

Copy And Paste Model

Biomedical Extension

Part 4 of 5

Claudio Serdar-Petkovski

Independent Researcher, Systems Geometry Group

Version 2.0, 10 September 2026

split from the single-document Version 1.4 of 10 August 2026

Abstract. This document applies the CAPM engine to allergic sensitization. The allergen is treated as a static blueprint and the immune network as the execution engine that copies, amplifies and maintains the response; allergy is then a false compatibility written into substrate memory. We state at the outset what the model cannot do: baseline severity is an input, never an outcome, so the framework describes the maintenance and modification of an established sensitization and not its acquisition. What follows is derived from that starting point: an inversion replacing the unobservable ambient encounter rate with the observable untreated baseline, which yields a reachability constraint capable of rejecting fitted parameter sets on internal-consistency grounds; a tolerance boundary through which treatment moves a patient rather than a wall behind which they sit; a polysensitization risk testable against existing repertoire-sequencing data; tolerance plasticity and a closed-form durable endpoint, without which no protocol can move the post-treatment equilibrium at all; a reinterpretation of the GRASS trial in which the arms differ in route-dependent induction efficiency rather than in patient properties; a sensitization threshold produced by the coupling matrix; and a pre-symptomatic acoustic signature resolved into a displacement core with a coherence-loss annulus. The document closes with a falsification ledger giving the current status of every prediction, including one already contradicted.

Developed in collaborative dialogue. Corrections to earlier versions are recorded in the Technical Appendix, not silently removed.

The CAPM Map

One framework, five documents. Read the one addressed to you; the others are referred to by name, never assumed.

Document	Its job	What you are asked to evaluate
P Philosophy	States the thesis (<i>the substrate changes; the equations do not</i>), the method, and the boundaries of the claim. Prose only.	Is the thesis coherent? Is the intuition-first method honest? Is the narrowed claim that survived testing the right one to make?
C Core Mathematical Apparatus	Defines the engine, the action, substrate memory and the population wave; reports what is known about their well-posedness; tests substrate invariance.	The normalisation pathology and its repair; the measure problem and the convergence bound; uniqueness of equilibrium; ρ_s independence; the invariance verdict.
G Cognitive Extension	Applies the engine to thought: geometry translated across substrates, and the detection of the substrate a pattern natively occupies.	Whether coherence across crossings predicts what is remembered, and whether substrate detection predicts anything its bridge costs did not contain.
► B Biomedical Extension	Applies the engine to allergic sensitization; derives predictions for immunotherapy, polysensitization and pre-symptomatic detection.	A falsification ledger: every prediction with its status, whether contradicted, pre-registered, testable now, or open.
A Technical Appendix and Bookkeeping	Numerical methods, verification runs, the corrections log, coherence bookkeeping, master equation table, V1.4 concordance.	Reproducibility. Every “earlier versions” statement lives here and nowhere else.

Anchor symbols. Shared across the documents; defined where indicated.

Symbol	Meaning	Defined in
$R(t)$	Replication probability: Boltzmann-weighted share of path weight reaching the nominated destination	C
$S[\gamma, t]$	Informational action along path γ ; $\sigma(t)$ is its temperature, in the position $\hbar$ occupies but with a real exponent	C
χ_j	Bridge-operator cost of the j -th substrate crossing; χ_{ij} its matrix form between subsystems	C
$c(s)$	Node compatibility: geometric $\vec{p}^\top M_s \vec{p}$ and thermodynamic $\rho_s / (\rho_s + e^{\Delta G_0 / k_B T})$, one quantity in two coordinate systems	C
$V(s, t, \rho_s)$	Novelty: $\rho_s \cdot \text{JSD}(P_{\text{pattern}} \ P_{\text{substrate}})$ over feature bins	C
$P_{\text{substrate}}, P_0, N, \bar{n}$	Substrate memory and its pristine baseline; execution-engine population and $\bar{n} = N/K$	C
$\alpha, \gamma_{\text{mem}}$	Learning rate and <i>forgetting</i> rate of the substrate; low γ_{mem} means long memory	C

Symbol	Meaning	Defined in
ρ_s	Structural resolution: how finely the substrate discriminates configurations	C, grounded in B
Θ, Mu	Propagation efficiency $(Z/Z_{\text{ideal}})^{1/\bar{n}}$; mutation–adaptation ratio $\eta^2 d/\gamma_{\text{relax}}$	C
$C_{\text{total}}, \omega_r$	Coherence across a set of substrate crossings; rate of recall attempts	G
κ, P_{tol}	Resensitization coefficient; regulatory substrate that makes γ_{mem} plastic	B
P^*, ν	Observable untreated baseline; ambient encounter rate inverted from it	B

Conventions. γ_{mem} is a forgetting rate throughout. Differences are oriented as $(P_{\text{substrate}} - P_0)$. The population enters memory laws as $\bar{n} = N/K$. Equation, section and theorem numbers carry the document letter, so Core Eq. (C.1) is the engine wherever it is cited and the letter alone tells you which document to open. Gray tags [v1.4(*n*)] give the number the same equation carried in the single-document Version 1.4.

Contents

B.1	Blueprint, Execution Engine, and the Parameter Map	5
B.1.1	Biomedical Lens Extension: Blueprint and Execution Engine	5
B.1.2	The Core Mapping	7
B.2	What the Model Cannot Do	8
B.2.1	Sensitization Cannot Occur	8
B.3	The Observable Baseline and the Ambient Exposure Inversion	9
B.4	Phase Space and the Tolerance Boundary	11
B.4.1	Why α and γ_{mem} Are Pathological in Allergy	11
B.4.2	The Phase Space and the Tolerance Boundary	12
B.5	Structural Resolution and the Three-Dimensional Phase Space	13
B.6	The Resonance Window	18
B.6.1	Derivation	18
B.6.2	The Counterintuitive CAPM Warning	19
B.6.3	The Optimal Waveform	19
B.7	Tolerance Plasticity and the Durable Endpoint	20
B.8	The GRASS Trial, Reinterpreted	23
B.8.1	Empirical Reassessment: The GRASS Trial	23
B.9	Compartments: Regulation, Rest State, Blockade, Ignition	24
B.9.1	The Regulatory Subsystem Absorbs Tolerance Plasticity	24
B.9.2	The Rest State Is Algebraic	24
B.9.3	What Constrains χ	24
B.9.4	An Unplanned Consequence: The Ignition Threshold	25
B.10	Why Allergic Memory Persists: The $\eta \rightarrow 0$ Limit	26
B.10.1	Why Allergy Is the $\eta \rightarrow 0$ Limit	26
B.11	Pre-Symptomatic Detection via Quantitative Ultrasound	26
B.11.1	The Pre-Symptomatic Window: Detecting the Execution Engine	27
B.11.2	What Quantitative Ultrasound Would Detect	27
B.11.3	The Pre-Symptomatic Time Window	30
B.12	Boundaries	30
B.12.1	Mutations and Evolving Pattern Identity	30
B.12.2	Parallel Immune Subsystems	31
B.12.3	Assumptions Adopted Without Derivation	32

B.12.4	Parameter Identifiability	32
B.13	Contributions and Falsification Ledger	33
B.13.1	The Ledger	35

Imported apparatus, defined in the Core document and restated here so this document is self-contained

This document uses the following Core objects unchanged. They are restated here so that the biomedical argument can be read without the Core document open; the numbered originals, with their derivations, are where the tags point.

The path-sum engine and the informational action:

$$R(t) = \frac{\sum_{\gamma \in \mathcal{P}(s_0 \rightarrow s_f; \Gamma_t)} \gamma_{\text{len}}^{|\gamma|} C(\gamma) e^{-S[\gamma, t]/\sigma(t)}}{\sum_{s_{\text{alt}} \in V_t} \sum_{\gamma' \in \mathcal{P}(s_0 \rightarrow s_{\text{alt}}; \Gamma_t)} \gamma_{\text{len}}^{|\gamma'|} C(\gamma') e^{-S[\gamma', t]/\sigma(t)}} \quad (\text{Core Eq. (C.1)})$$

$$S[\gamma, t] = \sum_{s \in \gamma} \frac{E(s)}{1 + \tanh(V(s, t) \cdot L(s) \cdot c(s))} + \sum_{j \in B_t} \chi_j \quad (\text{Core Eq. (C.2)})$$

Novelty, weighted by structural resolution:

$$V(s, t, \rho_s) = \rho_s \cdot \text{JSD}(P_{\text{pattern}} \parallel P_{\text{substrate}}(s, t)) \quad (\text{Core Eq. (C.4)})$$

The Memory-Dissipation Law and the driven population wave:

$$\frac{\partial P_{\text{substrate}}(s, t)}{\partial t} = \alpha \cdot N(t) \cdot P_{\text{pattern}}(s) - \gamma_{\text{mem}}(P_{\text{substrate}}(s, t) - P_0) \quad (\text{Core Eq. (C.11)})$$

$$\frac{dN}{dt} = \max\left(0, \omega_0 R(t) N \left(1 - \frac{N}{K}\right)\right) + \mu \nabla^2 N + \nu \left(1 - \frac{N}{K}\right) - \lambda N \quad (\text{Eq. (B.5)})$$

Everything in this document follows from these five equations together with the parameter map of Table B.2 and the biological role separation of §B.1.

B.1 Blueprint, Execution Engine, and the Parameter Map

B.1.1 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:

$$P_{\text{pattern}} \xrightarrow{\chi_{\text{immune}}} P_{\text{substrate}} \quad [\text{v1.4 (8)}] \quad (\text{B.1})$$

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:

$$P_{\text{pattern}} = P(s, d, r, c, t, H) \quad [\text{v1.4 (9)}] \quad (\text{B.2})$$

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:

$$\partial_t P_{\text{substrate}} = \alpha N(t) P_{\text{pattern}} - \gamma_{\text{mem}} (P_{\text{substrate}} - P_0). \quad [\text{v1.4 (10)}] \quad (\text{B.3})$$

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 (B.3) 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.

B.1.2 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 B.1 (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:

$$R(t) \big|_{s_f=s_{\text{attack}}} \approx 1 \tag{B.4}$$

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 B.2: Parameter correspondences with explicit biological role separation. The allergen functions as a static blueprint; the immune network is the active CAPM propagation engine.

B.2 What the Model Cannot Do

B.2.1 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 Core Remark C.2 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 Core Eq. (C.14) do what Core §C.6 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 Core §C.6 would otherwise be false, which is the failure mode Philosophy §P.6 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 B.3 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.

B.3 The Observable Baseline and the Ambient Exposure Inversion

Neither Core Eq. (C.11) nor its extended form Equation (B.19) 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:

$$\frac{dN}{dt} = \max\left(0, \omega_0 R(t)N\left(1 - \frac{N}{K}\right)\right) + \mu \nabla^2 N + \nu\left(1 - \frac{N}{K}\right) - \lambda N \quad [\text{v1.4 (40)}] \quad (\text{B.5})$$

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 B.1 (Ambient Exposure Inversion). *Let P^* denote the untreated equilibrium of the substrate memory law (Equation (B.19)) with $\zeta = 0$. Setting $\partial_t P_{\text{substrate}} = 0$ gives the engine density the baseline requires,*

$$\bar{n}^* = \left[\frac{\gamma_{\text{mem}}(P^* - P_0)}{\kappa(1 - P^*)} \right]^{1/q}, \quad [\text{v1.4 (41)}] \quad (\text{B.6})$$

and setting $dN/dt = 0$ in Equation (B.5) at $\bar{N}^* = K\bar{n}^*$ gives

$$\nu = \frac{\bar{N}^*[\lambda - \omega_0 R^*(1 - \bar{n}^*)]}{1 - \bar{n}^*}. \quad [v1.4 (42)] \quad (\text{B.7})$$

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 (B.6) states that $\kappa\bar{n}^{*q}$ equals

$$A = \frac{\gamma_{\text{mem}}(P^* - P_0)}{1 - P^*} \quad [v1.4 (43)] \quad (\text{B.8})$$

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 (B.8) into Equation (B.19) 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 B.8 is developed there.

Not every baseline is dynamically reachable. Equation (B.6) 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 B.8 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 (B.6) 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 (B.5) and its subsystem form Core Eq. (C.27) 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:

$$\Xi = \frac{g_N \bar{\zeta}}{\nu}, \quad \bar{\zeta} = \frac{\zeta \tau}{T} \quad [\text{v1.4 (44)}] \quad (\text{B.9})$$

Ξ is dimensionless, invariant under the schedule reparameterisations of Proposition B.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 (B.5) 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

$$\nu(t) = \bar{\nu} \left[1 + a \sin \left(2\pi \left(\frac{t}{T_{\text{yr}}} + \varphi \right) \right) \right]_+ \quad [\text{v1.4 (45)}] \quad (\text{B.10})$$

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

Under Equation (B.10) the untreated state is a limit cycle rather than a fixed point, and the closed form of Proposition 1 no longer applies: because Equation (B.19) 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.

B.4 Phase Space and the Tolerance Boundary

B.4.1 Why α and γ_{mem} Are Pathological in Allergy

In allergic individuals, Core Eq. (C.11) 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.

B.4.2 The Phase Space and the Tolerance Boundary

Definition B.2 (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:

$$(\alpha\zeta + \gamma_{\text{mem}})(P_{\text{substrate}} - P_0) = \kappa \cdot \bar{n}(t) \cdot (1 - P_{\text{substrate}}) \quad [\text{v1.4 (46)}] \quad (\text{B.11})$$

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 (B.20), (B.15) and (B.23) follow the same convention.

γ_{mem} here is the current value, not the constitutive one. Equation (B.11) is a balance among instantaneous rates, so every quantity in it is evaluated at the state the substrate currently occupies. Once Section B.7 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 B.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 (B.11) 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.

B.5 Structural Resolution and the Three-Dimensional Phase Space

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 Core Eq. (C.5) for a substrate trained to saturation on $\vec{w}$, and define

$$X = \left\langle \frac{1 - \text{JSD}(\vec{w}' \| \vec{b}^\rho)}{1 - \text{JSD}(\vec{w} \| \vec{b}^\rho)} \right\rangle_{\vec{w}'} \quad [\text{v1.4 (20)}] \quad (\text{B.12})$$

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. Core Remark C.5 requires w_0 to lie in an intermediate range, and without Equation (B.12) that requirement would amount to choosing a parameter so that a conclusion follows from it.

Definition B.3 (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:

$$\rho_s = f(D_{\text{TCR}}, H_{\text{MHC}}, F_{\text{DC}}) \quad [\text{v1.4 (21)}] \quad (\text{B.13})$$

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:

$$\rho_s \approx \frac{D_{\text{TCR}}}{D_{\text{TCR}}^{\text{ref}}} \quad [\text{v1.4 (22)}] \quad (\text{B.14})$$

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.

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 Core Eq. (C.4): as ρ_s increases, the effective novelty of any allergen encounter increases, which *lowers* the action $S[\gamma, t]$ along attack pathways and so raises their weight in the path sum, pushing the tolerance boundary outward. Formally, the resonance condition (Equation (B.20)) acquires a resolution correction:

$$(\alpha\bar{\zeta} + \gamma_{\text{mem}})(\bar{P}_{\text{substrate}} - P_0) > \kappa\bar{n}(1 - \bar{P}_{\text{substrate}}) \cdot \rho_s \quad [\text{v1.4 (47)}] \quad (\text{B.15})$$

Equation (B.15) is Equation (B.20) 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 param-

eter capable of explaining why two patients with identical α and γ_{mem} present with entirely different allergen profiles.

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.

Which channel carries the argument. Two channels connect ρ_s to the propagation of an attack response, and only one of them can support this prediction.

The compatibility channel cannot. Core Eq. (C.22) does make $c(s)$ increase monotonically with ρ_s at fixed ΔG_0 , and earlier versions of this work rested the polysensitization argument there. But Core Remark C.5 establishes that $c(s)$ depends on ρ_s and ΔG_0 only through the single combination $g_{\text{eff}} = \Delta G_0/k_B T - \ln \rho_s$. A high-resolution patient is therefore indistinguishable, in this channel, from a patient whose receptors simply bind more tightly. Any risk attributed to ρ_s through $c(s)$ alone is not attributable to resolution at all.

The novelty channel can. Under the epitope-feature construction of Core §C.3, ρ_s sets the width of the kernel through which the substrate registers a profile, and the same remark establishes that this dependence is not absorbed into g_{eff} . The polysensitization argument is therefore stated here through $V(s, t, \rho_s)$, and the $c(s)$ route is retained only as a consistency check.

Direction of the effect. Raising ρ_s raises V , and raising V *lowers* the action: the novelty term enters Core Eq. (C.2) through $E(s)/[1 + \tanh(VLc)]$, whose denominator grows with V . A lower action means a larger path weight and a larger $R(t)$, so a high-resolution substrate propagates the attack response more readily, across a wider range of environmental molecules. Earlier versions of this work asserted the opposite sign, stating that rising ρ_s increases the action; the conclusion was correct and the stated route was inverted. The corrected statement is recorded in Appendix §A.2.

Resolution is not free. Written as an unbounded parameter, ρ_s makes polysensitization risk grow without limit. It cannot, and the reason is contained in the framework’s own apparatus rather than imported from outside.

Core Eq. (C.22) does not distinguish foreign patterns from self patterns. A substrate that resolves finely enough to find every environmental molecule compatible will, by the same equation, find self compatible: the $\rho_s \rightarrow \infty$ limit noted in Core Remark C.4 is not only a statement about allergens. Vertebrate immune systems prevent this by deleting, during thymic negative selection, precisely those clones that supply the resolution. Resolution is therefore self-limiting, and the cost is paid in the same currency as the benefit.

Let $\hat{\rho}$ denote the nominal presentation breadth, set by the number of distinct MHC molecules an individual expresses, and let $g_{\text{self}} = \Delta G_{\text{self}}/k_B T$ be the baseline free energy of a representative

self peptide. Two consequences follow from Core Eq. (C.22) applied to self:

1. the number of self epitopes presented grows with $\hat{\rho}$, since each additional MHC molecule presents its own set;
2. the compatibility of a given clone with any one of them is $c_{\text{self}}(\hat{\rho}) = \hat{\rho}/(\hat{\rho} + e^{g_{\text{self}}})$, which also grows with $\hat{\rho}$.

Taking the deletion events as independent across presented self epitopes, the surviving fraction of the pre-selection repertoire is $\exp(-\beta \hat{\rho} c_{\text{self}})$ for a per-epitope deletion pressure $\beta > 0$, and the resolution actually realised is

$$\rho_s(\hat{\rho}) = \hat{\rho} \exp\left(-\beta \frac{\hat{\rho}^2}{\hat{\rho} + e^{g_{\text{self}}}}\right). \quad (\text{B.16})$$

Proposition B.2 (Polysensitization Risk, bounded by self-tolerance). *Let a patient occupy position $(\alpha, \gamma_{\text{mem}}, \hat{\rho})$, with realised resolution given by Equation (B.16). Then polysensitization risk is*

$$\Pi \propto \rho_s(\hat{\rho}) \cdot \frac{\alpha}{\gamma_{\text{mem}} + \epsilon}, \quad (\text{B.17})$$

with $\epsilon > 0$ preventing divergence at $\gamma_{\text{mem}} = 0$. Π is not monotone in presentation breadth. It rises, attains an interior maximum at the $\hat{\rho}^*$ solving

$$\beta \hat{\rho}^2 (\hat{\rho} + 2e^{g_{\text{self}}}) = (\hat{\rho} + e^{g_{\text{self}}})^2, \quad (\text{B.18})$$

and falls thereafter, approaching zero as deletion outruns presentation. In the regime $\hat{\rho} \gg e^{g_{\text{self}}}$ the optimum is $\hat{\rho}^* \simeq 1/\beta$: the optimal presentation breadth is set by the reciprocal of the stringency of negative selection.

Polysensitization is therefore not a separate pathological mechanism, and not a consequence of maximal resolution either. It is what happens to a substrate sitting near the resolution optimum while also learning fast and forgetting slowly.

Relation to existing theory. An interior optimum in individual MHC diversity is not a new proposal. Nowak et al. (1992) first argued that the benefit of presenting more antigens is offset by the cost of deleting more T-cell clones, and Woelfing et al. (2009) estimated the resulting optimum at 3 to 25 alleles; Borghans et al. (2003) dispute whether thymic depletion is the binding constraint, attributing the limit instead to self-reactivity. The contribution here is not the optimum but its derivation: Equation (B.16) follows from applying the framework's own compatibility function to self peptides, with no additional postulate beyond the independence of deletion events.

The contact is quantitative as well as structural. Solving Equation (B.18) for the β that reproduces an optimum of 3, 10 and 25 alleles gives per-epitope deletion pressures of roughly 0.7, 0.15 and 0.05 respectively at $g_{\text{self}} \in [2, 3]$, a physiologically unremarkable range. The

framework is not independently constrained to that range, so this is a consistency check rather than a prediction.

The empirical picture is unsettled in a way that matters for Definition B.3. Migalska et al. (2019) report that higher individual MHC class I diversity is associated with *smaller* TCR repertoires, the first direct test of the depletion hypothesis; a later study finds $CD4^+$ TCR richness *positively* correlated with MHC class II diversity (Migalska et al., 2023). Definition B.3 composes ρ_s from D_{TCR} and H_{MHC} as though both raise resolution together. Under the depletion result they oppose each other, at least for class I. Equation (B.16) is the minimal repair: $\hat{\rho}$ carries the MHC contribution, ρ_s is the realised repertoire measured after selection, and the two are related rather than summed. Whether the same relation holds for class II is open, and §B.12.3 records it as such.

What would falsify this. Equation (B.16) converts a prediction that required new data into one testable by re-analysis of data already published. Cohorts exist with HLA typed from whole-genome sequencing and IgE sensitization scored across allergen panels; Ledoux et al. (2023) report one of 797 individuals, in which polysensitization was modelled directly. That analysis asked which alleles a participant carries. The prediction here concerns how many: polysensitization prevalence, plotted against the count of distinct HLA class II alleles, should rise to an interior peak and fall, rather than increasing monotonically or showing no trend.

A monotone increase would refute Equation (B.16) while leaving the original unbounded form standing. A flat or absent relationship would refute both, and would indicate that ρ_s is carried by D_{TCR} alone with no presentation contribution, requiring Definition B.3 to be narrowed accordingly.

Status of the supporting evidence. No study has yet compared repertoire diversity between polysensitized and monosensitized patients at matched α and γ_{mem} , which remains the direct test. Two indirect results are consistent with the sign of the effect at resolutions below the optimum. Lee et al. (2023) find, in 67 children with B-ALL and a validation cohort of 113, that a more diverse and more rapidly changing TCR repertoire predicts clinical hypersensitivity to asparaginase. Chen et al. (2024) report higher TRAV diversity and CDR3 clonality in childhood atopic dermatitis, the more broadly sensitized phenotype, than in asthma or controls. Neither study measures polysensitization, and the first concerns drug hypersensitivity under chemotherapy; they are recorded here as directionally consistent, not as support.

A caution on Equation (B.16). The exponential survival form and the linear growth of presented self epitopes with $\hat{\rho}$ are modelling choices, not derivations. They are the same choices made in the classical treatment, and they are what produce an interior optimum; a different deletion law would move the optimum and could remove it. The conditions under which Equation (B.16) is claimed are therefore: deletion events independent across presented self epitopes, epitope count linear in $\hat{\rho}$, and a single representative g_{self} . This is exactly the class of unstated-condition failure that Philosophy §P.6 describes, and the conditions are stated here rather than

left for a later correction.

Remark B.1. 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.

B.6 The Resonance Window

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.

B.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):

$$\frac{\partial P_{\text{substrate}}}{\partial t} = \underbrace{\alpha \cdot \zeta(t) \cdot (P_0 - P_{\text{substrate}})}_{\text{tolerance rewriting}} + \underbrace{\gamma_{\text{mem}}(P_0 - P_{\text{substrate}})}_{\text{natural forgetting}} + \underbrace{\kappa \cdot N(t) \cdot (1 - P_{\text{substrate}})}_{\text{resensitization}} \quad [\text{v1.4 (50)}] \quad (\text{B.19})$$

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 B.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 $[t_k, t_k + T]$:*

$$(\alpha\bar{\zeta} + \gamma_{\text{mem}})(\bar{P}_{\text{substrate}} - P_0) > \kappa\bar{n}(1 - \bar{P}_{\text{substrate}}) \quad [\text{v1.4 (51)}] \quad (\text{B.20})$$

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 (B.19) 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. $\square$

B.6.2 The Counterintuitive CAPM Warning

Equation (B.20) 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 (B.20) together, and the protocol converges on doing nothing rather than on doing harm. Numerical integration of Equation (B.19) 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 (B.19) 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:

$$\left. \frac{\partial P_{\text{substrate}}}{\partial t} \right|_{\text{resens}} = \kappa \cdot \bar{n}(t)^q \cdot (1 - P_{\text{substrate}}), \quad q \geq 1 \quad [\text{v1.4 (52)}] \quad (\text{B.21})$$

At $q = 1$ this is Equation (B.19) 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.

B.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:

$$T^* = \frac{\ln(\kappa/(\alpha_p \zeta^* + \gamma_p))}{\lambda} \quad [\text{v1.4 (53)}] \quad (\text{B.22})$$

and the optimal dose:

$$\zeta^* = \frac{\kappa \bar{n}(1 - \bar{P}_{\text{substrate}})/(\bar{P}_{\text{substrate}} - P_0) - \gamma_p}{\alpha_p} \quad [\text{v1.4 (54)}] \quad (\text{B.23})$$

The difference is oriented as $(\bar{P}_{\text{substrate}} - P_0)$ in accordance with the convention fixed at Equation (B.11); 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 B.2 (The optimal dose is not identifiable as written). Equation (B.23) requires α_p as an input. But α and ζ enter Equation (B.19) 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.

B.7 Tolerance Plasticity and the Durable Endpoint

Core Eq. (C.11) and (B.19) treat γ_{mem} as a constant. That choice has a consequence which was not visible until the model was integrated numerically, and which invalidates the clinical prescription earlier versions of this work derived in the empirical-validation section.

Proposition B.4 (Treatment cannot move a fixed- γ_{mem} endpoint). *With γ_{mem} constant and ν determined by Equation (B.7), the equilibrium of Equation (B.19) 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:

$$\frac{dP_{\text{tol}}}{dt} = \alpha_{\text{tol}} \cdot \zeta(t) \cdot (1 - P_{\text{tol}}) - \gamma_{\text{tol}} \cdot P_{\text{tol}} \quad [\text{v1.4 (55)}] \quad (\text{B.24})$$

$$\gamma_{\text{mem}}(t) = \gamma_0 + (\gamma_{\text{max}} - \gamma_0) \cdot P_{\text{tol}} \quad [\text{v1.4 (56)}] \quad (\text{B.25})$$

Here γ_0 is the patient's constitutive forgetting rate the quantity ν was inverted from in Equation (B.7), 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 (B.24) is Core Eq. (C.11) with P_{tol} in place of $P_{\text{substrate}}$; no new dynamical form is introduced.

Proposition B.5 (Durable Endpoint). *With ν fixed by Equation (B.7) and hence $\kappa \bar{n}^q = A$ constant, the equilibrium of Equation (B.19) at a raised forgetting rate γ_f is*

$$P_{\text{end}} = \frac{\gamma_f P_0 + A}{\gamma_f + A} \quad [\text{v1.4 (57)}] \quad (\text{B.26})$$

with A given by Equation (B.8). 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 (B.24)–(B.26) extend the framework rather than replacing it.

Proposition B.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 B.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 (B.24) as written and requires an induction term with its own nonlinearity.

Proposition B.7 (Phase Invariance of the Endpoint). *Under Equations (B.24) and (B.10), 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 (B.24). 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 B.6, because pre-seasonal immunotherapy is established clinical practice and the comparison has been made.

The Test Selects Between Two Formulations

Proposition B.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} ,

$$\frac{dP_{\text{tol}}}{dt} = \alpha_{\text{tol}} \min(E, \bar{E})(1 - P_{\text{tol}}) - \gamma_{\text{tol}} P_{\text{tol}} \quad [\text{v1.4 (58)}] \quad (\text{B.27})$$

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 (B.24) 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 (B.24) the content of Proposition B.7 and 6.2×10^{-3} under Equation (B.27), 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 (B.27); its absence selects Equation (B.24). Either outcome removes an assumption that Section B.12.3 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 (B.26) 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.

B.8 The GRASS Trial, Reinterpreted

B.8.1 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 (B.19) 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 (B.6) 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 (B.7), $\kappa \bar{n}^q$ 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 (B.24)–(B.25) 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 (B.26) 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 B.3. 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.

B.9 Compartments: Regulation, Rest State, Blockade, Ignition

B.9.1 The Regulatory Subsystem Absorbs Tolerance Plasticity

Designate one subsystem regulatory. Its memory obeys Equation (B.24) rather than Core Eq. (C.28), and it sets the forgetting rate of every effector:

$$\gamma_i = \gamma_{0,i} + (\gamma_{\text{max},i} - \gamma_{0,i}) P_{\text{reg}} \quad [\text{v1.4 (62)}] \quad (\text{B.28})$$

Equations (B.24)–(B.25) 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.

B.9.2 The Rest State Is Algebraic

With $\zeta = 0$ the regulatory memory decays to zero, each effector runs at its constitutive $\gamma_{0,i}$, and Core Eq. (C.28) solves in closed form for a given engine density:

$$P_i^{\text{rest}} = \frac{\gamma_{0,i} P_0 + \kappa_i n_i^q}{\gamma_{0,i} + \kappa_i n_i^q} \quad [\text{v1.4 (63)}] \quad (\text{B.29})$$

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 (B.29) and will report equilibria it has not reached.

B.9.3 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 Fc ϵ RI, 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,

$$\chi = \psi^+ A^+ - \psi^- A^- \quad [\text{v1.4 (65)}] \quad (\text{B.30})$$

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

$$\Phi^+ = 1 - \bar{n}_{\text{eff}}|_{\psi^+=0} / \bar{n}_{\text{eff}} \quad (\text{upstream blockade}) \quad [\text{v1.4 (66)}] \quad (\text{B.31})$$

$$\Phi^- = \bar{n}_{\text{eff}}|_{\psi^-=0} / \bar{n}_{\text{eff}} - 1 \quad (\text{regulatory depletion}) \quad [\text{v1.4 (67)}] \quad (\text{B.32})$$

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 B.1); the dose-to-ambient ratio Ξ replaces the activation gain g_N (§B.3); the noise scale σ replaces the coherence attenuation constant (Cognitive §C.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 Core Proposition C.1 shows that the pattern is constrained from a different direction entirely.

B.9.4 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 B.9, 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.

B.10 Why Allergic Memory Persists: The $\eta \rightarrow 0$ Limit

B.10.1 Why Allergy Is the $\eta \rightarrow 0$ Limit

This resolves something the framework asserted but did not explain. Section B.1 separates the allergen, a static blueprint, from the immune network that propagates the response. Core Eq. (C.34) 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 the series. The Biomedical document is the $\text{Mu} \rightarrow 0$ theory. Extending them to oncology where Section B.12.1 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.

B.11 Pre-Symptomatic Detection via Quantitative Ultrasound

B.11.1 The Pre-Symptomatic Window: Detecting the Execution Engine

The stochastic seeding term (Core Eq. (C.14)) 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:

$$\lim_{P_{\text{substrate}} \rightarrow \emptyset} C(\gamma_{\text{anomaly}}) \rightarrow 0 \implies R(t)_{\text{equilibrium}} \rightarrow 0 \quad (\text{B.33})$$

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.

B.11.2 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 Core Eq. (C.4)

Earlier versions named the first signal V and described it as a novelty elevation. That collides with Core Eq. (C.4), 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 Core Eq. (C.4).

Numerically the point is unambiguous. At a focally activated site the Core Eq. (C.4) 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:

$$D(\mathbf{x}, t) = |P_{\text{substrate}}(\mathbf{x}, t) - P_0| \quad [\text{v1.4 (73)}] \quad (\text{B.34})$$

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:

$$L(\mathbf{x}) = \exp(-k_L \|\nabla n(\mathbf{x})\|) \quad \text{architecture is disrupted at interfaces} \quad [\text{v1.4 (74)}] \quad (\text{B.35})$$

$$c(\mathbf{x}) = \frac{\rho_s}{\rho_s + \exp(\Delta G_0/k_B T + k_c n(\mathbf{x}))} \quad \text{affinity shifts with local activity} \quad [\text{v1.4 (75)}] \quad (\text{B.36})$$

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 (B.35) 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 B.3: 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 B.4. This prediction is grounded in established QUS capabilities. Shear-wave elastogra-

phy 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.

B.11.3 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} :

$$\Delta t_{\text{pre}} \approx \frac{1}{\omega_0 R(t)} \ln \left(\frac{N_{\text{clinical}}}{N_{\text{QUS}}} \right) \quad [\text{v1.4 (76)}] \quad (\text{B.37})$$

This follows directly from the early exponential growth phase of Core Eq. (C.13) 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 of the pre-symptomatic cluster grows at the front velocity of Core Eq. (C.19), which applies only where ambient encounter vanishes and the empty state is unstable, giving a time-dependent spatial signature that would distinguish genuine early immune activation from measurement noise or tissue heterogeneity artifacts. genuine early immune activation from measurement noise or tissue heterogeneity artifacts.

B.12 Boundaries

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.

B.12.1 Mutations and Evolving Pattern Identity

Addressed in Core §C.11; 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,

$$\frac{d\vec{p}}{dt} = \eta \cdot \xi(t), \quad [\text{v1.4 (78)}] \quad (\text{B.38})$$

cannot be used as written, because it does not preserve the constraint $\|\vec{p}\|_2 = 1$ imposed by Definition 2. Core §C.11 supplies the projected form (Core Eq. (C.31)), 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. Core §C.11 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 (Core Eq. (C.1)) 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.

B.12.2 Parallel Immune Subsystems

Addressed in Core §C.10; 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 . Core §C.10 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 B.7 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 B.9.4 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 B.9.3 reduces the magnitudes to two scalars recoverable from cytokine-blockade and regulatory-depletion responses, and Core Proposition C.1 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 B.9 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.

B.12.3 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: Cognitive §G.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 (B.24); 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 B.7 now states the rival formulation (Equation (B.27)), 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 Core Eq. (C.27) 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 Core Eq. (C.18). 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.

B.12.4 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 (B.19) only as a product (Remark B.2).
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 (Core Remark C.5).
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 (§B.3).

Table B.4: 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.

B.13 Contributions and Falsification Ledger

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 (B.6)–(B.7)), 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 (B.24)–(B.25)) 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 B.6, that outcome depends on cycle-averaged dose rate alone and not on schedule, and Proposition B.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 resolution of the pre-symptomatic signature into a displacement core with a coherence-loss annulus (Section B.11). 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.
8. A prediction of a pre-symptomatic acoustic signature a
9. A prediction of a pre-symptomatic acoustic signature a co-localized cluster of D , L , and c anomalies driven by early immune execution quantitative ultrasound prior to clinical immune activation, with a quantitative time-window prediction (Equation (B.37)).

Future research will pursue five directions. (0) Closing the onset gap of Section B.2: 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 B.9.3, which reduces empirical determination of the coupling matrix χ_{ij} , which Core §C.10 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 B.9 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 Core Remark C.5.

B.13.1 The Ledger

The contributions above are claims. This table is what an immunologist is actually being asked to attack: each prediction, what would refute it, and where it currently stands. A prediction already contradicted is listed as such rather than removed, on the principle stated in Philosophy §P.6.

Table B.5: Falsification ledger. “Status” is as of this revision.

Prediction	What would refute it	Status	Where
Schedule invariance: outcome depends on cycle-averaged dose rate alone	Any trial varying dose and interval together at matched cumulative rate that finds a schedule effect	Contradicted. Rask et al. (2010) report a schedule effect at matched cumulative dose. The linear induction law (B.24) is refuted; an induction term carrying its own nonlinearity is required	Prop. B.6
Phase invariance: pre-seasonal and co-seasonal immunotherapy differ in symptom burden during treatment and not at all in durable outcome	A demonstrated durable advantage to pre-seasonal dosing, which would select the rival induction law (B.27) instead	Pre-registered, not yet testable. The predicted separation is $\sim 10^{-3}$ – 10^{-2} in model endpoint units; no units bridge to a clinical endpoint exists, and no published trial reports the required contrast	Prop. B.7

Prediction	What would refute it	Status	Where
Polysensitization risk, (B.17), non-monotone in presentation breadth with an interior optimum	Polysensitization prevalence increasing monotonically with HLA class II allele count (refutes the bounded form only), or showing no relation to it (refutes both forms, and narrows Definition B.3 to D_{TCR} alone)	Testable by re-analysis. Cohorts with HLA typed from whole-genome sequencing and polysensitization already scored exist (Ledoux et al., 2023); they tested allele identity, not allele count. The direct repertoire comparison at matched $\alpha, \gamma_{\text{mem}}$ remains unrun	Prop. B.2
Reachability: a patient with high γ_{mem} and low κ cannot sustain a high untreated baseline at any ambient exposure	A fitted parameter set that reproduces patient data while violating $\bar{n}^* < 1$	Applicable now to any published fit. Rejects roughly 63% of the $(\gamma_{\text{mem}}, \alpha)$ square at $P^* = 0.85$	Prop. B.1
GRASS reinterpretation: the arms differ in route-dependent induction efficiency α_{tol} , not in patient properties	Evidence that sublingual and subcutaneous routes induce regulatory apparatus with equal efficiency, or a constant- γ fit consistent with randomisation	Open. The constant- γ reading is internally inconsistent; the plastic reading is consistent but not independently confirmed	§B.8
Coupling magnitudes reduce to two scalars recoverable from blockade and depletion responses	A blockade/depletion pair that does not invert to a unique (ψ^+, ψ^-) in the physiological region	Protocol exists, not executed. Both interventions are routine; the measurement has not been made in these units	(B.31)–(B.32)
Ignition threshold: a sensitization threshold in ambient load, below which the network stays quiescent	A continuous dose–response across the predicted gap, or a patient baseline observed inside it	Open. The threshold <i>location</i> is robust to compartment count; the state above it is not, and the two should be reported separately	§B.9
Pre-symptomatic acoustic signature: a displacement core with a coherence-loss annulus, detectable before clinical activation	Absence of the annulus where the core is present, or co-localized maxima in all three channels, which would show the fields are not distinct functionals	Open. Requires no new instrumentation, only a new measurement target	§B.11

Prediction	What would refute it	Status	Where
Durable endpoint $P_{\text{end}} = (\gamma_f P_0 + A)/(\gamma_f + A)$	A post-treatment equilibrium that moves with protocol under constant γ_{mem} , or one that does not move under plasticity	Open. $1/\gamma_{\text{tol}}$ is the quantity to measure, from serial IgG4 or regulatory T-cell counts	Prop. B.5

On the first row. The refutation of Proposition [B.6](#) is recorded here rather than repaired, because repairing it means choosing a nonlinearity and no evidence currently selects one. What the contradiction establishes is narrower than it looks: it refutes the *linear* induction law $\alpha_{\text{tol}}\zeta(1 - P_{\text{tol}})$, not tolerance plasticity itself, since the durable endpoint ([B.26](#)) follows from γ_{mem} being a state variable and not from the particular law that drives it. The rival formulation ([B.27](#)), introduced for an unrelated reason, is nonlinear in exposure and is not excluded by the same data. Whether it survives is the first thing to check.

References

- José A. M. Borghans, André J. Noest, and Rob J. De Boer. Thymic selection does not limit the individual MHC diversity. *European Journal of Immunology*, 33(12):3353–3358, 2003.
- Ming-Han Chen, Yin-Ku Lin, Chih-Yung Chiu, et al. T-cell receptor diversity and allergen sensitivity in childhood asthma and atopic dermatitis. *Pediatric Allergy and Immunology*, 35(5):e14143, 2024.
- Isabella Marie Ledoux, Mohammad S. Ehlayel, et al. Associations between HLA class II alleles and IgE sensitization to allergens in the Qatar Biobank cohort. *Journal of Allergy and Clinical Immunology: Global*, 2(3):100120, 2023.
- Shawn H. R. Lee, Zhenhua Li, Evelyn H. Z. Lim, et al. Associations of T-cell receptor repertoire diversity with L-asparaginase allergy in childhood acute lymphoblastic leukemia. *Cancers*, 15(6):1829, 2023.
- Magdalena Migalska, Alvaro Sebastian, and Jacek Radwan. Major histocompatibility complex class I diversity limits the repertoire of T cell receptors. *Proceedings of the National Academy of Sciences*, 116(11):5021–5026, 2019.
- Magdalena Migalska et al. Evolutionary trade-offs constraining the MHC gene expansion: beyond simple TCR depletion model. *Frontiers in Immunology*, 14:1240723, 2023.
- Martin A. Nowak, Karl Tarczy-Hornoch, and Jonathan M. Austyn. The optimal number of major histocompatibility complex molecules in an individual. *Proceedings of the National Academy of Sciences*, 89(22):10896–10899, 1992.
- Benno Woelfing, Arne Traulsen, Manfred Milinski, and Thomas Boehm. Does intra-individual major histocompatibility complex diversity keep a golden mean? *Philosophical Transactions of the Royal Society B*, 364(1513):117–128, 2009.