A Covariant Postulation of Ring Bank Theory:
Constrained Concrescence, the Adjoint Clamp Identity, and Wake Holonomy

Brad Caldwell
caldwbr@gmail.com

June 8, 2026

Abstract

Ring Bank Theory (RBT) models the content of a conscious moment as the joint evolution of two geometric layers over a brief rolling window: a substrate *Attention Schema* that mints content at a thin *access manifold*, and a family of content-bearing *Paint Schemas* the substrate samples and projects. This paper restates the theory as a small set of postulates and sharpens its core formulas so that each is coordinate-correct on the relevant Lie group and, where possible, *derives* structure the prior presentation stated by hand. Eight refinements are advanced. (1) The clamp identity is rewritten covariantly in right-trivialized velocities, replacing a chart-level subtraction with an adjoint-transported equation valid on all of Sim(3); the hinge taxonomy (translate-, rotate-, scale-hinge) and the exactness of the scale law in logarithms then follow from the adjoint-orbit structure of $\mathfrak{sim}(3)$ rather than being posited. (2) The wake's "live factor" is identified as a *holonomy*—a path-ordered exponential of the realm's pose velocity—so that the *predicted* hysteresis of accumulated counter-pose is exactly path-dependence of a non-abelian transport. (3) Concrescence is cast as a constrained variational principle (a Gauss least-constraint projection), under which pose-forces are Lagrange multipliers, the coupling parameters (κ, α) are compliance weights and their partition, reference/displaced is a gauge choice, and the *matter-cannot-act* observation becomes a theorem. (4) The four printing modes collapse onto a two-parameter plane of compression ratio $\chi = \nu_{\text{ring}}/\nu_{\text{bin}}$ and sample rate ν_s, with extrusion the high-ν_s flicker-fusion regime rather than a high-persistence one. (5) Reversal-locked sampling is given a generative phase-locked-loop law. (6) Stitching is formalized as a connection along window-age with a self-coherence order parameter whose relaxation time reproduces the $\sim 1\,\text{s}$ onset of dissociation. (7) The mic is a rank-one selection whose salience functional unifies the companion paper's three pillars—geometric access, baseline dynamics, aboutness modulation—as additive source terms. (8) The sub-millisecond phenomenology of a kilohertz substrate is bounded by a population Fisher information (Cramér–Rao) estimate, turning a temporal puzzle into a representational one quantitatively. Throughout, each formal move is checked against the catalogued first-person phenomenology, and a concluding correspondence table ties every result to specific case files.

1 Introduction: what this paper sharpens

Ring Bank Theory treats conscious experience as a series of brief geometric events over a rolling window of roughly the last and next half-second. The architecture splits into an *Attention Schema* (the substrate machinery: the window $\mathcal{T}$, the running ringframe field ρ, the access manifold A at which content commits, the pose-state field ☧, the printing operator Π, the single-voice mic M, and the aboutness vector $\vec{C}$) and the *Paint Schemas* (the content-bearing surfaces directly intuited as experience: the real world R, the imaginal volume I, the bank of stored shapes B, and peripersonal space PPS). The companion *Journal of Ring Bank Phenomenology* catalogues the first-person case files on which the formalism rests; the present paper is a formal restatement that takes those constructs as fixed and asks a narrower question: *are the equations written in the strongest available form, and does that form give back the phenomenology?*

The prior formal core is already substantial. It supplies a three-factor decomposition of where a slice of content presents, a *clamp identity* governing what must move when arrivals at the access manifold disagree, a wake field with an explicit live factor, a gauge result removing the manifold's shape and content from the pose dynamics, and a system of pose- and rate-couplings parametrized by a strength κ and a transfer-symmetry α. The refinements below do not replace any of this. They tighten it.

The eight refinements, in one place. Each is developed in its own section and checked against the phenomenology in §14.

R1. The adjoint clamp identity (§4). The original clamp identity differentiates a composition of group elements but subtracts the results as if they lived in a vector space. We rewrite it in right-trivialized (spatial) velocities, so the left side is a genuine difference in $\mathfrak{sim}(3)$ and the right side is the printed body-jet transported by Ad. The hinge taxonomy and the logarithmic scale law become consequences of the adjoint-orbit structure of $\mathfrak{sim}(3)$.

R2. Wake holonomy (§5). The wake's live factor is a path-ordered exponential of the realm's pose velocity—a holonomy. The non-re-zeroing of counter-pose under a mid-gesture reference flip—a *prediction* of this form, not yet catalogued—would be exactly the path-dependence of a non-abelian transport.

R3. Constrained concrescence (§6). Concrescence is a Gauss least-constraint projection of the orchestrated trajectories onto the marriage constraint surface. Pose-forces are the Lagrange multipliers; (κ, α) are compliance weights and their partition; reference/displaced is a gauge/stiffness choice; and *matter-cannot-act* is the statement that the realized motion is the unique constrained projection, with no additional agentic term.

R4. The mode plane (§8). The four named printing modes are limits of a single compression ratio $\chi = \nu_{\text{ring}}/\nu_{\text{bin}}$ together with the sample rate ν_s; extrusion is the regime where ν_s climbs past flicker-fusion and the prints merge into a continuous extrusion, not a regime of elevated persistence.

R5. Reversal-locked sampling (§9). The anticipatory rate modulation that lands commit-events on kinematic keyframes is a phase-locked loop; the realized rate is the smoothest interpolation consistent with placing phase-zero at the next reversal.

R6. Stitching as a connection (§10). Across-window continuity is a connection along window-age; a self-coherence order parameter governs when the I-character is readable, and its relaxation time at full window depth reproduces the $\sim 1\,\text{s}$ timescales of dissociation onset and recovery.

R7. The selection functional (§11). The monophonic mic is a rank-one projector whose salience functional sums the three pillars—access, baseline, aboutness—so the companion paper's subtitle becomes three additive terms in one equation.

R8. Representational resolution (§12). The fineness of phenomenological time is bounded by the population Fisher information of the substrate, not by single-neuron firing rates; a Cramér–Rao estimate dissolves the sub-millisecond puzzle quantitatively.

Conventions. We represent Sim(3), the group of orientation-preserving similarities of $\mathbb{R}^3$ (rotation, translation, isotropic scale), by 4×4 matrices

$$g = \begin{pmatrix} sR & \mathbf{t} \\ \mathbf{0}^\top & 1 \end{pmatrix}, \qquad s \in \mathbb{R}_{>0},\ R \in \mathrm{SO}(3),\ \mathbf{t} \in \mathbb{R}^3,$$

so that group multiplication is matrix multiplication and the adjoint action is conjugation, $\mathrm{Ad}_g X = gXg^{-1}$. Its Lie algebra $\mathfrak{sim}(3)$ is seven-dimensional,

$$\mathfrak{sim}(3) \;=\; \underbrace{\mathfrak{so}(3)}_{\text{3 rot.}} \;\oplus\; \underbrace{\mathbb{R}^3}_{\text{3 transl.}} \;\oplus\; \underbrace{\mathbb{R}D}_{\text{1 scale}}\,, \tag{1}$$

with $D = \mathrm{diag}(1, 1, 1, 0)$ the dilation generator (so $\exp(uD) = \mathrm{diag}(e^u, e^u, e^u, 1)$). For a curve $g(t) \in \mathrm{Sim}(3)$ we write the *spatial* (right-trivialized) velocity $\omega_g := \dot{g}\, g^{-1} \in \mathfrak{sim}(3)$ and the *body* (left-trivialized) velocity $\xi_g := g^{-1}\dot{g}$. A subscripted dot, $\dot{g}_\varsigma$, always means $\mathrm{d}g_\varsigma/\mathrm{d}t$. We reserve t for clock-time, τ for window-age, $\tilde{\tau}$ for a body's intrinsic material ruler, and ℓ for within-event latency.

2 Primitives: the window, its axes, and the active set

We begin by fixing the substrate over which everything is defined, as a single postulate, and naming the objects that live on it.

Postulate 1 (Rolling window)**.** Conscious content at clock-time t is carried by fields over a bounded *rolling window* $\mathcal{T}$ of window-age $\tau \in [-\tau_{\max}, +\tau_{\max}]$ with $\tau_{\max} \sim 0.5\,\mathrm{s}$. The face $\tau = 0$ is the *bleeding edge* (the present), $\tau < 0$ the aging wake, and $\tau > 0$ the orchestrated near future. The window slides along clock-time: as $t \mapsto t + \mathrm{d}t$, the slice at $\tau = +\mathrm{d}t$ becomes the new bleeding edge.

Three time-like axes meet in this picture and must be kept distinct. Clock-time t is monotonic and unbounded; window-age τ is bounded and orthogonal to t; within-event latency ℓ is the per-event processing gradient from the sensory port (or baseline source) at $\ell \approx 0$ to fully-painted schema content at ℓ of tens to hundreds of milliseconds. A fourth object, the attention dwell-locus $\alpha(t)$, is a thin region in the (τ, t)-plane carrying its own τ-coordinate; in ordinary waking it sits at $\tau = 0$, but it may ride pastward ($\alpha_\tau < 0$, the postdictive "what just happened?" ride) or pin on orchestrated content ($\alpha_\tau > 0$). The printing operator stays at the bleeding edge regardless of where α wanders—a separation we will use in §9.

Definition 1 (State fields)**.** The *ringframe field* $\rho(\tau, t)$ has per-schema components $\rho_R, \rho_I, \rho_B, \rho_{\mathrm{PPS}}$ (and possible sub-component refinements), each running an independent trajectory through the window. The *pose-state field*

$$☧(\tau, t) = (☧^\rho_1, \ldots, ☧^\rho_{N_\rho(t)},\; ☧^\varsigma_1, \ldots, ☧^\varsigma_{N_\varsigma(t)})(\tau, t) \in \mathrm{Sim}(3)^{N(t)} \tag{2}$$

records the Sim(3) pose of every active component and every active Paint Schema at every point of the window, with $N(t) = N_\rho(t) + N_\varsigma(t)$ varying as objects join and leave the active set.

The active set is naturally a time-varying graph $G(t) = (V(t), E(t))$ whose vertices are the components and schemas of (2) and whose edges are the couplings of §6; births and deaths change $V(t)$, while couplings strengthening or failing change the weights on $E(t)$ without, in general, changing its topology. This graph picture will let us state dissociation precisely as an edge-weight collapse rather than an edge deletion.

The pose-state field carries a deliberate notational mnemonic. We write it with the Chi-Rho monogram ☧, which does two pieces of visual work at once (Figure 1). Read as a railroad-crossing signal, the X with a vertical post is the ringframe-tracks picture: ρ-component tracks diverging in $\pm\tau$ and converging at $\tau = 0$, with ☧ sitting at the convergence point as the pose-state the convergence resolves to. Read as a three-axis coordinate frame, the two diagonals plus the stem trace the Sim(3) frame—translations, rotations, and isotropic scale—at which a component is posed. Both readings are load-bearing and recur below: the first in the concrescence constraint of §4, the second in every appearance of a $\mathfrak{sim}(3)$-valued velocity or jet.

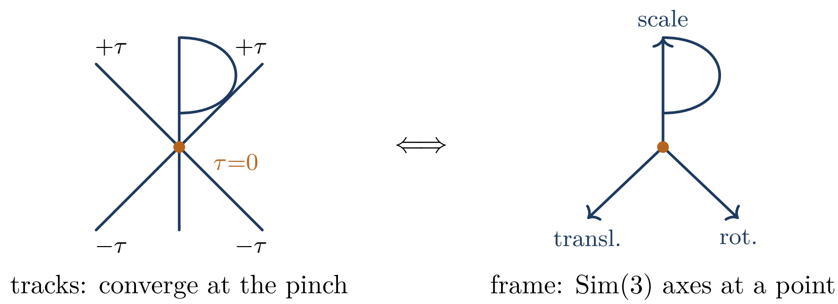

Figure 1: The Chi-Rho pose-state symbol ☧ under its two readings. Left: the converging-tracks reading, with components meeting at the pinch τ=0. Right: the coordinate-frame reading, the Sim(3) generators (translation, rotation, scale) meeting at a point. The same glyph names the convergence event and the pose it resolves to.

2.1 Paint Schemas, cameras, and the world anchor

The components ρ_X are not the felt content; they are the substrate tracks that the content rides. Each is rigidly glued, by the worldline-constant λ_X of Post. 2, to one *Paint Schema*—the surfaces directly intuited as experience. The real schema R carries the felt three-dimensional world of waking perception (roughly a room-scale volume about the body); the imaginal schema I carries thought-imagery, inner speech, and prospection (typically a smaller volume overlaying the front of the head); the bank B is a library of stored geometric shapes the substrate draws from; and peripersonal space PPS is the black embedding volume just inside and around the body, carrying the felt axes of "I am here." Two perspectival cameras render the visual content: a cyclopean camera Ξ_R for R and a narrower mind's-eye camera Ξ_I for I, whose axis settles onto Ξ_R's on eye-contact.

Underneath all of them sits the bare *world anchor* W: stationarity itself, carrying no content, the reference against which any schema's Sim(3) pose is specified. No schema owns stationarity; a schema is *reference* exactly when it is currently pose-locked to W, and which one that is reads off the case—R in ordinary waking (the Earth holds still while the avatar moves), PPS often under sedation (the body-frame holds still while the world counter-flows), and even PPS orbiting against an R-anchored world in dissociation. This is why the formalism never privileges a frame: W is a choice of section, and reference/displaced is which schema is gauge-fixed to it, a point made precise in §6. Each Paint Schema is itself time-extended—a schema-video carrying its own intrinsic ruler $\widetilde{\tau}$—so that the gluing of Post. 2 is a $\widetilde{\tau}$-by-$\widetilde{\tau}$ correspondence held for the pair's whole window-life, not a single point of attachment.

3 The three-factor decomposition and the printed jet

The pose at which a given slice of a component presents in modeling space factors into exactly three pieces, and isolating their clock dependence is what makes the dynamics tractable.

Postulate 2 (Three-factor presentation)**.** For each component–schema pair X, the pose at which the slice at material marker $\widetilde{\tau}$ of the body $\tilde{\rho}_X$ presents against the bare world anchor W is

$$g_X(\widetilde{\tau}, t) \;=\; g_{\varsigma_X}(t) \,\cdot\, \lambda_X \,\cdot\, m_X(\widetilde{\tau}), \tag{3}$$

where $m_X(\widetilde{\tau}) \in \mathrm{Sim}(3)$ is the *printing*—the voxel-marriage written along the body's intrinsic ruler, frozen at orchestration, so $\partial m_X/\partial t = 0$; $\lambda_X \in \mathrm{Sim}(3)$ is the *glue*—the rigid worldline-constant offset between the body and its schema-video; and $g_{\varsigma X}(t) \in \mathrm{Sim}(3)$ is the *realm*—the live pose of the whole schema against W, the only factor that can change with clock-time.

The argument lists carry the physics: m_X depends only on $\widetilde{\tau}$ because it is frozen into the body, $g_{\varsigma X}$ only on t because it is live, and λ_X on nothing because it never changes. The single live factor is what makes a re-pose of a schema carry its entire extruded body—wake included—at once, and it is why the zipper only ever has to act on ρ and on the manifold's pose at $\tau = 0$: the schemas follow by inheritance.

The quantity the dynamics will repeatedly need is the rate of change of the printing *along* the body. Because $m_X(\widetilde{\tau})$ is a curve in a non-abelian group, its "gradient" must be trivialized to land in the Lie algebra. We use the right-trivialization throughout.

Definition 2 (Printed jet)**.** The *printed jet* of the pair X is the right-trivialized derivative of the printing along the body,

$$J_X(\widetilde{\tau}) \;:=\; \frac{\partial m_X}{\partial \widetilde{\tau}}\, m_X(\widetilde{\tau})^{-1} \;\in\; \mathfrak{sim}(3). \tag{4}$$

It is a property of the pre-formed object alone—"freeze clock-time and walk along the frozen body, comparing the marriages of neighboring markers." Its three projections onto the summands of (1) are the per-marker increments of orientation, position, and size printed into the body.

Remark ("Jet" is used in a narrower sense here)*.* The word is borrowed from the companion journal's observational vocabulary but does *not* carry its full meaning. There, *jets* denote the entire stack of successive time-derivatives of a pose or shape variable—first, second, n^{th} (d/dt, $d^2/dt^2, \ldots$: "A is becoming less ring-like, and that becoming is accelerating"). The printed jet (4) is only the *first-order* term of such a stack—a single twist in $\mathfrak{sim}(3)$ (angular, linear, and scale rate bundled)—and the derivative is taken along the body's intrinsic ruler $\widetilde{\tau}$, not along clock-time t. Higher orders are the successive $\partial/\partial\widetilde{\tau}$ derivatives; they do not enter the first-order clamp identity (6), which relates one rate to another. Where the distinction matters, J_X may be read as the *printed twist*.

Writing the printed gradient as the algebra element (4) rather than as a bare $\partial m_X/\partial\widetilde{\tau}$ is the small move that makes the next section's identity coordinate-correct: a difference of group elements is not itself a group or algebra element, but (4) is unambiguously a vector in $\mathfrak{sim}(3)$, and the clamp identity will equate it (transported) to a difference of spatial velocities, which is also unambiguously in $\mathfrak{sim}(3)$.

4 Concrescence and the adjoint clamp identity

The architecture's central operation is *concrescence*: at every clock-instant, all $N(t)$ active bodies thread a single portal whose support is the access manifold A, their current slices forced into voxel-agreement there. Away from the bleeding edge the components owe each other nothing. The advance of clock-time *is* the pinching: what was at $\tau = +\mathrm{d}t$ becomes the new bleeding edge, the marriage acts on those arriving slices, and the previously pinched moment, now in $\tau < 0$, is released to drift.

Postulate 3 (Concrescence)**.** At each t there is a thin access manifold $A(t)$, a 0-to-2-dimensional support existing only at $\tau = 0$ and freshly constituted at every moment, carrying its own pose $g_A(t) \in \mathrm{Sim}(3)$ against W. For each active pair X let $\widetilde{\tau}^*_X(t)$ be the marker of $\tilde{\rho}_X$ currently at the portal. Concrescence is the constraint, holding in the zipped degrees of freedom,

$$g_A(t) \;=\; g_X(\widetilde{\tau}^*_X(t),\, t) \;=\; g_{\varsigma X}(t)\, \lambda_X\, m_X(\widetilde{\tau}^*_X(t)), \tag{5}$$

imposed simultaneously for all X in the active set.

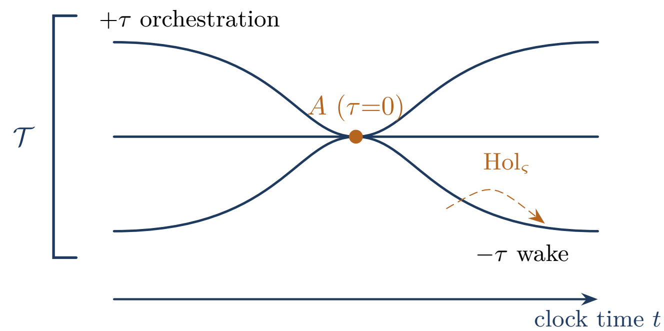

Figure 2: Concrescence at the bleeding edge. Per-component tracks run independent orchestrations in $+\tau$, are pinched into voxel-agreement at the access manifold A at τ=0, and unzip into the $-\tau$ wake, where each aged slice is carried by the realm holonomy Hol_ς of (10).

When the bodies' arrivals would disagree, something must give: a realm shifts pose, or A moves or deforms, or some mixture. The identity governing exactly what gives, and at what rate, is obtained by differentiating (5) in clock-time. The prior presentation wrote the result as $\dot g_A - \dot g_{\varsigma_X} = (\partial m_X/\partial\tilde\tau)\,(\mathrm{d}\tilde\tau^*_X/\mathrm{d}t)$ and noted that it must be "read per-generator" and is "exact in logarithms" only for pure scale. Both caveats are symptoms of subtracting group-element derivatives that do not live in a common vector space. The covariant form removes the caveats.

Proposition 1 (Adjoint clamp identity)**.** *Differentiating the concrescence constraint* (5) *and right-trivializing yields, for each active pair X,*

$$\boxed{\omega_A \;-\; \omega_{\varsigma_X} \;=\; \mathrm{Ad}_{g_{\varsigma_X}\lambda_X}\big(J_X(\tilde\tau^*_X)\big)\;\cdot\;\frac{\mathrm{d}\tilde\tau^*_X}{\mathrm{d}t}} \tag{6}$$

where $\omega_A = \dot g_A g_A^{-1}$ and $\omega_{\varsigma_X} = \dot g_{\varsigma_X} g_{\varsigma_X}^{-1}$ are the spatial velocities of the manifold and the realm, and J_X is the printed jet (4)*. Both sides are elements of $\mathfrak{sim}(3)$.*

Proof. Abbreviate $g_\varsigma \equiv g_{\varsigma_X}$, $m \equiv m_X(\tilde\tau^*_X(t))$, $\lambda \equiv \lambda_X$. By (5), $g_A = g_\varsigma\lambda m$. Differentiate, using $\partial m/\partial t = 0$ so that $\dot m = (\partial m/\partial\tilde\tau)\,\dot{\tilde\tau}^*_X$:

$$\dot g_A = \dot g_\varsigma \lambda m + g_\varsigma\lambda\,\frac{\partial m}{\partial\tilde\tau}\,\dot{\tilde\tau}^*_X.$$

Right-multiply by $g_A^{-1} = m^{-1}\lambda^{-1}g_\varsigma^{-1}$:

$$\omega_A = \dot g_A g_A^{-1} = \dot g_\varsigma g_\varsigma^{-1} + g_\varsigma\lambda\Big(\frac{\partial m}{\partial\tilde\tau}m^{-1}\Big)\lambda^{-1}g_\varsigma^{-1}\,\dot{\tilde\tau}^*_X = \omega_\varsigma + \mathrm{Ad}_{g_\varsigma\lambda}(J_X)\,\dot{\tilde\tau}^*_X,$$

which is (6). □

The reading is the same as before but now exact. The left side is the relative live motion of the manifold and the realm; the right side is *printed jet times feed rate*, parallel-transported by the adjoint of the realm-and-glue pose into the frame where the portal sits.

Remark (The prior identity is the linearization)*.* Evaluated near the group identity, where $g_{\varsigma_X}\lambda_X \approx \mathbf{1}$ so that $\mathrm{Ad}_{g_{\varsigma_X}\lambda_X} \to \mathrm{Id}$, and where the right-trivialized velocities $\omega_A = \dot g_A g_A^{-1}$, $\omega_{\varsigma_X} = \dot g_{\varsigma_X} g_{\varsigma_X}^{-1}$ reduce to the plain derivatives $\dot g_A, \dot g_{\varsigma_X}$, equation (6) collapses to $\dot g_A - \dot g_{\varsigma_X} = (\partial m_X/\partial\tilde\tau)\,(\mathrm{d}\tilde\tau^*_X/\mathrm{d}t)$—the prior chart-level clamp identity. The two caveats that accompanied it—"read per-generator" and "exact in logs only for scale"—are exactly the terms dropped in this linearization: the adjoint conjugation that mixes the rotation and translation generators away from the identity, and which the scale character is immune to. The covariant form (6) is thus a strict strengthening that agrees with the original wherever the original was meant to hold.

4.1 The two branches

Reference assignment decides how the left side of (6) splits. With A *as reference* ($\omega_A = 0$) the realm pays,

$$\omega_{\varsigma_X} \; = \; - \, \mathrm{Ad}_{g_{\varsigma_X}\lambda_X}(J_X) \, \frac{\mathrm{d}\widetilde{\tau}^*_X}{\mathrm{d}t}, \tag{7}$$

the minus sign being "something has to give"; the realm counter-moves at exactly the rate the printed change feeds through, and the wake counter-poses for free because every aged slice carries the same live factor $g_{\varsigma_X}(t)$ (§5). With *the component as reference* ($\omega_{\varsigma_X} = 0$) the manifold pays,

$$\omega_A \; = \; + \, \mathrm{Ad}_{g_{\varsigma_X}\lambda_X}(J_X) \, \frac{\mathrm{d}\widetilde{\tau}^*_X}{\mathrm{d}t}, \tag{8}$$

and A simply walks along the stationary body. Mixed cases divide the right side between the two terms. This yields a clean post hoc discriminator the case archive already uses: on the counter branch the wake warps while A's worldtube runs straight; on the traversal branch the wake sits undistorted in W while A's worldtube curves. Same printed body, same feed—the ledger balances, and which column was debited names the ground.

4.2 Why the scale law is exact, and the hinge taxonomy is forced

The single equation (6) now *derives* two facts the prior account stated separately. Both follow from the structure of the adjoint action on the decomposition (1).

Consider the *scale character* $\sigma : \mathfrak{sim}(3) \to \mathbb{R}$, the linear map reading off the dilation component (equivalently $\sigma(X) = \frac{1}{3}\,\mathrm{tr}$ of the upper 3×3 block of X). Because the scale factor $s : \mathrm{Sim}(3) \to \mathbb{R}_{>0}$ is a group homomorphism, $\sigma = \mathrm{d}\log s$ is a Lie-algebra homomorphism into the abelian $\mathbb{R}$, hence annihilates every commutator: $\sigma([\,\cdot\,,\,\cdot\,]) = 0$. Therefore σ is Ad-invariant, $\sigma(\mathrm{Ad}_g Y) = \sigma(Y)$ for all g, Y.

Corollary 1 (Exact logarithmic scale law)**.** *Applying σ to (6) kills the adjoint and gives, on the A-reference branch,*

$$\frac{\mathrm{d}\log s_{\varsigma_X}}{\mathrm{d}t} \; = \; - \frac{\partial \log s_{m_X}}{\partial \widetilde{\tau}} \, \frac{\mathrm{d}\widetilde{\tau}^*_X}{\mathrm{d}t}, \tag{9}$$

where s_{ς_X} and s_{m_X} are the scale parts of the realm pose and the printing. The scale law is frame-independent and exact; the previously noted "exact in logs" status is precisely the Ad*-invariance of σ.*

A body printed to grow 2% per marker, fed at three markers per second, shrinks its realm at exactly 6% per second—independent of where the manifold or realm sit. This is the dolly-zoom construction of a scale-hinge: camera approach and lens widening rate-matched so the subject holds constant size while the background warps.

The rotation and translation parts behave oppositely, and this is where the hinge taxonomy comes from. They span the derived subalgebra $\mathfrak{se}(3) = \mathfrak{so}(3) \oplus \mathbb{R}^3$ on which Ad acts *non*trivially, so their contribution to (6) is genuinely frame-dependent. The geometric signature of each generator is fixed by its fixed-point set under the affine action:

Corollary 2 (Hinge taxonomy)**.** *Decompose the transported jet* $\mathrm{Ad}_{g_{\varsigma_X}\lambda_X}(J_X) = \Theta + \mathbf{v} + \delta D$ *into rotation* $\Theta \in \mathfrak{so}(3)$*, translation* $\mathbf{v} \in \mathbb{R}^3$*, and scale* δD*. Each summand carries the relative portal velocity in one channel, producing a distinct wake signature dictated by its fixed-point structure:*

- *a* scale-hinge *(*$\delta \neq 0$*) has an isolated fixed* center *(the conjugation orbit of* D *under translations selects it), about which the wake dilates;*
- *a* rotate-hinge *(*$\Theta \neq 0$*) has a fixed* axis *(the screw axis of* Θ *through* A*), about which the wake revolves;*
- *a* translate-hinge *(*$\mathbf{v} \neq 0$*, others zero) has* no *fixed locus, since* $\mathbb{R}^3$ *is the abelian ideal of* $\mathfrak{sim}(3)$*; the wake simply streams.*

Combinations are free, and the deformation pattern of a wake reads back which generator was live and where A *sat.*

The taxonomy was previously a list of cases requiring, by hand, "a center," "an axis," or "no fixed locus." Here it is the orbit structure of Ad on (1): the center of the algebra (scale) gives the frame-free law, the abelian ideal (translation) gives the locus-free streaming, and the semisimple part (rotation) gives the axis. This *gives with* the phenomenology directly: the opposing-scale recession of the Unfolding Basketballs (one body upscaling, one downscaling through the shared portal) is two scale channels of opposite sign by (9); the apple-peel spiral is a rotate/scale-hinge whose curvature is written into A's worldtube on the traversal branch (8); and the Imaginal Head Turn is a rotate-hinge whose counter-turned pre-turn wake is the realm's ω_ς on the counter branch (7).

Remark (The clamp as instrument)*.* Equation (6) makes the phenomenology quantitative the way a voltage clamp does. Holding g_A fixed turns the realm's counter-velocity ω_{ς_X} into a readout of *printed jet times feed rate*, two quantities not separately introspectable. Watching how fast the surround contracts *is* measuring how fast the marriage-growth feeds through the portal.

5 The wake as holonomy

Once released into $\tau < 0$, each component drifts under its own dynamics, and the prior account gave the wake field as $\text{☧}_i^{\rho}(\tau, t) = [\, g_{\varsigma_i}(t)\, g_{\varsigma_i}(t - |\tau|)^{-1}\,]\, U_i(\tau)\, \text{☧}_i^{\rho}(0, t - |\tau|)$, with a bracketed "live factor" distinguished from an age-only operator U_i. The live factor has an exact identity: it is the holonomy of the realm's pose connection over the clock-time the slice has aged.

Proposition 2 (Wake live factor is a holonomy)**.** *Since* $\dot g_{\varsigma_i} = \omega_{\varsigma_i} g_{\varsigma_i}$*, the realm pose is the path-ordered exponential of its spatial velocity, and*

$$g_{\varsigma_i}(t)\, g_{\varsigma_i}(t - |\tau|)^{-1} \;=\; \mathcal{P}\mathrm{exp}\!\int_{t-|\tau|}^{t} \omega_{\varsigma_i}(t')\, \mathrm{d}t' \;=:\; \mathrm{Hol}_{\varsigma_i}[t - |\tau|,\, t]. \tag{10}$$

The full wake field is therefore

$$\text{☧}_i^{\rho}(\tau, t) \;=\; \mathrm{Hol}_{\varsigma_i}[t - |\tau|,\, t] \;\cdot\; U_i(\tau) \;\cdot\; \text{☧}_i^{\rho}(0,\, t - |\tau|), \qquad \tau < 0. \tag{11}$$

The two factors answer to different clocks, and now we can say exactly why no function of τ alone can absorb the live factor. The age-operator $U_i(\tau)$ depends on age alone—transit-revelation of the pre-formed body plus any structured deformation by the window's flow, all expressed in the realm's own frame, and itself factoring as $U_i = T_{\text{flow}}(\tau) \circ U_i^{\text{transit}}(\tau)$ with T_{flow} near-identity except where the flow field carries real structure. The holonomy (10) depends on the whole *path* of reference assignments between mint-time and now: on the traversal branch ($\omega_{\varsigma_i} = 0$) it is the identity and the wake sits undistorted in W; on the counter branch it is the accumulated ω_{ς_i} of (7), re-posing every aged slice at once regardless of age.

Corollary 3 (Hysteresis is path-dependence)**.** *Because* $\mathfrak{sim}(3)$ *is non-abelian, the path-ordered exponential* (10) *depends on the order in which reference assignments were held, not merely on net motion. If reference flips mid-gesture, the wake keeps the holonomy accumulated up to the flip and the newly displaced realm integrates onward from there; two episodes with identical kinematics but different reference histories leave differently-posed wakes. The wake is a record of* reference-assignment history*.*

This is the formal content of the journal's open observation target. Note also the consistency with Corollary 1: projecting the holonomy by the scale character gives an ordinary (path-independent) integral $\exp\int \mathrm{d}\log s_{\varsigma_i}$, so the *scale* of a wake slice never shows hysteresis; path-dependence lives entirely in the non-commuting rotation–translation interplay. The wake's information content—what makes it usable as short-term memory—is exactly the part no global holonomy can remove: the gauge-invariant relative poses $\text{☧}_i^{-1}\text{☧}_j$ among components in $\tau < 0$, i.e. their divergence from one another after the marriage released.

6 Constrained concrescence: pose-forces, couplings, and why matter cannot act

The clamp identity (6) is a first-order (velocity-level) constraint: it says what relative motion is admissible at the portal. It does not by itself say how the architecture *realizes* an update when several constraints compete, nor which component bends. The prior account supplied this through pose-forces (elements of $\mathfrak{sim}(3)$ applied at $\tau = 0$, propagating through each component's articulation like stress through a solid) and through couplings parametrized by a strength $\kappa \in [0, 1]$ and a transfer-symmetry $\alpha \in [0, 1]$. We show these are the pieces of a single variational principle, and that the principle is what the *matter-cannot-act* phenomenology is reporting.

6.1 A least-constraint principle

Each component arrives at the bleeding edge carrying a free (orchestrated) spatial velocity ω_i^{free}—the extension to $\tau = 0$ of its independent $+\tau$ trajectory. The marriage (5) imposes holonomic constraints $C_e(\text{☧}) = \text{☧}_i^{-1} T_{ij} \text{☧}_j = e$ on each coupled pair $e = (i, j) \in E(t)$. In general the free velocities violate these constraints; the architecture must choose a realized $\{\omega_i\}$ that satisfies them.

Postulate 4 (Constrained concrescence)**.** The realized spatial velocities are the unique minimizer of the weighted deviation from free motion subject to the marriage constraints:

$$\{\omega_i\} \;=\; \arg\max_{\{\omega_i\}} \Big(-G \Big), \qquad G = \tfrac{1}{2} \sum_{i \in V(t)} \kappa_i \left\| \omega_i - \omega_i^{\text{free}} \right\|_M^2 \;\; \text{s.t.} \;\; \dot{C}_e = 0 \;\forall e \in E(t), \tag{12}$$

where $\|\cdot\|_M$ is the gauge metric on $\mathfrak{sim}(3)$ (a weighted Killing-type form whose separate translation, rotation, and scale weights are set by the manifold's sensitivity, i.e. by the "being" γ) and $\kappa_i \geq 0$ are the coupling strengths.

This is Gauss's principle of least constraint transcribed to the pose group: the realized motion is the orthogonal (in the κM metric) projection of the free, orchestrated motion onto the constraint surface. Carrying out the constrained minimization in the standard way gives the realized velocities and the reactions that enforce them.

Proposition 3 (Pose-forces are Lagrange multipliers)**.** *The stationarity conditions of* (12) *read*

$$\kappa_i M \left(\omega_i - \omega_i^{\mathrm{free}}\right) \;=\; \sum_{e \ni i} \pm\, \Phi_e, \qquad \dot{C}_e = 0, \tag{13}$$

where $\Phi_e \in \mathfrak{sim}(3)$ *is the multiplier of constraint* e*. The* Φ_e *are exactly the* pose-forces*: each is the infinitesimal* $\mathfrak{sim}(3)$ *generator applied across the marriage edge, and the sign/incidence pattern* $\pm$ *on the right propagates it through the component's articulation hierarchy, reproducing the "stress through a solid" transmission.*

The realized update is thus $\omega_i = \omega_i^{\mathrm{free}} + \kappa_i^{-1} M^{-1} \sum_e (\pm\Phi_e)$: free orchestration plus the constraint reaction, scaled by the component's compliance κ_i^{-1}. Three readings of (κ, α) fall out immediately and match the catalogue.

- **Strength** κ is inverse compliance. A marriage edge at $\kappa \approx 1$ enforces its constraint rigidly; as $\kappa \to 0$ the edge's multiplier can no longer move the component and the constraint is effectively released—*partial dissociation is a* κ*-drop on an edge whose graph topology is unchanged*, precisely the journal's description.
- **Transfer-symmetry** α is the partition of an edge's reaction between its two endpoints. Writing the endpoint compliances as $(1-\alpha)$ and α, the value $\alpha = \frac{1}{2}$ splits the reaction evenly (symmetric marriage, both bend), $\alpha = 1$ loads it entirely onto the child (parent's transforms cascade, child does not back-drive), and $\alpha = 0$ the reverse. Within-component articulation sits at $\kappa \approx 1,\ \alpha \approx 1$; the marriage at A at $\kappa \approx 1,\ \alpha \approx \frac{1}{2}$.
- **Reference vs. displaced** is the limiting gauge choice $\kappa_i \to \infty$ on the reference component: an infinitely stiff vertex absorbs no reaction and is held against W, while the displaced vertices carry the full multiplier. The *sedation swap* is the exchange of which vertex carries $\kappa \to \infty$; the constraint surface $C_e = e$ is untouched, only the stiffness assignment flips, and the subjective signature—the world moves instead of the avatar—is the relabeled solution of the *same* (12).

Rate-coupling is the same principle on a parallel scalar axis: the constraint is a ratio of traversal velocities $\mathrm{d}\tilde{\tau}_i/\mathrm{d}t$ through stored content, with its own (κ, α), enforced simultaneously with the pose constraints at $\tau = 0$. A marriage can be pose-coupled but rate-decoupled, or conversely, because the two axes contribute independent constraint blocks to (12).

6.2 Matter cannot act

The variational form makes precise a phenomenological observation the journal records repeatedly—the *willless*, forced-progression quality of the cog-railway sessions, and the consciousness-as-servo case in which the body is laid down not by any force sweeping its atoms but by the manifold walking across voxels and marrying them.

Proposition 4 (Matter-cannot-act)**.** *Given the free orchestrations* $\{\omega_i^{\mathrm{free}}\}$*, the active graph* $G(t)$*, and the stiffness assignment* $\{\kappa_i, \alpha_e\}$*, the realized update* (13) *is* unique*. No additional agentic degree of freedom enters: the next configuration is the constrained projection of what was already orchestrated, and apparent agency is the felt character of this downhill flow on the constraint landscape.*

Sketch. G in (12) is strictly convex in $\{\omega_i\}$ on the affine constraint set $\{\dot{C}_e = 0\}$ (a positive-definite quadratic in the velocities under the metric κM, restricted to an affine subspace), so its minimizer exists and is unique whenever the active constraints are independent. The body's pose trajectory is then the time-integral of this uniquely determined velocity field; nothing in the data of the problem leaves room for an internal initiator. □

The body's posture is the integral of the manifold's walk-schedule across voxels; "laying yourself down" is the read-out of that committed sequence at the body-pose level. The clamp identity (6) is the velocity-level constraint inside (12); the pose-forces (13) are how it is met; and *matter-cannot-act* is the convexity that leaves the outcome with no free handle. The three are one mechanism seen at three levels.

7 Gauge: shape and being drop out

A constraint written over voxels in A should naively depend on the manifold's shape Γ (its morphology) and content γ (its voxel-level fill, loosely its "being"). It does not, and the reason is exactly the structure already in (12).

Proposition 5 (Shape/being gauge invariance)**.** *Provided A is confined to a single rigid segment of each component and is non-degenerate in the relevant degrees of freedom, the marriage Z_ρ pulled back to the pose group is the single group-relation*

$$\text{☧}_i(t) \;=\; T_{ij}\,\text{☧}_j(t), \tag{14}$$

in which the inter-component transforms T_{ij} depend only on the rigid attachment topology of A in each component. Different shapes Γ give different cost functions off *the constraint surface—different sensitivities to misalignment, i.e. different metrics M in* (12)*—but the same constraint surface $\{C_e = e\}$ in $\mathrm{Sim}(3)^{N(t)}$. The voxel content γ likewise parametrizes how the constraint is measured, not which poses satisfy it.*

Thus (Γ, γ) is a gauge choice: it enters (12) only through the metric $\|\cdot\|_M$ that weighs the projection, never through the constraint set onto which the projection lands. The physical content of the pose dynamics is the group-relation (14); required pose updates are shape-independent. This is what licenses treating the marriage as acting on ρ and on g_A alone: the explicit voxels drop out at the level of the equations of motion, surviving only as the sensitivity weighting that decides how sharply a misalignment is felt.

8 The printing operator: one mode plane

The ρ-zipper runs at $\sim$100% duty cycle while consciousness holds; what varies is which parts of A are *bright* (committed, introspectively available) and which are *dim*. The printing operator Π is the Attention Schema's sampling of the always-on zipper at the bleeding edge. It does not gate consciousness; it sets the intensity envelope. It factorizes into mode Π_ϕ, map Π_ψ, and rate Π_ω. The prior account named four discrete modes—phasic, compressed-phasic, interrupt, extrusion. They are not four mechanisms but regions of one two-parameter plane, and the second parameter is the *sample rate*, not persistence.

Definition 3 (Compression ratio and sample rate)**.** Let ν_{ring} be the tracer transit rate ($0 : 2\pi$ around A), ν_{bin} the enveloping bin rate, and ν_s the *sample/print rate* (the cadence at which fresh illuminations are committed). Define the *compression ratio* and the *tracer duty cycle*

$$\chi \;:=\; \frac{\nu_{\mathrm{ring}}}{\nu_{\mathrm{bin}}} \;\in [1, \infty), \qquad D \;:=\; \frac{1}{\chi} \;\in (0, 1]. \tag{15}$$

Per-voxel persistence (the wake's decay half-life) is *not* a mode parameter here: it is the ordinary τ-aging common to all four modes, not something elevated in any one of them.

Proposition 6 (The four modes are regions of (χ, ν_s))**.** *Below the flicker-fusion rate ν_{fus} the prints read as* discrete*, and χ alone names the mode:* phasic *is $\chi = 1$ (transit fills the bin, $D = 1$);* compressed-phasic *is $\chi > 1$ (a sprint of duration $D < 1$ then a wait);* interrupt *is the limit $\chi \to \infty$, $D \to 0$, the whole support flashing at one instant per bin with cadence read off ν_s.* Extrusion *is the regime $\nu_s \gtrsim \nu_{\text{fus}}$: the samples arrive fast enough to* fuse*, so the print reads as a continuously-present extruded volume with no gaps. There is no discrete bin or sweep to compress, so χ goes moot; persistence is unchanged.*

The continuity of extrusion is then flicker-fusion of a high but *finite* rate—plausibly only ∼20–40 Hz, not literally infinite—so the proposition carries a falsifiable edge: under the right conditions an extrusion should be resolvable back into discrete prints at its fusion rate, exactly as a high-refresh display or a fluorescent tube resolves into flicker when probed finely enough.

So the "turn-the-corner" transition from compressed-phasic to interrupt is simply $\chi \to \infty$ in the discrete band: sweeping faster and faster eventually stops reading as motion and starts reading as a single impulse. The complementary transition, interrupt→ extrusion, is ν_s rising past ν_{fus}: the gappy flashes close up into a continuous extrusion. The slow regime drift documented over hours in the long sedation brackets—clean phasic early, interrupt-like and frame-like late—is a monotone rise of χ with occupancy time, with the late shells additionally moving up the dimensionality of A. The "live-photo" pseudo-mode (content rate exceeding frame rate) is the off-diagonal region where the printing rate and the paint (content-update) rate decouple, which the plane accommodates without a new category: ν_{bin} for the screen, an independent content cadence for what plays on it.

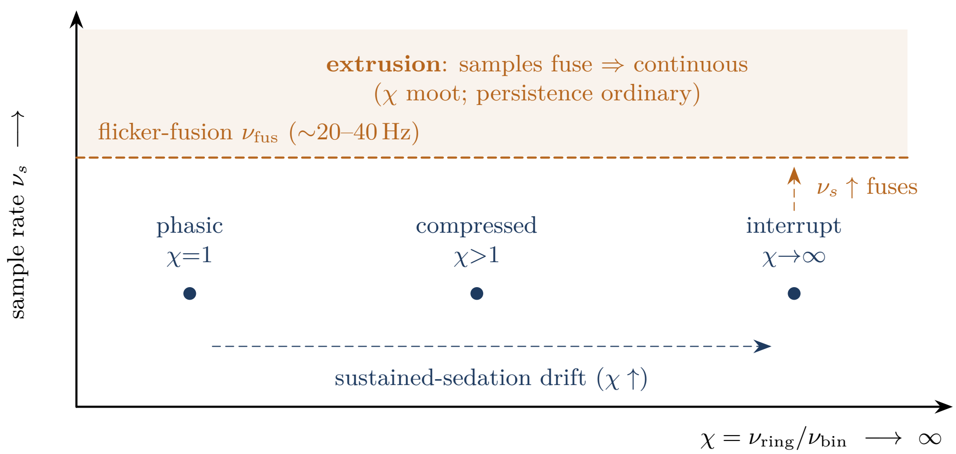

Figure 3: The four printing modes as regions of one (χ, ν_s) plane. Below flicker-fusion the prints are discrete and the compression ratio χ alone names the mode (phasic, compressed, interrupt, left to right); the dashed horizontal arrow is the slow phasic→interrupt drift of multi-hour sedation brackets. Above flicker-fusion the samples merge into a continuous *extrusion*—a high-but-finite sample rate (∼20–40 Hz), not elevated persistence—so the bin and sweep structure (and χ) wash out. The vertical arrow is the interrupt→extrusion fusion route.

9 Reversal-locked sampling as a phase-locked loop

A robust regularity in the catalogue is that commit-events land at *kinematic keyframes*—instants of velocity reversal where some Sim(3) channel of a ρ-component, of A, or of B has $\mathrm{d(pose)}/\mathrm{d}t = 0$. The trigger lives at the Attention-Schema layer, on the substrate's own machinery, not on the painted content. The prior account stated that the substrate achieves this not by adding an event on top of the regular cadence but by *anticipatorily modulating the rate itself*, so that phase-zero of the next commit just lands at the reversal, with the rate varying smoothly—no kink, no extra event. That is the behavior of a phase-locked loop, and naming it as one supplies a generative law and a testable prediction.

Define the *keyframe functional* on the spatial velocities of the Attention-Schema objects,

$$K(t) \;=\; \sum_a w_a \, \langle \omega_\bullet(t), e_a \rangle^2 \;=\; \|\omega_\bullet(t)\|_M^2, \tag{16}$$

whose zeros (in a given channel e_a) are the reversals; let $\{t_k\}$ be the upcoming reversal instants. Let the commit phase be $\varphi(t) = 2\pi \int^t \nu(t')\,\mathrm{d}t'$, with a commit-event at each $\varphi \in 2\pi\mathbb{Z}$.

Postulate 5 (Reversal lock)**.** The print rate $\nu(t)$ is governed by an Adler-type phase equation locking the commit phase to a reference phase Φ_{ref} built from the predicted reversal schedule,

$$\frac{\mathrm{d}\varphi}{\mathrm{d}t} \;=\; \omega_0(t) \;+\; K_c \, \sin\!\big(\Phi_{\mathrm{ref}}(t) - \varphi(t)\big), \qquad \Phi_{\mathrm{ref}}(t_k) = 2\pi k, \tag{17}$$

with ω_0 the relaxed cadence and K_c the lock gain. Between reversals the rate relaxes to ω_0; approaching a reversal it is pulled so that phase-zero arrives at t_k.

Proposition 7 (Smoothest-interpolation characterization)**.** *In the tight-lock limit the realized rate is equivalently the unique minimizer of* $\int (\dot\nu)^2\,\mathrm{d}t$ *subject to the interpolation constraints* $\int_{t_k}^{t_{k+1}} \nu\,\mathrm{d}t = 1$ *and* $\varphi(t_k) \in 2\pi\mathbb{Z}$*; i.e. exactly one cycle elapses between consecutive reversals and the rate profile is the smoothest (minimum-jerk) one consistent with hitting each keyframe at phase-zero. Hence "no kink, no extra event."*

The construction predicts a measurable micro-structure: the instantaneous print rate—and its EEG correlate—should dip or rise in advance of an upcoming movement reversal so that a commit lands on it, then relax. It *gives with* the squint-injection case (a $\sim$10 Hz cadence with one specialty ρ_I-dominant frame among ρ_R frames), with the two-handed gesture in which A object-switches between hands at successive max-abduction/adduction keyframes, and with the reported tempo phase-zero of auditory beat tracking, where the beat aligns with the velocity-zero of the substrate's own tracking ring at the envelope peak.

10 Stitching as a connection, and self as an order parameter

The marriage acts within a single τ-slice. Something else holds adjacent τ-slices of each component into one continuous fabric, so that "the last second" feels like one thing rather than a stack of disconnected slices. The prior account calls this *stitching*, notes that it is per-component, that it can tear independently of the marriage, that it operates at the window's full depth, and that self—the readable I-character—is what coherent multi-component stitching produces. We make the layer geometric.

Definition 4 (Stitch connection)**.** Each component i carries a *stitch connection* $\mathcal{A}_i$ along window-age, whose transport $S_i(\tau,\tau';t) = \mathcal{P}\exp\!\int_{\tau'}^{\tau} \mathcal{A}_i \,\mathrm{d}\tau''$ glues the slice at τ' to the slice at τ, modulated by a stitch-strength field $\sigma_i(\tau,t) \in [0,1]$. The marriage is the transverse ($\tau = 0$, cross-component) constraint; stitching is the longitudinal (along-τ, within-component) one.

That the two layers fail asymmetrically is now structural. The marriage is a constraint at a point in τ; stitching is a transport along τ. The marriage rarely loosens gracefully—its only failure mode is the global float-failure of extreme sedation, where the sample rate ν_s drops so low the whole process stops getting itself out of the water from one moment to the next. Stitching, by contrast, fails gracefully and locally: a single component's $\sigma_i \to 0$ peels it out of the fabric.

Proposition 8 (Window-depth timescale)**.** *If σ_i relaxes by a transport equation along window-age with characteristic speed set by $\mathcal{T}$'s flow—one window-depth $\tau_{\max}$ traversed in clock-time $\sim \tau_{\max}$—then a change in stitch-state propagates over the full depth and back in*

$$t_{\text{stitch}} \;\sim\; 2\,\tau_{\max} \;\approx\; 1\,s. \tag{18}$$

Dissociation onset and the return from a brief loss-of-consciousness blip are predicted to take of order one second, because the fabric of the second must tear or re-knit at full window depth.

This reproduces the catalogued ∼1 s onset and ∼1 s restitch of the dissociation triad without an additional fitted timescale: it is the round-trip time of the window itself.

Definition 5 (Self-coherence order parameter)**.** Let the cross-component stitch coherence be

$$\Sigma(t) \;=\; \frac{1}{|P|} \sum_{(i,j)\in P} \kappa^{\text{st}}_{ij}\, c(S_i, S_j), \tag{19}$$

where P ranges over co-active component pairs, κ^{st} is the inter-component stitch strength, and $c \in [0,1]$ measures compatibility of the two transports (unity when their holonomies agree, zero when they are unrelated). The readable I-character is available iff $\Sigma(t)$ exceeds a threshold Σ^*.

Proposition 9 (Self is a stitching effect)**.** *Self is not a primitive of the architecture. When $\Sigma > \Sigma^*$ the running story carries a coherent I; when a component's $\sigma_i \to 0$ and Σ falls below Σ^*, the zipper continues and content continues to commit, but the I-character drops out of readability. In the canonical dissociation case a body-control sub-component of ρ_R detaches over t_{stitch}; the marriage runs through with the remaining components (the missing role improvised, typically by I modeling R); no conscious servo controls the body though it continues to behave; restitching over another t_{stitch} restores Σ.*

The architecture thus has four dynamical layers, in increasing locality: stitching (along-τ, per-component, window-depth timescale); zippering (cross-component at $\tau = 0$, the marriage); pose-lock (Sim(3) couplings between components); and rate-lock (traversal couplings between $\widetilde{\tau}$-objects). Any can weaken or fail without forcing the others to—a statement now backed by the independence of the corresponding terms in (12) and the transverse/longitudinal split of Definition 4.

11 The mic: a rank-one selection unifying the three pillars

The mic M foregrounds exactly one high-level voice (one quale) at a time—exclusive and monophonic, like an old monophonic synthesizer sounding one note even when several keys are pressed. Its content may be sourced through any of the framework's three pillars, which are also the companion paper's subtitle: bottom-up *geometric access* (sensory flux), top-down *aboutness modulation* (imaginal, inner-voice, template responses), and *baseline dynamics* (intrinsic rhythm). We render the mic as a selection operator and the three pillars as additive source terms in its salience.

Definition 6 (Selection functional)**.** Let $\{\psi_k\}$ be the candidate voices at the bleeding edge. Their instantaneous salience is

$$s_k(t) \;=\; \beta_{\text{acc}}\, a_k(t) \;+\; \beta_{\text{about}}\, \langle \vec{C}(t),\, c_k \rangle \;+\; \beta_{\text{base}}\, b_k(t), \tag{20}$$

with a_k the bottom-up access drive, $\langle \vec{C}, c_k \rangle$ the alignment of voice k with the aboutness vector $\vec{C}(t)$, and b_k the baseline/intrinsic drive. The mic is the rank-one projector onto the winner,

$$\Pi_M(t) \;=\; |\psi_{k^*}\rangle\langle\psi_{k^*}|, \qquad k^* = \arg\max_k s_k(t), \tag{21}$$

its output seeding the ringframe that Paint Schema content then fills at latency ℓ.

The monophonic constraint is the rank-one form (21): a softmax over $\{s_k\}$ at low temperature collapses to a single winner, exactly the winner-take-all of a salience hub. The three pillars are not three mechanisms but three terms in one functional (20); whichever dominates at an instant sets the source of the current quale. A specialty injected frame is a transient in which $\beta_{\text{about}}\langle \vec{C}, c_k \rangle$ wins for one commit (the squint ρ_I-snapshot of §9); the $\sim$2 Hz flick of heavy sedation is b_k winning on a baseline rhythm; ordinary perception is a_k winning on sensory flux. The aboutness vector $\vec{C}(t)$ is the upstream control that biases (20) and, more broadly, coordinately modulates the downstream geometric variables—the "high-level orchestrated modulation of a myriad of variables by the switching of focused concept" of the companion paper. In the language of (12), $\vec{C}$ sets the orchestrated free trajectories ω_i^{free} that concrescence then projects onto the constraint surface.

12 Representational resolution: a Cramér–Rao bound

A standing objection to any substrate account of fine temporal phenomenology: human listeners discriminate interaural time differences of 20–200 μs, yet single neurons fire at most around 1 kHz and cortical representations of one physical event are staggered across regions by tens of milliseconds. The substrate looks unable to support the phenomenology. RBT's reply is that the brain is not *clocking* the event by firing a neuron at the right instant; it is *modeling* a rolling window $\rho(\tau, t)$, and the event is a feature of that model along the τ axis. We make the reply quantitative.

Proposition 10 (Parametric, not sampling, resolution)**.** *Let a population of M units with tuning $r_m(\tau)$ and noise variance v_m encode the window field along window-age. The achievable τ-resolution is bounded by the Cramér–Rao inequality*

$$\sigma_\tau \;\geq\; \frac{1}{\sqrt{I_F(\tau)}}, \qquad I_F(\tau) \;=\; \sum_{m=1}^{M} \frac{\big(r'_m(\tau)\big)^2}{v_m}, \tag{22}$$

which scales with the population Fisher information $I_F \propto M$, not with the single-unit rate. For $M \sim 10^4$ tuned units the bound gives σ_τ of tens of microseconds even when each unit is limited to $\nu_s \sim 1$ kHz.

The temporal mismatch between substrate firing rate and phenomenological time resolution is then no contradiction; it is the expected gap between a sampling rate and a model resolution, the same gap by which a learned continuous spatial representation has finer x-resolution than the neuron count would suggest. The reframe is precise: the old question—how does a $\sim$1 kHz substrate represent $\sim$200 μs phenomenology?—becomes how does the brain maintain a continuous $\rho(\tau, t)$ whose τ-axis carries Fisher information (22) of the required magnitude, together with a $\mathrm{Sim}(3)^{N(t)}$-valued ☧(t) updated by the marriage at $\tau = 0$. The same bound read forward recovers the journal's extrusion estimate: to resolve one part in a hundred of a 250 ms extruded surface one needs $\Delta\tau \approx 2.5$ ms, i.e. an effective $\nu \approx 400$ Hz of model fineness—a Fisher requirement on the representation, not a firing-rate claim about any neuron. The same logic underwrites the four-qualia-per-revolution time-division of the first catalogued ring: distinct θ-positions are distinct τ-features of one modeled revolution, resolvable to the precision (22) allows.

13 Worked reductions

To show the adjoint clamp identity is not merely tidier but computational, we run it through three canonical cases. Fix the matrix generators of $\mathfrak{sim}(3)$ in the 4×4 representation: the z-rotation L_z, the x-translation P_x, and the dilation D,

$$L_z = \begin{pmatrix} 0 & -1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad P_x = \begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad D = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \tag{23}$$

with brackets $[D, P_x] = P_x$, $[D, L_z] = 0$, $[L_z, P_x] = P_y$, and scale character $\sigma(L_z) = \sigma(P_x) = 0$, $\sigma(D) = 1$.

13.1 A scale-hinge: the dolly-zoom surround

Let the body be printed to grow uniformly, so its printed jet is pure dilation, $J_X = \kappa_g D$ with κ_g the fractional size increment per marker, and let it feed at $\mathrm{d}\tilde{\tau}^*_X/\mathrm{d}t = f$ markers per second. On the A-reference branch (7),

$$\omega_{\varsigma X} = -\operatorname{Ad}_{g_{\varsigma X}\lambda_X}(\kappa_g D)\, f.$$

The adjoint generally moves D off itself; for a pure translation by $\mathbf{b}$ along x, $g = \exp(b\,P_x)$, one computes

$$\operatorname{Ad}_{\exp(bP_x)}(D) = e^{\operatorname{ad}_{bP_x}} D = D + [bP_x, D] = D - b\,P_x, \tag{24}$$

the higher brackets vanishing since $[P_x, P_x] = 0$. The transported generator $D - bP_x$ is still a dilation, now centered at $\mathbf{b}$: this is the *center selection* of a scale-hinge made explicit, the choice of $\mathbf{b}$ being exactly the conjugating translation. Applying the scale character and using its Ad-invariance, $\sigma(D - bP_x) = \sigma(D) = 1$, so

$$\frac{\mathrm{d}\log s_{\varsigma X}}{\mathrm{d}t} = \sigma(\omega_{\varsigma X}) = -\kappa_g f.$$

With $\kappa_g = 0.02$ per marker and $f = 3$ markers/s the realm shrinks at 6% per second: $s_{\varsigma X}(t) = s_{\varsigma X}(0)\, e^{-0.06\,t}$. Because every wake slice shares the single live factor $g_{\varsigma X}(t)$, the whole surround—wake included—contracts uniformly at this rate while A holds clamp size and the incoming slices land at clamp size (each printed larger, each met by a realm already shrunk to match). All the motion is displaced into the background: the dolly-zoom, with A as the subject held fixed.

13.2 A rotate-hinge: the Imaginal Head Turn

Now let ρ_R's head be reference (world-fixed) and let ρ_I be printed to turn right, a pure rotation jet $J_I = \theta' L_z$ with θ' the angular increment per marker, fed at rate f. The marriage co-poses the two heads at A; the disagreement is resolved on the counter branch (7) by the imaginal realm paying,

$$\omega_{\varsigma I} = -\operatorname{Ad}_{g_{\varsigma I}\lambda_I}(\theta' L_z)\, f.$$

Conjugation by a similarity carries a rotation generator to a rotation generator of equal magnitude about the transported axis (the screw axis through A), so $\sigma(\omega_{\varsigma I}) = 0$—no size change—and the realm counter-rotates at angular rate $\theta' f$ about that axis. Integrating gives the wake holonomy (10),

$$\operatorname{Hol}_{\varsigma I}[t - |\tau|, t] = \mathcal{P}\exp\int_{t-|\tau|}^{t} \omega_{\varsigma I}\,\mathrm{d}t' = \exp\Big(-\int \theta' f\,\mathrm{d}t'\ \hat{L}\Big),$$

a rotation by exactly the executed head-turn angle about the A-axis $\hat{L}$, applied to every aged slice at once. This is the counter-turned pre-turn wake: the imaginal scene that was already laid down before the turn swings backward by the turn amount, because it rides the realm's live factor. A 30° orchestrated turn leaves a 30° counter-rotated wake, independent of slice age—the signature the case file reports.

13.3 A traversal: the apple-peel spiral

Finally hold the realm as reference, $\omega_{\varsigma_X} = 0$, with a mixed jet $J_X = \theta' L_z + v\, P_x + \kappa_g D$ (the manifold peeling around, advancing, and scaling into the fruit). The manifold pays, (8),

$$\omega_A = +\,\mathrm{Ad}_{g_{\varsigma_X}\lambda_X}(\theta' L_z + v\, P_x + \kappa_g D)\, f,$$

and integrating ω_A in clock-time traces a spiral worldtube for A through R. Here the holonomy (10) is the identity (the realm never moves), so the already-extruded tail stays fixed in W: the holographic prints of past tracer positions remain at their world locations as the bleeding edge spirals around the apple. The receipt is written as curvature in A's worldtube rather than as distortion of any realm—the clean discriminator promised by the two branches. The same three numbers (θ', v, κ_g) that on the counter branch would have warped a wake here bend a path instead; the ledger balances either way, and which column was debited names the ground.

14 Correspondence with the phenomenology

The test of each refinement is whether it gives back the catalogued first-person record. The following table maps every formal result above to the case files and recurring observations it is answerable to. The formalism is phenomenology-side throughout: it constrains what any substrate must do, and the substrate hooks (intrinsic-oscillation Π_ω dynamics, low-dimensional PCA rings as Π_ψ substrate, stored duration-objects behind the wake, and the phase-decoupling and voice-loss signatures of altered states) are correspondences, not claims of identity.

Result	Formal content	Phenomenology it gives with
Adjoint clamp identity (6)	relative spatial velocity = Ad-transported printed jet × feed rate	surround-contraction readout; dolly-zoom wake; counter vs. traversal branch discriminator
Exact scale law (9)	scale character is Ad-invariant	constant-rate realm shrink/grow; Unfolding Basketballs opposing scale flows
Hinge taxonomy (Cor. 2)	fixed-point structure of $\mathfrak{so}(3) \oplus \mathbb{R}^3 \oplus \mathbb{R}D$	apple-peel spiral (rotate/scale, curved A worldtube); translate-hinge streaming wake
Wake holonomy (11)	live factor $= \mathcal{P}\exp\int \omega_\varsigma$	Imaginal Head Turn counter-turned pre-turn wake; undistorted traversal wake
Hysteresis (Cor. 3)	path-dependence of non-abelian transport	reference-flip wake memory (open observation target)
Constrained concrescence (12)	Gauss projection of free orchestration	willless cog-railway; consciousness-as-servo body-laying
Pose-forces (13)	Lagrange multipliers, propagated as stress	neck counter-rotation; chest-to-shoulder cascade
(κ, α) as compliance	inverse stiffness and reaction partition	partial dissociation as κ-drop; parent→child articulation
Reference = gauge (§6)	$\kappa \to \infty$ stiffness choice	sedation swap (world moves vs. avatar)
Mode plane (Prop. 6)	four modes as (χ, ν_s) regions; extrusion = fusion	multi-hour phasic→interrupt drift; live-photo decoupling
Reversal-lock PLL (17)	phase lock to keyframe schedule	squint-injection specialty frame; two-handed object-switch; beat phase-zero
Stitch timescale (18)	round-trip window depth $\sim 1\,\mathrm{s}$	dissociation onset and restitch each $\sim 1\,\mathrm{s}$
Self order parameter (19)	threshold on cross-component coherence	I-character drops while zipper continues
Selection functional (20)	three pillars as additive salience	monophonic mic; baseline $2\,\mathrm{Hz}$ flick; aboutness specialty frames
Cramér–Rao bound (22)	population Fisher information $\propto M$	$20\,\mu\mathrm{s}$ binaural discrimination; $400\,\mathrm{Hz}$ extrusion estimate; four-qualia time-division

15 Summary: the postulated theory

The content of a conscious moment, under the sharpened theory, is a coupled pair of fields over a brief rolling window together with the constrained dynamics that act on them. The *statics* are a ringframe field $\rho(\tau, t)$ whose components run independent trajectories and are forced into voxel-agreement at the access manifold A at $\tau = 0$, and a pose-state field ☧(τ, t) valued in $\mathrm{Sim}(3)^{N(t)}$ over the active components and the schemas locked to them. The *dynamics* are the constrained projection of orchestrated free motion onto the marriage surface, with the following load-bearing identities, each now in coordinate-correct form.

- **Presentation** (Post. 2): $g_X(\tilde{\tau}, t) = g_{\varsigma_X}(t)\,\lambda_X\,m_X(\tilde{\tau})$—live realm, permanent glue, frozen printing.
- **Concrescence** (Post. 3): $g_A(t) = g_{\varsigma_X}(t)\lambda_X m_X(\tilde{\tau}^*_X)$ for all active X at once.
- **Adjoint clamp identity** (Prop. 1): $\omega_A - \omega_{\varsigma_X} = \mathrm{Ad}_{g_{\varsigma_X}\lambda_X}(J_X)\,\mathrm{d}\tilde{\tau}^*_X/\mathrm{d}t$, from which the exact logarithmic scale law and the hinge taxonomy follow as the adjoint-orbit structure of $\mathfrak{sim}(3)$.

- **Wake holonomy** (Prop. 2): the live factor is $\mathcal{P}\exp\int\omega_\varsigma$, making hysteresis the path-dependence of a non-abelian transport.
- **Constrained concrescence** (Post. 4): the realized update is the Gauss least-constraint projection of the free orchestration; pose-forces are its multipliers, (κ,α) its compliances, reference/displaced its gauge, and matter-cannot-act its convexity (Prop. 4).
- **Gauge** (Prop. 5): manifold shape Γ and being γ enter only the projection metric, never the constraint surface.
- **Printing** (Prop. 6): the four modes are regions of one compression ratio χ and the sample rate ν_s, extrusion being the flicker-fusion (high-ν_s) regime rather than a high-persistence one; reversal-locked sampling is a phase-locked loop (Post. 5).
- **Stitching** (Def. 4): a connection along window-age whose self-coherence order parameter Σ gates the I-character and whose round-trip timescale is the $\sim 1\,\mathrm{s}$ of dissociation.
- **Selection** (Def. 6): a rank-one mic whose salience sums geometric access, baseline dynamics, and aboutness modulation—the three pillars of conscious semiosis.
- **Resolution** (Prop. 10): phenomenological time-fineness is the population Fisher information of the model along τ, not the firing rate of the substrate.

The throughline is a single discipline: write each construct on the group where it actually lives, trivialize derivatives into $\mathfrak{sim}(3)$ before equating them, and let the algebra's own structure—its center, its abelian ideal, its semisimple part, its non-commutativity—carry the facts the prior account had to state by hand. The clamp identity becomes adjoint-covariant; the scale law's exactness becomes a homomorphism property; the hinge taxonomy becomes an orbit structure; hysteresis becomes a holonomy; pose-forces become multipliers; the willlessness of the moment becomes the convexity of a projection. None of this displaces the phenomenology that motivated it. It is offered as the form in which that phenomenology is least burdened by coordinates and most exposed to test.

What remains open. Three items are flagged for the archive rather than settled here. First, the flow operator $T_{\text{flow}}(\tau)$ of (11) is given only as a near-identity factor with occasional real structure; a measured functional form along τ (the anecdotal logarithmic slowing of the wake) would close it. Second, the mid-gesture reference flip predicted to leave a hysteretic wake (Cor. 3) has no clean catalogued instance yet and is a specific observation target. Third, the lock gain K_c and relaxed cadence ω_0 of (17) should be recoverable from the instantaneous-frequency micro-structure of EEG around movement reversals, turning Post. 5 into a quantitative substrate prediction.

Acknowledgments

The companion archive *Journal of Ring Bank Phenomenology* catalogues the first-person observations on which this formalism rests; readers wanting the empirical case files behind any construct above are pointed there. The compact expository statement of the unsharpened formalism appears in *Ring Bank Theory: Concrescence, Pose State, and the Printing Operator*; the broader program, including the three pillars of conscious semiosis, in *Ring Bank Theory of Conscious Semiosis: Geometric Access, Baseline Dynamics, and Aboutness Modulation.*