for static blueprints and recurring evasion for mutating ones, and the sole parameter distinguishing them is η .

This also delimits the biomedical claims of this paper. Sections 5–7 are the $\text{Mu} \rightarrow 0$ theory. Extending them to oncology where Section 11.1 correctly observes that the configuration changes as the disease spreads requires the full $\text{Mu} > 0$ treatment, in which $c(s)$ carries its own dynamics and the compatibility tensor becomes a stochastic process rather than a fixed property of the substrate.

9 Pre-Symptomatic Detection via Quantitative Ultrasound

9.1 The Pre-Symptomatic Window: Detecting the Execution Engine

The stochastic phase transition formalism (Equation (27)) predicts a detectable pre-symptomatic window. Before clinical immune activation, the allergen template has entered tissue but $N(t)$ the immune execution engine has not yet undergone full clonal expansion:

$$\lim_{P_{\text{substrate}} \rightarrow \emptyset} C(\gamma_{\text{anomaly}}) \rightarrow 0 \implies R(t)_{\text{equilibrium}} \rightarrow 0 \quad (72)$$

Critically, the pre-symptomatic signal the framework predicts is **not** the allergen itself diffusing through tissue. Under the biological role separation established in Section 5.1, the allergen is a static blueprint. What is detectable is the *earliest sub-clinical activity of the immune execution engine* $N(t)$ the moment dendritic cells first encounter the allergen template and begin local signaling, before full clonal expansion activates.

When immune cells first respond to the allergen template, they induce localized changes in tissue architecture that precede any macroscopic allergic reaction:

- Dendritic cell migration alters local cellular packing density and interstitial fluid distribution
- Early T-cell recruitment increases local vascular permeability at the microvascular level
- Nascent mast cell priming changes membrane tension states prior to degranulation

Each of these events alters the *mechanical properties of the local tissue matrix* its viscosity, elasticity, and acoustic impedance in spatially localized clusters, before any visible inflammatory response appears. These are precisely the quantities quantitative ultrasound (QUS) measures directly.

9.2 What Quantitative Ultrasound Would Detect

Under the CAPM framework, the pre-symptomatic acoustic signature is a spatial cluster of three micro-elastic tissue anomalies driven by early $N(t)$ activity. Two points must be settled before they can be stated, because earlier versions of this section were wrong on both.

The Displacement Field Is Not the Novelty of Equation (17)

Earlier versions named the first signal V and described it as a novelty elevation. That collides with Equation (17), which defines $V = \rho_s \text{JSD}(P_{\text{pattern}} \| P_{\text{substrate}})$ the divergence from the *pattern*. What this section means is the divergence of local tissue state from *its own baseline*. The two move in opposite directions, because immune activity drives $P_{\text{substrate}}$ toward P_{pattern} : the most active site is the one that has learned most, and therefore the site of *least* novelty in the sense of Equation (17).

Numerically the point is unambiguous. At a focally activated site the Equation (17) novelty is depressed by a factor of 0.72 relative to quiescent tissue, while displacement from baseline is elevated by 1.25; the two fields correlate at -0.9997 across the strip. We therefore give the quantity its own symbol:

$$D(\mathbf{x}, t) = |P_{\text{substrate}}(\mathbf{x}, t) - P_0| \quad (73)$$

D is what rises at an incipient lesion. V falls there.

L and c Require Spatial Forms, and Not Arbitrary Ones

The second problem is that this section asserts spatial structure the framework could not produce. $L(s) = e^{-\varepsilon(s)}$ is a per-pathway constant and $c(s)$ depends on ρ_s and ΔG_0 , both global; neither varies with $\mathbf{x}$ anywhere in Sections 2–8. Supplying spatial forms is therefore required, but *how* they are supplied decides whether the prediction of this section says anything at all.

The obvious construction makes both monotone in local engine density $n(\mathbf{x})$. It must be rejected: it yields $\text{corr}(L, c) = 1.0000$ exactly, the triple (D, L, c) has rank one, and the co-localization of three readings of a single scalar field is an identity rather than a prediction. We instead take the three to be different *functionals* of the state, which is also what the physical descriptions below actually say:

$$L(\mathbf{x}) = \exp(-k_L \|\nabla n(\mathbf{x})\|) \quad \text{architecture is disrupted at interfaces} \quad (74)$$

$$c(\mathbf{x}) = \frac{\rho_s}{\rho_s + \exp(\Delta G_0/k_B T + k_c n(\mathbf{x}))} \quad \text{affinity shifts with local activity} \quad (75)$$

D reads the value of P , L the gradient of n , and c the value of n . Measured pairwise correlations are -0.71 , -0.997 and 0.70 : no longer degenerate, though D and c remain closely related since both are near-monotone in activity at equilibrium. The triple has rank two rather than three, and that should be stated rather than

glossed.

These are **observables, not couplings**. They are computed from the state and fed back into nothing, so no dynamical result elsewhere in this paper depends on them. That is the appropriate status for a detection model: a detector should not perturb what it detects.

The Three Signals

- (1) **A $D(\mathbf{x}, t)$ elevation substrate displacement.** Local tissue state departs from its baseline as immune cell infiltration alters cellular compliance. In QUS terms: localized increases in acoustic impedance mismatch at cellular boundaries arising from dendritic and T-cell clustering, detectable via broadband ultrasonic attenuation spectroscopy before inflammatory edema is present.
- (2) **An $L(\mathbf{x})$ reduction structural coherence loss.** As early immune signaling disrupts normal cellular organization, tissue micro-architecture deviates from its baseline geometric regularity. Equation (74) locates this at the *boundary* of an activated region rather than its interior, since uniformly activated tissue is not architecturally disordered the interface between recruited and quiescent tissue is. Detectable as changes in local viscosity and shear-wave velocity via shear-wave elastography, which resolves changes in the Pascal range.
- (3) **A $c(\mathbf{x})$ shift free-energy landscape displacement.** The thermodynamic binding landscape shifts locally as early immune activation alters the biochemical environment around receptor sites, manifesting as altered backscatter coefficient patterns in high-frequency QUS at 20–50 MHz.

What the Fingerprint Is, and What It Is For

Because L responds to the interface while D and c respond to the interior, the signature is not a single blob. It is a **displacement core with a coherence-loss annulus**: D and c peak at the activation site while L is minimal on its flanks, at an offset set by the width of the recruited region. This is a sharper and more falsifiable claim than co-localization of three coincident maxima, and it follows from the equations rather than being asserted alongside them.

The requirement that all three agree is best understood as noise rejection rather than as evidence of three independent phenomena. The channels are read by separate instruments attenuation spectroscopy, elastography, and backscatter analysis so their measurement errors are independent even though the underlying fields are not.

Simulating this directly, with per-channel noise scaled to each field's spatial standard deviation, the rate of correctly locating an activated site is:

noise σ	single channel	all three required
1	98.1%	99.8%
2	80.1%	95.3%
3	53.9%	81.5%
5	29.1%	54.7%

Table 6: Detection rate for the activation site, one channel versus the co-localization criterion, under independent per-instrument noise.

At realistic noise the triple criterion locates the site roughly half again as often. The clinical recommendation of this section therefore stands, but for a different reason than earlier versions gave: not because three independent phenomena coincide, but because three noisy views of one process can be combined to reject artifacts. That is a weaker biological claim and a stronger methodological one.

Remark 15. *This prediction is grounded in established QUS capabilities. Shear-wave elastography currently resolves tissue stiffness changes at the level of individual lobules in breast tissue. High-frequency quantitative backscatter analysis resolves cellular-scale structural changes in thyroid and lymph node tissue. The CAPM prediction requires no new instrumentation it requires a new measurement target: early immune cell infiltration signatures rather than established inflammatory markers.*

9.3 The Pre-Symptomatic Time Window

The framework provides a quantitative prediction for how far ahead of clinical detection the acoustic signature should appear. The pre-symptomatic window Δt_{pre} is bounded by the time required for $N(t)$ to grow from the minimum acoustically detectable immune cell density N_{QUS} to the clinical detection threshold N_{clinical} :

$$\Delta t_{\text{pre}} \approx \frac{1}{\omega_0 R(t)} \ln \left(\frac{N_{\text{clinical}}}{N_{\text{QUS}}} \right) \quad (76)$$

This follows directly from the early exponential growth phase of Equation (26) when $N(t) \ll K$. The quantity N_{QUS} is a measurable, instrument-dependent physical threshold that future QUS calibration studies could establish, converting the qualitative prediction into a quantitative detection protocol. The Fisher-KPP propagation front (Equation (31)) provides an additional geometric prediction: the spatial extent

of the pre-symptomatic cluster grows at velocity $v = 2\sqrt{\mu \cdot \omega_0 R}$, giving a time-dependent spatial signature that would distinguish genuine early immune activation from measurement noise or tissue heterogeneity artifacts.

10 Cross-Substrate Parameter Calibration

10.1 The Invariance Claim, Tested

Section 1 makes substrate invariance the framework’s deepest commitment, and every quantitative result in this paper comes from the biomedical column of the table below. That makes the central thesis the least tested claim in the work. This subsection tests it, by running the equations under silicon parameters and asking which of them survive.

The distinguishing feature of silicon is not its labels but its timescales. A processor clock runs at order 10^9 Hz while garbage collection runs at order 1–100 Hz, so $\omega_0/\gamma_{\text{mem}} \sim 10^7$, against roughly 10 in tissue. Sweeping that separation gives a clean and immediate result.

The population wave transfers exactly. Setting $\nu = 0$ and integrating, the equilibrium density reproduces the logistic prediction $n^* = 1 - \lambda/(\omega_0 R)$ across three decades of clock rate: measured 0.1626, 0.5785, 0.8313, 0.9578, 0.9916 against predicted 0.1591, 0.5795, 0.8318, 0.9580, 0.9916. Equation (26) is substrate-invariant in the strong sense not merely relabelled but quantitatively correct in a domain it was not built for. This is the clearest positive evidence the framework has.

The ambient exposure inversion does not transfer, and the reason is structural. Proposition 1 requires $\nu \geq 0$, hence $\omega_0 R(1 - \bar{n}) \leq \lambda$: the substrate must be unable to sustain itself without external input. Tissue satisfies this remove the allergen and the response dies. Silicon does not: the clock keeps processes alive with no external input at all. The critical point sits at $\omega_0 \approx 0.98$ in these units, and above it the inversion returns no solution at any timescale separation tested from 10^2 to 10^7 .

This is a genuine limit on the framework rather than a defect in it. **The biomedical machinery of Sections 5–6 assumes a driven substrate; silicon is a self-sustaining one, and the two are opposite regimes.** P^* , the observable the inversion is built on, has no analogue in a system that does not relax to a baseline when its input is removed.

Substrate memory saturates and stops carrying information. Above $\omega_0/\lambda \approx 8$ the substrate state pins at $P \approx 0.975$ and ceases to respond: measured 0.9703, 0.9749, 0.9752 across two further decades. Read as cache occupancy this is arguably correct a busy processor’s cache is always full but it means Equation (24) discriminates

nothing in that regime, and any prediction resting on substrate memory is vacuous there.

The path sum requires a much smaller discount. Switch fabrics fan out far more widely than tissue. At branching factor 16 the ratio of Equation (29) is 5.0, and at 32 it is 10.1, so $R(t)$ is dominated by its longest paths unless γ_{len} falls to 0.18 or 0.09 respectively. The value used biomedically, 0.92, is inadmissible on silicon topology by two orders of magnitude.

10.2 The Multi-Agent Column: An Independent Confirmation

The silicon result of Section 10.1 is a single data point, and a single data point admits luck. The multi-agent column supplies a second and sharper test, because unlike silicon it can sit on *either* side of the driven/self-sustaining boundary: a message below the epidemic threshold requires continued seeding, and one above it spreads unaided.

In the framework's terms the basic reproduction number is

$$\mathcal{R}_0 = \frac{\omega_0 \Theta R}{\lambda} \quad (77)$$

and standard susceptible–infected–susceptible theory, which forms no part of this framework and against which nothing here was fitted, predicts a threshold at $\mathcal{R}_0 = 1$ with endemic prevalence $n^* = 1 - 1/\mathcal{R}_0$ above it.

ω_0	$\mathcal{R}_0$	inversion	adoption n^*	$1 - 1/\mathcal{R}_0$
0.8	0.577	applies	0.00000	0.00000
1.2	0.866	applies	0.00000	0.00000
1.6	1.155	undefined	0.13424	0.13386
3.0	2.165	undefined	0.53332	0.53806
8.0	5.773	undefined	0.82486	0.82677

Table 7: Multi-agent substrate against standard epidemic theory. Prevalence agrees to three decimal places; the ambient inversion applies below threshold and is undefined above it.

Bisecting for the point at which spread first appears locates it at $\mathcal{R}_0 = 0.971$, within 3% of the theoretical value of one.

This Agreement Is Algebraic, Not Empirical

It is tempting to read the table as independent confirmation of substrate invariance. It is not, and the distinction matters enough to state plainly. Setting Equation (26) with clearance to zero,

$$\omega_0 R (1 - n) = \lambda \implies n^* = 1 - \frac{\lambda}{\omega_0 R} = 1 - \frac{1}{\mathcal{R}_0},$$

which is the SIS endemic equilibrium in two lines. **The logistic equation with linear death is the SIS equation.** The agreement above is therefore an identity that had to hold, not a prediction that might have failed and the same objection applies to the silicon result of Section 10.1, where logistic growth was assumed and logistic saturation was recovered.

What the two exercises do establish is weaker but not empty: the parameters of this framework map onto those of two well-validated models without inconsistency, and the mapping is quantitative rather than nominal. That is a calibration result. It is not evidence that the framework's *additional* structure the path integral, the substrate memory, the tensor carries across, because none of that structure is exercised by an equilibrium calculation.

Two things do survive the objection. The first is the crossover itself: the ambient inversion applies precisely below threshold and is undefined precisely above it, which identifies what Proposition 1 encodes. It is not a fact about immunology but the condition $\mathcal{R}_0 < 1$, and that identification was not visible from the biomedical column alone. The second is the residual: the measured prevalence departs from $1 - 1/\mathcal{R}_0$ by up to 5.4×10^{-3} , because R and Θ drift with the state while plain SIS holds its parameters fixed. That departure is small, but it is the only part of the comparison that is not guaranteed in advance.

A Prediction That Requires Θ

Equation (77) contains Θ , and this is where the correction of Section 2.3 earns its place. Holding reach and audience fixed and degrading only the contextual match, $\mathcal{R}_0$ falls from 1.794 to 0.477 and the message crosses from spreading unaided to requiring continued seeding **a message can be made sub-critical purely by fitting its audience worse.**

Without Θ the same sweep produces the opposite: $\mathcal{R}_0$ would *rise* from 1.957 to 3.133 as the match degrades, so a poorly-targeted message would spread more readily than a well-targeted one. The pathology identified in Section 2.3 is not a technicality

of the biomedical reading; it inverts the central quantity of epidemic theory.

This is also where the framework says something SIS does not. In SIS, $\mathcal{R}_0$ is fixed by contact and recovery rates alone; the fit between message and audience has no representation. Here $\mathcal{R}_0$ carries a third factor, and it moves the threshold: at $\Delta G_0 = 0$ the same protocol gives $\mathcal{R}_0 = 1.277$ where plain SIS would give 1.957 regardless of fit. **The genuine contribution of the multi-agent column is not that the model reproduces epidemic theory it must but that it extends it with a fit-dependent reproduction number.** Whether real diffusion data show that dependence is an empirical question, and a sharper one than any equilibrium comparison.

Verdict

Invariance holds *in part*, and the partition is informative. Two substrates have now been tested rather than one. Their agreement on the population equation is algebraic rather than evidential (Section 10.2), but the boundary they locate is not: the ambient inversion transfers exactly where $\mathcal{R}_0 < 1$ and nowhere else, which identifies a condition that the biomedical column alone left implicit. The population dynamics carry across exactly. The path integral carries across with a topology-dependent reparameterisation that the framework now specifies. The substrate memory law carries across formally but degenerates. The ambient inversion does not carry across at all, because it encodes an assumption about the regime rather than about the mechanism.

The honest form of the claim is therefore narrower and more useful than the one Section 1 makes: **the propagation machinery is substrate-invariant; the biomedical apparatus built on top of it is not, and assumes a driven substrate.** Four entries in the table below already carry no counterpart outside the biomedical column, which is the same fact recorded at the level of individual parameters.

The table below consolidates all model parameters across the three deployment substrates silicon infrastructure, dynamic multi-agent networks, and the new biomedical substrate introduced by this refinement providing a unified reference for empirical calibration and cross-domain validation.

Var.	Bounds	Silicon Substrate	Multi-Agent Net- work	Biomedical (New)
$N(t)$	$[0, K]$	Active memory pro- cesses	Minds with active re- call	Activated immune cells
K	> 0	RAM capacity (bytes)	Total target popula- tion	Max immune cell count
ω_0	> 0	Processor clock (Hz)	Daily interaction rate	Immune cycle rate
$E(s)$	> 0	Transmission power (J/Tb)	Attention time (s)	Receptor binding en- ergy
$V(s, t)$	$[0, 1]$	Shannon diver- gence (bits)	Surprise response in- dex	Tissue acoustic impedance shift
$L(s)$	$[0, 1]$	Syntax validation rate	Logical coherence score	Cell structural regu- larity
$c(s)$	$[0, 1]$	Protocol matching score	Contextual rele- vance index	Quantum docking probability
α	$(0, 1)$	Cache update weight	Familiarity learning rate	IgE clonal expansion rate
γ_{mem}	$[0, 1]$	Garbage collection rate	Memory decay rate	IgE antibody half-life
μ	> 0	—	—	Molecular diffusion coefficient
κ	> 0	—	—	Resensitization coeffi- cient
ρ_s	> 0	Protocol version fi- delity	Audience discrimina- tion capacity	TCR diversity / HLA specificity

Table 8: Extended parameter reference frames including the biomedical substrate. Variable slots are mathematically open free parameters, whose specific numerical ranges are calibrated per domain during empirical testing.

11 Epistemological Boundaries and Limitations

The refined CAPM makes strong predictions. It is equally important to state clearly where the framework fails or requires further extension.

Three of the limitations recorded in earlier versions of this work have since been addressed in the body of the paper, and are retained here in reduced form both because the diagnosis remains the reason the corresponding machinery exists, and because each resolution left a narrower limitation behind it. Where that is the case, the section says so and states what still stands.

11.1 Mutations and Evolving Pattern Identity

Addressed in Section 8; a narrower limitation remains.

The framework as originally formulated assumed a fixed identity vector $\vec{p}$, while biological patterns evolve cancer especially changes its configuration as it spreads. The mutation term proposed for this purpose,

$$\frac{d\vec{p}}{dt} = \eta \cdot \xi(t), \quad (78)$$

cannot be used as written, because it does not preserve the constraint $\|\vec{p}\|_2 = 1$ imposed by Definition 2. Section 8 supplies the projected form (Equation (68)), which does, and derives the compatibility decay law, the escape time, and the mutation–adaptation ratio Mu from it.

What remains open is the coupling of $\text{Mu} > 0$ to the rest of the machinery. Section 8 treats $c(s)$ as decaying under mutation with the substrate held fixed; a full evolutionary treatment requires $c(s)$ to become a stochastic process entering the partition function (Equation (11)) and the population wave simultaneously, so that selection acts on $\vec{p}$ through $R(t)$ rather than mutation merely eroding recognition. Until that is done, the biomedical sections of this paper should be read as the $\text{Mu} \rightarrow 0$ theory, and the oncological application remains out of reach.

11.2 The Novelty Channel Is Structurally Weak

Addressed in Sections 2.6 and 7; narrower limitations remain.

Equation (17) exists to make novelty govern propagation: as a substrate learns a pattern, $V \rightarrow 0$, the action rises, and $R(t)$ falls. This feedback is the mechanism by which Section 3.1 claims novelty regenerates over time. Numerical evaluation shows the loop gain is close to zero.

Driving the substrate across its full range, from $P_{\text{substrate}} = 0.95$ to 0.1, moves $R(t)$ by 0.011 under Equation (12) and by 4×10^{-4} under Equation (14). The cause is not saturation of the hyperbolic tangent, as might be supposed: at the operating point the argument $V \cdot L \cdot c$ is approximately 0.3, well inside the linear region. The cause is that V enters as a quantity common to all candidate routes, and a common shift of every action cancels in the normalisation of Equation (11). Only the between-route *spread* in ΔS survives, and that spread is 0.481 under the raw action and 0.006 under the hybrid form.

Two partial repairs are available and both are reported here for completeness rather than adopted as claims.

Restoring the node index. Equation (17) indexes novelty by node, $V(s, t, \rho_s)$, but any well-mixed reduction collapses it to a single scalar. Allowing each route to read the substrate memory of the tissue it traverses raises the dynamic range from 2×10^{-4} to 2.1×10^{-3} , a factor of ten.

Unbounding the coupling. The factor $1/(1 + \tanh x)$ is bounded below by $1/2$, so novelty can never reduce the action by more than a factor of two however large V becomes. The substitution $1/(1 + \tanh x) \rightarrow e^{-x}$ agrees to first order in x identical wherever the original is linear and does not flatten. Combined with the node index it reaches 5.0×10^{-3} , and 1.8×10^{-2} under the raw action.

Neither repair is sufficient, and the reason is instructive. Both depend on different regions having learned different things, and in the present formulation they do not: no term diffuses $P_{\text{substrate}}$, so a spatially uniform execution population leaves a spatially uniform substrate. Simulated tissue is uniform to five decimal places under both quiescent and broad exposure, and even under focal exposure the candidate routes read novelty values within 0.011 of one another.

The conclusion is that novelty is close to inert in any *scalar* formulation of CAPM, and that no adjustment internal to the path integral repairs it. What such a formulation lacks is a structural reason for distinct compartments to hold distinct memories.

Two independent repairs resolve this in combination, and they address different halves of it.

Range comes from the tensor. With the compartments separated and each pathway reading the memory of the compartment it belongs to, the spread in V across effector subsystems rises from effectively zero to 3.9×10^{-2} under the scalar

divergence and to 6.9×10^{-2} under the epitope-feature form of Section 2.6 while the spread in R across a substrate sweep rises from 3.8×10^{-4} to 9.1×10^{-3} , roughly twenty-four fold. In a running simulation the per-pathway novelties separate rather than being identical: 0.075, 0.357 and 0.004 under the epitope-feature divergence of Section 2.6, or 0.133, 0.357 and 0.230 under the scalar form against which the figure was first measured.

Discrimination comes from the feature divergence. The tensor makes V differ between pathways but leaves it blind to *which* epitopes differ. Section 2.6 supplies that: two profiles with identical scalar means separate by 0.490 where the Bernoulli form gives exactly zero.

The diagnosis above is retained because it is the reason both constructions are prerequisites rather than refinements, not because the defect is still present.

Three narrower limitations survive. First, the partial repairs discussed earlier the node index and the unbounded coupling remain insufficient *on their own*. Second and this corrects a claim made in an earlier version of this section the apparent cost of the feature divergence in sensitivity is a scale factor rather than a loss of information. The Bernoulli form's V spans 0.620 across the learning range where Equation (19) spans 0.180, a ratio of 3.44; the shortfall in R was 3.5-fold. These are the same number. The novelty coupling constant k_V of Equation (12) is simply set too low for the narrower divergence, and raising it recovers the range: with the unbounded coupling (Section 11.2) at $k_V = 8$ the feature form reaches 9.98×10^{-3} against the scalar form's 9.07×10^{-3} , slightly ahead rather than behind. Critically, k_V leaves the resolution ratios of Remark 9 untouched at 0.85/0.159/0.013 for every value tested, so the discrimination and identifiability gains survive the rescaling intact. Nothing in the shortfall is irreducible, and the earlier assertion that some of it was should be disregarded. Third, no term diffuses $P_{\text{substrate}}$ (Section 11.6, item 4), so geometric heterogeneity in substrate memory cannot arise from the equations as written. Novelty is carried by compartment identity rather than by tissue location, which is a modelling choice and not a derived result.

11.3 Parallel Immune Subsystems

Addressed in Section 7; narrower limitations remain.

The immune network is not a simple directed graph. It contains parallel, redundant, and sometimes contradictory memory systems B cells, T cells, IgE antibodies, mast cells each requiring its own topology Γ_t . Section 7 implements the corresponding tensor: memory and engine density become vectors over subsystems, coupled by the

bridge operator χ_{ij} , with the regulatory compartment absorbing the plasticity of Section 6.4 as one of its entries.

Three limitations remain, and the first is the most consequential.

1. **χ_{ij} is unmeasured, but no longer unmeasurable.** Section 7.7 shows it is not bookkeeping: the coupling matrix determines the location of the sensitization threshold and therefore which baselines the model can represent at all. The values used in this paper remain plausible in sign and ordering rather than fitted. What has changed is that a protocol now exists: Section 7.6 reduces the magnitudes to two scalars recoverable from cytokine-blockade and regulatory-depletion responses, and Proposition 8 constrains the topology by an argument requiring no data at all. The limitation is therefore that the measurements have not been *made* here, not that they cannot be.
2. **The decomposition is a choice, and only partly vindicated.** Four compartments are used because four suffice to exhibit effector heterogeneity and regulatory suppression, not because the immune system has four. Section 7.6 now tests convergence in M and finds it topology-dependent: results converge under parallel activation and do not under a sequential relay. Since the matrix used here is a hybrid of the two, the quantities reported from it inherit that ambiguity, and only the threshold location not the state above it should be regarded as settled.
3. **Subsystems are well-mixed.** Each carries a scalar memory and a scalar density, so the tensor supplies compartmental but not spatial heterogeneity. A treatment in which each subsystem carries its own topology Γ_t in the full sense, rather than a single lumped node, remains undone.

11.4 Sensitization Cannot Occur

The framework describes a sensitized substrate in detail: how it is maintained, how it responds to exposure, how it can be driven toward tolerance and how far. It cannot describe how it came to be sensitized. **Baseline severity P^* is an input, never an outcome.**

The point is not subtle. Administering a single exposure to a patient at a low baseline and following them afterwards, over a 300-fold range of dose from $\zeta = 0.2$ to 60, every trajectory returns to within 2% of where it started. There is no dose at which the encounter leaves a durable mark. This is the same structural fact identified in Remark 6 a unique equilibrium admits no switching viewed from the clinical side

rather than the dynamical one, and it is the more consequential view: a paper about allergy contains no mechanism for becoming allergic.

Repairing this requires bistability, and specifically a low-severity state and a high-severity state coexisting at one environmental load, with a dose threshold separating them. The natural form is an Allee term, $\omega_0 R N(N - a)(1 - N/K)$, expressing that a clonal population below some precursor frequency cannot establish which is itself good immunology, and would additionally make Equation (27) do what Section 3.4 once claimed for it.

We have not made that change, and the reason is worth stating rather than burying. Every extension in this paper was forced by an internal inconsistency and reduces to the prior formulation in a limit: ν from P^* because ν is unobservable, plasticity because the clinical prescription had no mechanism, the subsystem tensor because novelty cancelled identically without it. An Allee term would not be forced in that way; it would be introduced because a claim in Section 3.4 would otherwise be false, which is the failure mode Section 11.5 describes. It also carries a specific cost: with two equilibria at one $\bar{\nu}$, the inversion of Proposition 1 becomes branch-dependent, and P^* ceases to be a single well-defined observable. That inversion is currently the cleanest result in the paper.

The correct order is therefore the reverse of the tempting one. Onset is a genuine gap and should be closed, but as an extension in its own right with the inversion re-derived on branches and the reachability constraint of Section 5.2 restated accordingly rather than as a patch that happens to rescue a sentence. Until then, this framework should be read as a model of *maintenance and modification* of an established sensitization, not of its acquisition.

11.5 A Recurring Failure Mode

The corrections recorded through this paper share a shape, and naming it is more useful than listing them again.

Almost none of the claims that failed under testing were simply false. They were *conditionally* true, with the conditions unstated. The front velocity of Remark 7 is correct where the empty state is unstable and the ambient rate vanishes, and inapplicable at the paper's own reference parameters, where neither holds. The tolerance boundary of Equation (46) is correct when evaluated at the current forgetting rate, and wrong by 0.31 in substrate memory when evaluated at the constitutive one. The co-localization criterion of Section 9.2 is informative when its three fields are distinct functionals, and an identity when they are not. The compatibility decay of Proposition 9 governs an expectation, and says nothing useful about a single

realisation. In each case the equation was sound and the scope was missing.

This is what one should expect of a framework built by carrying a geometric intuition across substrates, and it suggests where to look for the errors that remain. They are unlikely to be in the algebra. They are likely to be assumptions so natural to the original setting that they were never written down, and which fail silently once the same expression is used somewhere else. Stating the domain of validity alongside each result is therefore not editorial tidiness in a framework of this kind; it is the principal defence available.

11.6 Assumptions Adopted Without Derivation

Four modelling choices were required to make the system integrable and are not consequences of anything in the framework. A fifth, the attenuation constant of the cognitive coherence law, appeared in this list in earlier versions and has since been removed from it: Section 1.1 derives that constant as $1/\sigma$ rather than fitting it. Each changes results, and each is recorded here so that it can be argued with rather than inherited.

1. **Only the controlled dose induces tolerance now an explicit alternative rather than a bare assumption.** $\zeta(t)$ drives Equation (55); ambient exposure ν does not. The justification is negative if uncontrolled encounter induced tolerance efficiently, sensitization would resolve spontaneously in everyone and no principle in CAPM distinguishes the two sources. Section 6.4 now states the rival formulation (Equation (58)), shows the two are nested, and identifies the measurement that separates them. The assumption remains unresolved, but it is no longer unexamined.
2. **The regulatory substrate is systemic.** P_{tol} and P_{reg} are scalars rather than fields, on the grounds that the induced apparatus circulates. Focal exposure therefore yields systemic tolerance in this model.
3. **Exposure recruits the execution engine.** The gain g_N in Equation (59) has no counterpart in the original formulation, which has exposure rewriting the substrate without connecting it to the population that performs the rewriting. Its value is unconstrained by anything here.
4. **Substrate memory does not diffuse a decision, not an omission.** Only N carries a Laplacian in Equation (31). This was listed here in earlier versions as a gap, on the reasoning that spatial coupling might supply the heterogeneity the novelty channel needs. Numerical testing shows the opposite: adding a diffusion

term to $P_{\text{substrate}}$ lowers its spatial variance monotonically, from 6.00×10^{-3} at $\mu_P = 0$ to 5.00×10^{-3} at $\mu_P = 4$, because diffusion is a smoothing operator. It would weaken the heterogeneity rather than create it. The biological reading agrees: immunological memory resides in cell populations, which do not diffuse on the timescales modelled here. The absence of the term is therefore justified, and the earlier classification of it as a limitation was mistaken.

5. **Substrate memory does not diffuse.** Only N carries a Laplacian in Equation (31); $P_{\text{substrate}}$ has no spatial coupling. A spatially uniform engine therefore leaves a spatially uniform memory, which is why geometric heterogeneity alone could not revive the novelty channel and why Section 7 was necessary.

11.7 Parameter Identifiability

Several parameters of the framework cannot be separated from one another by any measurement of $P_{\text{substrate}}$. They are collected here because the distinction between a parameter and an identifiable combination governs which predictions can legitimately be fitted to data.

Confounded	Identifiable	Reason
α, ζ	$\alpha\bar{\zeta}$	They enter Equation (50) only as a product (Remark 12).
κ, γ_0	$A, \bar{n}^*$	Once ν is inverted, $\kappa\bar{n}^q = A$ identically and κ leaves the substrate dynamics. It survives only as an activated-cell density.
$\rho_s, \Delta G_0$	g_{eff}	$c(s)$ depends on the difference $\Delta G_0/k_B T - \ln \rho_s$ alone (Remark 9).
g_N, ν	Ξ	g_N appears only as $g_N\zeta$, added to ν in the same units. The identifiable quantity is the ratio of controlled to ambient encounter (Section 5.2).

Table 9: Confounded parameters and the combinations data can determine.

Separating any pair requires an observable outside the reactivity time course: α from direct clonal-expansion measurement, κ from activated-cell density, ρ_s from repertoire sequencing paired with independent calorimetric determination of ΔG_0 . Fitting procedures that estimate these individually from symptom or IgE trajectories return their initialisation rather than a measurement.

11.8 Self-Referential Predictability

The model cannot fully predict its own behavior when $P_{\text{substrate}}$ is altered by the measurement act itself analogous to the quantum measurement problem. Pre-symptomatic acoustic measurement that perturbs tissue mechanically may itself shift $P_{\text{substrate}}$ slightly. This boundary is not a defect but a fundamental epistemological limit that should be incorporated into experimental design.

A weaker but more immediately practical version of the same problem arises in the numerical treatment. Several quantities reported in this paper are averages over a dose cycle, and the act of sampling them interacts with the protocol being sampled: reading the state at a single instant makes every reported figure depend on where in the cycle it was read, which produced a non-monotone dependence of durable outcome on treatment duration before it was corrected. The resolution averaging over a complete cycle with the window boundary treated as an integration event is stated in the numerical remark of Section 3. The general point is that in a system driven by a periodic protocol, the observable is not the state but a functional of it, and the choice of functional is part of the model rather than a reporting convenience.

12 Conclusion and Future Directions

The refined Copy And Paste Model provides an integrated, mathematically formalized bridge connecting molecular-scale thermodynamic binding dynamics to macroscopic immunological population dynamics. The three refinements continuous field reformulation, thermodynamic diffusion augmentation, and statistical mechanics correction to the compatibility tensor extend the original CAPM from a network information model into a physically grounded framework applicable to biological substrates.

The principal novel contributions are:

1. A formal characterization of allergy as a substrate memory rewriting problem, with explicit biological role separation between the allergen as static blueprint and the immune network as active execution engine, and unified mathematical parameters α , γ_{mem} , and κ that individually locate each patient in a biologically meaningful phase space.
2. The tolerance boundary $\partial\mathcal{T}$ as a mathematically precise target for immunotherapy, replacing empirical trial-and-error with a trajectory optimization problem in phase space.
3. The ambient exposure inversion (Equations (41)–(42)), which replaces an unobservable environmental rate with the observable untreated baseline P^* , and which yields a reachability constraint capable of rejecting fitted parameter sets on internal-consistency grounds alone.
4. Tolerance plasticity (Equations (55)–(56)) and the durable endpoint $P_{\text{end}} = (\gamma_f P_0 + A)/(\gamma_f + A)$. With γ_{mem} constant no protocol can move the post-treatment equilibrium; making it a state variable written by the same law on a second substrate is what permits treatment duration to matter at all. Setting $\alpha_{\text{tol}} = 0$ recovers the earlier formulation exactly.
5. A reinterpretation of the GRASS trial in which the two arms share a constitutive forgetting rate and a common baseline and differ only in route-dependent induction efficiency a reading consistent with randomisation, which the constant- γ_{mem} fit is not.
6. The resonance window prediction: an individually optimized exposure protocol connected to the clinical concept of metronomic dosing together with Proposition 6, that outcome depends on cycle-averaged dose rate alone and not on schedule, and Proposition 7, that treatment phase relative to the pollen season

alters symptom burden but not durable outcome. The latter is testable against the existing pre-seasonal immunotherapy literature.

7. The reduction of the coupling matrix to two experimentally recoverable scalars (Section 7.6) and of the epitope kernel width to a cross-reactivity index (Section 2.6), completing a pattern in which every unobservable parameter of the framework the ambient exposure rate, the activation gain, the coherence attenuation constant, the coupling magnitudes, and the resolving width is expressed through a quantity an experiment yields. No free constant remains on which a result of this paper depends.
8. The subsystem tensor (Section 7), which is required rather than optional: with a single substrate memory the novelty term is common to all pathways and cancels identically, so Equation (17) carries no weight until the compartments are separated. The coupled form additionally produces a sensitization threshold (Section 7.7) that no scalar version exhibits, establishing that χ_{ij} carries dynamical content of its own.
9. The identification of Definition 2 and Equation (35) as one quantity in two coordinate systems (Section 4.4), with the change of variables $\Delta G_{\text{eff}}/k_B T = \ln(\rho_s(1-c)/c)$, the anchoring condition required for consistency between them, and the resulting bound on immune escape: because compatibility relaxes to $\bar{c} \geq b$ rather than to zero, a promiscuous receptor population retains residual recognition of an arbitrarily mutated pattern.
10. The mutation operator (Section 8), corrected to a diffusion on S^{d-1} so that $\|\vec{p}\| = 1$ is preserved, verified numerically against its own prediction to within 1–6% (Section 8.3), together with the compatibility decay law $\mathbb{E}[c](t) = \bar{c} + (c_0 - \bar{c})e^{-\eta^2 dt}$ and the mutation-adaptation ratio Mu. This supplies the explanation the blueprint/engine distinction of Section 1.4 previously asserted without one: allergic memory persists because $\eta \approx 0$, so the encoding is never invalidated by a moving target.
11. The resolution of the pre-symptomatic signature into a displacement core with a coherence-loss annulus (Section 9.2). Two of its three fields had no spatial form in the framework, and supplying them naively as monotone functions of local activity makes the co-localization criterion an identity rather than a prediction. Taking them as distinct functionals of the state instead gives the signature internal structure, and recasts the three-channel requirement as rejection of independent instrument noise, which is what it is quantitatively good for.

12. A prediction of a pre-symptomatic acoustic signature a co-localized cluster of V , L , and c anomalies driven by early immune execution engine activity $N(t)$ detectable by quantitative ultrasound prior to clinical immune activation, with a quantitative time-window prediction (Equation (76)).

Future research will pursue five directions. (0) Closing the onset gap of Section 11.4: the framework currently models the maintenance and modification of an established sensitization but contains no mechanism by which one is acquired, since a unique equilibrium returns every perturbation to its starting point. This is the largest single omission, and it is placed first because the repair bistability, most naturally through an Allee term requires the ambient exposure inversion to be re-derived on branches rather than on a single equilibrium, and should therefore be undertaken deliberately rather than as an adjustment. (i) Execution of the blockade protocol of Section 7.6, which reduces empirical determination of the coupling matrix χ_{ij} , which Section 7 shows is not bookkeeping: it sets the location of the sensitization threshold and therefore which baselines the model can represent at all. Cytokine-mediated cross-regulation between the Th2, B-cell, mast and regulatory compartments is measurable, and χ is the framework's point of contact with that literature. Section 7.6 narrows what must be measured to two numbers: the admissible topologies are decided by a convergence argument requiring no data, and the magnitudes within a topology reduce to the scalar pair (ψ^+, ψ^-) recoverable from two response ratios. The experimental question is no longer a matrix but a pair of measurements that have, in other contexts, already been made. (ii) Determination of γ_{tol} from serial IgG₄ or regulatory T-cell measurement, which decides whether the durable endpoint constitutes clinical cure or a slowly decaying remission. (iii) Experimental collaboration with QUS researchers to characterize tissue micro-elastic signatures driven by early $N(t)$ activation preceding clinical sensitization. (iv) Extension of the free-energy compatibility correction toward a full computational immunology formalism, integrating calorimetric ΔG_0 measurements to constrain $c(s)$ independently of ρ_s and so resolve the confound of Remark 9.

A note on method. Several claims in earlier versions of this work did not survive numerical evaluation of their own equations, and are corrected rather than quietly removed: the direction of the interval warning in the Counterintuitive Warning subsection, the identifiability of the optimal dose, the independence of ρ_s , the constant- γ_{mem} reading of the GRASS fits, the unprojected form of the mutation operator, and the direction of Equation (7), and the use of the symbol V in Section 9.2 for a quantity anticorrelated with Equation (17), the description of Equation (27) as producing a phase transition when the system it perturbs has a unique equilibrium,

and the claim that Equation (28) is formally identical to the Feynman path integral when its measure is unspecified. In each case the correction is recorded in place rather than the claim removed, and Section 11.5 states what they have in common.

Two items moved the other way. The absence of diffusion in substrate memory was listed as a limitation until numerical testing showed that adding it would reduce the spatial heterogeneity it was supposed to supply; it is now recorded as a justified modelling decision. And the attenuation constant of the cognitive coherence law was introduced as an empirical fit and is now derived (Section 1.1): it is $1/\sigma$, and the four published coherence values agree on that σ to within 0.4%. A framework that takes substrate invariance seriously should expect its constants to turn out to be the same constants, and here one did. An intuition-first framework earns its standing by surviving the attempt to break it, and a correction recorded is worth more than a claim retained.

The framework proposed here is a lens, not a final answer. Its value lies not in replacing existing immunology but in providing a unified mathematical language one that connects the geometry of information propagation to the dynamics of biological substrate memory, and in doing so, reveals structure that neither field could see alone.

A Summary of Key Equations

The following table provides a consolidated master reference for every named equation in the refined CAPM, including both the original formulation and the three biomedical extensions.

Table 10: Master equation reference for the refined CAPM.

Equation	Name
$R(t) = \sum_{\gamma} \gamma_{\text{len}}^{ \gamma } C(\gamma) e^{-S/\sigma} / \mathcal{Z}$	CSP Probability Engine
$S[\gamma, t] = \sum_s E(s) / [1 + \tanh(V L c)] + \sum_j \chi_j$	Action Functional
$c(s) = \max(0, \vec{p}^T M_s \vec{p})$	Compatibility Tensor
$\partial_t P_{\text{substrate}} = \alpha N P_{\text{pattern}} - \gamma_{\text{mem}} (P_{\text{substrate}} - P_0)$	Memory-Dissipation Law
$\sigma(t) = \sigma_0 + (\sigma_{\text{max}} - \sigma_0) \tanh(\beta N / K)$	Congestion Noise
$dN/dt = \max(0, \omega_0 R N (1 - N/K)) - \lambda N$	Population Wave
$R(t) = \mathcal{Z}^{-1} \int \mathcal{D}\gamma e^{-S[\gamma, t]/\sigma(t)}$	Continuous Field Extension
$dN/dt = [\dots] + \mu \nabla^2 N - \lambda N$	Reaction-Diffusion Augmentation
$V(s, t, \rho_s) = \rho_s \cdot \text{JSD}(P_{\text{pattern}} \ P_{\text{substrate}})$	Resolution-Weighted Novelty
$c(s) = \rho_s / (\rho_s + \exp(\Delta G_0 / k_B T))$	Stat. Mech. + Resolution Correction
$\Pi \propto \rho_s \cdot \alpha / (\gamma_{\text{mem}} + \epsilon)$	Polysensitization Risk
$T^* = \ln(\kappa / (\alpha_p \zeta^* + \gamma_p)) / \lambda$	Optimal Exposure Interval
<i>Extensions introduced in this revision</i>	
$S_{\text{hybrid}} = S[\gamma, t] / \gamma + \lambda \gamma $	Hybrid Action
$\sigma(t) \geq \sigma_{\text{min}}$	Noise Floor
$dN/dt = [\dots] + \nu(1 - N/K) - \lambda N$	Ambient Exposure
$\bar{n}^* = [\gamma_{\text{mem}} (P^* - P_0) / (\kappa(1 - P^*))]^{1/q}$	Baseline Engine Density

(Table A.1 continued)

$\nu = \bar{N}^*[\lambda - \omega_0 R^*(1 - \bar{n}^*)]/(1 - \bar{n}^*)$	Ambient Exposure Inversion
$A = \gamma_{\text{mem}}(P^* - P_0)/(1 - P^*)$	Environmental Pressure
$\kappa \bar{n}^q(1 - P_{\text{substrate}})$	Superlinear Resensitization
$dP_{\text{tol}}/dt = \alpha_{\text{tol}}\zeta(1 - P_{\text{tol}}) - \gamma_{\text{tol}}P_{\text{tol}}$	Tolerance Plasticity
$\gamma_{\text{mem}} = \gamma_0 + (\gamma_{\text{max}} - \gamma_0)P_{\text{tol}}$	Plastic Forgetting Rate
$P_{\text{end}} = (\gamma_f P_0 + A)/(\gamma_f + A)$	Durable Endpoint
$c(s) = [1 + \exp(g_{\text{eff}})]^{-1}$	Confounded Compatibility
$\nu(t) = \bar{\nu}[1 + a \sin(2\pi(t/T_{\text{yr}} + \varphi))]_+$	Seasonal Ambient Exposure
<i>Subsystem tensor</i>	
$dn_i/dt = \omega_i R_i n_i(1 - n_i) - \lambda_i n_i + (\dots + \sum_j \chi_{ij} n_j)(1 - n_i)$	Coupled Engine Densities
$dP_i/dt = \alpha_i \zeta(P_0 - P_i) + \gamma_i(P_0 - P_i) + \kappa_i n_i^q(1 - P_i)$	Subsystem Memory
$V_i = \rho_i \text{JSD}(P_{\text{pattern}} \ P_i)$	Per-Subsystem Novelty
$V_\gamma = \rho_{\text{sub}(\gamma)} \text{JSD}(\vec{w} \ \vec{q}^{\text{sub}(\gamma)})$	Pathway–Compartment Binding
$\chi = \psi^+ A^+ - \psi^- A^-$	Coupling Magnitude Scaling
$\Phi^+ = 1 - \bar{n}_{\text{eff}} _{\psi^+=0}/\bar{n}_{\text{eff}}$	Blockade Response
$\Phi^- = \bar{n}_{\text{eff}} _{\psi^-=0}/\bar{n}_{\text{eff}} - 1$	Depletion Response
$X = \langle (1 - \text{JSD}(\vec{w}' \ \vec{b}^\rho)) / (1 - \text{JSD}(\vec{w} \ \vec{b}^\rho)) \rangle$	Cross-Reactivity Index
$D(\mathbf{x}, t) = P_{\text{substrate}}(\mathbf{x}, t) - P_0 $	Substrate Displacement
$L(\mathbf{x}) = \exp(-k_L \ \nabla n\)$	Coherence Field
$c(\mathbf{x}) = \rho_s / (\rho_s + e^{\Delta G_0/k_B T + k_c n(\mathbf{x})})$	Affinity Field
$dP_{\text{tol}}/dt = \alpha_{\text{tol}} \min(E, \bar{E})(1 - P_{\text{tol}}) - \gamma_{\text{tol}}P_{\text{tol}}$	Regularity Induction (rival)
$\gamma_i = \gamma_{0,i} + (\gamma_{\text{max},i} - \gamma_{0,i})P_{\text{reg}}$	Regulatory Modulation

(Table A.1 continued)

$P_i^{\text{rest}} = (\gamma_{0,i}P_0 + \kappa_i n_i^q)/(\gamma_{0,i} + \kappa_i n_i^q)$	Algebraic Rest State
<i>Pattern mutation</i>	
$d\vec{p} = \eta(\mathbb{I} - \vec{p}\vec{p}^T) \circ d\vec{W}$	Mutation on the Sphere
$\mathbb{E}[c](t) = \bar{c} + (c_0 - \bar{c})e^{-\eta^2 dt}$	Compatibility Decay
$\tau_{\text{escape}} = (\eta^2 d)^{-1} \ln[(c_0 - \bar{c})/(c_{\text{thr}} - \bar{c})]$	Escape Time
$\text{Mu} = \eta^2 d/(\gamma_{\text{mem}} + A)$	Mutation–Adaptation Ratio
$\Delta G_{\text{eff}}/k_B T = \ln(\rho_s(1 - c)/c)$	Geometric–Thermodynamic Bridge
$a + b = [1 + \exp(\Delta G_0/k_B T - \ln \rho_s)]^{-1}$	Anchoring Condition
$\Delta G_{\text{eff}}^{\text{max}}/k_B T = \ln(\rho_s(1 - \bar{c})/\bar{c})$	Escape Ceiling
$\Xi = g_N \bar{\zeta}/\nu$	Dose to Ambient Ratio
$b \gamma_{\text{len}} c_{\text{max}} e^{-E_{\text{min}}/\sigma} < 1$	Path-Sum Convergence
$\Theta = (\mathcal{Z}/\mathcal{Z}_{\text{ideal}})^{1/\bar{n}}$	Propagation Efficiency
$\vec{b}^\rho = \mathcal{N}[\vec{w} * G_{w_0/\rho_s}]$	Resolvable Profile
$q_i = P b_i^\rho + (1 - P)/B$	Feature-Space Novelty
$C_{\text{total}} = \exp(-\sum_j \chi_j/\sigma)$	Coherence Across Crossings
$\gamma_{\text{mem}} = \omega_r(1 - C_{\text{total}})$	Coherence to Forgetting Rate