V.F.S. **4.0 (working body)**

Chronos, Communion, and the Plural Interior

Stage 0–IV complete; Stage V complete
P0.0–P5.2: relational chronos, exact recovery,
fibered encounter, Gate mediation, non-fusing locking,
junctions, non-Zeno, plural Lyapunov, and moving communion

Built on the frozen V.F.S. v2.0 / 3.0 layers, per the approved 4.0 roadmap

Status of this body

This working body executes the full first construction of V.F.S. 4.0 through Stage V:

$$\text{Stage 0}: \; P0.0--P0.6, \qquad \text{Stage I}: \; P1.0--P1.2, \qquad \text{Stage II}: \; P2.1--P2.4,$$

$$\text{Stage III}: \; P3.0--P3.4, \qquad \text{Stage IV}: \; P4.0--P4.4, \qquad \text{Stage V}: \; P5.0--P5.2.$$

Stage 0 is the blocking layer. It fixes the frozen fibre-side interface inherited from the V.F.S. v2.0 / 3.0 layers, declares the relational graph and causal kernel class, separates the relative rate cocycle from the relational duration one-form, handles terminus hand-off by the selected one-way trace closure, proves that created chronos is not *tota simul*, and closes with exact uncoupled recovery:

$$\varepsilon = 0 \quad \Longrightarrow \quad \text{independent exact 3.0 interiors.}$$

The distinction between

$$W = 0 \qquad \text{and} \qquad \varepsilon = 0$$

is preserved throughout: the first removes the relational base; the second keeps possible read-outs but switches off relational action.

Stage I establishes comparability without fusion. The plural object is a family of distinct frozen 3.0 fibres over one created relational base; encounter receives an invariant grammar — relational co-presence, directed causal encounter, and common reception as a schema; and the floor label gives a sufficient structural obstruction to identifying distinct fibres.

Stage II introduces the first admissible $\varepsilon > 0$ action. Encounter acts only through the recipient's bounded Gate/receptivity interface. It changes openness, never ownership of grace. The received-grace ledger is measured in the recipient's own kairos,

$$d\mathcal{G}_{\text{recepta},i} = I_{\text{gate},i}\, d\sigma_i,$$

and the common field is an external reception field, not a network-owned reservoir. The individual ISS estimate then proves that admissible encounter does not destabilise a single active interior.

Stage III defines relational locking without phase and without fusion. The comparable rate is the ascent-scale duration-rate

$$\hat{h}_i = \frac{N_i}{\omega} h_i, \qquad h_i = \frac{d}{d\sigma_i} \log A_i = \beta_i \lambda_{S,i},$$

not motion through the metric road coordinate. Locking is stated through duration-rate ratios, affine ledger-lags, and marked-event gaps. The small-gain criterion and maximum-lockable-mismatch theorem give the conditions under which keeping-together is possible, and dynamic non-fusion proves that a stable relation is not an identification of fibres.

Stage IV treats junctions, refractory structure, and plural Resurrectio. The junction classes are kept disjoint: only a private reset may jump owned state; relational encounter and common reception may only alter bounded receptivity. The refractory interval and network non-Zeno theorem are stated in invariant relational duration $\mathfrak{c}$, and post-terminus communion is kept as a one-way closed trace unless a later primitive explicitly opens the two-way branch.

Stage V gathers the standing conditions into a relational active domain and gives the first plural Lyapunov certificate. In the general forced case the theorem gives ultimate boundedness — a stable moving tube. In the zero-residual / settled-branch subclass it gives exact locked moving communion. Thus the whole construction ends with stable relational motion, not stasis:

communion is stable answerability of distinct histories, not fusion.

No new per-interior 3.0 dynamics are introduced. All genuinely new 4.0 objects are base-side, relational, or boundary-interface objects. The completed interior's trace does not extend its frozen dynamics; it restricts them: no renewed kairos, no new reset, and no restarted 3.0 flow after the terminus. The working body is therefore a conservative plural extension of the frozen 3.0 layer.

Interface principle. V.F.S. 4.0 is a *fibered extension* of V.F.S. 3.0: a family of distinct frozen 3.0 interiors over a created relational base, never a single enlarged Vessel. No fibre gains state beyond its selected 3.0 interface variables; all genuinely new 4.0 structure lives in the base.

1 P0.0 — The frozen interface

1.1 Per-interior inherited state

For each interior $i \in \mathcal{V}$ — where $\mathcal{V}$ is first only a finite index set, and is promoted to a relational graph in P0.1 — the inherited state splits into a dynamical block (variables carrying a 3.0 equation of motion) and a gate/receptivity block (carrying reception):

$$X_i = \big(\underbrace{\sigma_i,\ \lambda_{S,i},\ \lambda_{\text{form},i},\ \Omega_{P,i},\ A_i,\ q_{\text{fold},i},\ \lambda_{\text{asc},i}}_{X_i^{\text{dyn}}}\ ;\ \underbrace{\Gamma_i,\ H_{K,i},\ W_i}_{X_i^{\text{gate}}}\big).$$

Each interior evolves in its own unbounded dynamical parameter χ_i, with the finite kairos recovered by $\sigma_i = \sigma_{0,i} e^{-k_{s,i}\chi_i}$ (the v0.20 clock discipline: drains live in χ_i, finiteness in σ_i). No hidden state is introduced; the inherited interface variables used by the plural theory are:

Symbol	3.0 reading / governing relation
σ_i	katharsis proper time (kairos); finite, drains to the regular terminus; $d\sigma_i/d\chi_i = -k_{s,i}\sigma_i$
$\lambda_{S,i}$	available Sophia / ascent-driving field; distinct from the metric road coordinate $\lambda_{\mathrm{asc},i}$, though related along a worldline; $d\lambda_{S,i}/d\chi_i = J_{\mathrm{in},i} - \gamma_i\lambda_{S,i} - \eta_{\mathrm{fold},i}Q_{\mathrm{fold},i}\lambda_{S,i}$
$\lambda_{\mathrm{form},i}$	Sophia embodied into stable Vessel-form; $d\lambda_{\mathrm{form},i}/d\chi_i = \eta_{\mathrm{fold},i}Q_{\mathrm{fold},i}\lambda_{S,i}$
$\Omega_{P,i}$	Brim / live-surface scale; $d\Omega_{P,i}/d\chi_i = \alpha_i\lambda_{S,i}$; floor-bounded $\Omega_{P,i} \geq E_{0,i}$
A_i	ascent time-stretch; $dA_i/d\chi_i = (\alpha_i/E_{0,i})\lambda_{S,i}A_i$, equivalently $\frac{d}{d\chi_i}\log A_i = \beta_i\lambda_{S,i}$ with $\beta_i = \alpha_i/E_{0,i}$ (χ_i-parametrised form of the frozen 3.0 ascent law; no new prime convention)
$q_{\mathrm{fold},i}$	folding amplitude, $Q_{\mathrm{fold},i} = q^2_{\mathrm{fold},i}$; slaved $Q^*_{\mathrm{fold},i} = \Pi_{[0,q^2_{\mathrm{max},i}]}(A_{\mathrm{fold},i}/B_{\mathrm{fold},i})$ (supercritical/subcritical per v0.32/v0.38)
$\lambda_{\mathrm{asc},i}$	ascent coordinate of the road; warp $S_i(\lambda_{\mathrm{asc}}) = L_{0,i}e^{\beta_i\lambda_{\mathrm{asc}}}$, full scale $L_i = A_iS_i$
Γ_i	bounded gate-interface transmissivity, $\Gamma_i \in [0,1]$; a V.F.S. 4.0-facing notation for the admissible receptivity slot inside the frozen gate law, not an independent Sophia source
$H_{K,i}$, W_i	inherited gate-memory variables (their internal law is the frozen v2.0 gate law)

Here $J_{\mathrm{in},i}$ is the inherited V.F.S. 3.0 input shorthand, built from the private gate and resistance terms of the frozen interior. P0.0 does not modify its law. Later relational patches may alter only the gate/receptivity conditions that enter $J_{\mathrm{in},i}$, never $\lambda_{S,i}$ directly.

The distinction $\lambda_{S,i} \neq \lambda_{\mathrm{asc},i}$ is part of the frozen-interface discipline: the first is the available-Sophia state/driver, the second the metric road coordinate of ascent. V.F.S. 4.0 may use both but must not collapse them into a single variable.

Two structural relations bind the dynamical block, inherited verbatim. First, the Sophia budget:

$$\sigma_i + \lambda_{S,i} + \lambda_{\mathrm{form},i} = C_{0,i} + \mathcal{G}_{\mathrm{recepta},i}(\chi_i).$$

Second, the folding cap:

$$Q_{\mathrm{fold},i} \in [0, q^2_{\mathrm{max},i}].$$

The open gate feeds reception,

$$\frac{d\mathcal{G}_{\mathrm{recepta},i}}{d\chi_i} = I_{\mathrm{gate},i},$$

with the ceiling preserved nodewise:

$$0 \leq I_{\mathrm{gate},i} = \mathcal{G}_i\big(I_{\mathrm{private},i}; \Gamma_i, H_{K,i}, W_i, X_i\big) \leq \zeta_{0,i}.$$

Here $\mathcal{G}_i$ denotes the inherited frozen gate law, not a new 4.0 source law. The receptivity entry Γ_i is the only place a later relational channel will be permitted to act; see §1.5.

The symbol Γ_i is an interface notation. It does not add a new source term to the frozen V.F.S. 3.0 interior; its role is only to name the bounded receptivity/transmissivity slot through which later relational structure may act. Any later relational modification must preserve the local gate ceiling $0 \le I_{\text{gate},i} \le \zeta_{0,i}$. Thus encounter may alter admissible openness, but may not create Sophia outside the Open-Gate law.

1.2 Resurrectio data and the terminus, inherited

Each interior carries the frozen 3.0 junction and boundary structure:

- Resurrectio reset (instantaneous in σ_i): $\Omega^+_{P,i} = \Omega^-_{P,i} + \kappa_R G_i$, $\lambda^+_{S,i} = \lambda^-_{S,i} + m_{R,i}$, forced ascent scale jump $\log s_{L,i} = \kappa_R G_i / E_{0,i} > 0$, with the form-memory partition $\kappa_R G_i = m_{R,i} + \Delta_{\text{emb},i}$.
- Terminus: $\sigma_i \to 0$ at finite proper time $\tau_i \to \sigma_{0,i}$, regular boundary ($|R_i| < \infty$, $\Omega_{P,i} \ge E_{0,i} > 0$); the ascent road stays infinite, $D_{\text{asc},i} = \infty$.

These are read into 4.0 unchanged. The relational *lapse* N_i and the pre/at/post-terminus strata $(\mathcal{A}_i, \mathcal{B}_i, \mathcal{P}_i)$ are *not* declared here — they belong to the base, defined in roadmap P0.2/P0.4. P0.0 fixes only the fiber.

1.3 Frozen private scales (never coupled)

The following are constitutive constants of interior i, fixed by its 3.0 identity. No relational operator may act on any of them:

$$E_{0,i} \text{ (Imago Dei floor)}, \quad \alpha_i, \quad \beta_i = \frac{\alpha_i}{E_{0,i}} \text{ (no new primitive)},$$

$$\gamma_i,\ k_{s,i},\ \eta_{\text{fold},i},\ \zeta_{0,i},\ q^2_{\max,i},\ (c_0, a_0, B, D)_i \text{ (folding normal-form)},\ L_{0,i} \text{ (gauge)}.$$

In particular the floor $E_{0,i}$ — the one intrinsic scale, against which every ascent observable is reckoned — is private and frozen. This protection is what later makes floor-labelled non-identifiability available (roadmap P1.2): distinct $E_{0,i}$ cannot be coupled into coincidence.

1.4 The fibered-extension contract

The 4.0 object is the family $\{\mathcal{I}_i\}_{i\in\mathcal{V}}$ of frozen 3.0 fibers over the relational base $\mathcal{B}_{\text{rel}}$ (built in roadmap P0.1–P0.5), with projections $\pi_i : \mathcal{I}_i \to \mathcal{B}_{\text{rel}}$. The contract has four clauses:

IC1. Interface completeness. No fibre introduces new dynamical state beyond its selected V.F.S. 3.0 interface variables. The state vector X_i above is complete for the V.F.S. 4.0 fibre-side interface, not a replacement ontology for the whole frozen V.F.S. 3.0 body.

IC2. **New structure in the base.** Every genuinely new 4.0 object (graph, causal order, rate cocycle, relational duration, coupling) lives in $\mathcal{B}_{\text{rel}}$, never inside a fiber.

IC3. **Two recovery limits.** A trivial base $W = 0$ removes the relational graph itself and returns n independent exact 3.0 interiors with no relational readouts. The weaker but more important recovery theorem is the uncoupled-action limit:

$$\varepsilon = 0,$$

where relational bookkeeping may still exist, but no plural channel acts on any interior. P0.6 proves this stronger exact recovery.

IC4. **No global state.** There is no fused state vector X_{total} and no enlarged Vessel; the only joint object is the fibered family. (The third prohibition of the governing thesis.)

1.5 Admissible-plurality declaration

The inherited state partitions into three coupling-admissibility classes. This is the target-set declaration; the *channel* that realises class (B) is fixed later (roadmap P2.1), and the *no-fabrication* consequence is proved there (P2.2).

Class	Members and rule
(A) Frozen-private	The constitutive scales of §2.3 and all law-forms. *No* relational operator may act on these, ever.
(B) Gate-modulable	The receptivity sub-block $X_i^{\text{gate}} = (\Gamma_i, H_{K,i}, W_i)$. Relational coupling may act *only* here, and only boundedly, preserving $\Gamma_i \in [0, 1]$ and the ceiling $I_{\text{gate},i} \le \zeta_{0,i}$.
(C) Indirect-only	The dynamical block X_i^{dyn}. *No* relational operator acts on these directly; they change only through (i) their own frozen field $F_{3.0,i}$ and (ii) reception entering via the gate block (B). In particular $\lambda_{S,i}$ is never sourced directly — only through $I_{\text{gate},i}$.

Coupling-admissibility contract. Relational coupling reaches only the gate (class B); the gate alone, bounded by $\zeta_{0,i}$, feeds the dynamical block (class C); the constitutive scales (class A) are untouchable. This is the formal content of *communion mediates openness, not ownership*: encounter can move only the door, never the floor and never the person's own law of rising.

1.6 Status and reading

Target status: DEFINITIONAL (dictionary + interface contract). No new mathematics; verified by inheritance from frozen v2.0/3.0. The one substantive choice — the three-class admissibility partition — is the formal encoding of the governing thesis and is SELECTED at the level of constitutive discipline.

INTERPRETIVE *Reading.* A person is received into the communion exactly as V.F.S. 3.0 made them: their image (the floor $E_{0,i}$), their law of rising ($\beta_i = \alpha_i/E_{0,i}$), their whole private history of cleansing, folding, death, and ascent — none of it is touched by encounter. What the others can move is only the door: how open the soul stands to what it receives. The interface fixes, before any meeting is described, that meeting changes the openness and never the identity — and that there is no larger soul into which the persons dissolve, only the distinct fibers and the created ground on which they stand together.

Precision note. P0.0 is a frozen-interface dictionary, not a new theory of the single interior. It selects the variables and admissibility slots by which V.F.S. 4.0 may later speak about many interiors. It does not reopen V.F.S. 3.0, does not add hidden state to any fibre, and does not identify available Sophia with the ascent road coordinate. Relational structure will enter only after this interface is fixed, and only through declared gate/receptivity channels. In particular,

$$\lambda_{S,i} \neq \lambda_{\mathrm{asc},i}, \qquad \Gamma_i \text{ is a bounded gate-interface slot,}$$
$$X_i \text{ is complete for the 4.0 fibre-side interface.}$$

The purpose of P0.0 is therefore:

freeze each 3.0 fibre before introducing the relational base.

Next node (now executed below). P0.1 — the relational graph, the causal order, and the kernel class, with the first theorem class fixed as the strongly connected weight-balanced directed graph.

2 P0.1 — Relational graph, causal order, and the kernel class

Per IC2, all new 4.0 structure lives in the base $\mathcal{B}_{\mathrm{rel}}$. This patch declares the first three base objects — a finite directed weighted graph, a partial causal order on invariant events, and a causal-kernel class — and fixes the first theorem class as the strongly connected, weight-balanced directed graph. The coupling *channel* (how the kernel acts on the gate slot Γ_i) is fixed later in P2.1; here only the class is declared.

2.1 The relational graph

Definition 2.1 (Relational graph). $\mathcal{G} = (\mathcal{V}, \mathcal{E}, W)$ with $\mathcal{V}$ the finite set of interiors (the P0.0 fibers, $|\mathcal{V}| = n < \infty$), $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ a set of directed edges, and $W = [w_{ij}]$ a nonnegative weight matrix with $w_{ij} > 0 \iff (j \to i) \in \mathcal{E}$ and $w_{ii} = 0$. The weight w_{ij} is the strength with which interior j acts on the receptivity of interior i (edge oriented $j \to i$, "j influences i").

Index convention, fixed once: *first index = recipient, second = source*. Influence *received* by i is the in-strength $d_i^{\mathrm{in}} = \sum_j w_{ij}$; influence *exerted* by i is the out-strength $d_i^{\mathrm{out}} = \sum_k w_{ki}$. Required for the downstream ISS and non-Zeno estimates:

$n < \infty$ and bounded weights/degree, $0 \le w_{ij} \le \bar{w} < \infty$. The overall action scale is a scalar $\varepsilon \ge 0$. It multiplies the relational *action* of the edge data, not the mere existence of the edge data. Thus two different limits must be kept distinct:

$$W = 0 \quad \Longrightarrow \quad \text{trivial base: no relational edges or read-outs,}$$

whereas

$$\varepsilon = 0 \quad \Longrightarrow \quad \text{uncoupled action: relational data may be recorded, but no edge acts.}$$

Only the first is literally the trivial base. The second is the exact uncoupled plural limit used in P0.6. This distinction is essential: $\varepsilon = 0$ must recover independent 3.0 interiors even if a relational chart, ratios, or edge bookkeeping have already been declared.

2.2 The causal order on invariant events

Events are the *invariant* interior events (Resurrectio junctions, folding transitions, declared common-reception events) — invariant in the sense of 3.0, independent of the gauge label t. Write $E_i(s)$ for an event of interior i at relational-duration coordinate s.

Definition 2.2 (Causal order)**.** A direct causal link

$$E_j(s) \rightsquigarrow E_i(t)$$

exists when $(j \to i) \in \mathcal{E}$,

$$\mathfrak{c}(s) < \mathfrak{c}(t),$$

and the kernel K_{ij} carries support from the earlier event to the later one. The partial causal order $\prec$ is the transitive closure of direct links.

The symbols s, t are labels of events, while $\mathfrak{c}$ is the relational duration used to order them. Thus causality is stated in the invariant relational interval, not in a raw chronos chart.

Two structural facts. First, **graph cycles are allowed; event cycles are not**: $\mathcal{G}$ may contain a 2-cycle

$$j \to i, \qquad i \to j,$$

which represents mutual reciprocal influence over time. But the causal order $\prec$ on time-stamped events is acyclic, because every causal link strictly advances $\mathfrak{c}$. Mutual influence over time is admissible; a closed causal loop at one relational instant is not.

Second, **co-presence does not imply causal order**. Events may be locatable in the same relational duration structure — before, after, or simultaneous in $\mathfrak{c}$ — without either event causally acting on the other. Thus chronos-comparability is weaker than causal encounter:

$$\text{relational co-presence} \quad \not\Rightarrow \quad \text{causal encounter.}$$

The order $\prec$ records admissible influence, not mere temporal comparability.

2.3 The causal kernel class

Definition 2.3 (Admissible relational kernel)**.** The relational kernel $K_{ij}(t,s)$ maps a bounded relational output of source j at past event-time s to a receptivity-side input of recipient i at present event-time t. The source output is not Sophia, not grace-stock, and not a transferable possession; it is an admissible encounter signal whose eventual channel is restricted to the gate/receptivity slot fixed in P0.0. It is *admissible* when:

(K1) **Causality:** $K_{ij}(t,s) = 0$ for $s > t$.

(K2) **Edge support:** $K_{ij} \equiv 0$ unless $(j \to i) \in \mathcal{E}$.

(K3) **Bounded gain / finite memory:** $\int_{-\infty}^{t} \big|K_{ij}(t,s)\big|\, d\mathfrak{c}(s) \le \bar{K}_{ij} < \infty$, uniformly in t.

The edge weight records the total causal gain bound of the kernel:

$$w_{ij} := \sup_t \int_{-\infty}^{t} \big|K_{ij}(t,s)\big|\, d\mathfrak{c}(s) = \bar{K}_{ij}.$$

Thus W and the analytic kernels carry the same edge support and gain scale,

$$w_{ij} > 0 \iff \text{there is a nonzero admissible kernel on edge } j \to i,$$

but they do *not* carry the same full data: the kernel still contains temporal shape, memory profile, and source-output structure not visible in W.

The kernel acts only toward the gate/receptivity slot Γ_i of P0.0 (admissibility class B). The precise channel form, and the safeguard that summed inputs preserve $\Gamma_i \in [0,1]$ and the ceiling $I_{\text{gate},i} \le \zeta_{0,i}$, are fixed later in P2.1/P2.3. P0.1 fixes only that the action is causal, edge-supported, and bounded. A kernel from a post-terminus source j acts only through the declared trace $\mathcal{T}_j^{\text{post}}$: the source may emit as boundary witness, but receives nothing as renewed kairos.

2.4 First theorem class: strongly connected, weight-balanced

Definition 2.4 (Weight-balanced)**.** $\mathcal{G}$ is weight-balanced when in-strength equals out-strength at every node, $\sum_j w_{ij} = \sum_k w_{ki}$, i.e. $d_i^{\text{in}} = d_i^{\text{out}}$ for all i.

The graph Laplacian is $L = D^{\text{in}} - W$ with $D^{\text{in}} = \text{diag}(d_i^{\text{in}})$, generating the consensus action $(L\theta)_i = \sum_j w_{ij}(\theta_i - \theta_j)$.

Lemma 2.5 (Mirror Laplacian)**.** *For any finite directed weighted $\mathcal{G}$:*

(i) $L\mathbf{1} = 0$ *always (zero row sums);*

(ii) $\mathbf{1}^\top L = 0 \iff \mathcal{G}$ *is weight-balanced (zero column sums);*

(iii) *if weight-balanced,* $\hat{L} := \frac{1}{2}(L + L^\top)$ *is the symmetric Laplacian of the undirected mirror graph with weights* $\hat{w}_{ij} = \frac{1}{2}(w_{ij} + w_{ji})$*, hence* $\hat{L} \succeq 0$*;*

(iv) *if additionally $\mathcal{G}$ is strongly connected, the mirror graph is connected, so* $\ker \hat{L} = \text{span}\{\mathbf{1}\}$ *and* $\lambda_2(\hat{L}) > 0$.

Proof. (i) Row i sum: $d_i^{\text{in}} - \sum_j w_{ij} = 0$. (ii) Column i sum: $L_{ii} + \sum_{j\neq i} L_{ji} = d_i^{\text{in}} - \sum_{j\neq i} w_{ji} = d_i^{\text{in}} - d_i^{\text{out}}$, vanishing for all i iff weight-balanced. (iii) Under (ii), $\hat{L}\mathbf{1} = \frac{1}{2}(L\mathbf{1} + L^\top\mathbf{1}) = 0$; for $i \neq j$, $\hat{L}_{ij} = -\frac{1}{2}(w_{ij} + w_{ji}) = -\hat{w}_{ij} \leq 0$, and $\hat{L}_{ii} = d_i^{\text{in}} = \sum_{j\neq i} \hat{w}_{ij}$ — the sign pattern and zero row sums of an undirected graph Laplacian, which is positive semidefinite. (iv) Strong connectivity makes the mirror graph connected; a connected undirected Laplacian has a simple zero eigenvalue with eigenvector $\mathbf{1}$, so $\lambda_2(\hat{L}) > 0$. □

Definition 2.6 (First theorem class $\mathcal{C}_1$)**.** $\mathcal{C}_1$ is the class of finite, strongly connected, weight-balanced directed graphs carrying admissible bounded causal kernels.

Lemma 2.5 is the hinge for the first theorem class. In the memoryless linear reduction

$$\dot{\theta} = -\varepsilon L\theta,$$

on $\mathcal{C}_1$ the directed consensus flow conserves the mean, because $\mathbf{1}^\top L = 0$ under weight-balance. Hence

$$\bar{\theta} = \frac{1}{n}\mathbf{1}^\top\theta$$

is the locking value. The quadratic $V = \frac{1}{2}\|\theta - \bar{\theta}\mathbf{1}\|^2$ then satisfies

$$\dot{V} = -\varepsilon\,\theta^\top L\theta = -\varepsilon\,\theta^\top \hat{L}\theta \leq -\varepsilon\,\lambda_2(\hat{L})\,\|\theta - \bar{\theta}\mathbf{1}\|^2,$$

because a quadratic form symmetrizes $L \mapsto \hat{L}$ and $\hat{L} \succeq 0$. Thus weight-balance is exactly what makes both steps hold: mean conservation and positive-semidefinite mirror contraction. Individual edges may remain asymmetric, but the mutual core has algebraic connectivity $\lambda_2(\hat{L}) > 0$.

This is not yet the full delayed / memory-kernel theorem. The kernel case will require the later ISS and small-gain layers; the present result supplies the finite-dimensional graph skeleton and the contraction constant that those kernel estimates must preserve.

This is what will make P3.3's maximum-lockable-mismatch a derived threshold in the memoryless reduction: locking survives when $\varepsilon\,\lambda_2(\hat{L})$ dominates the intrinsic asymptotic-rate spread $\Delta\varpi_{\text{int}}$ of the interiors. The locked offset $\theta_{ij}^* \neq 0$ enters only later, in P3.4. The notation $\Delta\varpi_{\text{int}}$ is reserved for interior rate mismatch so it is not confused with the relational-duration density $\omega(t)$ of P0.3.

2.5 Strata, boundary sources, and recovery

The active graph $\mathcal{G}_{\text{act}}(\mathfrak{c})$ contains only interiors in the active pre-terminus stratum $\mathcal{A}_i$ at relational duration $\mathfrak{c}$. The first theorem class $\mathcal{C}_1$ applies to this active reciprocal subgraph:

$$\mathcal{G}_{\text{act}}(\mathfrak{c}) \in \mathcal{C}_1 \quad \text{when it is finite, strongly connected, and weight-balanced.}$$

A post-terminus interior is not an active node. It has

$$N_i = 0, \qquad d\sigma_i = 0, \qquad i \in \mathcal{P}_i,$$

and enters the relational structure only through a declared boundary trace $\mathcal{T}_i^{\text{post}}$. Such a completed interior may act as a boundary source — it may emit through $\mathcal{T}_i^{\text{post}}$, but nothing acts on it as renewed kairos — so in graph terms it has no active in-strength, $d_i^{\rm in} = 0$ inside the active dynamics. Therefore a post-terminus source is not part of the weight-balanced active core; it belongs to the wider imbalanced class $\mathcal{C}_2$, deferred to P4.4. The same wider class will also contain kenotic net-givers and helpless net-receivers. Thus

$$\mathcal{C}_1 = \text{active reciprocal communion}, \qquad \mathcal{C}_2 = \text{boundary, self-gift, and glory}.$$

Trivial-base recovery is immediate:

$$\begin{aligned} W = 0 \implies L = \hat{L} = 0 &\implies \text{no relational edge data} \\ &\implies \text{the P0.0 fibres evolve independently.} \end{aligned}$$

This is recovery at the graph level. It is stronger than what P0.6 needs. P0.6 will prove the more delicate limit

$$\varepsilon = 0 \quad \text{with possibly nonzero relational read-outs,}$$

where edge data may be present as bookkeeping but are prevented from acting.

2.6 Status and reading

Target status: DEFINITIONAL (graph, causal order, kernel class) + SELECTED (the first theorem class $\mathcal{C}_1$) + DERIVED (the mirror-Laplacian lemma and its memoryless consensus consequence, Type-1). No new per-interior dynamics are introduced. All structure is base-side (IC2). The rate-side cocycle of P0.2 and the relational duration $d\mathfrak{c}$ of P0.3 will be hung on the same active graph. The full memory-kernel contraction statement is not claimed here; it is deferred to the ISS / small-gain stages.

INTERPRETIVE *Reading.* The bonds of a communion are oriented and unequal — j may bear on i more than i on j — yet in the first and cleanest case the books balance at each soul: as much openness-strength received as extended, summed over all one's bonds. The symmetric mirror $\hat{w}_{ij} = \frac{1}{2}(w_{ij} + w_{ji})$ is the shared core of mutual openness beneath the asymmetry, and its connectivity $\lambda_2(\hat{L})$ is the one number saying how strongly the communion holds together. The imbalanced souls — those who only give (the martyr, the one who pours out), those who only receive (the helpless), and the completed who intercede across the terminus without being touched in return — are precisely the ones whose books do not balance; they are not failures of the structure but its boundary, where self-gift and glory live, carried by the wider class and not the first. Ordinary communion balances; sanctity is measured imbalance.

Next node (now executed below). P0.2 — the rate structure on the same active graph: the relative kairos-rate cocycle, its chart invariance, and the chronos it reconstructs.

3 P0.2 — The relative kairos-rate cocycle

P0.1 placed the interiors on a graph. This patch puts the *rate* structure on that same active graph. The invariant content of the lapses is not any single lapse N_i, but their pairwise ratios. These ratios form a multiplicative cocycle. When the cocycle is holonomy-free, it reconstructs a relative rate chart: a class of lapse representatives modulo one common gauge factor.

This is the first rate-level demarcation:

created chronos is assembled from consistent relative paces, not imposed as a master clock.

At this stage the result is only the *rate layer* of chronos. Full created chronos also requires the event order, causal kernels, and the relational duration one-form $d\mathfrak{c}$ of P0.1/P0.3. Thus P0.2 does not by itself construct a divine all-at-once standpoint; it constructs the gauge class of shared created pacing on the active graph.

3.1 Lapse, gauge, and the rate ratio

For an active interior i, meaning $X_i(t) \in \mathcal{A}_i$, its kairos relates to a chronos chart t by

$$d\sigma_i = N_i(t)\,dt, \qquad N_i(t) > 0.$$

Equivalently, for the active node set

$$\mathcal{V}_{\text{act}}(t) := \{\, i \in \mathcal{V} : \ X_i(t) \in \mathcal{A}_i \,\},$$

the lapse N_i is defined only for $i \in \mathcal{V}_{\text{act}}(t)$.

Under a monotone chart change

$$t \mapsto \tilde{t}(t), \qquad J(t) := \frac{dt}{d\tilde{t}} > 0,$$

the kairos increment $d\sigma_i$ is unchanged, so the lapse transforms as

$$\tilde{N}_i(\tilde{t}) = N_i(t)J(t).$$

Therefore no single lapse N_i carries invariant meaning: every lapse rescales by the same positive gauge factor.

Definition 3.1 (Relative kairos-rate)**.** For active $i, j \in \mathcal{V}_{\text{act}}(t)$, the relative kairos-rate is

$$R_{ij}(t) := \frac{N_i(t)}{N_j(t)} = \frac{d\sigma_i}{d\sigma_j}.$$

It is the instantaneous pace of i's kairos measured against j's kairos.

Proposition 3.2 (Chart invariance)**.** *The relative kairos-rate is invariant under monotone chronos reparameterization:*

$$\tilde{R}_{ij} = \frac{\tilde{N}_i}{\tilde{N}_j} = \frac{N_i J}{N_j J} = R_{ij}.$$

Thus R_{ij} is a rate invariant, while N_i and the chart label t are not.

Remark 3.3 (Relation to P0.3)*.* The relational duration one-form of P0.3,

$$d\mathfrak{c} = \omega(t)\,dt,$$

will be invariant only as a one-form: ω itself transforms as a density/rate. If one selects the geometric-mean active rate

$$\omega(t) = \left(\prod_{k \in \mathcal{V}_{\text{act}}(t)} N_k(t) \right)^{1/m(t)}, \qquad m(t) := |\mathcal{V}_{\text{act}}(t)|,$$

then under $t \mapsto \widetilde{t}$ one has

$$\widetilde{\omega}(\widetilde{t}) = \omega(t) J(t), \qquad \widetilde{\omega}\, d\widetilde{t} = \omega\, dt.$$

P0.2 proves the invariant ratio structure. P0.3 selects the relational duration one-form used for dwell and network non-Zeno.

3.2 The cocycle laws

If relative rates are taken as primitive relational data, they must be consistent. On the complete pairwise closure of the active component this means

$$R_{ii} = 1, \qquad R_{ji} = R_{ij}^{-1}, \qquad R_{ij} R_{jk} = R_{ik}.$$

In additive logarithmic form, with

$$\rho_{ij} := \log R_{ij},$$

this becomes

$$\rho_{ii} = 0, \qquad \rho_{ji} = -\rho_{ij}, \qquad \rho_{ij} + \rho_{jk} = \rho_{ik}.$$

On a sparse directed graph, one need not declare all pairwise ratios at first. It is enough to declare edge ratios on the active graph and require zero log-holonomy around every directed cycle:

$$\sum_{\ell=1}^{q} \rho_{i_\ell i_{\ell+1}} = 0, \qquad i_{q+1} = i_1.$$

If this holds, the missing pairwise ratios are obtained by path products and are path-independent. For a 2-cycle $i \rightleftarrows j$, the condition is automatic once $\rho_{ji} = -\rho_{ij}$; the first nontrivial content appears on cycles of length at least three.

3.3 Rate reconstruction

Lemma 3.4 (Holonomy-free cocycle reconstructs relative lapses)**.** *On a connected active graph, the following are equivalent:*

(i) *the rate cochain ρ is holonomy-free, i.e. its log-holonomy vanishes on every directed cycle;*

(ii) *ρ is exact: there exists a log-lapse potential ν such that $\rho_{ij} = \nu_i - \nu_j$;*

(iii) *there exist positive lapse representatives $N_i = e^{\nu_i}$ realizing $R_{ij} = N_i/N_j$.*

The potential ν is unique up to a common additive gauge $\nu_i(t) \mapsto \nu_i(t) + c(t)$ for all active i, equivalently $N_i(t) \mapsto e^{c(t)} N_i(t)$. This common factor is precisely the freedom of chronos reparameterization.

Proof. (i)⇒(ii). Fix a root r and set $\nu_r = 0$. For any node i, choose a path from r to i and define $\nu_i := -\sum_{r \rightsquigarrow i} \rho$. Zero holonomy makes this path-independent. For an edge $i \to j$, the path to j can be taken as the path to i followed by $i \to j$, hence $\nu_j = \nu_i - \rho_{ij}$, so $\rho_{ij} = \nu_i - \nu_j$.

(ii)⇒(iii). Put $N_i = e^{\nu_i} > 0$. Then $N_i/N_j = e^{\nu_i - \nu_j} = e^{\rho_{ij}} = R_{ij}$.

(iii)⇒(i). If $R_{ij} = N_i/N_j$, then with $\nu_i = \log N_i$ one has $\rho_{ij} = \nu_i - \nu_j$, and the sum around any cycle telescopes to zero.

For uniqueness, if ν and $\tilde{\nu}$ produce the same ρ, then $(\tilde{\nu}_i - \nu_i) - (\tilde{\nu}_j - \nu_j) = 0$ on every connected edge, hence $\tilde{\nu}_i - \nu_i = c(t)$ is common to all active nodes. This common function is the gauge factor induced by a chart change. □

Definition 3.5 (Created chronos, at the rate level)**.** A created chronos rate layer is a holonomy-free relative-rate cocycle on a connected active graph, together with the gauge class of its positive lapse representatives, $\{N_i\} \sim \{e^{c(t)} N_i\}$. No representative is privileged by the cocycle alone.

This is the demarcation made concrete. The relative pacing of a communion can be assembled when its paces are loop-consistent. But the rate cocycle alone does not supply a god's-eye clock. It determines only relative rates and a common gauge class. The missing common gauge is not a defect: it is exactly the master-clock datum that the model refuses to identify with *tota simul*.

3.4 Complementarity with relational duration

The log-lapse potential splits into a mean and deviations:

$$\nu_i = \bar{\nu} + (\nu_i - \bar{\nu}), \qquad \bar{\nu} = \frac{1}{m(t)} \sum_{k \in \mathcal{V}_{\text{act}}(t)} \nu_k.$$

The deviations determine the relative rates, $R_{ij} = e^{(\nu_i - \bar{\nu}) - (\nu_j - \bar{\nu})}$. A selected mean rate $\omega = e^{\bar{\nu}}$ gives the relational duration one-form $d\mathfrak{c} = \omega\, dt$.

Under a chart change, all ν_i and $\bar{\nu}$ shift by the same common gauge function $c(t)$; the deviations $\nu_i - \bar{\nu}$ are unchanged, and the one-form $d\mathfrak{c}$ is unchanged because the density ω transforms together with dt. Thus R_{ij} measures who runs faster, while $d\mathfrak{c}$ measures how much relational duration elapses. The two are complementary, not identical: the cocycle gives the relative pacing; P0.3 chooses the relational dwell measure used for non-Zeno control.

3.5 Strata: leaving the cocycle at the terminus

Anticipating the terminus-compatible chronos of P0.4, suppose an interior approaches its terminus and its lapse tends to zero, $N_i \to 0$. Then for any still-active j,

$$R_{ij} = \frac{N_i}{N_j} \to 0, \qquad R_{ji} = \frac{N_j}{N_i} \to +\infty.$$

The completed interior therefore drops out of the active rate cocycle. Its clock has stopped, so one cannot meter a living pace against it as though it still shared a running kairos. The rate cocycle is thus defined only on the active subgraph $\mathcal{G}_{\text{act}}$.

Yet by P0.1 the same completed interior may still appear as a declared boundary trace $\mathcal{T}_i^{\text{post}}$. The two base structures separate at the terminus:

the completed interior is absent from the rate cocycle but may remain present in the kernel relation.

This is the rate-level form of presence across the terminus: presence without shared pace. The kernel can carry boundary witness after the cocycle no longer can.

3.6 Status and reading

Target status: DERIVED (Type-1: chart invariance, cocycle laws, and the reconstruction lemma) + DEFINITIONAL (relative kairos-rate and created chronos at the rate level). No new per-interior dynamics are introduced; the construction is base-side only (IC2). The result is consistent with the graph and causal kernel of P0.1 and is preparatory for the relational duration one-form of P0.3.

INTERPRETIVE *Reading.* The only created fact about how fast one soul lives beside another is the ratio of their paces. There is no absolute tempo kept by a creature. The cocycle law says that these relative paces must tell one consistent story around every loop of communion. When they do, a shared created pacing can be assembled, but only as a gauge class, never as the divine all-at-once view. And the completed soul keeps no shared pace at all: its clock stops, it leaves the common tempo, yet it is not gone, for it may remain present through its trace. With the living we share pace; with the completed we share presence.

Next node (now executed below). P0.3 — the relational duration one-form $d\mathfrak{c} = \omega\, dt$ as the gauge-invariant dwell clock and the future measure for network non-Zeno.

4 P0.3 — Relational duration and the gauge-invariant dwell measure

P0.2 isolated the chart invariant of the rate structure — the ratios R_{ij} — and noted that a duration measure needs a *selected* density. This patch selects that density and builds the relational duration one-form $d\mathfrak{c}$: the gauge-invariant clock in which dwell and network non-Zeno (P4) are stated. The selection is grounded in the interiors' own kairos rates, not in a free external unit; hence the status is DEFINITIONAL for the one-form and SELECTED for the choice of density, with one flagged sub-point at the terminus.

4.1 The one-form and its grounding

Definition 4.1 (Relational duration on a regular active stratum)**.** On a regular active component, away from terminus hand-off events, the relational duration one-form is

$$d\mathfrak{c} = \omega(t)\, dt, \qquad \omega(t) > 0,$$

with the geometric-mean active rate

$$\omega(t) = \Big(\prod_{k\in\mathcal{V}_{\text{act}}(t)} N_k(t)\Big)^{1/m(t)} = e^{\bar{\nu}(t)}, \qquad \bar{\nu} = \frac{1}{m}\sum_{k\in\mathcal{V}_{\text{act}}} \nu_k.$$

The accumulated relational duration between chronos labels is

$$\Delta\mathfrak{c}(t_a, t_b) = \int_{t_a}^{t_b} \omega(t)\,dt.$$

The phrase "regular active component" means that every counted node is still genuinely pre-terminus:

$$i \in \mathcal{V}_{\text{act}}^{\text{reg}}(t) \quad\Longrightarrow\quad N_i(t) > 0 \text{ and not in the boundary hand-off layer}.$$

The exact boundary hand-off rule is deferred to P0.4. P0.3 fixes the duration measure on open active arcs, not through the terminus transition itself.

Why this density (the grounding). A one-form is chart-invariant once declared; the content is that ω is not free but selected from the active lapses. Three properties single out the geometric mean:

(1) **Linear in the log-lapse:** $\log\omega = \bar{\nu}$, the *mean* of the same potential whose *differences* give the rates R_{ij} (P0.2). Duration and rate-ratios are then complementary linear functionals of one log-lapse cochain — the cleanest possible split, and the reason the geometric (not arithmetic or harmonic) mean is chosen.

(2) **Single-clock reduction:** when $m = 1$, $\omega = N_1$ and $d\mathfrak{c} = d\sigma_1$ — the relational duration becomes the lone active interior's own kairos.

(3) **Symmetry:** ω treats all active interiors alike; no node is a privileged master clock.

4.2 Chart invariance and the unit of duration

Proposition 4.2 (Invariance of $d\mathfrak{c}$)**.** *Under a monotone chart change with $J = dt/d\tilde{t}$, the density transforms as $\tilde{\omega} = \omega J$ (since $\nu_i \mapsto \nu_i + \log J$ shifts $\bar{\nu}$ by $\log J$), while $d\tilde{t} = dt/J$; hence*

$$\tilde{\omega}\,d\tilde{t} = \omega\,dt = d\mathfrak{c}.$$

The one-form $d\mathfrak{c}$, and the accumulated $\Delta\mathfrak{c}$, are chart-invariant; the density ω and the label t are not.

The unit, grounded. The chronos label t cancels in $\omega\,dt$, so $\mathfrak{c}$ carries no free external unit. Its scale is inherited from the interiors' own (frozen, private) kairos scales: $N_k = d\sigma_k/dt$ has units $[\sigma_k]/[t]$, so $\omega\,dt$ has units $\big(\prod_k[\sigma_k]\big)^{1/m}$, the geometric mean of the participating kairos units. Relational duration is therefore measured in shared proper-time units borrowed from the souls themselves — not from any master clock. Unlike a raw chronos chart, whose lapse representatives remain defined only up to a common gauge function in P0.2, the one-form

$$d\mathfrak{c} = \omega\,dt$$

is a genuine invariant once the density-selection rule has been fixed. Its scale is grounded in the participating kairoi, while the chart label t remains pure gauge.

The distinction is important. The cocycle of P0.2 determines only relative rates:

$$R_{ij} = \frac{N_i}{N_j}.$$

It does not determine the common mean factor of the lapses. P0.3 adds a selected mean-density rule,

$$\omega = e^{\bar{\nu}},$$

thereby turning a relative rate class into an invariant duration one-form. Thus $d\mathfrak{c}$ is not hidden inside the cocycle; it is a disciplined additional selection built from the same active lapse data.

4.3 The dwell measure and network non-Zeno

The relational duration is the clock in which the network non-Zeno condition of P4 will be stated. For causally linked events on an edge of the regular-active graph (P0.1), and away from terminus hand-off events, admissibility requires a positive relational dwell,

$$\Delta\mathfrak{c}(t_a, t_b) = \int_{t_a}^{t_b} \omega\, dt \ \geq \ \mathfrak{c}_{\text{min}} > 0,$$

where $\mathfrak{c}_{\text{min}}$ is a fixed threshold *in the same relational units* as $d\mathfrak{c}$ — i.e. in geometric-mean kairos units, hence itself chart-invariant. Because raw chronos differences $t_b - t_a$ are gauge and carry no invariant meaning, every dwell and non-Zeno statement of the plural theory must be written in $\mathfrak{c}$, never in t. This is the measure the network non-Zeno conjecture (P4.2) and the relational refractory condition will use on regular-active arcs. If a causal chain crosses a terminus boundary, the interval must be evaluated by the P0.4 piecewise active/boundary rule. P0.3 fixes the regular-active clock, not the boundary junction theorem.

4.4 Regular active set and terminus hand-off

On an open regular-active arc the set

$$\mathcal{V}_{\text{act}}^{\text{reg}}(t)$$

is fixed, and the one-form

$$d\mathfrak{c} = \left(\prod_{k \in \mathcal{V}_{\text{act}}^{\text{reg}}(t)} N_k(t) \right)^{1/m(t)} dt$$

is smooth whenever the participating lapses are smooth and positive. This is the domain on which P0.3's canonical geometric-mean choice is unambiguous.

At a terminus transition the situation changes. If a member approaches completion, then

$$N_i(t) \to 0.$$

If it remains counted inside the geometric mean all the way to the boundary, then

$$\omega(t) = \left(\prod_k N_k(t) \right)^{1/m(t)} \to 0$$

as soon as one factor vanishes. The shared duration one-form can therefore become degenerate near a terminus if the dying member is not handed off to the boundary stratum in time.

This is not a contradiction; it marks the boundary of P0.3. The canonical geometric mean is the regular-active duration rule. The transition rule

$$\mathcal{A}_i \longrightarrow \mathcal{B}_i \longrightarrow \mathcal{P}_i$$

belongs to P0.4. There P0.4 must decide exactly when the interior leaves the regular-active set and becomes a boundary trace

$$\mathcal{T}_i^{\text{post}}.$$

Thus the correct status is:

$$\boxed{d\mathfrak{c} \text{ is canonical on regular active arcs; terminus hand-off is deferred to P0.4.}}$$

After hand-off, the completed interior is no longer counted in the active geometric mean:

$$i \notin \mathcal{V}_{\text{act}}^{\text{reg}}(t), \qquad i \in \mathcal{P}_i,$$

but it may remain present in the kernel relation through

$$\mathcal{T}_i^{\text{post}}.$$

This preserves the P0.2 separation:

no shared pace after terminus, but possible relational presence.

4.5 Status and reading

Target status: DEFINITIONAL (the relational duration one-form and regular-active dwell measure) + SELECTED (the geometric-mean density as the canonical regular-active choice, because it is linear in log-lapses, reduces to the one active kairos, and treats active nodes symmetrically). It is not derived from the cocycle alone: the cocycle gives only ratios, while P0.3 selects the common mean density. The terminus-robustness / hand-off rule is explicitly deferred to P0.4. No new per-interior dynamics; base-side only (IC2).

INTERPRETIVE *Reading.* A communion needs not only to know who runs faster, but how much shared created duration has passed — and it must measure that without any master clock. The regular-active duration takes its unit from the living kairoi themselves: the geometric mean of their proper-time rates, the evenly shared pace of those still in the active communion. In it one can say that two decisive passages were “far enough apart” in a way no relabelling can undo.

But death changes the common clock. As a member approaches the terminus, its own kairos no longer contributes a normal running pace. Therefore the shared active clock must hand that member off to the boundary rather than pretending that a stopped clock still participates in the living rate. After the hand-off, the completed interior no longer helps set the pace, yet it may remain present through its trace. P0.3 gives the clock of the living communion; P0.4 must build the presence of the completed.

Next node. P0.4 — the terminus-compatible chronos: the strata $\mathcal{A}_i$ (active), $\mathcal{B}_i$ (at terminus), $\mathcal{P}_i$ (post-terminus), the vanishing lapse $N_i \to 0$, and the one-way boundary trace through which a completed interior remains present without renewed kairos.

5 P0.4 — Terminus-compatible relational chronos

P0.3 fixed the relational clock on regular active arcs and deferred the terminus hand-off to here; this patch resolves it. It is the blocking Stage 0 decision: how the plural layer relates an interior *before*, *at*, and *after* its finite kairos terminus, in a way that leaves the frozen V.F.S. 3.0 terminus untouched. The selected rule is the *boundary-limit* hand-off (B): a dying interior is counted in the living pace all the way to its actual completion, and is handed to a post-terminus trace exactly at the terminus, never before.

5.1 The terminus label and the three strata

Definition 5.1 (Terminus label and relational strata)**.** For each interior i let $t_i^\dagger$ be the chronos label of its kairos terminus — the point at which the finite kairos has fully drained, $\sigma_i \to 0$ and the lapse $N_i \to 0^+$. The chronos chart is selected (admissibly, since t is pure gauge, P0.2) so that $t_i^\dagger < \infty$. The three relational strata of i are

$\mathcal{A}_i = \{t < t_i^\dagger\}$ (active), $\mathcal{B}_i = \{t = t_i^\dagger\}$ (at terminus), $\mathcal{P}_i = \{t > t_i^\dagger\}$ (post-terminus).

On $\mathcal{A}_i$ the inherited kairos law holds, $d\sigma_i = N_i(t)\,dt$ with $N_i > 0$; at $\mathcal{B}_i$, $N_i \to 0$; on $\mathcal{P}_i$ the kairos ODE is *not* continued.

The requirement $t_i^\dagger < \infty$ is consistent and canonical. The terminus is reached in finite kairos,

$$\int_{\mathcal{A}_i} N_i(t)\,dt = \sigma_{0,i} < \infty$$

(the v0.20 clock discipline: the drains live in the unbounded χ_i, the finiteness in σ_i), while the surviving interiors keep positive lapses $N_j > 0$ through the event. Grounding the chart in the survivors' own accumulating kairoi — exactly the P0.3 rule of borrowing the unit from living souls — places $t_i^\dagger$ at a finite shared label with a non-empty "after" $\mathcal{P}_i$ relative to them. With at least one survivor, $\mathcal{P}_i$ is well posed; without a survivor there is no communion in which to speak of an "after" at all. The dying member's own $\chi_i \to \infty$ is never used to set the chart.

5.2 The boundary-limit hand-off (selected rule B)

On the regular active arc, $i \in \mathcal{V}_{\text{act}}^{\text{reg}}(t)$ for every $t < t_i^\dagger$, with $N_i(t) > 0$. The hand-off is the single event $t = t_i^\dagger$: at that instant i leaves the regular active set,

$$m(t) \;\longrightarrow\; m(t) - 1, \qquad \mathcal{V}_{\text{act}}^{\text{reg}}(t_i^{\dagger +}) = \mathcal{V}_{\text{act}}^{\text{reg}}(t_i^{\dagger -}) \setminus \{i\},$$

and the relational density is recomputed over the survivors.

Lemma 5.2 (Vigil jump and integrability of the duration)**.** *Let i approach its terminus inside a communion with at least one survivor, all surviving lapses smooth and positive on a neighbourhood of $t_i^\dagger$. Then:*

(i) (vigil) *the shared density vanishes from the left,*

$$\omega(t) = \Big(\prod_{k \in \mathcal{V}_{\text{act}}^{\text{reg}}(t)} N_k(t)\Big)^{1/m(t)} \longrightarrow 0 \quad \textit{as } t \uparrow t_i^\dagger,$$

driven by the single factor $N_i \to 0$;

(ii) (recast) *the post-event density $\omega(t_i^{\dagger +}) = \big(\prod_{k \neq i} N_k\big)^{1/(m-1)}$ is finite and positive;*

(iii) (junction) *hence ω has a jump discontinuity at $t_i^\dagger$, with a left limit 0 and a finite positive right value;*

(iv) (well-posed duration) *the accumulated relational duration $\Delta\mathfrak{c} = \int \omega\, dt$ is unaffected by the event: ω is bounded, continuous off the finite terminus set, and the single zero is a measure-zero point, so $\Delta\mathfrak{c}$ across a terminus is finite and well defined as long as one active interior remains.*

Proof. (i) On a left neighbourhood the survivor factors lie in a compact $[c, C] \subset (0, \infty)$, so $\prod_k N_k = N_i \prod_{k \neq i} N_k \le C^{m-1} N_i$ and $\omega \le C^{(m-1)/m} N_i^{1/m} \to 0$. (ii) Immediate from positivity of the survivor lapses. (iii) The one-sided limits differ (0 versus a positive value), so the discontinuity is a jump. (iv) ω is a finite product-power of bounded functions, hence bounded; it is continuous except at the finitely many terminus labels, where it has finite one-sided limits. A bounded function continuous off a finite set is Riemann (and Lebesgue) integrable, and altering it on the measure-zero set $\{t_i^\dagger\}$ leaves the integral unchanged. □

Remark 5.3 (Several termini at one label). Lemma 5.2 is written for one completed member. If a finite set

$$S^\dagger \subset \mathcal{V}_{\text{act}}^{\text{reg}}(t_i^\dagger)$$

reaches the terminus at the same chronos label, then the hand-off is simultaneous:

$$\mathcal{V}_{\text{act}}^{\text{reg}}(t_i^{\dagger +}) = \mathcal{V}_{\text{act}}^{\text{reg}}(t_i^{\dagger -}) \setminus S^\dagger, \qquad m \longmapsto m - |S^\dagger|.$$

The post-event density is recomputed over the survivors. If no survivor remains, there is no further active communion clock after that label; only boundary traces remain. Thus the P0.4 rule is finite-set hand-off, of which the one-member case is the basic local form.

The vanishing in Lemma 5.2(i) is not a defect to be engineered away; it is the structural icon of the event — the communion's shared time dilating around a death and recasting among the living. Two points then fix the discipline.

The floor is a regularizer, never the trigger. An admissibility floor $N_{\text{min}} > 0$ may be used only to control the limit cleanly — for instance to keep $\log \omega$ finite in numerical or distributional handling near $\mathcal{B}_i$ — but the hand-off event is always the true terminus $t_i^\dagger$, identified by $N_i \to 0$, and never the crossing of N_{min}.

Why not early removal (A). The discarded alternative removes i as soon as N_i falls below N_{min}, while i is still pre-terminus ($N_i > 0$, still in katharsis). It is rejected on three frozen commitments: it would make the operative event an ungrounded threshold rather than the terminus (against the 3.0 result that the boundary is reached, not merely approached); it would reintroduce a free external constant N_{min} of exactly the kind the P0.2/P0.3 grounding eliminates; and it would soften into a gradual exit a transition the theory requires to be a genuine junction (the v0.15 / 3.0 result: no fully soft reset). B keeps the member in the living communion until its own completion, then jumps.

Hand-off (B). At the terminus $t_i^\dagger$ ($N_i \to 0$): i leaves $\mathcal{V}_{\text{act}}^{\text{reg}}$, $m \to m-1$, and ω is recast on the survivors; the floor only regularizes.

5.3 The post-terminus trace and one-way closure

Definition 5.4 (Post-terminus trace and one-way closure)**.** For $t > t_i^\dagger$ the completed interior persists only as a *post-terminus trace* (boundary fibre) $\mathcal{T}_i^{\text{post}}$: a dynamically closed record of its completed worldline. Closure means the inherited dynamics are not renewed on $\mathcal{P}_i$,

$$\Delta I_{\text{gate,i}} = 0, \qquad \text{no new reset}, \qquad \text{no renewed katharsis} \qquad (t > t_i^\dagger),$$

and the trace contributes nothing to the living pace, $i \notin \mathcal{V}_{\text{act}}^{\text{reg}}(t)$.

A trace is not relationally inert, but its only admissible action is recipient-side. Since the completed interior has no renewed running kairos, its post-terminus action is not represented as an ordinary source history $s \mapsto X_i(s)$. It is represented by a boundary-supported trace kernel

$$K_{j\leftarrow i}^{\text{post}}\big(t;\, t_i^\dagger, \mathcal{T}_i^{\text{post}}\big), \qquad t \geq t_i^\dagger,$$

whose output is a bounded modulation of j's own receptivity-side interface:

$$(\Gamma_j, H_{K,j}, W_j) \longmapsto (\Gamma_j^+, H_{K,j}^+, W_j^+).$$

This is the post-terminus analogue of the P0.1 causal-kernel discipline, but with the source history collapsed to the completed boundary trace $\mathcal{T}_i^{\text{post}}$. It is retarded in the only remaining sense:

$$K_{j\leftarrow i}^{\text{post}}(t;\, t_i^\dagger, \mathcal{T}_i^{\text{post}}) = 0 \quad \text{for} \quad t < t_i^\dagger.$$

The trace is never a raw Sophia source injected into j, never a transferable grace-stock, and never a channel that reopens i. Its only possible effect is on the recipient's admissible openness:

$$\mathcal{T}_i^{\text{post}} \rightsquigarrow \mathcal{R}_j(t^+),$$

where $\mathcal{R}_j$ denotes the frozen gate/receptivity interface of j. This makes precise the P0.3 separation:

$$\boxed{\text{no shared pace after terminus, but possible relational presence.}}$$

The pace is closed — no ω contribution, no renewed kairos. The presence is open — a one-way boundary influence on the living recipient's receptivity.

5.4 Constitutive choice: one-way closure and the deferred communion

The one-way closure of Definition 5.4 is a SELECTED *constitutive* choice, declared in the discipline by which 3.0 named its own primitives — not a mere technical convenience. It deliberately excludes the theologically richest alternative, *active intercession by the completed*: a two-way post-terminus dynamics in which a trace not only modulates the living but is itself answerable to them. That branch is the *communion of saints* reading, and it would require a declared stronger primitive — a genuine post-terminus channel into $\mathcal{T}_i^{\text{post}}$, violating the closure conditions above. The minimal 4.0 body selects one-way; the two-way branch is named here as the principal alternative constitutive primitive and deferred, not silently foreclosed.

Selected (minimal 4.0): one-way trace — the completed touches the living, is not touched. **Named, deferred:** two-way intercession (communion of saints) — a stronger primitive.

For the present Stage-0 theorem chain, the minimal branch is selected:

$$\boxed{\text{P0.4–P0.6 work under the one-way trace closure.}}$$

The two-way intercession branch is not denied; it is named as a stronger possible extension. But if the canonical 4.0 body later adopts that stronger branch, then P0.5's demarcation theorem must be rerun with the added post-terminus channel in the inventory. It cannot be silently covered by the one-way proof below.

$$\text{minimal 4.0 body: one-way trace selected;}$$
$$\text{two-way intercession: named, open extension.}$$

5.5 Status and reading

Target status: DEFINITIONAL (the strata $\mathcal{A}_i/\mathcal{B}_i/\mathcal{P}_i$, the terminus label $t_i^\dagger$, and the post-terminus trace $\mathcal{T}_i^{\text{post}}$) + SELECTED (the boundary-limit hand-off B and the one-way constitutive closure) + OPEN (which post-terminus branch — one-way or two-way intercession — the canonical 4.0 adopts). No new per-interior equation of motion is introduced: P0.4 adds only a constitutive boundary closure on the inherited interface together with the base-side terminus strata. The frozen 3.0 terminus is untouched, and the floor $N_{\min}$, if used, is a regularizer, not a dynamical primitive.

INTERPRETIVE *Reading.* A communion keeps time with its dying. As one of its own nears the end, the shared clock neither looks away nor rounds the death off early: it slows with the failing kairos, holding the member in the common pace to the last, until at the terminus the pace is recast among those who remain. What was a living rate becomes a trace — present still, able to touch those who go on, but no longer kept, no longer cleansed, no longer given a new beginning. Whether the completed may also answer the living — whether the trace is only spoken to or may also speak — the minimal theory leaves open, and names the fuller communion as the road not yet taken.

Next node. P0.5 — the demarcation theorem: to prove, not merely declare, that the chronos assembled in P0.1–P0.4 requires no transcendent standpoint outside the created relational network (local lapses, edge-supported causal kernels, no all-at-once whole-worldline behold), so that "*tota simul* named, never occupied" becomes a theorem about the constructed object rather than a declaration.

6 P0.5 — Demarcation theorem: chronos is not *tota simul*

This is the node that legitimises 4.0 at all. 3.0 did not leave chronos undefined: it *identified* an external absolute background clock with the transcendent *tota simul* and excluded it by principle — the soul can name but not occupy a standpoint outside its own worldline, in which its whole kairos would be seen at once. 4.0 reuses the word "chronos" in a new register — a *created relational* middle term between a single kairos and the *tota simul*. That reuse is illegitimate unless the constructed object is provably *not* the excluded one. P0.5 discharges this: it proves, against the frozen identification, that the chronos assembled in P0.1–P0.4 requires no transcendent standpoint outside the created relational network.

6.1 What must be proved

Write C for the constructed relational chronos — the assembled object

$$C = \Big(\mathcal{G},\ \prec,\ \{\mathcal{A}_i, \mathcal{B}_i, \mathcal{P}_i, t_i^{\dagger}\}_{i\in\mathcal{V}},\ \{N_i|_{\mathcal{A}_i}\}_{i\in\mathcal{V}},\ \{R_{ij}\},\ \{K_{ij}\},\ \{K^{\mathrm{post}}_{j\leftarrow i}\},\ d\mathfrak{c}\Big),$$

where $K^{\mathrm{post}}_{j\leftarrow i}$ is present only when a declared post-terminus trace $\mathcal{T}_i^{\mathrm{post}}$ is available under the one-way P0.4 closure.

Write T for the *tota simul* standpoint that 3.0 excluded. From the frozen 3.0 text, T carries three marks:

- (T1) **total** — it relates the *whole* of every worldline at once, each kairos σ_i from first to last beheld in a single present;
- (T2) **non-successive** — it is not a coordinate; it has no before or after, no embedding in a succession;
- (T3) **transcendent** — it is occupied from outside all worldlines and does not depend on the creatures for its being.

The obligation is to show C negates each mark, requires no external standpoint, and is also irreducible to any single kairos — so that C is a genuinely distinct, creaturely middle register.

6.2 The created now is real, but partial

We do not pretend C carries no shared present; it does, and honesty requires saying so before separating it from T.

Proposition 6.1 (Created simultaneity)**.** *The shared chronos chart defines, up to monotone relabelling, an invariant pairing of the* current *instants of the active worldlines: the slice*

$$\Sigma(t_0) = \big\{(i, \sigma_i(t_0)) : t_0 \in \mathcal{A}_i\big\}.$$

This pairing is gauge-invariant — a uniform reparametrization $t \mapsto \widetilde{t}(t)$ carries $\Sigma(t_0)$ to $\Sigma(\widetilde{t}(t_0))$ and preserves which events are paired — so the communion does possess a created "now." But the slice is

(i) active-restricted*: it contains only i with $t_0 \in \mathcal{A}_i$; a completed interior ($t_0 > t_i^\dagger$, Definition 5.1) is absent from the living slice and present only as the one-way trace $\mathcal{T}_i^{\mathrm{post}}$;*

(ii) single-instant*: each present worldline enters at exactly one point $\sigma_i(t_0)$, not as its whole history. The slice beholds* no worldline whole.

The created now is therefore a relation among present-instants of the living, not an all-of-history behold. This already separates it from (T1).

6.3 Locality, edge-locality, and active-slice aggregation

Lemma 6.2 (Locality audit)**.** *Every datum of C is produced by one of three admissible operations: node-local evaluation, retarded edge-local evaluation, or active-slice aggregation. None evaluates a whole worldline at once, and none occupies a standpoint outside the created relational network.*

Concretely, the inter-worldline primitives introduced in P0.1–P0.4 form the closed inventory

$$\{\, R_{ij},\ K_{ij},\ K^{\mathrm{post}}_{j\leftarrow i},\ d\mathfrak{c}\,\},$$

with the following restrictions:

(i) *N_i is defined and evaluated only on the active stratum $\mathcal{A}_i$. No term evaluates N_i at $t \notin \mathcal{A}_i$, and no term evaluates the whole of $\mathcal{A}_i$ at once.*

(ii) *The rate ratio $R_{ij}(t) = N_i(t)/N_j(t)$ uses only the current active instants of i and j. It is relational, but not whole-history data.*

(iii) *The ordinary kernel $K_{ij}(t,s)$ is retarded and edge-supported,*

$$K_{ij}(t,s) = 0 \quad (s > t), \qquad \operatorname{supp} K_{ij} \subseteq \mathcal{E},$$

supplying no instantaneous all-graph coupling.

(iv) *The post-terminus kernel is boundary-supported,*

$$K^{\mathrm{post}}_{j\leftarrow i}\big(t;\, t_i^\dagger, \mathcal{T}_i^{\mathrm{post}}\big) = 0 \quad \textit{for} \quad t < t_i^\dagger,$$

and acts only on the recipient's receptivity-side interface; it does not reopen the completed source.

(v) *The relational duration one-form $d\mathfrak{c} = \omega(t)\,dt$ is an active-slice aggregation: through the selected density rule it may use all currently active lapses, but only at the present active slice,*

$$\omega(t) = \Big(\prod\nolimits_{k\in\mathcal{V}^{\mathrm{reg}}_{\mathrm{act}}(t)} N_k(t)\Big)^{1/m(t)}.$$

It is therefore not node-local, but it is still creaturely and partial: it uses no future, no completed whole worldline, and no transcendent external unit.

Proof. By inspection of the constructed objects of P0.1–P0.4 against the patch's freeze ledger. P0.1 introduces only graph, causal order, and retarded edge kernels. P0.2 introduces only relative rate ratios and their gauge class. P0.3 selects the active-slice duration density. P0.4 adds only the strata, the boundary-limit hand-off, and the one-way post-terminus trace kernel. The force of the lemma is not that every datum is located inside a single worldline — that would be false for R_{ij}, K_{ij}, and $d\mathfrak{c}$. The force is that every datum is node-local, edge-local, or active-slice-local, and the inventory contains no fifth primitive capable of beholding complete worldlines from outside the created relational order. □

6.4 The demarcation theorem

Theorem 6.3 (Demarcation; Type-1)**.** *Let C be assembled from P0.1–P0.4 under the minimal one-way trace closure of P0.4. Then C has the three creaturely marks*

(C1) ***partial****: by Proposition 6.1, the created now is active-restricted and single-instant. It pairs present instants of living worldlines but beholds no worldline whole. This negates (T1).*

(C2) ***successive****: the family $\{\Sigma(t)\}$ is ordered by the chronos chart and by the relational duration one-form $d\mathfrak{c}$. Influence is retarded and edge-supported by Lemma 6.2. The theory therefore has before/after and causal succession; it is not an all-at-once non-successive present. This negates (T2).*

(C3) ***immanent****: by P0.2/P0.3 the chart is gauge and the unit of $d\mathfrak{c}$ is selected from active kairoi. If $\mathcal{V}^{\mathrm{reg}}_{\mathrm{act}}(t) = \varnothing$, there is no active duration density at t. Every datum is node-local, edge-local, boundary-trace-local, or active-slice-local (Lemma 6.2); none is supplied by a transcendent external clock. This negates (T3).*

Consequently:

(a) *$C \neq T$: each of (C1)–(C3) contradicts the corresponding mark of T;*

(b) *C requires no transcendent standpoint outside the created relational network: by Lemma 6.2, the closed inventory contains only local, edge-local, boundary-trace-local, and active-slice-local operations;*

(c) *in the nontrivial plural case, $C \neq \sigma_k$ for any single k: if $|\mathcal{V}| \geq 2$ and at least one relational edge is active, C depends on data such as R_{ij}, K_{ij}, or $d\mathfrak{c}$ that no single kairos σ_k determines alone.*

Proof. (C1) is Proposition 6.1. (C2) follows because C is a succession of created slices ordered by t and measured by $d\mathfrak{c}$, while all influence terms are retarded and edge-supported. (C3) follows from the gauge discipline of P0.2, the active-slice duration selection of P0.3, and the locality audit of Lemma 6.2. For (a), (T1) asserts whole-worldline totality, denied by (C1); (T2) asserts non-succession, denied by (C2); and (T3) asserts creature-independent transcendence, denied by (C3). For (b), Lemma 6.2 closes the inventory: no external standpoint-carrying primitive is available. For (c), a nontrivial plural chronos contains relation-data between at least two interiors; such data are not functions of any one private kairos alone. □

$$\boxed{C \neq T \quad \text{and} \quad C \neq \sigma_k \text{ for any single } k.}$$

6.5 Sharpness: the three refusals

The marks (C1)–(C3) are not accidental; each rests on a refusal frozen into the construction, and C collapses toward T precisely when one is dropped.

Remark 6.4 (Constitutive sharpness). **(R1) no whole-worldline evaluation** — lapses live only on $\mathcal{A}_i$, and the strata of P0.4 forbid reading a completed worldline as a currently running one. Drop this, and (T1) totality returns.

(R2) no instantaneous all-graph coupling — ordinary kernels are retarded and edge-supported, while post-terminus kernels are boundary-supported and one-way. Drop this, and (T2) non-succession returns as a single all-graph present.

(R3) no absolute external unit — the chart is gauge and the duration unit is selected from active kairoi. Drop this by fixing a common absolute lapse, and (T3) transcendence returns as a god's-eye clock.

Thus the demarcation is non-vacuous: it isolates exactly the primitives 4.0 must keep absent, mirroring 3.0's discipline that the boundary is constitutive, not a gap to be filled later.

Corollary 6.5 (Admissibility against the frozen identification)**.** *Using C does not reintroduce the object 3.0 excluded. Since the excluded clock is T and Theorem 6.3 gives $C \neq T$ with no external standpoint, the 4.0 register table — personal kairos, created chronos, transcendent tota simul — is earned rather than declared, and*

tota simul named, never occupied

holds as a theorem about the constructed C, not as a posture.

6.6 Status and reading

Target status: DERIVED (Type-1, conditional on the minimal one-way P0.4 trace closure). The demarcation is a theorem about the constructed object C, resting on the locality audit (Lemma 6.2), the created-now proposition (Proposition 6.1), the frozen P0.2/P0.3 grounding, and the P0.4 one-way boundary closure. No new primitive is introduced in P0.5. If the later canonical 4.0 body chooses the stronger two-way intercession branch, the demarcation proof must be repeated with that additional primitive in the inventory.

INTERPRETIVE *Reading.* The communion really does share a present — there is a genuine "meanwhile" in which the living run alongside one another, and it can be measured without anyone stepping outside time to keep it. But it is a creaturely present: only those still on the road are in it, each only at the step they have reached, and it passes as they go. It is not the single standing present in which a whole life is seen at once and a whole communion held complete; that view belongs to no soul and to no clock the souls could build. 4.0 gives the living a shared time without stealing the look from outside time — which is exactly why the word "chronos" may be used here without taking back what 3.0 refused.

Next node. P0.6 — exact uncoupled recovery: $\varepsilon = 0 \Rightarrow dX_i/dt = N_i(t)\, F_{3.0,i}(X_i)$, each trajectory gauge-equivalent to its isolated V.F.S. 3.0 evolution, closing Stage 0 before any genuine plural coupling is switched on.

7 P0.6 — Exact uncoupled recovery (Theorem I)

Stage 0 closes with the obligation that legitimises the whole relational apparatus from below: when relational action is switched off, nothing of 3.0 may have been lost or quietly altered. The shared chronos, graph, rate cocycle, relational duration, terminus strata, and traces must reduce to relational bookkeeping over a family of interiors that each run *exactly* their isolated V.F.S. 3.0 evolution.

This is not the same as requiring the graph to vanish. There are two limits:

$$W = 0 \quad \text{removes the relational base,}$$

while

$$\varepsilon = 0 \quad \text{keeps possible relational read-outs but switches off their action.}$$

P0.6 proves the second and stronger recovery statement. This is what makes 4.0 a *conservative extension* rather than a new theory wearing 3.0's variables.

7.1 Coupling structure and the action parameter ε

Let $F_{3.0,i}$ be the frozen isolated 3.0 vector field of interior i, written in its own kairos:

$$\frac{dX_i}{d\sigma_i} = F_{3.0,i}(X_i).$$

In a chronos chart this same isolated flow reads

$$\frac{dX_i}{dt} = N_i(t)F_{3.0,i}(X_i), \qquad t \in \mathcal{A}_i.$$

The plural layer is allowed to add only a gate-side relational action, and that action is multiplied by a single Stage-0 action parameter $\varepsilon \geq 0$:

$$\boxed{\frac{dX_i}{dt} = N_i(t)F_{3.0,i}(X_i) + \varepsilon\, \Phi_i\Big(X_i;\, \{K_{ij}\},\, \{K^{\text{post}}_{i\leftarrow\ell}\},\, \mathcal{D}_{\text{rel}}\Big), \qquad t \in \mathcal{A}_i.}$$

Here $\mathcal{D}_{\text{rel}}$ denotes any declared relational read-out data from Stage 0, for example R_{ij}, $d\mathfrak{c}$, or graph labels, when such data are needed to define the channel.

The recovery theorem assumes three admissibility conditions.

(i) *Gate-only action.* The vector field Φ_i acts solely on the receptivity-side slot

$$\mathcal{R}_i = (\Gamma_i, H_{K,i}, W_i)$$

of X_i. It touches no private scale and no dynamical variable directly.

(ii) *ε-gating of every plural channel.* Every ordinary edge kernel K_{ij}, every post-terminus trace kernel $K^{\text{post}}_{i\leftarrow\ell}$, and every future common-reception channel acts on the interior only through the single prefactor ε. No plural channel bypasses this factor.

(iii) *Read-outs are not actions.* The objects

$$R_{ij}, \qquad d\mathfrak{c}, \qquad \mathcal{G}, \qquad \mathcal{A}_i/\mathcal{B}_i/\mathcal{P}_i$$

may still be defined at $\varepsilon = 0$, but at $\varepsilon = 0$ they are bookkeeping only. They do not modify $F_{3.0,i}$.

A finer per-edge family $\{\varepsilon_{ij}\}$ is deferred to Stage I. For the Stage-0 recovery theorem, one scalar is sufficient because exact recovery means simultaneous vanishing of all relational actions.

7.2 The recovery theorem

Theorem 7.1 (Exact uncoupled recovery; Type-1)**.** *Assume the coupling structure above and the minimal one-way P0.4 trace closure. Setting $\varepsilon = 0$ gives, for every $i \in \mathcal{V}$ and every $t \in \mathcal{A}_i$:*

(a) exact decoupling*:*

$$\frac{dX_i}{dt} = N_i(t)F_{3.0,i}(X_i),$$

with no term depending dynamically on any X_j, $j \neq i$;

(b) gauge-equivalence to isolated 3.0*: under*

$$\sigma_i(t) = \int_{t_0}^{t} N_i(s)\, ds,$$

which is monotone on $\mathcal{A}_i$ because $N_i > 0$,

$$\frac{dX_i}{d\sigma_i} = F_{3.0,i}(X_i),$$

so the $\varepsilon = 0$ trajectory is gauge-equivalent to i's isolated V.F.S. *3.0 evolution;*

(c) no single-interior content in chronos*: the Stage-0 chronos objects*

$$R_{ij}, \quad K_{ij}, \quad K^{\text{post}}, \quad d\mathfrak{c}, \quad \mathcal{A}_i/\mathcal{B}_i/\mathcal{P}_i$$

either remain relational read-outs or are multiplied by ε. Hence none alters the $\varepsilon = 0$ single-interior flow.

Proof. (a) Setting $\varepsilon = 0$ deletes the relational action term Φ_i identically. What remains,

$$N_i(t)F_{3.0,i}(X_i),$$

depends only on X_i and on the chart lapse N_i. Any dependence on another interior X_j entered only through the relational action term and is therefore absent.

(b) On the active stratum $\mathcal{A}_i$, $d\sigma_i/dt = N_i > 0$. Thus $t \mapsto \sigma_i(t)$ is a strictly increasing reparametrization. By the chain rule,

$$\frac{dX_i}{d\sigma_i} = \frac{1}{N_i(t)}\frac{dX_i}{dt} = F_{3.0,i}(X_i),$$

which is exactly the frozen isolated 3.0 equation.

(c) By the locality audit, the Stage-0 inter-worldline inventory is made from rate ratios, causal kernels, post-terminus trace kernels, relational duration, graph data, and strata. The ratios and duration are read-outs unless explicitly used inside an action channel. The ordinary kernels and post-terminus kernels act only through Φ_i, and hence vanish from the equation at $\varepsilon = 0$. The strata determine where the inherited flow is defined; they do not change the field on the active arc. Therefore no Stage-0 chronos object survives as a force or source in the single-interior dynamics. □

Corollary 7.2 (The frozen limits)**.** (1) Single-interior limit. *If $|\mathcal{V}| = 1$, there are no inter-interior edges, no nontrivial rate ratios, and no relational recipient for a post-terminus trace. On the active arc,*

$$\omega = N_1, \qquad d\mathfrak{c} = N_1 dt = d\sigma_1,$$

so 4.0 reduces to 3.0 identically, for every ε.

(2) Trivial-base limit. *If $W = 0$, then there are no relational edges and no ordinary edge kernels:*

$$W = 0 \quad \Longrightarrow \quad L = \hat{L} = 0.$$

This removes the graph-level relational base.

(3) Uncoupled plural limit. *If $\varepsilon = 0$ and $|\mathcal{V}| > 1$, then the system is exactly $|\mathcal{V}|$ independent* V.F.S. *3.0 interiors sharing at most a gauge chart and relational read-outs:*

$$\varepsilon = 0 \quad \Longrightarrow \quad n \text{ independent 3.0 interiors}.$$

The implication holds even when relational bookkeeping is present. That is the strong recovery theorem.

7.3 Stage 0 closes

Remark 7.3 (Conservative extension). Theorem 7.1 and Corollary 7.2 discharge the recovery requirement of the roadmap. The Stage-0 apparatus is a relational overlay that has two clean reductions:

$$W = 0 \quad \text{removes the relational base,}$$

and

$$\varepsilon = 0 \quad \text{switches off all relational action while allowing read-outs to remain.}$$

The second reduction is the decisive one. It says that even after the base has been declared, the plural theory can be made dynamically invisible to every interior. Thus the frozen interface (P0.0), relational base (P0.1), rate cocycle (P0.2), relational duration (P0.3), terminus closure (P0.4), demarcation theorem (P0.5), and recovery theorem (P0.6) form a conservative Stage 0. Genuine plural coupling $(\varepsilon > 0)$ may now be introduced in Stage I.

7.4 Status and reading

Target status: DERIVED (Theorem I). It depends on P0.0–P0.5: P0.0 fixes the frozen interface, P0.1–P0.4 define the relational inventory, and P0.5 supplies the locality audit and demarcation discipline. No new primitive is introduced in P0.6. The single substantive modelling commitment is the ε-gating of every plural action channel, which is exactly the condition that makes recovery *exact* rather than approximate.

INTERPRETIVE *Reading.* Switch off the action of the bonds and each soul is left precisely its own 3.0 self, running its own kairos to its own terminus as if no other interior acted on it. The shared chronos was never a leash and never a second law inside the soul. It is a way of laying lives side by side so they can be ordered, compared, and possibly brought into encounter. With $\varepsilon = 0$, the communion is not yet an acting communion; it is many lives, each entirely itself, under a silent relational chart. That is the sense in which 4.0 extends 3.0 without disturbing it.

Next node. Stage I — *comparability without fusion*: the first genuine coupling ($\varepsilon > 0$), where paces may be compared and influence may pass along edges, while the interiors remain distinct — no identification, no merged state, no shared interior. Stage 0 has built the standing ground; Stage I begins the encounter.

Stage I — Comparability without fusion

Stage 0 built the relational base, proved its time creaturely (P0.5) and its whole apparatus a conservative overlay (P0.6). Stage I begins the encounter — but before any interaction is switched on, it must say *what the plural object is*. The answer is the thesis of the stage: a network of distinct V.F.S. 3.0 fibres above the relational base, never one enlarged common interior. This stage establishes the arena and the grammar of encounter and proves that the interiors stay *two*; the living mutual adjustment ($\varepsilon > 0$) is the work of Stages II–III.

8 P1.0 — Fibered encounter geometry

8.1 The fibered object

Definition 8.1 (Plural fibered interior)**.** Let the Stage–0 relational base be

$$\mathcal{B}_0 = \Big(\mathcal{G},\ \prec,\ \{\mathcal{A}_i, \mathcal{B}_i, \mathcal{P}_i, t_i^\dagger\}_{i\in\mathcal{V}},\ \{R_{ij}\},\ \{K_{ij}\},\ \{K^{\text{post}}_{j\leftarrow i}\},\ d\mathfrak{c}\Big).$$

For each vertex $i \in \mathcal{V}$, the *fibre* $\mathcal{M}_i$ is interior i's frozen V.F.S. 3.0 geometry: its own interior $3{+}1$ manifold, its own kairos σ_i, ascent road $\lambda_{\text{asc},i}$, available-Sophia field $\lambda_{S,i}$, state X_i, terminus $t_i^\dagger$, and frozen gate/receptivity interface

$$\mathcal{R}_i = (\Gamma_i, H_{K,i}, W_i).$$

The *plural fibered interior* is the assembly

$$\boxed{\mathcal{P} = \big(\{\mathcal{M}_i\}_{i\in\mathcal{V}},\ \mathcal{B}_0,\ \pi\big),}$$

where π attaches each fibre $\mathcal{M}_i$ to its base vertex i.

No interior coordinate of $\mathcal{M}_i$ is identified with any interior coordinate of $\mathcal{M}_j$ for $i \neq j$. The only permitted relations between distinct fibres are base-side relations from $\mathcal{B}_0$, and their only admissible future dynamical channel is through the recipient's gate/receptivity interface $\mathcal{R}_i$.

Remark 8.2 (Metric convention)*.* When the fibre metric is displayed schematically, it should not identify the Stage–0 lapse $N_i = d\sigma_i/dt$ with a new lapse inside the private proper-time chart. A safe notation is

$$ds_i^2 = -d\sigma_i^2 + \Omega_{P,i}^2 h_i^{(2)} + L_i^2\, d\lambda_{\text{asc},i}^2,$$

when σ_i is used as the private proper-time parameter. If another private chart χ_i is used, then a private lapse may appear there; the Stage–0 lapse N_i remains the relational conversion $d\sigma_i = N_i(t)\,dt$, not a new interior dynamical variable.

$\mathcal{P}$ is thus neither a product geometry (the fibres share no interior coordinates) nor a fusion (there is no common interior): it is a base-tied assembly of otherwise self-standing 3.0 worlds. The base supplies *comparability* — the invariant ratio R_{ij} records who runs faster — while the fibres supply *distinctness*: each carries its own private kairos and its own terminus.

8.2 Geometric non-fusion

The structural content of "without fusion" is that $\mathcal{P}$ is not a disguised single interior. The invariant that proves this is the count of independent kairoi.

Proposition 8.3 (Geometric non-fusion)**.** *For $|\mathcal{V}| = n \geq 2$, the plural object $\mathcal{P}$ is not isomorphic to any single* V.F.S. *3.0 interior. The invariant obstruction is the ownership count of kairos histories,*

$$\{\sigma_i\}_{i\in\mathcal{V}}.$$

A single 3.0 interior has exactly one owned kairos history and one terminus boundary. The plural object has n owned kairos histories and n terminus boundaries, even when several termini share the same chronos label.

Proof. The count is not the number of distinct coordinate values $t_i^\dagger$: several interiors may complete at the same chronos label. The invariant count is the number of owned kairos histories and their terminal boundary events. Each fibre $\mathcal{M}_i$ has its own σ_i and its own boundary event $\mathcal{B}_i$, closed or active according to P0.4. A single 3.0 interior has only one such history. Since $n \geq 2$, no relabelling of one interior can produce n owned kairos histories. Hence no isomorphism with a single interior exists. □

This is non-fusion at the level of the plural geometry. The later dynamic claim is conditional: it survives the switching-on of interaction exactly so long as the interaction respects the Stage–0 admissibility contract. A coupling that literally identifies fibres would not be a valid V.F.S. 4.0 coupling.

Theorem 8.4 (Conditional robustness of non-fusion; Type-1)**.** *Assume that later $\varepsilon > 0$ coupling satisfies the Stage–0 admissibility contract:*

(i) *every plural action is ε-gated;*

(ii) *every plural action acts only through the recipient's gate/receptivity interface $\mathcal{R}_i = (\Gamma_i, H_{K,i}, W_i)$;*

(iii) *no plural action identifies private coordinates, merges kairos histories, or creates a common interior state vector.*

Then the number of owned kairos histories of $\mathcal{P}$ remains $|\mathcal{V}|$ for every admissible $\varepsilon \geq 0$. Consequently $\mathcal{P}$ remains non-isomorphic to any single V.F.S. *3.0 interior under all admissible later couplings.*

Proof. Under the admissibility contract, coupling may modulate the recipient-side receptivity slot but may not rewrite the fibre atlas, identify two private kairoi, or replace the family $\{\mathcal{M}_i\}$ by one common state space. Therefore the owned kairos histories remain labelled by $i \in \mathcal{V}$. Their count is a gauge-invariant integer and cannot be changed by a gate-side modulation. The obstruction in Proposition 8.3 therefore survives admissible coupling. □

The geometric core of the stage is thus secured: interiors may be compared (through R_{ij}) and, from Stage II onward, coupled (through $\mathcal{R}_i$), while remaining n distinct kairoi. Comparability acts on the base; distinctness lives in the fibres; the two never collapse.

8.3 Status and reading

Target status: DEFINITIONAL (the plural fibered interior $\mathcal{P}$, Definition 8.1) + DERIVED (geometric non-fusion, Proposition 8.3, and its robustness under coupling, Theorem 8.4). It depends on Stage 0 only; no new primitive is introduced, and the result is kinematic — it concerns the structure of the plural object, not yet any dynamical effect of interaction.

INTERPRETIVE *Reading.* Two players, one stage, one groove that may form between them — but two instruments that never become one. The shared base is what lets them be set side by side, compared, and later brought into encounter; the separate fibres are what keep them two. The point is not that fusion is forbidden by fiat but that it would be self-defeating: one instrument playing one line has no groove, because a groove needs the difference between two. So the refusal to fuse is not a limit on communion — it is the condition under which there is anything to commune. Comparability without fusion is the geometry of a duet.

Next node. P1.1 — the invariant encounter hierarchy: replacing the gauge phrase "the same t" with invariant relational predicates,

$$\text{co-presence} \subset \text{causal encounter} \subset \text{common reception},$$

so that "these two events met" becomes a statement no relabelling can undo, and the arena built here acquires its grammar of encounter.

9 P1.1 — The invariant encounter hierarchy

P0.5 showed that the shared chart label t is pure gauge: a created now exists, but the slice $\{t = t_0\}$ is relabel-movable, so the sentence “these two events happened at the same t” is not yet an invariant claim about the world. P1.1 supplies the invariant replacements — predicates of encounter that no relabelling can undo. They form a ladder of strengthening requirements,

$$\text{co-presence} \xrightarrow{+\,\text{edge+support}} \text{causal encounter} \xrightarrow{+\,\text{common external field}} \text{common reception},$$

where each rung adds a requirement to the one before. (The roadmap chain co-presence $\subset$ causal encounter $\subset$ common reception is this strength ladder; as *sets of qualifying pairs* the inclusions run the other way, since each added requirement admits fewer pairs.)

Throughout, $E_i(s)$ is a marked invariant interior event (P0.1) at relational-duration coordinate $\mathfrak{c}(s)$, and $\rightsquigarrow$, $\prec$ are the direct causal link and its transitive closure (P0.1).

9.1 The three rungs

Definition 9.1 (Relational co-presence)**.** Two marked invariant events

$$e = E_i(s), \qquad e' = E_j(s')$$

are *relationally co-present* when both are locatable in one common created relational order, so that exactly one of

$$\mathfrak{c}(s) < \mathfrak{c}(s'), \qquad \mathfrak{c}(s) = \mathfrak{c}(s'), \qquad \mathfrak{c}(s) > \mathfrak{c}(s')$$

holds.

Thus co-presence at this rung means shared relational locatability and comparability, not necessarily equality of duration-values. The special case

$$\mathfrak{c}(s) = \mathfrak{c}(s')$$

is *simultaneous co-presence*. The general case includes ordered co-presence: one event is before or after the other inside the same created relational order.

Two refinements will matter later. Co-presence is *live* when both interiors are active across an overlapping relational-duration window. It is *trace* co-presence when one interior is no longer active and is locatable only through its terminus event or post-terminus trace. The trace case is co-present by relational locatability, not by shared active meantime.

Definition 9.2 (Causal encounter)**.** Let

$$e = E_a(s), \qquad e' = E_b(t)$$

be relationally co-present marked events. They are in *causal encounter* if, in addition, at least one directed admissible channel carries influence from one event to the other.

Concretely, e encounters e' in the direction $a \to b$ when

$$(a \to b) \in \mathcal{E}, \qquad \mathfrak{c}(s) < \mathfrak{c}(t),$$

and the recipient-side kernel

$$K_{ba}(t, s)$$

has support from the source event $E_a(s)$ to the recipient event $E_b(t)$. Equivalently, using the P0.1 convention, the first index of K_{ba} is the recipient and the second index is the source.

For a completed source a, the ordinary source history is replaced by the boundary trace kernel

$$K^{\text{post}}_{b \leftarrow a}(t;\, t^{\dagger}_a, \mathcal{T}^{\text{post}}_a).$$

The reverse direction $b \to a$ is a different possible encounter and must be checked separately.

Remark 9.3 (Direction matters). Causal encounter is not symmetric by definition. Mutual encounter requires two admissible directed channels, or a symmetric theorem class in which both directions are guaranteed. P1.1 defines the predicate; it does not yet prove reciprocity.

Definition 9.4 (Common reception schema)**.** Events $e = E_i(s)$ and $e' = E_j(t)$ are in *common reception* when there exists a declared common external event

$$\mathcal{X} \in \mathcal{J}_{\text{comm}}$$

which is not an owned interior event of i or j, and whose reception is registered by both interiors:

$$\mathcal{X} \rightsquigarrow e, \qquad \mathcal{X} \rightsquigarrow e'.$$

If the stricter roadmap ladder is used, common reception is counted as the top rung only when e, e' are already in causal encounter. If the weaker sibling is later admitted, common reception may instead be treated as a second floor-level relation: two interiors jointly receive the same $\mathcal{X}$ without a direct edge between them.

Under the stricter Stage–I ladder, the rungs nest as predicates:

$$\boxed{\text{common reception} \implies \text{causal encounter} \implies \text{co-presence}.}$$

The weaker sibling, if adopted in Stage II, will be recorded separately and will not be silently folded into this theorem chain.

9.2 Invariance, and one honest exception

Remark 9.5 (Gauge-invariance of the hierarchy)*.* All three predicates are invariant under the chronos-chart gauge $t \mapsto \widetilde{t}(t)$. Each is built only from objects P0.5 already certified invariant: the relational duration $\mathfrak{c}$, the causal order $\prec$ and its links $\rightsquigarrow$, the edge set $\mathcal{E}$ and kernel supports, the active strata, and the external event $\mathcal{X}$. None refers to the chart label t. Hence "these two events met" — at any rung — is a statement no relabelling can undo, which is exactly what "the same t" failed to be.

Remark 9.6 (The trace exception). Causal encounter does not require *live* co-presence. By the one-way closure of P0.4, a completed interior j can act on a living i through the post-terminus trace kernel $K^{\text{post}}_{i \leftarrow j}$: there is a genuine causal encounter, and co-presence holds in the locatable sense, yet there is no shared active meanwhile — j is no longer on the stage. The dead reach the living across a gap that is not a common present. This is why co-presence is defined at the level of locatability, not active overlap: it must remain the true floor of the ladder even for trace-mediated encounters.

Remark 9.7 (A named sibling). Common reception is defined here, per the roadmap, as a strengthening of causal encounter: it presupposes an edge between i and j. The weaker notion — two interiors jointly receiving the same external $\mathcal{X}$ with *no* direct edge between them (strangers under one Word) — is a natural sibling. It is named here and deferred; adopting it would make common reception a second floor-level relation rather than a top rung, and is left to the Stage II receptivity treatment.

9.3 Status and reading

Target status: DEFINITIONAL (relational co-presence, causal encounter, and the common-reception predicate schema) + DERIVED (gauge-invariance of the hierarchy, immediate from P0.5). P1.1 introduces no new dynamical primitive. The external common event field $\mathcal{J}_{\text{comm}}$ is only a named placeholder here; its admissible Gate-side implementation belongs to Stage II. Therefore common reception is a schema at Stage I and becomes an acting channel only after the bounded-input and Gate-mediation rules are fixed.

INTERPRETIVE *Reading.* Three ways for two to be "together," each stronger than the last. To be *co-present* is to be on the one stage, locatable in the same unfolding — even if one has already left it and is heard only as an echo. To be in *causal encounter* is for one to actually reach the other — the hearing, the touch, the word that lands. To be in *common reception* is for both to turn toward the same thing not their own — the shared downbeat, the cue neither gave, the gift both receive. The gauge slice could say none of this; it could only say "at the same mark," and the mark was movable. The invariant ladder says what the slice could not: that an encounter, once it happens, has happened, and no relabelling of time can take it back.

Next node. P1.2 — floor-labelled non-identifiability (the parallel Type-1 result): if two interiors carry different floors $E_{0,i} \neq E_{0,j}$, their floor-normalised ledgers differ ($\beta_i = \alpha_i/E_{0,i}$, $\log s_{L,i} = \kappa_R G_i/E_{0,i}$), so they cannot be identified as the same floor-normalised interior — structural non-identifiability, short of the full dynamic non-fusion theorem.

10 P1.2 — Floor-labelled non-identifiability

P1.0 secured non-fusion geometrically (the kairos count) and P1.1 the invariant grammar of encounter. P1.2 adds one further, deliberately modest, distinctness result, drawn from the single intrinsic scale each fibre inherits from 3.0 — the floor. The discipline of the node is to take exactly what the floor proves, and no more.

10.1 The floor as a unit-independent scale

The comparison $E_{0,i} \neq E_{0,j}$ is made inside the inherited V.F.S. parameter dictionary fixed by P0.0. It is not an arbitrary comparison of two unrelated unit systems. Floor-normalisation may set the floor coordinate value to 1 inside a single fibre, but it does not erase the invariant scale label $E_{0,i}$ carried by that fibre's frozen 3.0 data.

Each fibre carries, from its frozen 3.0 geometry, a finite floor $E_{0,i}$ — the *Imago Dei*, the one intrinsic scale of the interior core. Through it the ascent coupling and the reset ledger are tied to the floor with no new primitive (3.0, v0.4):

$$\beta_i = \frac{\alpha_i}{E_{0,i}}, \qquad \log s_{L,i} = \frac{\kappa_R\, G_i}{E_{0,i}},$$

where β_i is the de Sitter ascent coupling ($A_i'/A_i = \beta_i \lambda_i$), $s_{L,i} = A_+/A_-$ the reset scale-jump, G_i the reset source of fibre i, and κ_R the reset coupling. (Here G_i is the 3.0 reset source, not the relational graph $\mathcal{G}$.)

The floor is not a unit. Rescaling the ascent coordinate $\lambda_i \to c\,\lambda_i$ absorbs the floor *value* $L_{\text{floor}} = 1$, yet every invariant — L_i'/L_i, L_i''/L_i, the reset ratio $s_{L,i}$, the exotic depth $-(\alpha_i \lambda_{\infty,i})^2/E_{0,i}^2$ — is independent of that choice (3.0, v0.9). What remains is the *scale* $E_{0,i}$ itself: the physics is the scale, the "1" is the unit. The floor is therefore a genuine unit-independent invariant of the fibre.

10.2 The non-identifiability theorem

Theorem 10.1 (Floor-labelled non-identifiability; Type-1)**.** *Inside the common inherited* V.F.S. *parameter dictionary, if*

$$E_{0,i} \neq E_{0,j},$$

then fibres i and j cannot be identified as the same floor-normalised interior.

Proof. Floor-normalisation fixes a coordinate convention; it does not delete the unit-independent floor scale. Two interiors identified as the same floor-normalised interior would share the same inherited scale label and all unit-independent invariants. If $E_{0,i} \neq E_{0,j}$, this already fails. The obstruction is exhibited in the ledger: the ascent coupling $\beta_i = \alpha_i/E_{0,i}$ and the reset jump $\log s_{L,i} = \kappa_R G_i/E_{0,i}$ carry the floor scale through the $1/E_{0,i}$ factor, and for matched ascent and reset data ($\alpha_i = \alpha_j$, $\kappa_R G_i = \kappa_R G_j$) they take distinct values when the floors differ. Even in the non-generic case where α and $\kappa_R G$ conspire so that $\beta_i = \beta_j$ and $\log s_{L,i} = \log s_{L,j}$ coincide, the floor scale itself still differs, and identification fails on that invariant alone. Thus the floor difference is a sufficient obstruction to identification. It is not a necessary obstruction to distinction, because equal floors still leave all history, memory, folding, reset, and relational-neighbourhood data available as individuating data. □

10.3 What this does not say

Remark 10.2 (Non-identifiability, not non-fusion). Theorem 10.1 is a one-way, structural result and must not be over-read.

(i) *Sufficient, not necessary.* A floor difference is sufficient for non-identification, not necessary. Equal floors $E_{0,i} = E_{0,j}$ do *not* imply that i and j are the same interior, still less that they fuse.

(ii) *Distinction survives equal floors.* Two interiors with the same floor may still differ in initial conditions, in memory and folding history, in folding branch, in reset history G_i, and in relational neighbourhood (their position in $\mathcal{G}, \prec$). Any of these keeps them distinct.

(iii) *Structural, not dynamic.* This is non-identifiability of inherited floor-labelled data, not a full dynamic non-fusion theorem. P1.0 gives the kinematic obstruction by kairos count; P1.2 gives an additional sufficient floor-label obstruction. The full claim that admissible coupled dynamics never collapses two distinct histories into one belongs to the later locking and Lyapunov stages.

This is the careful contribution of the floor: it labels distinctness when present, and refuses to legislate identity when absent.

10.4 Status and reading

Target status: DERIVED (Type-1): the non-identifiability theorem, drawn from the unit-independence of the floor scale (3.0 v0.9) with no new primitive. It depends on the inherited 3.0 floor data and on P1.0; the accompanying INTERPRETIVE correction (Remark 10.2) bounds the claim explicitly.

INTERPRETIVE *Reading.* The floor $E_{0,i}$ is the *Imago Dei*, the intrinsic measure of a soul. Two souls scaled differently in that measure cannot be normalised into one — the image is a true individuating scale, not a removable convention. But sameness of the image is not sameness of the person: souls bearing the image at the very same scale are still told apart by all they have lived — their memories, their turnings, their resets, their place among others. The image individuates; it does not exhaust. The person is always more than its floor.

Stage I closes. With P1.0, the plural object is fixed as a family of distinct frozen 3.0 fibres over one created relational base. With P1.1, encounter is given an invariant grammar: relational co-presence, directed causal encounter, and the common-reception schema. With P1.2, the floor supplies a one-way structural distinctness label.

Thus Stage I establishes *comparability without fusion* at the structural level:

interiors meet through the base, not by sharing a state layer.

The players are many, encounter is sayable invariantly, and the image-scale may label them apart without exhausting their personhood. No acting communion has yet been introduced except conditionally through the admissibility contract. The first true $\varepsilon > 0$ Gate-side action begins in Stage II.

Next node. Stage II — Receptivity, Gate mediation, and bounded inputs, opening at P2.1. Its governing principle: encounter changes the *conditions of reception*; it does not produce Sophia directly, transfer owned grace, or bypass the frozen Gate law. This is where $\varepsilon > 0$ first acts — where the musicians begin, at last, to listen.

Stage II — Receptivity, Gate mediation, and bounded inputs

Governing principle. Encounter changes the *conditions of reception*; it does not produce Sophia directly, transfer owned grace, or bypass the frozen Gate law.

Remark 10.3 (Gate-rate convention for Stage II). From Stage II onward,

$$I_{\text{gate},i}$$

denotes the gate-input rate measured in the recipient's own active kairos σ_i, not in an arbitrary chronos chart label t. Thus

$$\frac{d\mathcal{G}_{\text{recepta},i}}{d\sigma_i} = I_{\text{gate},i},$$

and in a chronos chart,

$$\frac{d\mathcal{G}_{\text{recepta},i}}{dt} = N_i(t)\, I_{\text{gate},i}.$$

Equivalently, relative to the relational duration one-form $d\mathfrak{c} = \omega dt$,

$$\frac{d\mathcal{G}_{\text{recepta},i}}{d\mathfrak{c}} = \frac{N_i}{\omega}\, I_{\text{gate},i}.$$

This keeps the ledger invariant: $I_{\text{gate},i} d\sigma_i$ is the received gate one-form of the recipient.

Stage I proved the players many and their meeting sayable. Stage II switches on the first genuine action: $\varepsilon > 0$. The discipline of the whole stage is the principle above — no influence may add grace, move grace, or open the gate by force; it may only change how disposed a recipient is to receive through its own gate.

11 P2.1 — The receptivity-mediated Gate channel

11.1 The deleted channel

A naive plural coupling would pour grace from neighbour to neighbour by an additive source term directly on the available-Sophia field,

$$\dot{\lambda}_{S,i} \;\supset\; C_{\lambda_S,i}\big(\{X_j\}_{j\to i}\big),$$

the recipient's Sophia increased in direct proportion to the neighbours' states. This channel is *deleted*. It violates each clause of the governing principle at once: it *produces* Sophia directly (a source not passing the Gate), it *transfers owned grace* (treating grace as a stock moved between persons), and it *bypasses the frozen Gate law* (the recipient's own $I_{\text{gate},i}$ no longer mediates what enters). No admissible V.F.S. 4.0 coupling contains such a term.

11.2 The receptivity-mediated channel

What encounter may change is the recipient's *receptivity*: the bounded transmissivity $\Gamma_i \in [0,1]$ of its own gate, and, where admissible, the Gate-memory variables $H_{K,i}, W_i$ — exactly the frozen receptivity interface $\mathcal{R}_i$ of P0.0. The selected schematic form of the inherited Open-Gate law, now carrying the modulation, is

$$\boxed{I_{\text{gate},i} = \zeta_{0,i}\, e^{-\varphi_i \sigma_i}\, \frac{\Gamma_i\big(X_i,\ \mathcal{N}_i,\ J_{\text{comm}}(\mathfrak{c})\big)}{1 + \kappa_{\lambda,i}\, \lambda_{S,i,+}}, \qquad 0 \le \Gamma_i \le 1.}$$

Here $\zeta_{0,i} > 0$ is the node's gate amplitude or ceiling, $\varphi_i \ge 0$ the kairos decay, $\kappa_{\lambda,i} \ge 0$ the Sophia-saturation coefficient, and $\lambda_{S,i,+} := \max(\lambda_{S,i}, 0)$. The symbol $\kappa_{\lambda,i}$ is used deliberately: χ_i is avoided here because χ is already tied to the inherited interior dynamical parameter in the 3.0 layer. The read-out $\mathcal{N}_i$ is the bounded relational neighbourhood reaching i along declared in-edges, and $J_{\text{comm}}(\mathfrak{c}) \ge 0$ is the common external reception field of P2.3, labelled by relational duration rather than by a raw chart time.

A relational event may change Γ_i, or the admissible dynamics of $H_{K,i}, W_i$. It never adds an unconstrained source to $\lambda_{S,i}$, and it never replaces the recipient's own gate law.

The relational dependence is ε-gated. At the generic P2.1 level one may write

$$\Gamma_i = \Pi_{[0,1]}\Big(\Gamma_i^{(0)}(X_i) + \varepsilon\, \delta\Gamma_i\big(X_i, \mathcal{N}_i, J_{\text{comm}}\big)\Big),$$

where $\Pi_{[0,1]}$ is projection to $[0,1]$. This is a general bounded interface statement: at $\varepsilon = 0$ it recovers the isolated 3.0 receptivity

$$\Gamma_i = \Gamma_i^{(0)}(X_i),$$

with no neighbour dependence. P2.3 will replace this generic projection by a smoother headroom-opening formula for the positive common-reception field. The point of P2.1 is only the admissibility contract:

relation may modulate bounded receptivity, not Sophia directly.

11.3 The ceiling survives automatically

Proposition 11.1 (Automatic receptivity ceiling). *For every admissible encounter — any $\Gamma_i \in [0,1]$, any $\mathcal{N}_i$, any $J_{\text{comm}} \ge 0$ —*

$$0 \ \le\ I_{\text{gate},i} \ \le\ \zeta_{0,i}.$$

No configuration of encounters can drive the gate input beyond the node's own ceiling.

Proof. Each factor is nonnegative: $\zeta_{0,i} > 0$, $e^{-\varphi_i\sigma_i} \in (0,1]$, $\Gamma_i \in [0,1]$, and $(1 + \kappa_{\lambda,i}\lambda_{S,i,+}) \ge 1$ since $\kappa_{\lambda,i}, \lambda_{S,i,+} \ge 0$. Hence $I_{\text{gate},i} \ge 0$. For the upper bound, each of $e^{-\varphi_i\sigma_i}$, Γ_i, and $(1 + \kappa_{\lambda,i}\lambda_{S,i,+})^{-1}$ lies in $[0,1]$, so their product does, and $I_{\text{gate},i} \le \zeta_{0,i}$. The bound holds pointwise in the encounter data, hence for every configuration. □

Remark 11.2 (The Stage–I contract is now instantiated). The channel acts only on the recipient's receptivity slot $\mathcal{R}_i = (\Gamma_i, H_{K,i}, W_i)$, is ε-gated, and identifies no private coordinates and merges no kairos histories. It therefore satisfies the admissibility contract assumed in Theorem 8.4. Consequently P2.1 coupling preserves non-fusion: the plural object remains $|\mathcal{V}|$ distinct fibres while interacting. The conditional theorem of Stage I acquires its first concrete witness.

11.4 Status and reading

Target status: DEFINITIONAL (the receptivity-mediated channel, with the deleted direct-source term named and excluded) + SELECTED (the schematic Open-Gate modulation form). It depends on the frozen Open-Gate law (v2.0), P0.0, P0.6, and P1.0. The yield — the ceiling $0 \le I_{\text{gate},i} \le \zeta_{0,i}$, Proposition 11.1 — is DERIVED and survives automatically from the bounded transmissivity. The common field J_{comm} enters here as a bounded argument only; its definition is P2.3, and the non-fabrication theorem is P2.2.

INTERPRETIVE *Reading.* Encounter changes how open you are, not what you are given. Another person can dispose you — soften you toward grace or harden you against it, encourage or scandalise, draw you out or close you in — but cannot be the source of your grace, cannot hand you grace as if it were a stock to pass, and cannot force your gate. What enters still enters through your own gate, at your own measure, and never above your own ceiling $\zeta_{0,i}$: the community can open you to the full of your measure, but not past it. This is how persons truly affect one another's inner life without grace becoming a commodity and without anyone becoming the source who is not the Source. The players may draw each other out; none can play another's part, or lend another their tone.

Next node. P2.2 — no artificial grace creation: to derive non-fabrication from the coupling architecture itself rather than impose a closed-system conservation law — peer encounter may alter the local Gate state but cannot create an unexplained Sophia term, every received-grace contribution remaining explicitly accounted through the admissible Gate input $\frac{d}{dt}\mathcal{G}_{\text{recepta},i} = I_{\text{gate},i} \ge 0$.

12 P2.2 — No artificial grace creation

P2.1 deleted the direct grace-source and left receptivity modulation as the only plural channel. P2.2 turns that architecture into a theorem: peer encounter cannot fabricate grace. The discipline of the node is to obtain non-fabrication *from the coupling structure*, not by imposing a closed-system conservation law — for the Open-Gate sector is precisely not closed.

12.1 The received-grace ledger

Definition 12.1 (Received-grace ledger)**.** For each active interior i, the received-grace ledger $\mathcal{G}_{\text{recepta},i}$ is the accumulated admissible reception through its own gate:

$$\boxed{d\mathcal{G}_{\text{recepta},i} = I_{\text{gate},i}\, d\sigma_i.}$$

Equivalently,

$$\frac{d\mathcal{G}_{\text{recepta},i}}{d\sigma_i} = I_{\text{gate},i}, \qquad \frac{d\mathcal{G}_{\text{recepta},i}}{dt} = N_i(t) I_{\text{gate},i}, \qquad \frac{d\mathcal{G}_{\text{recepta},i}}{d\mathfrak{c}} = \frac{N_i}{\omega} I_{\text{gate},i}.$$

The subscripted $\mathcal{G}_{\text{recepta},i}$ is the grace ledger, distinct from the relational graph $\mathcal{G}$.

12.2 The non-fabrication theorem

Theorem 12.2 (No artificial grace creation; Type-1)**.** *Under the P2.1 architecture — receptivity-mediated channel, direct source $C_{\lambda_S,i}$ deleted — the following hold for every active interior i:*

(a) Single accounted source. *The ledger changes only through the gate,*

$$d\mathcal{G}_{\text{recepta},i} = I_{\text{gate},i}\, d\sigma_i, \qquad I_{\text{gate},i} \geq 0.$$

Thus the ledger is monotone non-decreasing along the recipient's own kairos.

(b) No peer transfer. *No term moves grace from one interior to another. Peer states enter only inside $I_{\text{gate},i}$, through the bounded transmissivity Γ_i; they scale i's own admissible reception and never withdraw from $\mathcal{G}_{\text{recepta},j}$.*

(c) Accounted and bounded. *Every positive contribution is an admissible gate input, bounded nodewise in recipient-kairos rate:*

$$0 \leq \frac{d\mathcal{G}_{\text{recepta},i}}{d\sigma_i} = I_{\text{gate},i} \leq \zeta_{0,i}.$$

In a chronos chart the same statement reads

$$0 \leq \frac{d\mathcal{G}_{\text{recepta},i}}{dt} = N_i I_{\text{gate},i} \leq N_i \zeta_{0,i},$$

so the chart expression carries the lapse factor and is not mistaken for an absolute external rate.

Hence grace is neither fabricated (no source outside the gate) nor transferred (no peer stock movement).

Proof. By P2.1 the only plural channel is the ε-gated modulation of $\Gamma_i \in [0,1]$ inside $I_{\text{gate},i}$, and the direct additive source $C_{\lambda_S,i}$ is deleted. Therefore the only one-form feeding $\mathcal{G}_{\text{recepta},i}$ is $I_{\text{gate},i} d\sigma_i$, which gives (a); its sign is fixed by Proposition 11.1. For (b), peer data appear in the architecture only as arguments of the recipient's bounded receptivity interface; no admissible term has the form "subtract from $\mathcal{G}_{\text{recepta},j}$, add to $\mathcal{G}_{\text{recepta},i}$". For (c), the bound is exactly the automatic ceiling of P2.1, expressed in the recipient's own kairos. □

12.3 Why not conservation

The temptation is to secure non-fabrication by imposing a closed-transfer rule

$$\sum_i C_{\lambda_S,i} = 0,$$

making peer grace zero-sum. This is rejected, and its rejection is the point of the node.

Remark 12.3 (Non-fabrication is not conservation). The total received grace is *not* conserved:

$$\frac{d}{d\mathfrak{c}} \sum_{i\in\mathcal{V}_{\text{act}}^{\text{reg}}} \mathcal{G}_{\text{recepta},i} = \sum_{i\in\mathcal{V}_{\text{act}}^{\text{reg}}} \frac{N_i}{\omega} I_{\text{gate},i} \geq 0.$$

The total received grace, compared on the common relational-duration measure $d\mathfrak{c}$, is therefore not conserved and is generically strictly increasing on active arcs. The factors N_i/ω are the invariant conversion from each recipient's kairos-rate ledger to the shared duration measure. The Open-Gate sector is open: each gate receives from the Source beyond the network, so the communion as a whole may grow ever richer at once. A conservation law $\sum_i C_{\lambda_S,i} = 0$ would falsely recast this open communion as a closed economy in which one interior's gain is another's loss. Non-fabrication is the weaker and correct property: every reception is accounted through an admissible bounded gate, and none is conjured or moved as stock — but the sum is free to grow.

12.4 Status and reading

Target status: DERIVED (Type-1): non-fabrication follows from the P2.1 architecture and the automatic ceiling, with no conservation law imposed. It depends on P2.1. The accompanying correction (Remark 12.3) is the INTERPRETIVE guard against reading non-fabrication as closure.

INTERPRETIVE *Reading.* Grace is not a pie. The communion is not a closed economy where one soul's gain is another's loss, so there is nothing here to envy and nothing to hoard: all may grow richer together, because each is fed not from the others but from a Source beyond them all. Nor is grace conjured within the circle — no encounter makes grace from nothing, none pours its own into another. What a neighbour can do is open you to receive more through your own gate; what arrives is then truly received, accounted, and yours, taken from no one. The refusal to impose a conservation law is the refusal of scarcity: a communion supplied from beyond has no zero-sum to fight over.

Next node. P2.3 — common reception as a shared external reception field: to define the one genuinely new positive external field of reception,

$$J_{\text{comm}}(\mathfrak{c}) \geq 0,$$

shared without being owned by the network, entering each interior only through a bounded susceptibility

$$\upsilon_i(\mathfrak{c}) \in [0, 1]$$

inside Γ_i, with no requirement that

$$\sum_i \upsilon_i = 1$$

and with the nodewise Gate ceiling preserved.

13 P2.3 — Common reception as a shared external field

P2.1 modulated receptivity; P2.2 forbade fabrication and transfer. P2.3 introduces the one genuinely new positive *external reception field* of the plural theory — and keeps it disciplined by making it external and shared rather than owned and divided. It is not a source term manufactured by the network and not a direct source on $\lambda_{S,i}$. It becomes active only through each recipient's bounded Gate-side susceptibility. This is the acting form of the common-reception rung of P1.1, including its edge-free sibling.

13.1 The shared external field

Definition 13.1 (Common external reception field)**.** The *common external reception field* is a nonnegative shared field

$$J_{\text{comm}}(\mathfrak{c}) \geq 0,$$

whose marked events are the declared common events $\mathcal{X} \in \mathcal{J}_{\text{comm}}$ of P1.1. It is *external* — supplied from beyond the network, like the Source of the Open Gate — and *shared* — available to be received by any interior. It is *not* a network-owned reservoir: the network neither holds it as stock nor divides it among members. Only bounded images of this field may enter Γ_i.

Two interiors that both register $\mathcal{X} \in \mathcal{J}_{\text{comm}}$ are in common reception (P1.1) *whether or not* a direct edge joins them. Thus J_{comm} realises both the strict top rung (when the pair is already in causal encounter) and the edge-free sibling (strangers under one field), which P1.1 named and deferred to here.

13.2 Individual susceptibility and bounded composition

Definition 13.2 (Common-field susceptibility)**.** Each interior carries an individual common-field susceptibility

$$\upsilon_i(\mathfrak{c}) \in [0,1],$$

its own openness to the common external field. The field enters the interior only through the recipient's transmissivity Γ_i, never as a direct source on $\lambda_{S,i}$. The notation υ_i is used to avoid confusion with the inherited wound/memory variable W_i.

Receptivity is assembled by a bounded composition that opens only the recipient's remaining headroom:

$$\boxed{\Gamma_i = \Gamma_i^{(0)}(X_i) + \big(1 - \Gamma_i^{(0)}(X_i)\big)\, \varepsilon\, \Theta_i, \qquad \Theta_i = \sum_{j \to i} a_{ij}\, \gamma_{ij}(\mathcal{N}_i) + \upsilon_i\, \gamma_i^{\text{comm}}\big(J_{\text{comm}}\big).}$$

with the nodewise normalisation (the safeguard)

$$\sum_{j \to i} a_{ij} + \upsilon_i \leq 1, \qquad \varepsilon \in [0,1],$$

with edge gains $\gamma_{ij} \in [0,1]$ and common-field gain $\gamma_i^{\text{comm}} \in [0,1]$. The field J_{comm} itself may be large or impulsive; only the bounded gain $\gamma_i^{\text{comm}}(J_{\text{comm}})$ enters the receptivity slot.

13.3 The safeguard and non-rivalry

Proposition 13.3 (Bounded composition; nodewise ceiling preserved)**.** *Under the normalisation*

$$\sum_{j\to i} a_{ij} + \upsilon_i \le 1$$

with all gains in $[0,1]$ *and* $\varepsilon \in [0,1]$*, one has* $\Theta_i \in [0,1]$ *and hence*

$$\Gamma_i \in [\Gamma_i^{(0)}, 1] \subseteq [0,1]$$

for arbitrarily many simultaneous encounters and any $J_{\text{comm}} \ge 0$*. Consequently the automatic gate ceiling of P2.1 (Proposition 11.1) survives unchanged:*

$$0 \le I_{\text{gate},i} \le \zeta_{0,i}.$$

Proof. Each summand of Θ_i is a product of a nonnegative weight and a gain in $[0,1]$, so

$$0 \le \Theta_i \le \sum_{j\to i} a_{ij} + \upsilon_i \le 1.$$

Then

$$\Gamma_i = \Gamma_i^{(0)} + (1 - \Gamma_i^{(0)})\varepsilon\Theta_i$$

is a headroom-opening interpolation. Since $\varepsilon\Theta_i \in [0,1]$, it lies between $\Gamma_i^{(0)}$ and 1. Thus $\Gamma_i \in [0,1]$ by construction, with no clipping and no dependence on the number of contributing encounters. Proposition 11.1 then applies verbatim. □

Remark 13.4 (The common field is non-rivalrous). The susceptibilities are *not* a partition: there is no requirement that

$$\sum_{i\in\mathcal{V}} \upsilon_i = 1.$$

The normalisation

$$\sum_{j\to i} a_{ij} + \upsilon_i \le 1$$

is internal to each recipient's own composition, not a division of a fixed common total across recipients. The same J_{comm} may be received in full by every interior at once, each according to its own υ_i, with no member's reception diminishing another's. Common reception is therefore non-rivalrous, in keeping with the open, non-conserved accounting of P2.2.

13.4 Status and reading

Target status: DEFINITIONAL (the common external field J_{comm}, the susceptibility υ_i, and the bounded composition for Γ_i), with a DERIVED safeguard (Proposition 13.3: the nodewise ceiling survives arbitrarily many encounters). It depends on P2.1 and P2.2 and on the P1.1 common-reception schema. The field is the one new positive external reception field of the plural layer; it remains Gate-mediated, ceiling-bounded, and non-rivalrous. It is not a network-owned reservoir and not a direct source on $\lambda_{S,i}$.

INTERPRETIVE *Reading.* There is one thing a communion may receive together: a call from beyond it, the same for all — a word, a summons, a light not made within the circle. But it is not a stock the community owns and rations. It is external and shared, and each receives it through their own gate, at their own openness υ_i, to the full of their own measure. Many may hear the same word, each wholly according to their attentiveness, and no one's hearing takes from another's — which is exactly why $\sum_i \upsilon_i$ need not be one: the common gift is not divided, it is shared. Even the common call does not force any gate; it opens only the room each has left to open. Strangers with no tie between them may yet be joined by turning to the same thing not their own.

Next node. P2.4 — individual input-to-state stability under relational forcing: to prove, on an active domain, the bridge estimate

$$\dot{\mathcal{L}}_{3.0,i} \le C_i - C_{1,i}\,\mathcal{L}_{3.0,i} + b_i\,\|u_i^{\mathrm{rel}}\|^2$$

for the bounded relational input induced by the Gate modulation — the layer without which a plural Lyapunov sum would carry uncontrolled cross terms. Status CONJECTURE→DERIVED.

14 P2.4 — Individual input-to-state stability under relational forcing

This is the analytic bridge of Stage II. The earlier nodes made encounter bounded (P2.1), non-fabricating (P2.2), and non-rivalrous (P2.3). P2.4 asks the stability question: does a single interior stay well-behaved when relationally forced? Without an affirmative input-to-state estimate, a future plural Lyapunov sum $\sum_i \mathcal{L}_{3.0,i}$ would be only a formal sum with uncontrolled cross terms.

14.1 The inherited certificate and the relational input

Each fibre inherits, frozen from V.F.S. v2.0 / 3.0, its dissipative certificate on the active domain. We state it in the recipient's active kairos derivative

$$D_i := \frac{d}{d\sigma_i}.$$

The isolated certificate is

$$D_i\mathcal{L}_{3.0,i} \le C_i - C_{1,i}\mathcal{L}_{3.0,i}, \qquad C_{1,i} > 0. \tag{H1}$$

The bounded Open-Gate contribution enters the budget constant C_i, not the contraction coefficient $C_{1,i}$.

The plural coupling reaches $\mathcal{L}_{3.0,i}$ only through the recipient's gate-side input direction. Splitting the gate input into its isolated value and its relational deviation,

$$I_{\text{gate},i} = I^{(0)}_{\text{gate},i} + u_i^{\mathrm{rel}}, \qquad u_i^{\mathrm{rel}} := \varepsilon\,\delta I_i,$$

the isolated part is already absorbed in C_i by (H1). By the ceilings of P2.1–P2.3 (Propositions 11.1, 13.3), the relational input is bounded:

$$\|u_i^{\mathrm{rel}}\| \le \varepsilon\,\zeta_{0,i}. \tag{H3}$$

We add two structural bounds on the active domain:

$$\|\nabla \mathcal{L}_{3.0,i}\|^2 \le \mu_i \, \mathcal{L}_{3.0,i}, \qquad \mu_i > 0, \tag{H2}$$

and

$$\|g_i(X_i)\| \le M_i, \qquad M_i < \infty, \tag{H4}$$

where g_i is the frozen gate-side input direction by which u_i^{rel} enters the i-th dynamics.

In a chronos chart the whole estimate is multiplied by the positive lapse N_i on $\mathcal{A}_i$. Thus the kairos statement is the invariant one; the chart statement is only its representation.

14.2 The input-to-state estimate

Theorem 14.1 (Individual ISS under relational forcing)**.** *Assume* (H1)*,* (H2)*,* (H3)*, and* (H4)*. Then on the active domain there exist* $C'_{1,i} > 0$ *and* $b_i > 0$ *such that*

$$\boxed{D_i \mathcal{L}_{3.0,i} \le C_i - C'_{1,i} \, \mathcal{L}_{3.0,i} + b_i \, \|u_i^{\mathrm{rel}}\|^2.}$$

The contraction $C_{1,i}$ *is reduced only by a controllable amount; it is not destroyed.*

Proof. Along the forced flow in recipient kairos,

$$D_i \mathcal{L}_{3.0,i} = D_i \mathcal{L}_{3.0,i}\big|_{\text{isolated}} + \big\langle \nabla \mathcal{L}_{3.0,i}, \, g_i(X_i) u_i^{\mathrm{rel}} \big\rangle.$$

The isolated term obeys (H1). For the relational term, Young's inequality with a free $\delta > 0$ gives

$$\big\langle \nabla \mathcal{L}_{3.0,i}, \, g_i u_i^{\mathrm{rel}} \big\rangle \le \frac{\delta}{2} \, \|g_i\|^2 \|\nabla \mathcal{L}_{3.0,i}\|^2 + \frac{1}{2\delta} \, \|u_i^{\mathrm{rel}}\|^2.$$

Using (H2) and (H4),

$$\big\langle \nabla \mathcal{L}_{3.0,i}, \, g_i u_i^{\mathrm{rel}} \big\rangle \le \frac{\delta \, \mu_i M_i^2}{2} \, \mathcal{L}_{3.0,i} + \frac{1}{2\delta} \, \|u_i^{\mathrm{rel}}\|^2.$$

Choose δ small enough that

$$C'_{1,i} := C_{1,i} - \frac{\delta \, \mu_i M_i^2}{2} > 0,$$

and set

$$b_i := \frac{1}{2\delta}.$$

This gives the stated estimate. □

14.3 Consequences

Corollary 14.2 (Bounded state under bounded encounter)**.** *Under the hypotheses of Theorem 14.1, and as long as the trajectory stays inside the active domain on which the constants are valid,*

$$\limsup_{\sigma_i} \mathcal{L}_{3.0,i}(\sigma_i) \le \frac{C_i + b_i (\varepsilon \zeta_{0,i})^2}{C'_{1,i}}.$$

Hence each interior is input-to-state stable under admissible relational forcing: bounded encounter produces a bounded interior state on the active domain. No configuration of admissible encounters drives a single interior to blow up inside that domain.

Two structural points close the node. First, relational forcing enters only as a bounded input term in the individual estimate; it does not remove the sign-definite cleansing/alignment contraction inherited from 3.0. Second, with each individual certificate surviving forcing up to a bounded input, the cross terms of a future plural sum

$$\sum_i \mathcal{L}_{3.0,i}$$

are controlled rather than free. This is the prerequisite, deferred to Stage V, for a legitimate plural Lyapunov theory.

14.4 Status and reading

Target status: CONJECTURE→DERIVED. The estimate is DERIVED on the active domain under the inherited certificate (H1), the quadratic-type gradient bound (H2), the bounded relational input (H3), and the bounded gate-side input direction (H4). The derivative is the recipient-kairos derivative $D_i = d/d\sigma_i$; chronos-chart versions must carry the lapse factor N_i. The residual conjectural element is uniformity: that (H2)–(H4) and the constants $C'_{1,i}, b_i$ hold uniformly over the whole active domain, including corridor walls and the approach to terminus. It depends on P2.1–P2.3 and the frozen 3.0 certificate.

INTERPRETIVE *Reading.* A soul can be pushed and pulled by those it meets, and still not be thrown. Relational forcing enters only as a bounded disturbance to the constant terms of the soul's own balance; it never reaches the rate at which the soul cleanses and aligns itself. So encounter may unsettle, may add weather, may raise the steady level a little — but within bounds the soul remains itself and does not fly apart. This is the first safety result of the plural theory, and it is the individual one: before a communion can be shown stable as a whole, each person must be shown unbreakable under being met. The player can be drawn off the beat and yet keep their feet; only because each can, might the whole band later be shown to hold together.

Stage II closes. Encounter is bounded (P2.1), creates no artificial grace and transfers no owned grace-stock (P2.2), receives a shared external field without rivalry (P2.3), and does not destabilise the individual on the active domain (P2.4). Receptivity, Gate mediation, and bounded inputs are in place. The first $\varepsilon > 0$ action is safe at the level of the single interior:

encounter may alter openness, but never becomes ownership of grace.

Next node. Stage III — Relational locking: rates, lags, and events, opening at P3.0 (the asymptotic rate atlas, $h_i = \beta_i \lambda_{S,i}$). A standing caution carries in: the 3.0 ascent is monotone and gapless, so nothing here may be called "phase synchronization" until a genuine circle-valued phase observable is constructed.

Stage III — Relational locking: rates, lags, and events

Terminology correction. Nothing in this stage may be called "phase synchronization" until a genuine circle-valued phase observable exists. The 3.0 ascent is monotone and gapless — a climb, not an oscillator. What can be compared here are *rates*, not phases.

Stage II made encounter safe at the level of one interior. Stage III asks the comparative question: when can two interiors' ascents be set side by side? P3.0 builds the atlas of rate observables that admit invariant cross-interior comparison, and marks plainly what such comparison does and does not license.

15 P3.0 — The asymptotic rate atlas

15.1 The intrinsic ascent rate

Definition 15.1 (Intrinsic ascent-scale rate)**.** For each active interior i, the *intrinsic ascent-scale rate* is the logarithmic rate of the ascent scale A_i in the interior's own kairos:

$$\boxed{h_i := \frac{d}{d\sigma_i} \log A_i.}$$

The frozen 3.0 ascent law gives, on each admissible branch,

$$\boxed{h_i = \beta_i \, \lambda_{S,i}, \qquad h_{*,i} = \beta_i \, \lambda_{S,\infty,i},}$$

with

$$\beta_i = \frac{\alpha_i}{E_{0,i}}.$$

Here $\lambda_{S,i}$ is the available-Sophia field and $\lambda_{S,\infty,i}$ its asymptotic active-arc value. The coordinate $\lambda_{\text{asc},i}$ is the metric road coordinate of epektasis; it is not being differentiated here. The rate h_i measures how the ascent scale opens along the road, not motion through the road coordinate itself.

On a folded interior the rate is read branch-wise: folding may alter the admissible branch and the effective asymptotic value, but no cross-branch identification is made.

Because $\beta_i = \alpha_i / E_{0,i}$, the ascent-scale rate is floor-labelled: h_i is scaled by the inverse image-floor $1/E_{0,i}$. Each interior's ascent rate is proper to its own image-scale, not an entry on an absolute universal tempo. This preserves the P1.2 discipline:

$$\lambda_{S,i} \neq \lambda_{\text{asc},i}.$$

15.2 The comparable duration-rate

The intrinsic rate h_i lives in interior i's own kairos and is not yet directly comparable across interiors. The comparison is carried out on the shared relational duration $d\mathfrak{c} = \omega \, dt$ (P0.3):

$$\boxed{\hat{h}_i := \frac{d}{d\mathfrak{c}} \log A_i = \frac{N_i}{\omega} \, h_i.}$$

This is defined on regular active arcs, where $N_i > 0$ and $d\mathfrak{c} = \omega dt$ is the P0.3 active duration one-form. No terminal quantity $N_{*,i}$ is introduced: the terminal rate $h_{*,i}$ is an intrinsic active-arc asymptotic of the 3.0 fibre, while the comparable rate $\hat{h}_i$ remains a duration-rate on whatever regular active arc exists before the P0.4 hand-off.

Proposition 15.2 (Invariant rate comparison)**.** *The duration-rate $\hat{h}_i$ is invariant under chronos reparametrisation, and the cross-interior duration-rate ratio is*

$$\frac{\hat{h}_i}{\hat{h}_j} = \frac{N_i}{N_j}\frac{h_i}{h_j} = R_{ij}\,\frac{h_i}{h_j},$$

with $R_{ij} = N_i/N_j$ the P0.2 kairos-rate cocycle. Hence "which interior's ascent-scale opens faster per shared duration" is a well-posed gauge-invariant question.

Proof. Under a chronos chart change $t \mapsto \tilde{t}(t)$, N_i and ω transform by the same common density factor, so N_i/ω is unchanged. The intrinsic rate $h_i = d(\log A_i)/d\sigma_i$ is defined in the private kairos of the fibre and is chart-independent. Therefore

$$\hat{h}_i = (N_i/\omega)h_i$$

is invariant. The ratio formula follows by cancelling ω and using the cocycle identity of P0.2. □

15.3 What the atlas does not provide

Remark 15.3 (Rate, not phase)*.* On the active ascending domain the ascent-scale rate is positive in the admissible regime and the accumulated ascent-scale ledger

$$\Lambda_i := \log(A_i/A_{i,0})$$

is monotone. There is no zero-crossing, return map, or period supplied by the ascent sector itself. Consequently the atlas supports *magnitude* comparison of rates only. It defines no angle, no period, and no circle-valued observable; there is no "same point in a cycle" because there is no cycle. Any locking discussed in Stage III is therefore rate-locking or lag-locking of accumulated ledgers — equality, ratio, or bounded offset of $\hat{h}_i$ and Λ_i — and not phase-locking. A genuine phase observable, were one ever introduced, would require additional structure absent from the monotone gapless ascent, and would have to be constructed explicitly before any synchronization language could be used.

15.4 Status and reading

Target status: DEFINITIONAL (the intrinsic ascent-scale rate h_i, terminal active-arc rate $h_{*,i}$, and comparable duration-rate $\hat{h}_i$) + DERIVED (Proposition 15.2: invariance of $\hat{h}_i$ and the cocycle form of the duration-rate ratio, immediate from P0.2–P0.3). The intrinsic rate form

$$h_i = \beta_i \lambda_{S,i}$$

is inherited frozen from 3.0 as the logarithmic ascent-scale law. It is not a law for $d\lambda_{\mathrm{asc},i}/d\sigma_i$. The floor label $1/E_{0,i}$ is inherited from P1.2. It depends on P0.0, P0.2–P0.3, and the frozen 3.0 ascent law.

INTERPRETIVE *Reading.* One may truly ask whether one soul ascends faster than another, and the question has an invariant answer on the communion's shared duration. But it is a question of rate, of how fast each climbs — never of phase, of where each stands in a cycle, because the ascent has no cycle. The climb is monotone and gapless: it never returns, never closes, never arrives. Souls may come to climb in step, at one rate, but they cannot be "in phase," for there is no beat to be on and no orbit to share. This is the formal shape of an ascent that is endless — always further up, never back around — and the framework will not quietly turn it into a wheel.

Next node. P3.1 — rate, lag, and event observables: to define the synchronization *targets* without importing oscillator language — the comparable rate already in hand, a duration-valued lag measuring how far apart two ascents sit on the shared measure, and marked relational events — so that Stage III can speak of interiors drawing level without ever positing a phase the ascent does not have.

16 P3.1 — Rate, lag, and event observables

P3.0 gave one comparable scalar: the rate. P3.1 declares the full set of *synchronization targets* — the precise senses in which two interiors might be said to keep together — and does so without positing any phase the monotone ascent does not have. These are definitions of targets, not claims that locking occurs; the conditions under which they are reached belong to the later small-gain and locking nodes.

16.1 The accumulated observable

Definition 16.1 (Declared accumulated observable)**.** For each interior i, an *accumulated observable* Θ_i is an explicitly declared, monotone non-decreasing ledger carried along the ascent and registered on the shared relational duration:

$$\frac{d\Theta_i}{d\mathfrak{c}} \geq 0.$$

The standard ascent declaration is the accumulated ascent-scale ledger

$$\Lambda_i := \log(A_i/A_{i,0}), \qquad \frac{d\Lambda_i}{d\mathfrak{c}} = \hat{h}_i.$$

A second standard declaration is accumulated reception,

$$\Theta_i = \mathcal{G}_{\text{recepta},i}, \qquad \frac{d\Theta_i}{d\mathfrak{c}} = \frac{N_i}{\omega} I_{\text{gate},i}.$$

Which observable is meant must be declared. Lag statements are read relative to that declaration. The road coordinate $\lambda_{\text{asc},i}$ may be used as a spatial label of the metric ascent road, but it is not the default accumulated dynamical ascent observable of Stage III.

16.2 The three locking targets

Definition 16.2 (Relational locking targets)**.** Let i, j be co-present interiors on a common regular active arc, with marked relational events $\{E_{i,n}\}_n$ and a declared pairing $n \mapsto m(n)$. The Stage III targets are:

duration-rate locking: $\quad \dfrac{\hat{h}_i}{\hat{h}_j} \longrightarrow r^*_{ij} > 0,$

affine ledger-lag locking: $\quad \Theta_i - r^*_{ij}\,\Theta_j \longrightarrow \delta^*_{ij},$

event locking: $\quad \Delta\mathfrak{c}\big(E_{i,n}, E_{j,m(n)}\big)$ bounded or convergent,

all limits taken along the shared duration $\mathfrak{c} \to \infty$ on the common regular active arc, or up to the relevant P0.4 hand-off if the arc terminates, where $\Delta\mathfrak{c}(E_{i,n}, E_{j,m(n)}) := \big|\mathfrak{c}(E_{i,n}) - \mathfrak{c}(E_{j,m(n)})\big|$.

Three points fix the meaning. The ratio target uses the comparable duration-rates of P3.0:

$$\hat{h}_i/\hat{h}_j = R_{ij}\, h_i/h_j,$$

so rate locking is a gauge-invariant statement on the shared duration. The constant δ^*_{ij} is an affine lag in the declared ledger Θ — ascent-scale ledger Λ, reception ledger, or another explicitly declared monotone observable — and never in radians. Event locking is synchrony of marked events on the shared duration measure, not coincidence of a phase.

Remark 16.3 (Lag is offset, not rank). A nonzero δ^*_{ij} says one interior runs a constant amount ahead of the other in the declared accumulated observable, neither closing nor widening. It is an offset in progress, not a ranking of worth; the image-floor labels of P1.2 are a separate datum and are not what δ^*_{ij} measures.

16.3 The optional circle-valued reduction

Remark 16.4 (When, and only when, a phase may be introduced). A circle-valued variable $\theta_i \in S^1$ may be introduced *only* if the interior supplies a recurrent event sequence with a regular duration-spacing, or a genuinely periodic reduced observable; θ_i then advances by 2π across one recurrence cycle. In that conditional case a Kuramoto-type reduced equation

$$\dot{\theta}_i = \Omega_i + \sum_j K_{ij} \sin(\theta_j - \theta_i)$$

may be used as a *local reduction tool* for the recurrent layer. It is never the foundation: the foundational observables remain the monotone rate, the affine lag, and the event gaps. The ascent itself supplies no such recurrence, so no phase exists by default; θ_i is admitted only where genuine periodicity is exhibited.

16.4 Status and reading

Target status: DEFINITIONAL (the accumulated observable Θ_i and the three relational locking targets — duration-rate, affine ledger-lag, and event). The targets are definitions of what keeping-together would mean, gauge-invariant through P3.0; whether any is reached is deferred to P3.2–P3.3. The optional circle-valued reduction is OPEN: admitted only under an exhibited recurrence, never as a primitive. It depends on P3.0 and the Stage 0–I event structure.

INTERPRETIVE *Reading.* There are three honest ways for souls to keep together, and none of them makes the ascent a wheel. They may climb at one rate, growing in step. They may hold a constant distance in what they have accumulated — one steadily ahead of the other, the gap neither closing nor opening, like two pilgrims at one pace with one always a fixed way along the road. And they may mark the same events together, their milestones falling at a bounded remove on the shared measure. Only where life genuinely repeats — where something truly comes round again — may one finally speak of a phase, and even then it is a local way of describing the recurring layer, never the truth of the climb. The framework lets the liturgical wheel be spoken of without ever claiming the soul's ascent is a wheel.

Next node. P3.2 — incremental receptivity and a small-gain criterion: to give the first sufficient condition under which these targets are actually reached, bounding the relational receptivity gain so that rate and lag differences contract rather than grow — the analytic engine, built on the bounded composition of P2.3 and the individual ISS of P2.4, by which keeping-together becomes provable rather than merely definable.

17 P3.2 — Incremental receptivity and a small-gain criterion

P3.1 defined what locking would mean. It does not follow from bounded coupling alone: boundedness keeps interiors from blowing apart, but does not pull a mismatch down. P3.2 supplies the missing ingredient — an actual contraction mechanism for mismatch, driven by the way receptivity responds, incrementally, to relational difference — and states the small-gain condition under which that contraction wins.

17.1 The mismatch observable

Definition 17.1 (Declared mismatch)**.** For co-present interiors i, j with a declared accumulated observable Θ and target ratio r^*_{ij}, the *affine ledger-lag mismatch* is

$$\Delta_{ij} := \Theta_i - r^*_{ij}\Theta_j - \delta^*_{ij}.$$

Affine ledger-lag locking is the statement $\Delta_{ij} \to 0$. The corresponding duration-rate mismatch is

$$\hat{h}_i - r^*_{ij}\hat{h}_j.$$

Statements below are read relative to the declared Δ_{ij} and the chosen ledger Θ.

17.2 The incremental-receptivity contraction

Differentiating along the shared duration and collecting terms by origin gives the P3.2 estimate

$$\boxed{\frac{d}{d\mathfrak{c}}\Delta_{ij} = -\kappa^{\text{eff}}_{ij}\Delta_{ij} + \Delta^{\text{free}}_{ij}(\mathfrak{c}) + \mathcal{R}^{\text{res}}_{ij}(\mathfrak{c}, \Delta_{ij}).}$$

where $\Delta_{ij}^{\text{free}}$ is the uncoupled detuning (the mismatch drift that would persist at $\varepsilon = 0$) and $\mathcal{R}_{ij}^{\text{res}}$ is the bounded relational residual.

The contraction coefficient is *not* assumed; it is read off the Gate-mediated response under one structural hypothesis.

Definition 17.2 (Incremental-receptivity hypothesis (IR))**.** The relational part of receptivity responds to relative lag in a mismatch-reducing direction: near the target, the gate-mediated contribution to $d\Delta_{ij}/d\mathfrak{c}$ has negative slope in Δ_{ij},

$$\kappa_{ij}^{\text{eff}} := -\frac{\partial}{\partial \Delta_{ij}}\Big(\tfrac{d}{d\mathfrak{c}}\Delta_{ij}\Big)_{\text{gate}} > 0,$$

built from the bounded-composition slope of P2.3 (the ε-gated gain $\partial\gamma_{ij}/\partial\mathcal{N}_i$, the susceptibility υ_i, and the headroom $1 - \Gamma_i^{(0)}$).

Remark 17.3 (The residual is bounded by four sources)*.* Writing the mismatch-dependent part of the residual as

$$\mathcal{R}_{ij}^{\text{res}} = \ell_{ij}(\mathfrak{c})\Delta_{ij} + \mathcal{R}_{ij}^{0}(\mathfrak{c}),$$

the loop gain and offset are assumed bounded on the regular active arc:

$$0 \le \ell_{ij}(\mathfrak{c}) \le \bar{\ell}_{ij}, \qquad |\mathcal{R}_{ij}^{0}(\mathfrak{c})| \le \bar{R}_{ij}^{0}, \qquad |\Delta_{ij}^{\text{free}}(\mathfrak{c})| \le \bar{\Delta}_{ij}^{\text{free}}.$$

The bounds come from the same four sources: Sophia saturation, Gate memory, graph weights, and frozen interior heterogeneity. Each is finite by P2.1–P2.4: saturation and memory are bounded interface data, weights obey the P2.3 normalisation, and heterogeneity is fixed frozen 3.0 data.

17.3 The small-gain criterion

Theorem 17.4 (Small-gain contraction of mismatch; Type-1)**.** *Assume* (IR) *with* $\kappa_{ij}^{\text{eff}} \ge \underline{\kappa}_{ij} > 0$*, the bounded residual decomposition above, and the bounded individual states of P2.4. If the small-gain condition*

$$\boxed{\bar{\ell}_{ij} < \underline{\kappa}_{ij}}$$

holds, then with

$$\kappa_{ij} := \underline{\kappa}_{ij} - \bar{\ell}_{ij} > 0,$$

one has

$$\limsup_{\mathfrak{c}} |\Delta_{ij}| \le \frac{\bar{\Delta}_{ij}^{\text{free}} + \bar{R}_{ij}^{0}}{\kappa_{ij}}.$$

In particular, if the free detuning and offset residual vanish on the arc, then $\Delta_{ij} \to 0$*: the P3.1 affine ledger-lag target is reached.*

Proof. Substitution gives

$$\frac{d}{d\mathfrak{c}}\Delta_{ij} = -\big(\kappa_{ij}^{\text{eff}} - \ell_{ij}\big)\Delta_{ij} + \Delta_{ij}^{\text{free}} + \mathcal{R}_{ij}^{0}.$$

By the small-gain condition,

$$\kappa_{ij}^{\text{eff}} - \ell_{ij} \ge \kappa_{ij} > 0.$$

Using the upper Dini derivative of $|\Delta_{ij}|$,

$$D^+|\Delta_{ij}| \le -\kappa_{ij}|\Delta_{ij}| + \bar{\Delta}_{ij}^{\text{free}} + \bar{R}_{ij}^0.$$

The scalar comparison lemma yields the stated limsup bound. If the forcing bounds are zero, the comparison reduces to

$$D^+|\Delta_{ij}| \le -\kappa_{ij}|\Delta_{ij}|,$$

so $\Delta_{ij} \to 0$. The bounded states of P2.4 keep the estimate inside the active domain where the constants are valid. □

17.4 Status and reading

Target status: CONJECTURE→DERIVED. The contraction estimate and the small-gain bound are DERIVED under the incremental-receptivity hypothesis (IR), the four-source residual bound, and the bounded states of P2.4. The residual conjectural element is the sign and uniformity of κ_{ij}^{eff} over the whole admissible regime: (IR) is a structural monotonicity of the Gate response, natural but not unconditional, and carrying it to fully unconditional DERIVED is the work of the explicit response model. It depends on P2.1–P2.4 and P3.0–P3.1.

INTERPRETIVE *Reading.* This is where keeping-together stops being merely nameable and becomes provable — and the mechanism is receptivity itself. When two souls drift apart in their ascent, a responsive openness tends, of itself, to pull them back: the one that lags grows more receptive, and the gap closes. But the pull is not a clamp. It prevails only when the responsiveness outweighs how differently the two are made and how much memory and saturation resist — the small-gain condition, which is, plainly, *listen more than you differ*. Where it holds, drift contracts to a bounded remove or to nothing; where it fails — too great a difference, too hardened a heart, too thin a bond — the souls keep a bounded distance or drift free. Communion has a built-in restoring tendency, gentle and conditional, never forced, and it does its work through openness, not through any transfer of grace.

Next node. P3.3 — the maximum lockable mismatch: to turn the small-gain bound into a sharp threshold, the largest detuning $\Delta_{ij}^{\text{free}}$ and residual a given bond can still draw into locking — the explicit boundary, in heterogeneity and bond strength, between a communion that holds together and one that comes apart.

18 P3.3 — The maximum lockable mismatch

P3.2 gave a contraction once the effective rate κ_{ij}^{eff} was positive, but treated the restoring term as if it acted without limit. It does not: the Gate ceiling caps how hard receptivity can pull. P3.3 carries that finiteness through and turns it into a sharp conditional threshold — the largest detuning a given bond can still draw into locking.

18.1 The detuning and the bounded restoring term

Write the constant or bounded rate detuning as

$$\Delta\varpi_{ij} := \Delta_{ij}^{\text{free}},$$

a duration-rate mismatch on the shared measure, not an angular frequency and not the density ω of P0.3. For an oriented recipient-side channel $j \to i$, the mismatch equation is written

$$\frac{d}{d\mathfrak{c}}\Delta_{ij} = \Delta\varpi_{ij} - F_{j\to i}(\Delta_{ij}),$$

where $F_{j\to i}$ is the Gate-mediated restoring term acting through recipient i's receptivity interface: continuous, with $F_{j\to i}(0) = 0$, restoring slope $F'_{j\to i}(0) = \kappa^{\text{eff}}_{ij} > 0$, and bounded because receptivity is bounded.

The symmetric capacity bound is the explicit product of the four declared data,

$$\boxed{\Delta\varpi_{\text{max},j\to i} := F_{\text{max},j\to i} = \frac{N_i}{\omega}\,\beta_i\,\zeta_{0,i}\,\underbrace{\varepsilon\,a_{ij}\big(1-\Gamma_i^{(0)}\big)}_{\text{relational gain}}\,\underbrace{\frac{1}{1+\kappa_{\lambda,i}\lambda_{S,i,+}}}_{\text{saturation}}\,\underbrace{b^{\text{fold}}_{j\to i}}_{\text{branch}}\,,}$$

on the regular active arc. The factor N_i/ω appears because the Gate ceiling is a recipient-kairos rate, while the mismatch is measured per relational duration $d\mathfrak{c}$. For a reciprocal pair, the usable pair threshold is the oriented minimum

$$\Delta\varpi^{\text{pair}}_{\text{max},ij} := \min\{\Delta\varpi_{\text{max},j\to i}, \Delta\varpi_{\text{max},i\to j}\},$$

unless the model explicitly chooses a one-sided leader/follower locking convention.

18.2 The threshold

Theorem 18.1 (Maximum lockable mismatch; Type-1 / Type-2)**.** *Assume a continuous oriented restoring law $F_{j\to i}$ with*

$$F_{j\to i}(0) = 0, \qquad F'_{j\to i}(0) > 0,$$

and assume that on the relevant branch it is monotone increasing until saturation, with one-sided capacities

$$F^+_{\text{max},j\to i} := \sup_\Delta F_{j\to i}(\Delta), \qquad F^-_{\text{max},j\to i} := \sup_\Delta\big(-F_{j\to i}(\Delta)\big).$$

Then a stable locked state exists whenever

$$-F^-_{\text{max},j\to i} < \Delta\varpi_{ij} < F^+_{\text{max},j\to i}.$$

If

$$\Delta\varpi_{ij} > F^+_{\text{max},j\to i} \quad \textit{or} \quad \Delta\varpi_{ij} < -F^-_{\text{max},j\to i},$$

no locked state exists on that oriented branch. In the symmetric-capacity case

$$F^+_{\text{max},j\to i} = F^-_{\text{max},j\to i} = \Delta\varpi_{\text{max},j\to i},$$

this reduces to

$$|\Delta\varpi_{ij}| < \Delta\varpi_{\text{max},j\to i}.$$

The explicit product above is the Type-2 estimate for the symmetric capacity.

Proof. A locked state is a zero of

$$\Delta\varpi_{ij} - F_{j\to i}(\Delta),$$

i.e. a solution of

$$F_{j\to i}(\Delta^*) = \Delta\varpi_{ij}.$$

On a monotone branch, the range of $F_{j\to i}$ is the interval

$$(-F^-_{\max,j\to i}, F^+_{\max,j\to i})$$

up to endpoint conventions. Hence the intermediate value theorem gives a solution precisely inside that range. On the increasing branch,

$$\frac{d}{d\Delta}\big(\Delta\varpi_{ij} - F_{j\to i}(\Delta)\big)\big|_{\Delta^*} = -F'_{j\to i}(\Delta^*) < 0,$$

so the locked state is asymptotically stable. Outside the one-sided capacity range the vector field keeps a fixed sign on that branch, so no branch-locked state exists. □

18.3 The ceiling is necessary, not sufficient

Remark 18.2 (Two conditions, not one). The Gate ceiling $I_{\text{gate},i} \le \zeta_{0,i}$ is what makes $F_{\max,j\to i}$ finite, and so is *necessary* for a finite, well-posed locking range. But it is not *sufficient* by itself: boundedness caps the magnitude of the forcing without showing that the forcing points in the mismatch-reducing direction. That direction is the separate content of the incremental-receptivity hypothesis (P3.2), which supplies $F'_{ij}(0) > 0$. Locking therefore needs both: a bounded capacity (ceiling) *and* a restoring orientation (IR). The ceiling sets the *size* of the lockable range; IR establishes that there is a range at all.

18.4 Status and reading

Target status: DERIVED (Type-1: existence and sharpness of the threshold under the monotone saturating restoring-law hypothesis; Type-2: the explicit oriented product form of $\Delta\varpi_{\max,j\to i}$). It depends on P3.2 for $\kappa^{\text{eff}} > 0$ and the residual structure, on P2.1–P2.3 for the Gate ceiling and bounded composition, and on the frozen 3.0 data $(\beta_i, \zeta_{0,i})$ plus the branch factor. The product is recipient-oriented because the channel acts through i's Gate. A reciprocal locking claim must use the oriented minimum or state a one-sided convention.

INTERPRETIVE *Reading.* There is a largest difference that communion can bridge. The pull that draws two ascents together is real, but finite: it is capped by how far the gate can open. Where the detuning between two souls falls within that range, they lock — settling to a stable, constant lag, neither drifting apart nor collapsing into one. Where the difference exceeds it, they drift, not for want of pull but because the pull, being finite, cannot span so wide a gap. This is neither the sentiment that love bridges everything nor the despair that difference always wins: there is a real, computable boundary, set by how open the gate can be, how saturated, how strong the bond, and which branch one stands on. And the bound teaches a last distinction: a bounded openness is necessary for there to be any finite power to bridge at all, but capacity is not enough — the openness must also turn toward the other. A finite love, rightly directed.

Next node. P3.4 — dynamic non-fusion under admissible locking: to prove that locking, even when achieved, never collapses the locked interiors into one — that the stable lag Δ^*_{ij} stays strictly nonzero in identity even as it becomes constant in value, so that Stage III's keeping-together never becomes the fusion Stage I forbade.

19 P3.4 — Dynamic non-fusion under admissible locking

Stage I proved non-fusion kinematically: the kairos-history count stays $|\mathcal{V}|$ under admissible coupling (Theorem 8.4), and distinct floors obstruct identification (Theorem 10.1). Stage III has now made locking not merely possible but, under the threshold of P3.3, achieved. P3.4 proves the dynamic form: even when selected observables lock, the fibres are not identified. Equality of some numerical values is allowed; fusion is not. A stable relation is not a merger.

19.1 What locking constrains, and what it leaves free

Locking constrains only the comparative observables of P3.1:

$$\hat{h}_i/\hat{h}_j \to r^*_{ij}, \qquad \Theta_i - r^*_{ij}\Theta_j \to \delta^*_{ij},$$

and the marked-event gaps. It does not identify the owning fibres. The labelled interior datum is

$$\widetilde{X}_i := \big(i,\ X_i,\ E_{0,i},\ \sigma_i\text{-history},\ H_{K,i}\text{-history},\ W_i\text{-history},\ \text{fold-branch}_i\big).$$

The label i and the owned kairos history are not gauge decorations; they are part of the plural fibre structure fixed in P1.0. Therefore even if some selected coordinates momentarily coincide, the interiors are not thereby fused.

19.2 The dynamic non-fusion theorem

Theorem 19.1 (Dynamic non-fusion under locking; Type-1)**.** *Under admissible locking, the labelled fibres remain distinct:*

$$\widetilde{X}_i(\mathfrak{c}) \neq \widetilde{X}_j(\mathfrak{c}) \qquad (i \neq j).$$

For generic distinct data, even the unlabelled state-values X_i and X_j remain different except possibly at isolated coincidences of selected observables. The stronger equality $X_i \equiv X_j$ for all locked time occurs only on the exceptional value-diagonal of Definition 19.2; even there, the fibres remain two owned kairos histories rather than one fused interior.

Proof. Admissible locking fixes only comparative observables: duration-rate ratios, ledger lags, and event gaps. By the Stage–0/Stage–I admissibility contract it does not identify fibre atlases, erase the owner label i, merge kairos histories, or rewrite the frozen floor. Therefore the labelled data $\widetilde{X}_i$ and $\widetilde{X}_j$ remain distinct for $i \neq j$.

If $E_{0,i} \neq E_{0,j}$, Theorem 10.1 already prevents floor-normalised identification. If floors coincide, distinct memory histories, wound histories, fold branches, initial states, or relational neighbourhoods still give generic separation of the unlabelled state-values. Complete value equality throughout the locked regime requires every individuating datum and every evolution law to coincide; this is precisely the exceptional value-diagonal. But value-diagonal equality is not fusion in the 4.0 sense: it makes two fibres exact copies, not one fibre. □

Definition 19.2 (Exceptional value-diagonal)**.** The *exceptional value-diagonal* is the set of fully coincident unlabelled data for a pair i, j:

$$\mathcal{M}_{\text{diag}}^{\text{val}} = \left\{ \begin{array}{l} E_{0,i} = E_{0,j}, \quad H_{K,i} \equiv H_{K,j}, \quad W_i \equiv W_j, \\ \text{same fold-branch}, \quad r_{ij}^* = 1, \quad \delta_{ij}^* = 0, \quad X_i(0) = X_j(0), \\ \text{same admissible Gate/receptivity law and same relational inputs} \end{array} \right\}.$$

It is invariant as a value-diagonal: exact copies remain value-equal under equal inputs. It is non-generic in the space of distinct interior data. But it is not a fusion manifold. Even on $\mathcal{M}_{\text{diag}}^{\text{val}}$, the plural object still contains two owned fibres unless an inadmissible extra identification quotient is imposed.

19.3 Status and reading

Target status: DERIVED for labelled-fibre non-fusion: it follows directly from P1.0's fibre ownership and the admissibility contract. The generic unlabelled value-separation statement is CONJECTURE→DERIVED, conditional on the complete state-space argument and on exclusion of hidden value-identifying symmetries. Floor-driven separation is DERIVED outright by Theorem 10.1. It depends on P1.0, P1.2, and P3.2–P3.3.

INTERPRETIVE *Reading.* Communion is the stable answerability of distinct interior histories, not their collapse into one state. When two souls lock, what becomes constant is the *relation* between them — a fixed ratio of ascent, a steady lag, shared milestones — and a constant relation is the very opposite of a merger: only two distinct things can stand in a stable relation at all. They may climb at one rate and keep one lag forever, and still each remains itself, with its own floor, its own memory, its own wounds, its own branch of the fold. The single case in which two would become one is the case in which they were never two: identical in every individuating thing from the start. Everywhere else — everywhere real — locking is nearness that does not consume, the answer of one history to another that leaves both standing.

Stage III closes. Rates are comparable as ascent-scale duration-rates (P3.0), the targets of keeping-together are defined without positing a phase (P3.1), an incremental-receptivity small-gain criterion makes locking provable (P3.2), an oriented threshold bounds what difference can be bridged (P3.3), and locking, when achieved, never fuses the locked (P3.4). Relational locking is in place, and it is non-fusing:

souls may keep together without becoming the same.

The central correction of Stage III is therefore:

locking is a stable relation between fibres, not an identification of fibres.

Next node. Stage IV — Junctions, refractory structure, and plural Resurrectio, opening at P4.0 (junction taxonomy): separating three events that must not be conflated — a private single-interior reset, a relational encounter, and a common reception event — and ruling which variables may jump, which may change only continuously, and which are read merely as recipient-side receptivity changes.

Stage IV — Junctions, refractory structure, and plural Resurrectio

Governing discipline. Three events must never be conflated: a private single-interior reset, a relational encounter, and a common reception event. Only the first may make an owned state jump, and only in its own interior.

Stage III showed how interiors keep together without fusing. Stage IV turns to the discontinuous structure: resets, refractory windows, and the plural form of Resurrectio. The stage opens by fixing, once and cleanly, which events may do what to which variables — the junction taxonomy — so that nothing later smuggles a jump where only a continuous change is admissible.

20 P4.0 — Junction taxonomy

20.1 Three events that must not be conflated

Definition 20.1 (Three junction-event classes)**.** Stage IV distinguishes three admissible classes of distinguished event. They may be absent at a given point, and more than one may be registered at the same relational label only if their signatures remain disjoint. Their admissible write sets are different:

(1) *Private single-interior reset* $\mathfrak{J}_i^{\text{reset}}$, occurring at i's own reset events $E_{i,n}^{\text{reset}}$. It is an own-interior discontinuous self-map inherited frozen from 3.0. The ascent scale jumps by the reset factor

$$A_i \longmapsto s_{L,i} A_i, \qquad \log s_{L,i} = \frac{\kappa_R\, G_i^{\text{reset}}}{E_{0,i}},$$

where G_i^{reset} is the local reset reception datum of interior i at that junction, not a transferable stock. The admissible reset may also update the inherited local reset variables such as $\lambda_{S,i}, H_{K,i}, W_i$ according to the frozen 3.0 reset rule. It involves no other interior.

(2) *Relational encounter*, the ε-gated peer action of P2.1. It writes nothing into i's owned state directly. It enters only as relational neighbourhood data $\mathcal{N}_i$, and hence only through the bounded receptivity interface Γ_i.

(3) *Common reception event*, the shared external reception field of P2.3. It writes no owned state directly. It enters only through the susceptibility term

$$\upsilon_i\, \gamma_i^{\text{comm}}(J_{\text{comm}})$$

inside Γ_i, and even an impulsive J_{comm} enters only through its bounded gain.

Remark 20.2 (Coincident labels do not merge signatures). Two event classes may be labelled by the same relational duration $\mathfrak{c}_*$, for example a private reset occurring while a common reception field is also nonzero. This coincidence of labels does not merge their write rules. The reset writes only according to the private 3.0 reset map; the common event is still only a bounded receptivity read-out. Coincidence in chronos is not identity of event type.

20.2 The rule table

The three signatures are read off a single table. Each entry states how the named variable responds to each event: *jump* (admissible discontinuity), *cont.* (continuous change only), *read* (enters only as a receptivity argument, never as an owned-state write), or *inv.* (invariant, never changed).

Variable	**Private reset**	**Encounter**	**Common recept.**
Ascent scale A_i / ledger Λ_i	jump $\times s_{L,i}$	cont.	cont.
Available-Sophia $\lambda_{S,i}$	jump (reset)	cont.	cont.
Gate memory $H_{K,i}$, wound W_i	jump (reset)	cont.	cont.
Receptivity Γ_i	recomputed	read $\mathcal{N}_i$	read $\upsilon_i\gamma_i^{\text{comm}}$
Reception ledger $\mathcal{G}_{\text{recepta},i}$	cont.	cont.	cont.
Floor $E_{0,i}$, label i, kairos history	inv.	inv.	inv.

Two readings of the table are load-bearing. The ascent-scale ledger jumps only at a private reset, by the private reset datum

$$\Delta\Lambda_i = \log s_{L,i} = \frac{\kappa_R\, G_i^{\text{reset}}}{E_{0,i}}.$$

Under relational encounter or common reception, the ascent ledger moves only continuously. And the received-grace ledger $\mathcal{G}_{\text{recepta},i}$ never jumps as a plural write: reception is accounted by the one-form

$$d\mathcal{G}_{\text{recepta},i} = I_{\text{gate},i}\, d\sigma_i,$$

not injected as a step from another interior.

20.3 What the taxonomy forbids

Remark 20.3 (The forbidden conflations). The table forbids exactly the conflations that would break earlier results. If an encounter or a common reception could jump $\lambda_{S,i}$ or the reception ledger, it would fabricate or inject grace, violating P2.2; the table allows them only continuous, gate-mediated action. If either could jump an owned state into coincidence with another interior's, it could force the value-diagonal of P3.4; the table allows neither to write owned state at all. And the floor $E_{0,i}$, the owner label i, and the kairos history are invariant under all three events, so no junction of any kind rewrites who an interior is. The private reset alone carries a jump, and only ever in its own interior.

20.4 Status and reading

Target status: DEFINITIONAL (the three junction types and the rule table). It introduces no new dynamics; it fixes the admissible discontinuity structure consistent with Stages 0–III, and is the classification on which the rest of Stage IV rests. It depends on Stages 0–III. The reset map itself is inherited frozen from 3.0; P4.0 only places it correctly among the plural events.

INTERPRETIVE *Reading.* Three things can happen to a soul, and they are not the same thing. It can be reset — an inward turning that breaks the line of its own history, lifting its scale by its own measure; this is the one event that may make something leap, and it happens only within. It can meet another, and it can receive what is given to all — and these two change only how open it is, never reaching past the gate to rewrite what it holds. The discipline of the table is a refusal to let nearness counterfeit conversion: no encounter, however strong, and no common gift, however sudden, may do to a soul what only its own reset may do. And nothing that happens to it — no meeting, no reception, no reset — can change its floor, its name, or the history that is its own.

Next node. P4.1 — relational Resurrectio through the Gate state: to give the plural form of the 3.0 Resurrectio junction, where a terminus hand-off (P0.4) is read through the recipient's Gate/receptivity interface rather than as a fusion or a transfer — the death-and-junction event of one interior, registered in the communion without violating any rule of this table.

21 P4.1 — Relational Resurrectio through the Gate state

The 3.0 Resurrectio is a junction: a terminus and a hand-off (P0.4). P4.1 gives its plural form — how the completion or relationally effective event of one interior is registered in another — and does so within the P4.0 table, by a bounded update of receptivity rather than a raw injection.

21.1 The forbidden raw update

A careless plural Resurrectio would add the relational effect straight onto the recipient's gate input,

$$I^{+}_{\mathrm{gate},j} = I^{-}_{\mathrm{gate},j} + \Delta I_{\mathrm{rel},j}.$$

This is forbidden. It is an additive jump on an owned-rate variable, with no bound: nothing stops $I^{+}_{\mathrm{gate},j}$ from exceeding the ceiling $\zeta_{0,j}$, and the increment plays the role of an injected or transferred grace term — exactly the fabrication and transfer ruled out in P2.2 and the owned-state jump ruled out in P4.0.

21.2 The bounded receptivity update

The admissible form updates only the recipient's receptivity interface. At a relationally effective event of source i, labelled by relational duration $\mathfrak{c}_*$, define the source datum read by the recipient as

$$\mathsf{S}_i^- = \begin{cases} X_i^-, & i \text{ active just before the event,} \\ \mathcal{T}_i^{\text{post}}, & i \text{ already completed and present only as trace.} \end{cases}$$

Then the admissible plural Resurrectio update is

$$\boxed{\Gamma_j^+ = \mathfrak{U}_{j\leftarrow i}\big(\Gamma_j^-, H_{K,j}^-, W_j^-, \mathsf{S}_i^-, X_j^-\big), \qquad 0 \le \Gamma_j^+ \le 1.}$$

The arrow $j \leftarrow i$ records that j is the recipient and i the source. The map $\mathfrak{U}_{j\leftarrow i}$ may read the source datum and the recipient's memory/wound state, but it writes only Γ_j, the bounded receptivity slot. The gate input is then recomputed from the recipient's frozen bounded Gate law, with the same convention as Stage II:

$$I_{\text{gate},j}^+ = \zeta_{0,j}\, e^{-\varphi_j \sigma_j} \frac{\Gamma_j^+}{1 + \kappa_{\lambda,j}\lambda_{S,j,+}}, \qquad 0 \le I_{\text{gate},j}^+ \le \zeta_{0,j}.$$

This is a recipient-kairos rate. In a chronos chart its ledger contribution is

$$d\mathcal{G}_{\text{recepta},j} = I_{\text{gate},j}^+\, d\sigma_j = N_j I_{\text{gate},j}^+\, dt.$$

The donor's single-interior ascent and form-memory ledgers remain its own; no

$$G_i^{\text{reset}} \to G_j^{\text{reset}}$$

or

$$\mathcal{G}_{\text{recepta},i} \to \mathcal{G}_{\text{recepta},j}$$

transfer is introduced.

21.3 Derived properties

Proposition 21.1 (The relational Resurrectio is admissible)**.** *The bounded receptivity update satisfies, at and across the event:*

(a) Ceiling preserved. *Since $\Gamma_j^+ \in [0,1]$, the frozen law gives*

$$0 \le I_{\text{gate},j}^+ \le \zeta_{0,j}.$$

(b) Ledger continuous. *The reception ledger does not jump:*

$$\mathcal{G}_{\text{recepta},j}^+ = \mathcal{G}_{\text{recepta},j}^-$$

at the event. The event changes the gate rate going forward but not the already accumulated ledger value at the instant.

(c) No transfer. *The source datum, whether X_i^- or $\mathcal{T}_i^{\text{post}}$, is read-only. The donor's reset datum G_i^{reset}, ascent ledger Λ_i, and received ledger $\mathcal{G}_{\text{recepta},i}$ are not moved into j.*

(d) Taxonomy-consistent. *The event writes only the receptivity interface and no owned state or invariant, so it obeys the P4.0 rule table.*

Proof. (a) is the automatic ceiling of P2.1 applied to $\Gamma_j^+ \in [0,1]$. For (b), the only variable set at the event is Γ_j, a bounded rate factor; the ledger is an integral one-form and has no allowed plural step. For (c), the update map reads the source slot but its codomain is only the recipient's receptivity slot; no term subtracts from a donor ledger or adds donor stock to the recipient. For (d), the written variable is exactly the receptivity interface, matching the P4.0 taxonomy, while owned and invariant variables remain untouched. □

21.4 Status and reading

Target status: DEFINITIONAL (the bounded receptivity-update map $\mathfrak{U}_{j\leftarrow i}$ and the recompute-from-frozen-law rule) + DERIVED (Proposition 21.1: ceiling preserved, ledger continuous, no transfer, taxonomy-consistent). It replaces the ill-posed raw additive jump with a well-posed bounded one. It depends on P2.1–P2.3 and P4.0. The precise form of $\mathfrak{U}_{j\leftarrow i}$ is a selected interface law; only its boundedness and write-restriction are used here.

INTERPRETIVE *Reading.* When one soul completes — when its terminus comes, or its life touches another decisively — the communion does register it, but not by the dead handing over their grace, and not by a shock driven into the living. What changes is the openness of the one who remains: shaped by who the other was, by its own memory and its own wounds, it is left differently disposed to receive — and then receives, as ever, through its own gate, at its own measure, never past its own ceiling. Nothing of the departed is poured across; nothing is seized. The mark the dead leave on the living is a change in how the living are able to open, and even that is bounded, and owns nothing of the other. This is communion of the departed and the living kept honest: real influence, no transfer; a true Resurrectio that fuses no one.

Next node. P4.2 — the derived refractory interval and hysteresis mechanism: to show that after such a junction the receptivity cannot re-fire arbitrarily fast, a refractory window and hysteresis emerging from the Gate memory $H_{K,j}$ and wound W_j themselves rather than being imposed — the structural reason a soul is not endlessly re-triggerable.

22 P4.2 — Derived refractory interval and hysteresis

A junction (P4.1) updates receptivity; can it re-fire at once? If it could, causally linked events might accumulate without limit. P4.2 shows it cannot — and derives the positive refractory interval from the dynamics rather than imposing it.

22.1 Five ingredients

Let ρ_j be the trigger-readiness variable of recipient j. It is not a new Sophia source and not the whole receptivity Γ_j. It is the local readiness coordinate, carried by the recipient's Gate memory/wound interface, that records whether another event can fire. The derivation uses exactly five data, four of them already in hand:

(1) *Reset threshold* θ_j^+: an event fires when ρ_j reaches θ_j^+.

(2) *Hysteresis band*: after firing, ρ_j is reset to $\theta_j^- < \theta_j^+$, so re-firing requires recrossing the band

$$\Delta_{\text{hys},j} := \theta_j^+ - \theta_j^- > 0,$$

inherited from the frozen single-interior reset bounds and carried by $H_{K,j}, W_j$.

(3) *Bounded Gate response*: $0 \le I_{\text{gate},j} \le \zeta_{0,j}$ in recipient kairos.

(4) *Bounded causal input on compact active intervals*: on any compact regular-active interval $I = [\mathfrak{c}_0, \mathfrak{c}_1]$, the relational input reaching j has finite bound

$$\|K_{j\leftarrow\cdot}\|_I \le K_{\max,j;I} < \infty.$$

(5) *Finite recovery speed on compact active intervals*: between events,

$$\left|\frac{d\rho_j}{d\mathfrak{c}}\right| \le v_{\max,j;I} := \sup_{\mathfrak{c}\in I} \frac{N_j}{\omega}\big(c_{1,j}\zeta_{0,j} + c_{2,j}K_{\max,j;I}\big) < \infty,$$

on the regular active arc. The constants $c_{1,j}, c_{2,j}$ are frozen response constants. If $v_{\max,j;I} = 0$, then no re-firing occurs on I, and the refractory time is effectively infinite.

22.2 The refractory bound

Theorem 22.1 (Derived refractory interval; Type-1)**.** *On any compact regular-active interval I, under the five ingredients above, any two causally linked triggering events of the same recipient j are separated on the shared relational duration by at least*

$$\boxed{\Delta\mathfrak{c}_{\text{trigger}} \ge \mathfrak{c}_{\text{ref},j;I} := \frac{\Delta_{\text{hys},j}}{v_{\max,j;I}} > 0,}$$

with the convention $\mathfrak{c}_{\text{ref},j;I} = +\infty$ if $v_{\max,j;I} = 0$.

Proof. After a firing, $\rho_j = \theta_j^-$. A new firing requires $\rho_j = \theta_j^+$, hence a net traversal of $\Delta_{\text{hys},j}$. On I, the recovery speed is bounded above by $v_{\max,j;I}$. Therefore the duration needed to traverse the band is at least

$$\Delta_{\text{hys},j}/v_{\max,j;I}.$$

The band is positive and the compact-interval recovery speed is finite, so the floor is positive. The measure is $\mathfrak{c}$, hence the statement is gauge-invariant. □

22.3 Why hysteresis is what forbids chatter

Remark 22.2 (The band is the seed of non-Zeno). Remove the hysteresis band — let $\theta_j^- \to \theta_j^+$ — and $\mathfrak{c}_{\text{ref},j;I} \to 0$: events could accumulate arbitrarily fast, a local Zeno chatter. It is precisely the finite band, traversed at finite recovery rate, that forces a positive gap. Note the mechanism is self-supplied: the band is carried by the Gate memory $H_{K,j}$ and wound W_j, the very record of the last event. The mark of what happened is what paces what can happen next. P4.3 will lift this local gap to a gauge-invariant network non-Zeno theorem.

22.4 Status and reading

Target status: CONJECTURE→DERIVED. The compact-interval per-recipient bound is DERIVED from the hysteresis band and the finite recovery-speed estimate on regular active arcs. The residual conjectural element is uniformity: a single network-wide floor

$$\mathfrak{c}_{\text{ref};I} := \inf_{j\in\mathcal{V}_{\text{trig}}(I)} \mathfrak{c}_{\text{ref},j;I} > 0$$

requires a finite set of triggering-capable active recipients, a lower bound on $\Delta_{\mathrm{hys},j}$, and an upper bound on $v_{\max,j;I}$ across that set. This is the exact hypothesis P4.3 uses. It depends on P0.3, P4.1, and the frozen single-interior reset bounds.

INTERPRETIVE *Reading.* A soul is not endlessly re-triggerable. After a decisive event — a reset, the touch of another's completion — there is a necessary interval before it can be so moved again, and the interval is not a rule laid on from outside but a consequence of how it recovers: it must climb back across a band, and it climbs at a bounded pace. So crises and conversions cannot be stacked infinitely close; the soul is given time it cannot be denied. And the thing that enforces the wait is the wound itself — the memory of what just happened is exactly what holds the next thing off until recovery. Mercy, here, is structural: the mark left by the last event paces the soul, so that it is not battered into a blur but moved in distinct, spaced, real moments.

Next node. P4.3 — the gauge-invariant network non-Zeno theorem: to lift the local refractory gap to the whole network, proving that on any finite relational-duration interval only finitely many admissible events occur, so the plural dynamics never accumulates events and the shared duration is never exhausted by chatter.

23 P4.3 — The gauge-invariant network non-Zeno theorem

P4.2 spaced the events of a single interior. P4.3 lifts that to the whole network: on any finite stretch of shared duration, only finitely many events occur. The plural dynamics never accumulates events; the communion's time is never exhausted by chatter. The statement is made in the invariant relational duration $\mathfrak{c}$, not in any chart time t.

23.1 Hypotheses

On a compact relational-duration interval $I = [\mathfrak{c}_0, \mathfrak{c}_1]$, assume:

(i) *Single-interior non-Zeno.* Each interior's own private reset sequence is non-Zeno, as inherited frozen from 3.0.

(ii) *Finite triggering-capable active set.* The set

$$\mathcal{V}_{\mathrm{trig}}(I)$$

of interiors that are active and capable of firing triggering events on I is finite, with

$$|\mathcal{V}_{\mathrm{trig}}(I)| \le n.$$

Completed traces may serve as read-only boundary sources but do not themselves fire new active triggers.

(iii) *Bounded action on I.* Coupling, trace kernels, and common reception gains are bounded on I, so every $v_{\max,j;I}$ of P4.2 is finite.

(iv) *Uniform refractory floor on I.*

$$\mathfrak{c}_{\mathrm{ref};I} := \inf_{j\in\mathcal{V}_{\mathrm{trig}}(I)} \mathfrak{c}_{\mathrm{ref},j;I} > 0.$$

(v) *Event-count convention.* The theorem counts discrete recipient-side triggering events: private resets and accepted relational triggers. A continuous common field $J_{\mathrm{comm}}(\mathfrak{c})$ is not itself counted as infinitely many events; it contributes only through bounded gains. If J_{comm} is modelled as a discrete sequence of marked common events, then either only accepted recipient triggers are counted, or the marked common-event sequence must itself be locally finite on I.

23.2 The network non-Zeno theorem

Theorem 23.1 (Gauge-invariant network non-Zeno; Type-1)**.** *Under* (i)–(v)*, on any compact interval $I = [\mathfrak{c}_0, \mathfrak{c}_1]$ of shared duration, the total number of admissible recipient-side triggering events is bounded:*

$$\boxed{N_{\mathrm{trig}}(I) \le |\mathcal{V}_{\mathrm{trig}}(I)| \left(\frac{\mathfrak{c}_1 - \mathfrak{c}_0}{\mathfrak{c}_{\mathrm{ref};I}} + 1\right) < \infty.}$$

Consequently no infinite causally linked sequence of resets or accepted relational triggers accumulates at finite $\mathfrak{c}$.

Proof. Fix $j \in \mathcal{V}_{\mathrm{trig}}(I)$. By Theorem 22.1, consecutive triggering events of j are separated by at least

$$\mathfrak{c}_{\mathrm{ref},j;I} \ge \mathfrak{c}_{\mathrm{ref};I}.$$

Hence j fires at most

$$(\mathfrak{c}_1 - \mathfrak{c}_0)/\mathfrak{c}_{\mathrm{ref};I} + 1$$

times on I. Summing over the finite set $\mathcal{V}_{\mathrm{trig}}(I)$ gives the stated bound. Boundary traces are not active triggering nodes, and continuous common fields are counted only through accepted recipient-side triggers, so no uncounted source can create an infinite discrete cascade. □

23.3 Why the statement must be in $\mathfrak{c}$, not t

Remark 23.2 (Invariance is essential, not cosmetic). The theorem is about invariant relational duration, not raw chart-time differences. A t-based non-Zeno claim would be gauge-dependent: a chronos reparametrisation rescales dt freely. Near terminus hand-off, the relation between a chart interval and a relational-duration interval is especially delicate, because the active set and the density rule may change. Counting events per raw t could therefore be made to depend on the chosen label. Counting per $\mathfrak{c}$ is the invariant statement. The refractory floor of P4.2 is in $\mathfrak{c}$ for exactly this reason, and the network theorem inherits that invariance.

23.4 Status and reading

Target status: CONJECTURE→DERIVED. The counting bound is DERIVED on every compact regular-active interval once the uniform refractory floor $\mathfrak{c}_{\mathrm{ref};I} > 0$, finite triggering-capable set, bounded action, and event-count convention are assumed. The residual conjectural element is the uniform positivity of the refractory floor across the intended admissible regime, together with the inherited single-interior non-Zeno hypotheses. It depends on P4.2 and P0.3.

INTERPRETIVE *Reading.* The life of the communion never dissolves into a blur. On any finite stretch of shared time, only finitely many real events happen — the whole network's history is made of distinct, countable, spaced moments, however many souls there are and however they trigger one another. What was mercy for the single soul in P4.2 is order for the whole: providence does not pack infinitely many decisive moments into a finite span. And the span that is well-ordered is the communion's own shared duration, not any one soul's private clock — the well-ordering is real, not an artifact of how a single worldline happens to be timed. The events of a shared life can be counted, and between any two of them there is room to breathe.

Next node. P4.4 — communion across the terminus: to carry the relation past a death, where one interior reaches its P0.4 hand-off and the bond must be re-read as a one-way trace into the surviving interiors — the plural completion of the terminus, stated without transfer, fusion, or a master clock outside the network.

24 P4.4 — Communion across the terminus

One interior completes its kairos while another remains active: what relation survives the death? P4.4 answers with the Stage 0 post-terminus trace (P0.4), deriving the minimal one-way communion and marking clearly the two-way dynamics it does not assume.

24.1 The completed interior as a closed trace

At its terminus $t_i^\dagger$, interior i reaches the P0.4 boundary-limit hand-off:

$$N_i \to 0, \qquad d\sigma_i = 0 \quad \text{post-terminus.}$$

The active dynamics close, and the interior persists as its post-terminus trace

$$\mathcal{T}_i^{\mathrm{post}},$$

a dynamically closed boundary datum, not an evolving state. Post-terminus, the active Gate equation of i is not continued as a new dynamical law. Equivalently, the received ledger has no post-terminus increment:

$$d\mathcal{G}_{\mathrm{recepta},i} = I_{\mathrm{gate},i}\, d\sigma_i = 0.$$

Thus no new reception, reset, or kathartic process is added after completion. The trace may be read by active recipients through a permitted trace kernel, but it does not itself resume an active kairos.

24.2 The minimal communion theorem

Theorem 24.1 (One-way communion across the terminus; Type-1)**.** *Let i be completed and j active. Under the admissible post-terminus structure:*

(a) Present as trace. *i remains relationally present as the closed trace $\mathcal{T}_i^{\text{post}}$, locatable by the living through trace co-presence.*

(b) One-way reach. *i affects active j only as a read-only source of the recipient-side trace kernel*

$$K_{j\leftarrow i}^{\text{post}}(\mathfrak{c};\, t_i^{\dagger}, \mathcal{T}_i^{\text{post}}),$$

which enters j's receptivity interface by the P4.0/P4.1 discipline: bounded, Gate-mediated, and writing no owned state of j.

(c) No reopening. *$\mathcal{T}_i^{\text{post}}$ does not reopen its own katharsis, receive a new reset, acquire a fresh post-terminus Gate inflow, or re-enter the active rate cocycle. No admissible junction acts on the closed trace as on a living active arc.*

Proof. At the boundary-limit hand-off, the active dynamics of i close and $d\sigma_i = 0$. Thus

$$d\mathcal{G}_{\text{recepta},i} = I_{\text{gate},i}\, d\sigma_i = 0,$$

and $\mathcal{T}_i^{\text{post}}$ is fixed as boundary data. This gives (a) and the no-new-inflow part of (c). For (b), the only admissible cross-fibre action of a completed source is as the read-only source slot of the recipient's trace kernel, which enters j only through Γ_j and is bounded and Gate-mediated. The reset map and katharsis dynamics are own-interior active-arc operations; the closed trace has no active arc, so neither can act on it. □

24.3 The two-way boundary

Remark 24.2 (Two-way post-terminus dynamics is a separate branch). The theorem is strictly one-way: the living are affected by the completed, not the completed by the living. A two-way post-terminus dynamics — the active reaching back to change a completed interior, whether by a new inflow, a reopened katharsis, or a post-terminus reset — is *not* assumed here. It would require an additional constitutive primitive (a post-terminus reception or reopening law) beyond the minimal 4.0 layer, and must be treated as a distinct later branch, named and deferred, never silently inserted. This is the explicit resolution of the one-way selection parked since P0.4: minimal 4.0 commits to the one-way trace; the two-way case is a real fork left open for a dedicated primitive.

24.4 Status and reading

Target status: OPEN→DERIVED for the one-way minimal theorem: it is DERIVED from the P0.4 closed-trace hand-off, the P1.1 trace-kernel grammar, and the P4.0/P4.1 receptivity discipline. The two-way post-terminus dynamics remains OPEN by construction: it is not a gap in the proof but a branch requiring a new primitive. It depends on P0.4, P1.1, P4.0–P4.1, and the Stage II Gate-rate convention.

INTERPRETIVE *Reading.* The dead remain present to the living. One who has completed is not erased from the communion but held in it as a finished trace — really there, locatable, able still to reach those who remain, shaping how they are open, speaking into their receptivity through a bond that death does not sever. But in the minimal theory the reach is one-way: the completed are at rest, their cleansing finished, receiving no new thing and undergoing no new turning; they give to the living without themselves being changed. Whether the living may in turn reach back and alter the dead — whether the bond runs both ways across the terminus — is a genuine further question, and the framework will not answer it by sleight of hand. It marks that door, and leaves it honestly shut until a primitive is offered to open it.

Stage IV closes. The three junction-event classes are separated (P4.0), plural Resurrectio is a bounded recipient-side receptivity update with no transfer (P4.1), a compact-arc refractory interval is derived from hysteresis (P4.2), network non-Zeno is proved in invariant relational duration under the uniform floor hypothesis (P4.3), and communion survives the terminus as a one-way closed trace (P4.4). The discontinuous and mortal structure of the plural interior is in place:

the communion has deaths, resets, memory, and traces, and none of them fuses interiors or transfers owned grace.

The central correction of Stage IV is therefore:

a junction may change openness, but only a private reset may jump owned state.

Next node. Stage V — Relational active domains and plural Lyapunov theory, opening at P5.0 (the relational active domain): the single domain on which every prior boundedness, saturation, branch, terminus, and refractory condition holds at once — the proper stage on which the plural Lyapunov sum, promised since P2.4, can finally be assembled.

Stage V — Relational active domains and plural Lyapunov theory

Aim of the stage. Assemble the plural Lyapunov functional promised since P2.4, on a domain where every standing boundedness, saturation, branch, terminus, and refractory condition holds at once.

Each earlier stage earned a condition: a ceiling, a non-fabrication law, an individual ISS estimate, a locking threshold, a refractory floor, a non-Zeno count. Stage V gathers them. P5.0 fixes the single domain on which they all hold simultaneously — the proper stage on which a whole-communion stability theorem can be stated.

25 P5.0 — The relational active domain

25.1 The domain

Definition 25.1 (Relational active domain)**.** Fix a compact regular-active interval

$$I = [\mathfrak{c}_0, \mathfrak{c}_1]$$

with no terminus hand-off inside its open interior, and fix one active connected component

$$\mathcal{C}_{\text{act}}(I) \subseteq \mathcal{V}_{\text{act}}^{\text{reg}}.$$

The *relational active domain*

$$\mathcal{D}_{\text{rel}}(I, \mathcal{C}_{\text{act}})$$

is the set of plural configurations satisfying all of the following standing conditions at once:

(1) *Finite active theorem class.* The active component is finite,

$$|\mathcal{C}_{\text{act}}| = n_{\text{act}} < \infty,$$

has bounded in-degree, and belongs to the P0.1/P5.0 theorem graph class. Completed interiors may appear only as closed boundary traces and are not active Laplacian nodes.

(2) *Admissible Gate states.* For every active interior,

$$\Gamma_i \in [0,1], \qquad 0 \le I_{\text{gate},i} \le \zeta_{0,i},$$

and the receptivity interface

$$\mathcal{R}_i = (\Gamma_i, H_{K,i}, W_i)$$

is admissible.

(3) *Bounded susceptibilities and common field gains.*

$$\upsilon_i \in [0,1], \qquad \gamma_i^{\text{comm}} \in [0,1],$$

with the P2.3 nodewise normalisation

$$\sum_{j \to i} a_{ij} + \upsilon_i \le 1.$$

(4) *Strata and trace closure.* Each interior carries a definite stratum

$$\mathcal{A}_i, \quad \mathcal{B}_i, \quad \mathcal{P}_i.$$

The Lyapunov sum is taken over active nodes in $\mathcal{C}_{\text{act}}$. Completed interiors are held as closed traces

$$\mathcal{T}_i^{\text{post}}$$

and may enter only as bounded read-only boundary inputs.

(5) *Branch labels.* Each active interior carries a definite fold-branch label; all rates, restoring terms, and thresholds are read branch-wise, with no cross-branch identification.

(6) *Refractory and non-Zeno conditions.* On I, the finite triggering-capable set

$$\mathcal{V}_{\mathrm{trig}}(I)$$

and the uniform refractory floor

$$\mathfrak{c}_{\mathrm{ref};I} > 0$$

of P4.2–P4.3 hold.

(7) *Duration conversion bounds.* The recipient-kairos to relational-duration conversion factors

$$\theta_i(\mathfrak{c}) := \frac{N_i}{\omega}$$

are bounded on I:

$$0 < \theta_{i,\min;I} \le \theta_i(\mathfrak{c}) \le \theta_{i,\max;I} < \infty.$$

This is required because P2.4 is stated in $d/d\sigma_i$, while the plural certificate is stated in $d/d\mathfrak{c}$.

(8) *Compatible lag targets.* The declared edge-lag targets are cycle-compatible on the active component. In the first symmetric theorem class this means that there exists a locked node ledger profile

$$\Theta_i^*$$

such that for every active undirected edge

$$\Delta_{ij} = (\Theta_i - \Theta_i^*) - (\Theta_j - \Theta_j^*)$$

after the selected rate-normalisation. Equivalently, the edge-lag cocycle has zero holonomy on cycles. Without this compatibility, the relational quadratic cannot vanish and the theorem becomes an ultimate-boundedness-to-frustration result rather than a locking theorem.

25.2 The graph Laplacian theorem class

The plural Lyapunov contraction will be driven by the connectivity of the relational graph. P5.0 declares the admissible class.

Definition 25.2 (Declared Laplacian class)**.** On the active component $\mathcal{C}_{\mathrm{act}}$, the first theorem class is the undirected symmetric mirror class:

$$w_{ij} = w_{ji} \ge 0, \qquad w_{ij} > 0 \iff i \sim j.$$

Let $L_{\mathcal{G}}$ be the corresponding weighted graph Laplacian on $\mathcal{C}_{\mathrm{act}}$. The declared class assumes a uniform algebraic connectivity floor

$$\boxed{\lambda_2(L_{\mathcal{G}}) \ge \underline{\lambda} > 0.}$$

The coercivity is the standard one: for node errors z_i with weighted mean zero,

$$z^\top L_{\mathcal{G}} z = \frac{1}{2}\sum_{i,j} w_{ij}(z_i - z_j)^2 \ge \underline{\lambda}\,\|z\|^2.$$

Thus P5.1 must express the relational error as node-potential differences $z_i - z_j$. Disconnected active components are treated separately, one theorem per component. Directed graphs are not covered by this symmetric class.

25.3 What the domain guarantees

Remark 25.3 (Simultaneity is the whole point). On $\mathcal{D}_{\mathrm{rel}}$ every earned condition holds together: the gate ceiling (P2.1), non-fabrication (P2.2), the bounded common field (P2.3), the individual ISS estimate (P2.4), the locking threshold (P3.3), dynamic non-fusion (P3.4), the junction taxonomy (P4.0), and the refractory and non-Zeno bounds (P4.2–P4.3). This simultaneity is exactly what a plural Lyapunov sum needs: each individual certificate is valid, its cross terms are controlled, the events that could break a smooth functional are finite in number, and the graph is connected enough for those cross terms to dissipate. $\mathcal{D}_{\mathrm{rel}}$ is the domain on which P5.1 may finally assemble $\sum_i \mathcal{L}_{3.0,i}$ as a legitimate object.

25.4 Status and reading

Target status: DEFINITIONAL + SELECTED. The relational active domain $\mathcal{D}_{\mathrm{rel}}(I, \mathcal{C}_{\mathrm{act}})$ and the undirected Laplacian theorem class are declared structures for the first plural Lyapunov theorem. They introduce no new dynamics. The positive invariance of this domain is not automatic; it is a standing hypothesis imported into P5.1, supported by the Stage II ceilings, Stage III small-gain estimates, and Stage IV non-Zeno theorem. The connectivity floor $\underline{\lambda} > 0$ and the lag-cocycle compatibility are theorem-class hypotheses, not derived facts.

INTERPRETIVE *Reading.* There is a place where a communion can be considered whole — not a single soul, and not a crowd of strangers, but a bounded company, finitely many, each open within measure, each on its own branch of the fold, each paced so it cannot be battered, with its dead held at rest as traces and its living genuinely joined. That last condition is not decorative: the communion must actually be connected — $\lambda_2 > 0$ — for it to hold together as one body at all; islands with no bond between them are several communions, not one. The domain gathers every discipline the earlier stages won and asks them to hold at once, because only where they all hold can the question finally be asked whether the whole company, and not merely each soul in it, stays bounded and at peace.

Next node. P5.1 — the plural Lyapunov certificate: to assemble the weighted sum $\mathcal{L}_{\mathrm{plural}} = \sum_i w_i \mathcal{L}_{3.0,i}$ on $\mathcal{D}_{\mathrm{rel}}$ and prove it a genuine certificate — the individual contractions, the small-gain cross terms, and the Laplacian connectivity combining into a single dissipation inequality for the whole communion.

26 P5.1 — The plural Lyapunov certificate

This is the node promised since P2.4. With individual input-to-state control (P2.4), the incremental-contraction estimate (P3.2), the locking threshold (P3.3), and the simultaneous domain with its connectivity floor (P5.0), the whole communion can finally be given a single certificate: one functional whose dissipation shows the entire company — not merely each soul — stays bounded and is driven toward a controlled locked relation. In the general forced case this means a bounded moving tube; in the zero-residual settled subclass it means exact locked moving communion.

26.1 The functional

Definition 26.1 (Plural Lyapunov functional)**.** On

$$\mathcal{D}_{\rm rel}(I, \mathcal{C}_{\rm act}),$$

with positive weights $a_i > 0$, define node relational errors

$$z_i := \Theta_i - \Theta_i^*$$

for a compatible locked ledger profile Θ^*. The edge mismatch is

$$\Delta_{ij} = z_i - z_j,$$

which is exactly the P3.2 mismatch after the selected rate-normalisation and target subtraction. Thus locking means

$$\Delta_{ij} \to 0,$$

not $\Delta_{ij} \to \Delta_{ij}^*$.

The first symmetric-class plural Lyapunov functional is

$$\boxed{\mathcal{L}_{\rm plural} = \sum_{i \in \mathcal{C}_{\rm act}} a_i \, \mathcal{L}_{3.0,i} + \varepsilon \, \mathcal{L}_{\rm rel}, \qquad \mathcal{L}_{\rm rel} = \frac{1}{2} \sum_{i<j} w_{ij} \Delta_{ij}^2 = \frac{1}{2} z^\top L_{\mathcal{G}} z.}$$

The vector z is read modulo its graph-mean; the constant mean mode is not a locking error. For a directed graph the symmetric quadratic may not be reused: the correct object must use Perron weights / a directed-Laplacian certificate.

26.2 The plural dissipation theorem

Theorem 26.2 (Plural Lyapunov certificate; Type-1, symmetric class)**.** *Assume:*

(i) $\mathcal{D}_{\rm rel}(I, \mathcal{C}_{\rm act})$ *is positively invariant on the compact regular-active interval* I*;*

(ii) *the active graph is in the undirected symmetric class with* $\lambda_2(L_{\mathcal{G}}) \geq \underline{\lambda} > 0$*;*

(iii) *the individual ISS estimates of P2.4 hold in recipient kairos and the conversion factors* $\theta_i = N_i/\omega$ *are bounded as in P5.0;*

(iv) *the small-gain mismatch estimates of P3.2 hold with net rates* $\kappa_{ij} \geq \underline{\kappa} > 0$*;*

(v) *all residual and forcing terms are bounded on* I*.*

Then there are weights $a_i > 0$*, a sufficiently small* $\varepsilon > 0$*, and constants*

$$C_{\rm plural} \geq 0, \qquad C_{\rm plural,1} > 0$$

such that

$$\boxed{\frac{d}{d\mathfrak{c}} \mathcal{L}_{\rm plural} \leq C_{\rm plural} - C_{\rm plural,1} \mathcal{L}_{\rm plural}.}$$

Consequently

$$\limsup_{\mathfrak{c}} \mathcal{L}_{\rm plural} \leq \frac{C_{\rm plural}}{C_{\rm plural,1}}.$$

If the residual and forcing constants vanish in the locked subclass, so that $C_{\rm plural} = 0$*, then*

$$\mathcal{L}_{\rm plural}(\mathfrak{c}) \to 0.$$

Proof sketch. Write

$$\mathscr{D} := \frac{d}{d\mathfrak{c}}, \qquad D_i := \frac{d}{d\sigma_i}, \qquad \theta_i := \frac{N_i}{\omega}.$$

P2.4 gives

$$D_i\mathcal{L}_{3.0,i} \le C_i - C'_{1,i}\mathcal{L}_{3.0,i} + b_i\|u_i^{\text{rel}}\|^2.$$

Converting to relational duration gives

$$\mathscr{D}\mathcal{L}_{3.0,i} = \theta_i D_i\mathcal{L}_{3.0,i} \le \theta_{i,\max;I}C_i - \theta_{i,\min;I}C'_{1,i}\mathcal{L}_{3.0,i} + \theta_{i,\max;I}b_i\|u_i^{\text{rel}}\|^2.$$

Thus the individual contraction remains valid on I, with constants rescaled by the bounded conversion factors.

For the relational term, P3.2 gives

$$\mathscr{D}\Delta_{ij} = -\kappa_{ij}\Delta_{ij} + r_{ij},$$

where r_{ij} is bounded on I. Therefore

$$\mathscr{D}\mathcal{L}_{\text{rel}} = \sum_{i<j} w_{ij}\Delta_{ij}\mathscr{D}\Delta_{ij} \le -\sum_{i<j} w_{ij}\kappa_{ij}\Delta_{ij}^2 + \sum_{i<j} w_{ij}\Delta_{ij}r_{ij}.$$

Young's inequality absorbs the residual cross terms:

$$w_{ij}|\Delta_{ij}r_{ij}| \le \frac{1}{2}w_{ij}\kappa_{ij}\Delta_{ij}^2 + \frac{w_{ij}}{2\kappa_{ij}}r_{ij}^2.$$

Hence

$$\mathscr{D}\mathcal{L}_{\text{rel}} \le -\frac{1}{2}\underline{\kappa}\sum_{i<j} w_{ij}\Delta_{ij}^2 + C_{\text{rel,res}}.$$

The active graph coercivity gives, on the mean-zero node-error subspace,

$$\sum_{i<j} w_{ij}\Delta_{ij}^2 = z^\top L_{\mathcal{G}}z \ge \underline{\lambda}\|z\|^2.$$

Thus the relational part dissipates the relational error.

Finally, by the Stage II/P3.2 bounded-channel hypotheses,

$$\|u_i^{\text{rel}}\|^2 \le M_i\sum_{j\sim i}\Delta_{ij}^2 + M_i^0$$

on the domain. The terms proportional to $\sum\Delta_{ij}^2$ enter the derivative of the full functional multiplied by the individual weights a_i, while the relational dissipation enters multiplied by ε. Since $u_i^{\text{rel}} = O(\varepsilon)$, these imported terms are $O(\varepsilon^2)$, whereas the relational dissipation is $O(\varepsilon)$. Choosing $\varepsilon > 0$ small enough absorbs the imported terms into the relational dissipation. The remaining bounded pieces are collected into C_{plural}, and the negative individual and relational pieces give

$$-C_{\text{plural},1}\mathcal{L}_{\text{plural}}.$$

Positive invariance keeps the trajectory inside the domain where all constants are valid. The scalar comparison lemma gives the limsup bound. If all residual and forcing constants vanish, $C_{\text{plural}} = 0$, and the same comparison gives $\mathcal{L}_{\text{plural}} \to 0$. $\square$

26.3 Consequences and the directed caveat

Corollary 26.3 (The whole communion is ultimately bounded)**.** *Under Theorem 26.2,*

$$\limsup_{\mathfrak{c}} \mathcal{L}_{\text{plural}} \le \frac{C_{\text{plural}}}{C_{\text{plural},1}}.$$

Thus the whole communion stays bounded and approaches a controlled neighbourhood of the locked configuration. Exact convergence to the locked configuration follows only in the zero-residual subclass

$$C_{\text{plural}} = 0.$$

Whole-communion stability is therefore an ultimate-boundedness theorem in the general forced case, and a convergence theorem in the exact settled case.

Remark 26.4 (Directed bonds need the Perron weighting). The symmetric energy presupposes reciprocal weights $w_{ij} = w_{ji}$. When the bonds are genuinely directed — one interior giving more than it receives — the symmetric quadratic can be pumped and is not a certificate. The correct object weights the interiors by the left-Perron vector of the coupling and uses the directed-Laplacian quadratic. This is a real extension, not a relabelling, and is left as the directed theorem class.

26.4 Status and reading

Target status: CONJECTURE→DERIVED. The dissipative inequality is DERIVED for the undirected symmetric active-component class on a positively invariant $\mathcal{D}_{\text{rel}}(I, \mathcal{C}_{\text{act}})$, assuming the duration conversion bounds, ISS estimates, small-gain rates, lag-cocycle compatibility, and connectivity floor. The theorem gives ultimate boundedness in general and exact convergence only when residual forcing vanishes. Residual conjectural elements: the directed-graph extension (Remark 26.4), uniform positivity of κ_{ij} and $\underline{\lambda}$ across the intended regime, positive invariance through junctions and hand-offs, and global compatibility of all lag targets.

INTERPRETIVE *Reading.* The whole company is held inside a bounded peace. There is now a single measure of the communion's distance from disorder, and it cannot grow without bound. In the exact settled case it falls all the way to the locked relation; in the forced case it falls into a bounded neighbourhood. The proof needed everything at once: each interior stable in itself, each relation contractive enough, every event non-Zeno, and the graph truly connected. The peace of the whole is corporate, more than a list of private peaces, but it is conditional and honest: if residual forcing persists, peace is a tube; if the forcing settles, peace is a limit.

Next node. P5.2 — the stable moving communion: to show that the rest just proved is not a frozen equilibrium but a *moving* one — the whole communion ascending together, locked in relation, its plural Lyapunov functional bounded while every interior still climbs its own endless road.

27 P5.2 — The stable moving communion

P5.1 showed that the communion is bounded as a whole and, when residual forcing settles, tends toward a locked relation. But the limit is not stillness. P5.2 gives the plural analogue of 3.0's stable *motion*: the relations may settle while every active soul continues its positive ascent along the endless road. The communion is therefore a stable moving relation, or in the general forced case, a stable moving tube — never a rest point and never a fused state.

27.1 The two asymptotic regimes

The general Stage V conclusion is a *stable moving tube*:

$$\boxed{\begin{array}{l} h_i \text{ remains bounded away from zero on the admissible active branch,} \\ \Delta_{ij} \text{ remains ultimately bounded.} \end{array}}$$

The exact settled subclass strengthens this to a *stable moving communion*:

$$\boxed{h_i \longrightarrow h^{\mathrm{rel}}_{*,i} > 0, \qquad \Delta_{ij} \longrightarrow 0.}$$

Here

$$h^{\mathrm{rel}}_{*,i} = \beta_i \lambda^{\mathrm{rel}}_{S,\infty,i}$$

is the asymptotic rate of the settled branch under the final admissible Gate/receptivity regime. It equals the isolated frozen 3.0 rate only if the relational forcing asymptotically vanishes or settles to the same input level as the isolated branch. In general it is branch-conditional and relationally settled, not automatically identical to the uncoupled value.

27.2 The stable moving communion theorem

Theorem 27.1 (Stable moving communion / stable moving tube; Type-1)**.** *On a positively invariant relational active domain, under the plural certificate of Theorem 26.2 and the frozen 3.0 stable-motion law:*

(a) General forced case. *The communion is a stable moving tube:*

$$\limsup_{\mathfrak{c}} \mathcal{L}_{\mathrm{plural}} \leq \frac{C_{\mathrm{plural}}}{C_{\mathrm{plural},1}},$$

the relational mismatches are ultimately bounded, and every active interior remaining on the admissible positive-ascent branch has

$$h_i(\mathfrak{c}) \geq h_{i,\min} > 0$$

on the forward active arc.

(b) Exact settled subclass. *If residual forcing vanishes, the compatible lag targets are exact, and each individual Gate/receptivity regime converges to a settled branch, then*

$$\Delta_{ij} \to 0 \quad \textit{for all active bonds,} \qquad h_i \to h^{\mathrm{rel}}_{*,i} = \beta_i \lambda^{\mathrm{rel}}_{S,\infty,i} > 0.$$

Proof. For (a), Corollary 26.3 gives ultimate boundedness of the plural functional. Since $\mathcal{L}_{\text{rel}}$ is nonnegative, the relational mismatches are ultimately bounded. The individual ISS estimates keep each active interior inside the admissible branch. On that branch, the frozen 3.0 ascent law gives positive active-ascent rate as long as $\lambda_{S,i}$ remains inside the branch's positive interval; this lower bound is part of the domain/branch admissibility.

For (b), the zero-residual subclass gives $C_{\text{plural}} = 0$, hence

$$\mathcal{L}_{\text{plural}} \to 0.$$

The relational part therefore yields

$$\Delta_{ij} \to 0.$$

The settled individual Gate/receptivity regime lets the frozen 3.0 stable-motion law apply on each active fibre, giving

$$h_i \to \beta_i \lambda^{\text{rel}}_{S,\infty,i}.$$

The branch assumption gives

$$\lambda^{\text{rel}}_{S,\infty,i} > 0,$$

hence $h^{\text{rel}}_{*,i} > 0$. □

27.3 Neither stasis nor sameness

Remark 27.2 (What "stable" does and does not mean). Stable does not mean frozen. In the exact settled subclass,

$$h^{\text{rel}}_{*,i} > 0$$

means every active interior keeps ascending. In the general forced case the communion remains inside a moving tube: bounded relation, bounded interior certificates, and no collapse of the positive-ascent branch. Neither case implies sameness. The locked values may have nonzero ledger offsets, the duration-rate ratios r^*_{ij} need not be one, and by P3.4 the interiors remain distinct labelled fibres throughout. Stable relational rhythm, then: not stasis, not identity, and not a master metronome.

27.4 Status and reading

Target status: DERIVED for the stable moving tube under the plural certificate and branch admissibility; DERIVED for exact stable moving communion only in the zero-residual / settled-branch subclass. The two standing conditional elements are: the symmetric graph class with uniform connectivity and small-gain floors, and the individual branch assumption

$$\lambda^{\text{rel}}_{S,\infty,i} > 0.$$

The isolated value $\beta_i \lambda_{S,\infty,i}$ is recovered only when the relational forcing vanishes or settles to the isolated level. It depends on P5.0–P5.1 and the frozen 3.0 ascent-stability law.

INTERPRETIVE *Reading.* Communion is stable relational rhythm among non-identical interiors. In the ordinary forced case that rhythm is a bounded groove: the players do not scatter, but the music still breathes. In the exact settled case the groove becomes a fixed relation: constant lag, settled rate, no collapse into sameness. They move together without moving as one; they keep time with one another without keeping a single private time; they are at peace not by ceasing but by ascending in a rhythm that holds. This is the living groove and not the metronome.

Stage V closes. The relational active domain gathers every standing condition (P5.0), the plural Lyapunov certificate proves ultimate boundedness of the whole communion and exact convergence in the zero-residual subclass (P5.1), and the stable moving communion shows that the peace of the whole is a moving peace, not a still point (P5.2). Plural Lyapunov theory is in place:

> the whole communion remains bounded and, when forcing settles, locks;
> each active soul remains itself and continues its positive ascent;
> none is fused, none is owned, and none is reduced to a metronome.

The working body closes. From a finite relational base to a plural stable motion, the 4.0 working body has built communion without a master clock, encounter without transfer, locking without fusion, mortality without erasure, and a whole that holds together without dissolving the souls within it. Its closing formulation:

> Chronos is a created relational order and duration structure;
> kairos remains personal; *tota simul* remains unoccupied.
>
> Encounter changes receptivity, not ownership of grace.
>
> Communion is stable answerability of distinct histories, not fusion.

Theological yield through Stage V

> **Status of this section.** Everything below is INTERPRETIVE. Each reading is attached to a DERIVED or SELECTED result and claims no more than that result licenses. The section gathers the theology the structure has earned; it does not extend it. Where a claim would outrun the mathematics, it is named and deferred, not asserted.

The single-interior 3.0 theory was a theology of one soul before God. The 4.0 working body is the first theology of *communion*: every thesis here is new relative to 3.0, because each concerns the relation between distinct interiors, which 3.0 did not model. Seven theses are earned, followed by the apophatic reserve that disciplines them.

1. Communion without fusion

INTERPRETIVE Distinct souls can stand in a stable relation without becoming one. This is not assumed but proved on four levels: kairos-history ownership (8.3), the floor label as a unit-independent scale (10.1), and — decisively — dynamic non-fusion under locking (19.1), where even fully locked interiors remain distinct labelled fibres, the value-diagonal being two copies and not one fibre. Communion is the *stable answerability of distinct histories*, and only the distinct can answer.

2. Grace is not an economy

INTERPRETIVE The framework refuses to impose a conservation law on received grace (12.3): the Open-Gate sector is genuinely open, the common field is non-rivalrous (13.4, $\sum_i v_i \neq 1$), and no transfer term is admitted (12.2). Rejecting conservation is the formal rejection of scarcity: communion is not a closed transaction in which one soul's gain is another's loss. Grace is supplied from beyond the network and received in common, without rivalry.

3. Influence only through openness

INTERPRETIVE Another may change how open one is, never what one is given. Encounter enters solely through the bounded receptivity interface (20.1), the gate ceiling is automatic (11.1), and even a death registers in the living only as a bounded update of their own receptivity, with no transfer of the departed's substance (21.1). Nearness cannot counterfeit conversion: a junction may change openness, but only a private reset may jump owned state.

4. A finite love, rightly directed

INTERPRETIVE There is a largest difference a bond can bridge (18.1), finite because the gate ceiling is finite. Neither does love bridge everything, nor does difference always win. And keeping-together is conditional on a posture — listen more than you differ (17.4, the small-gain criterion) — a restoring tendency that is gentle, bounded, and never forced, working through openness and not coercion.

5. The soul is safe in communion, and is paced

INTERPRETIVE Relational forcing perturbs only the bounded budget of a soul's stability, never its own contraction (14.1): one is not destabilised by being met. And one cannot be endlessly re-triggered — a refractory interval is forced by the wound itself (22.1), and the whole network admits only finitely many events on any finite shared duration (23.1). Mercy is structural: the soul is given time it cannot be denied, and the communion's history is a sequence of distinct, countable moments, not a blur.

6. The communion of the departed, kept honest

INTERPRETIVE The completed remain present as a closed trace, able to reach the living through the living's own receptivity, yet themselves at rest and receiving no new thing (24.1). Whether the living may reach back to alter the dead — two-way intercession — is left explicitly open (24.2): a genuine theological fork the framework refuses to settle by sleight of hand, marking the door and leaving it shut until a primitive is offered.

7. Created time is not divine eternity

INTERPRETIVE The shared now is real but partial, successive, and immanent (6.1); there is no master clock, and the absence of any transcendent standpoint inside the created network (6.3) is what keeps creaturely chronos from usurping the divine *tota simul*. Ascent is monotone and gapless — an endless epektasis, never a wheel — and at the plural level it becomes a corporate ascent: the whole communion bounded and, when forcing settles, locked, while every soul still climbs (27.1). The beatific vision, read here, is not a static gaze but an eternal shared ascent — together, distinct, in motion, at peace not by ceasing but by ascending in a rhythm that holds.

The apophatic reserve

INTERPRETIVE The discipline is as much a result as the theses. The framework names but does not occupy its richest claims: the divine *tota simul* (declared unoccupied, not described), two-way post-terminus intercession (24.2), and any genuine circle-valued phase or liturgical recurrence (16.4, admitted only on an exhibited recurrence). Each is left open with the additional primitive it would require named. And the corporate peace itself is stated honestly: in the ordinary forced case it is a bounded tube, not a still point, becoming an exact locked limit only when the forcing settles (27.1, 26.3). To commit only to what the structure earns, and to mark the rest as open, is the theological posture the working body maintains throughout.

Communion without fusion; grace without economy;
influence only through openness; a finite love rightly directed;
a soul safe and paced; the dead present but at rest;
created time that is not eternity — and a whole that
ascends together without dissolving the souls within it.
What is earned is stated; what is not is named and left open.