Formal → Coupling Hypothesis

Mathematical and Operational Admissibility of Operators Post-$_{4b}$

1 Scope and Positioning

This document formalizes the → *coupled pipeline* used in Chamber XXXV. is treated as a *global selection layer* acting on an ensemble of discrete structures; is treated as a *geometric / stability layer* acting on the -filtered ensemble. The goal is to define (i) a mathematically testable coupling hypothesis and (ii) a class of -operators admissible *after* the canonical $_{4b}$ selection.

Canonical upstream fact (assumed). $_{4b}$ is the R_Λ-*aligned* bandpass selector: it filters structures by proximity to a shared baseline target V_{target} and yields a coherent acceptance band while preserving protected macro-invariants (within a fixed guardrail). Chamber XXXV takes this as given and does not redefine .

2 Objects, Ensemble, and Baseline Metrics

2.1 Ensemble

Let $\mathcal{S}$ denote the space of finite structures (e.g., graphs) generated under a controlled generator family. An experimental run constructs a finite ensemble

$$E = \{S_1, \ldots, S_M\} \subset \mathcal{S}.$$

2.2 Baseline scalar proxy V and residual R_Λ

Each structure S is assigned scalar observables (examples used upstream):

$$Z_0(S) \in [0,1], \qquad \text{Gap}(S) > 0, \qquad \text{CycleT}(S) \geq 0.$$

Define the vacuum proxy

$$V(S) \;=\; \alpha\, Z_0(S) \;+\; \beta\, \frac{1}{\text{Gap}(S) + \varepsilon} \;+\; \gamma\, \text{CycleT}(S),$$

with fixed weights $\alpha, \beta, \gamma \geq 0$ and small $\varepsilon > 0$.

Let the baseline target be a functional of the baseline ensemble (e.g., the median):

$$V_{\text{target}} \;=\; \text{median}\{V(S) : S \in E\}.$$

Define the ensemble mean

$$\overline{V}(E) \;=\; \frac{1}{|E|} \sum_{S \in E} V(S),$$

and the residual proxy

$$R_\Lambda(E) \;=\; \frac{|\overline{V}(E) - V_{\text{target}}|}{\max(|V_{\text{target}}|, \varepsilon)}.$$

2.3 Protected macro-invariant vector $I(S)$

Let $I : \mathcal{S} \to \mathbb{R}^d$ denote protected observables ("-protected macro invariants") computed per structure and then ensemble-averaged. Write

$$I(S) = \big(I_1(S), \ldots, I_d(S)\big), \qquad \overline{I}(E) = \frac{1}{|E|} \sum_{S \in E} I(S).$$

A canonical example set is:

$$I_1(S) = \frac{\rho(A_S)}{n_S}, \quad I_2(S) = \frac{\text{Tr}(L_S)}{n_S}, \quad I_3(S) = H(\text{degree distribution of } S),$$

but the admissibility framework below does not depend on the specific choice, only that I is fixed and guarded.

3 $_{4b}$ Selection as a Fixed Upstream Map

$_{4b}$ is an acceptance mask defined on E by proximity to the shared baseline target V_{target}. Let $k = \lfloor fM \rfloor$ where $f \in (0, 1)$ is the keep fraction. Define the score

$$\text{score}(S) = -|V(S) - V_{\text{target}}|.$$

Let $E_\Omega \subset E$ be the subset of size k maximizing the score (i.e., closest to V_{target}):

$$E_\Omega = \Omega_{4b}(E) = \text{Top-}k \text{ structures by score}(S).$$

The selection is *canonical* in the sense that it is deterministic given E, V_{target}, and f.

4 The → Coupling Hypothesis

4.1 as a stabilizing operator family

A -operator is a map acting on structures:

$$\tau_\theta : \mathcal{S} \to \mathcal{S},$$

parameterized by θ (which may include step size, iteration count, smoothing strength, etc.). We extend τ_θ to ensembles pointwise:

$$\tau_\theta(E_\Omega) = \{\tau_\theta(S) : S \in E_\Omega\}.$$

4.2 Two core requirements: (i) stationarity improvement, (ii) invariant preservation

Define the residual after :

$$R_\Lambda^\tau(\theta; E) := R_\Lambda\big(\tau_\theta(E_\Omega)\big), \quad \text{where } E_\Omega = \Omega_{4b}(E).$$

Define relative drift of protected invariants:

$$\text{drift}_j(\theta; E) := \frac{\big|\,\overline{I}_j(\tau_\theta(E_\Omega)) - \overline{I}_j(E)\,\big|}{\max\big(\big|\overline{I}_j(E)\big|\,, \eta\big)}, \qquad j = 1, \ldots, d,$$

with a small floor $\eta > 0$ to prevent denominator blow-up.

4.3 Coupling hypothesis (mathematical statement)

Hypothesis H(→): Existence of a -locking regime. There exists a nonempty parameter region Θ_{lock} such that for a nontrivial class of generators (families of E), for all $\theta \in \Theta_{\text{lock}}$,

$$\text{(H1) Residual contraction:} \quad R_\Lambda^\tau(\theta; E) \;\le\; (1-\kappa)\, R_\Lambda(E_\Omega) \quad \text{for some } \kappa > 0, \tag{1}$$

$$\text{(H2) -protection:} \quad \max_{1\le j\le d} \text{drift}_j(\theta; E) \;\le\; \delta, \tag{2}$$

$$\text{(H3) Nondegeneracy:} \quad \frac{|E_\Omega|}{|E|} \in [a, b] \subset (0, 1) \quad \text{(coherent acceptance band, not 0\% or 100\%).} \tag{3}$$

Interpretation. $_{4b}$ creates a coherent band E_Ω in which can act *contractively* on stationarity residual while leaving protected macro-invariants stable. The coupling is *not* " always improves things"; it is the claim that a *regime exists* where becomes an admissible stabilizer *because has already selected the right band.*

4.4 Operational hypothesis (testable form)

For a run family $\{E^{(r)}\}_{r=1}^{R}$ (different seeds, same configuration class), define:

$$\mathrm{Improve}_r(\theta) = \mathbf{1}\Big(R_\Lambda^\tau(\theta; E^{(r)}) < R_\Lambda(E_\Omega^{(r)}) - \Delta\Big),$$

$$\mathrm{Stable}_r(\theta) = \mathbf{1}\Big(\max_j \mathrm{drift}_j(\theta; E^{(r)}) \le \delta\Big),$$

$$\mathrm{Band}_r = \mathbf{1}\Big(|E_\Omega^{(r)}|/|E^{(r)}| \in [a, b]\Big).$$

Then the coupling hypothesis is supported if there exists θ such that

$$\frac{1}{R}\sum_{r=1}^{R} \mathrm{Improve}_r(\theta)\,\mathrm{Stable}_r(\theta)\,\mathrm{Band}_r \;\ge\; p_0,$$

for a chosen confidence level p_0 (e.g., $p_0 = 0.8$).

5 Admissibility Framework for Operators (Post-$_{4b}$)

5.1 Definition (Admissible operator)

Fix guardrails $\Delta > 0$ (minimum residual improvement), $\delta > 0$ (maximum invariant drift), and $[a, b] \subset (0, 1)$ (coherent band window). A -operator family $\{\tau_\theta\}_{\theta\in\Theta}$ is *admissible post-*$_{4b}$ if there exists $\theta \in \Theta$ such that, across a prescribed multi-seed run set,

1. **Residual improvement (-conditioned):**
$$R_\Lambda\big(\tau_\theta(E_\Omega)\big) \le R_\Lambda(E_\Omega) - \Delta.$$
2. **-protection:**
$$\max_j \mathrm{drift}_j(\theta; E) \le \delta.$$
3. **No -backreaction:** τ_θ does not modify V_{target} and does not re-run internally. (is upstream and fixed.)
4. **Structural well-posedness:** $\tau_\theta(S) \in \mathcal{S}$ for all S, and the map is deterministic given (S, θ).

5.2 Important non-admissible behaviors

A operator is *not* admissible post-$_{4b}$ if it:

- trivially drives R_Λ down by redefining V_{target} or modifying the residual formula;
- masks instability by collapsing protected invariants via averaging artifacts;
- induces acceptance degeneracy (effectively requiring Ω to accept nearly all or nearly none to pass);
- is non-deterministic unless its randomness is treated as part of the operator definition and audited.

6 Canonical -Operator Families Admissible *in Principle*

Below are -operator families designed to be admissible *after* $_{4b}$. Each is expressed abstractly so it can be instantiated on graphs, lattices, or other discrete structures.

6.1 $_A$: Laplacian heat-flow smoothing (local regularizer)

Let A_S be adjacency, L_S the (combinatorial) Laplacian. Define a softened edge-weight matrix $W_S(t)$ evolving by a discrete heat step:

$$W_S(t+1) = \Pi_{\mathcal{W}}\Big(W_S(t) - \lambda\, L(W_S(t))\Big),$$

where $L(\cdot)$ is a Laplacian-like operator on weights, $\lambda > 0$ is step size, and $\Pi_{\mathcal{W}}$ projects back to an admissible weight space (e.g., nonnegative, symmetric, bounded). Then τ_A outputs a new structure $\tau_A(S)$ by thresholding or sampling from $W_S(T)$ at time T.

Admissibility intuition. Local smoothing can reduce high-frequency irregularity (a proxy for curvature noise) while keeping global macro-invariants stable if λ and T are small.

6.2 $_B$: Spectral band-limiter (mode trimming without target shift)

Let $\{\lambda_i, u_i\}$ be Laplacian eigenpairs (or an approximation). Define a band-pass reconstruction using modes $i \in \mathcal{B}$:

$$\widetilde{L}_S = \sum_{i \in \mathcal{B}} \lambda_i \, u_i u_i^\top,$$

and reconstruct $\tau_B(S)$ as the closest structure (in a chosen metric) whose Laplacian matches $\widetilde{L}_S$ within tolerance. Operationally this can be done by iterative edge rewiring minimizing $\|L_{S'} - \widetilde{L}_S\|$.

Admissibility intuition. $_{4b}$ already selects near a global stationarity band; $_B$ should only trim unstable modes *within* that band (no moving the target).

6.3 $_C$: Divergence-minimizing rewiring (least-divergence step)

Define a structure-level divergence functional $D(S)$ (distinct from V), e.g. a flux inconsistency measure computed from local constraints. Define as a constrained descent step:

$$\tau_C(S) = \arg \min_{S' \in \mathcal{N}(S)} D(S') \quad \text{subject to} \quad |V(S') - V(S)| \leq \epsilon_V,$$

where $\mathcal{N}(S)$ is a local neighborhood (single-edge rewires, small edits) and ϵ_V is a small tolerance to prevent -target backreaction.

Admissibility intuition. This is a pure "stability step" that cannot cheat by shifting V far away from -selected values.

6.4 $_D$: Curvature proxy equalization (local curvature flattening with invariant locks)

Assume a local curvature proxy $\kappa_S(v)$ defined on nodes/regions (e.g., degree-based curvature, cycle tension density, or Ollivier-style curvature if available). Define as a local equalization with hard locks on protected totals:

$$\kappa_S \mapsto \kappa_S - \mu \nabla \mathcal{E}(\kappa_S), \quad \text{with constraints } \overline{I} \text{ fixed within tolerance.}$$

This is implemented by local edits that reduce curvature variance while enforcing constraints.

Admissibility intuition. If is truly "geometric stabilization", it should appear as curvature-variance reduction under strict macro locks.

6.5 $_E$: Multi-scale (coarse-grain then refine, no target shift)

Define a coarse-graining operator C and a refinement operator R. Let

$$\tau_E(S) = R\big(\tau_{\text{coarse}}(C(S))\big),$$

where τ_{coarse} is any admissible candidate on the coarse structure, and refinement preserves protected invariants (within tolerance) by construction.

Admissibility intuition. If is resolution-stable, a properly constrained multi-scale should pass across n and across generator families.

7 Admissibility Tests and Verdict Logic for Chamber XXXV

7.1 Inputs and fixed constants

For each run:

- ensemble size M, node scale n, generator family, seed set;
- $_{4b}$ keep fraction f and fixed upstream weights (α, β, γ);
- guardrails: Δ (residual improvement), δ (max drift), $[a, b]$ (acceptance band), floors ε, η.

7.2 Core pipeline (operational algorithm)

1. Generate E.
2. Compute $V(S)$, V_{target}, $R_\Lambda(E)$, and $\overline{I}(E)$.
3. Apply $_{4b}$ to get $E_\Omega = \Omega_{4b}(E)$ and compute $R_\Lambda(E_\Omega)$.
4. Apply candidate: $E_{\Omega,\tau} = \tau_\theta(E_\Omega)$.
5. Compute $R_\Lambda(E_{\Omega,\tau})$ and drifts drift_j.
6. Verdict for (E, θ) is PASS iff:

$$R_\Lambda(E_{\Omega,\tau}) \le R_\Lambda(E_\Omega) - \Delta, \qquad \max_j \text{drift}_j \le \delta, \qquad |E_\Omega|/|E| \in [a, b].$$

7.3 Coupling signatures to report (beyond PASS/FAIL)

A candidate that passes should also be reported with:

- residual contraction ratio

$$\mathrm{CR}(\theta) = \frac{R_\Lambda(E_{\Omega,\tau})}{R_\Lambda(E_\Omega)} \in (0,1);$$

- protected drift profile $(\mathrm{drift}_1, \ldots, \mathrm{drift}_d)$;
- sensitivity curves over θ showing a stable lock window Θ_{lock} (not a single tuned point).

8 Minimal -Admissibility Checklist (Post-$_{4b}$)

A operator family is admissible post-$_{4b}$ only if it satisfies all of:

1. **-fixedness:** does not redefine V_{target} and does not internally re-run .
2. **Residual contraction:** R_Λ strictly improves beyond Δ on the -filtered ensemble.
3. **Macro stability:** protected drift stays below δ across seeds.
4. **Lock window:** there exists a nontrivial Θ_{lock} interval (robustness).
5. **Cross-family sanity:** does not only pass on one degenerate generator regime.

9 Admissible τ Operators Post-Ω_{4b}

This section defines when a τ-operator is admissible after canonical Ω_{4b} selection. Throughout, Ω_{4b} is treated as fixed and upstream.

9.1 Setting

Let $E = \{S_1, \ldots, S_M\}$ be a finite ensemble of structures. Let

$$E_\Omega := \Omega_{4b}(E)$$

denote the Ω_{4b}-filtered sub-ensemble with acceptance ratio

$$\frac{|E_\Omega|}{|E|} \in (0,1).$$

Each structure S carries:

- a scalar stationarity proxy $V(S)$,
- a protected macro-invariant vector $I(S) \in \mathbb{R}^d$.

Define ensemble averages

$$\overline{V}(E) = \frac{1}{|E|} \sum_{S \in E} V(S), \qquad \overline{I}(E) = \frac{1}{|E|} \sum_{S \in E} I(S),$$

and the residual proxy

$$R_\Lambda(E) = \frac{|\overline{V}(E) - V_{\text{target}}|}{\max(|V_{\text{target}}|, \varepsilon)}.$$

9.2 Definition (-Operator)

A τ-operator is a deterministic map

$$\tau_\theta : \mathcal{S} \to \mathcal{S},$$

parameterized by θ, extended pointwise to ensembles:

$$\tau_\theta(E_\Omega) = \{\tau_\theta(S) : S \in E_\Omega\}.$$

9.3 Definition (Admissibility Post-Ω_{4b})

Fix guardrails:

$\Delta > 0$ (minimum residual improvement), $\delta > 0$ (maximum invariant drift), $[a, b] \subset (0, 1$

A τ-operator family $\{\tau_\theta\}_{\theta \in \Theta}$ is *admissible post-*Ω_{4b} if there exists $\theta \in \Theta$ such that all of the following hold:

(A1) -fixedness. τ_θ does not modify V_{target} and does not internally reapply or emulate any Ω selection.

(A2) Residual contraction.

$$R_\Lambda\big(\tau_\theta(E_\Omega)\big) \le R_\Lambda(E_\Omega) - \Delta.$$

(A3) Macro-invariant protection. For each protected component I_j,

$$\frac{\big|\overline{I}_j\big(\tau_\theta(E_\Omega)\big) - \overline{I}_j(E)\big|}{\max(|\overline{I}_j(E)|, \eta)} \le \delta, \quad j = 1, \dots, d,$$

with fixed denominator floor $\eta > 0$.

(A4) Non-degeneracy. The Ω acceptance band remains coherent:

$$\frac{|E_\Omega|}{|E|} \in [a, b].$$

(A5) Structural well-posedness. For all $S \in E_\Omega$, $\tau_\theta(S) \in \mathcal{S}$ and the map $S \mapsto \tau_\theta(S)$ is deterministic given θ.

9.4 Definition (-Locking Regime)

A parameter subset $\Theta_{\text{lock}} \subset \Theta$ is called a *τ-locking regime* if admissibility conditions (A1)–(A5) hold uniformly for all $\theta \in \Theta_{\text{lock}}$ across a prescribed multi-seed run family.

Existence of a nontrivial Θ_{lock} is interpreted as evidence that τ acts as a genuine stabilizing completion conditioned on Ω_{4b}, rather than as a tuned or degenerate post-processing step.

9.5 Non-Admissible Behaviors

A τ operator is *not* admissible post-Ω_{4b} if it:

- reduces R_Λ by redefining V, V_{target}, or the residual itself;
- violates macro-invariant protection beyond δ;
- requires near-total acceptance or near-total rejection to pass;
- relies on uncontrolled randomness not included in θ.

10 Necessary (Not Sufficient) Conditions for τ-Admissibility Post-Ω_{4b}

This section derives conditions that *must* hold for a τ-operator to be admissible post-Ω_{4b}, but which alone do not guarantee admissibility.

10.1 Notation

Let E be a baseline ensemble and $E_\Omega = \Omega_{4b}(E)$ the fixed Ω_{4b}-filtered subensemble. For a τ operator τ_θ, define

$$E_{\Omega,\tau} := \tau_\theta(E_\Omega).$$

Write

$$\overline{V}(X) = \frac{1}{|X|}\sum_{S\in X} V(S), \qquad \overline{I}(X) = \frac{1}{|X|}\sum_{S\in X} I(S),$$

and

$$R_\Lambda(X) = \frac{\left|\overline{V}(X) - V_{\text{target}}\right|}{\max(|V_{\text{target}}|, \varepsilon)}.$$

10.2 NC0: Well-posedness and -fixedness (structural necessities)

NC0a (Well-posedness). A necessary condition is:

$$\tau_\theta(S) \in \mathcal{S} \quad \forall S \in E_\Omega,$$

and τ_θ must be deterministic given θ (or randomness must be explicitly parameterized and audited as part of θ).

NC0b (-fixedness). A necessary condition is that τ_θ does *not* redefine V, does *not* redefine V_{target}, and does *not* reapply or emulate Ω selection internally. Otherwise any improvement in R_Λ is not attributable to τ as a post-Ω_{4b} operator.

10.3 NC1: Residual improvement implies a mean-shift inequality

Admissibility requires an improvement margin $\Delta > 0$:

$$R_\Lambda(E_{\Omega,\tau}) \leq R_\Lambda(E_\Omega) - \Delta.$$

Let $D := \max(|V_{\text{target}}|, \varepsilon)$ and define

$$d_\Omega := \left|\overline{V}(E_\Omega) - V_{\text{target}}\right|, \qquad d_{\Omega,\tau} := \left|\overline{V}(E_{\Omega,\tau}) - V_{\text{target}}\right|.$$

Then admissibility implies the necessary inequality

$$d_{\Omega,\tau} \leq d_\Omega - \Delta D.$$

Equivalently,

$$\left|\overline{V}(E_{\Omega,\tau}) - V_{\text{target}}\right| \leq \left|\overline{V}(E_\Omega) - V_{\text{target}}\right| - \Delta \max(|V_{\text{target}}|, \varepsilon).$$

Interpretation. A post-Ω_{4b} admissible τ must, at minimum, move the Ω-filtered mean $\overline{V}$ *closer* to the fixed target by an amount large enough to clear the improvement margin.

10.4 NC2: If $\overline{V}$ is unchanged, admissibility is impossible

If τ_θ preserves the ensemble mean of V on E_Ω, i.e.

$$\overline{V}(E_{\Omega,\tau}) = \overline{V}(E_\Omega),$$

then

$$R_\Lambda(E_{\Omega,\tau}) = R_\Lambda(E_\Omega),$$

so the strict improvement requirement cannot hold for any $\Delta > 0$. Hence a necessary condition for admissibility (when $\Delta > 0$) is:

$$\overline{V}(E_{\Omega,\tau}) \neq \overline{V}(E_\Omega),$$

and moreover it must shift in the direction that reduces $|\overline{V} - V_{\text{target}}|$.

10.5 NC3: Macro-invariant protection implies componentwise drift bounds

Fix drift tolerance $\delta > 0$ and denominator floor $\eta > 0$. Admissibility requires for each protected component I_j:

$$\frac{\left|\overline{I}_j(E_{\Omega,\tau}) - \overline{I}_j(E)\right|}{\max(|\overline{I}_j(E)|, \eta)} \leq \delta.$$

Therefore a necessary condition is the componentwise absolute bound

$$\left|\overline{I}_j(E_{\Omega,\tau}) - \overline{I}_j(E)\right| \leq \delta \max(|\overline{I}_j(E)|, \eta), \qquad j = 1, \ldots, d.$$

Interpretation. No matter how strongly τ improves R_Λ, if it forces any protected macro observable outside this envelope, it cannot be admissible.

10.6 NC4: Admissibility forces a compatibility window between V-shift and invariant drift

Let $\Delta\overline{V} := \overline{V}(E_{\Omega,\tau}) - \overline{V}(E_\Omega)$ and $\Delta\overline{I} := \overline{I}(E_{\Omega,\tau}) - \overline{I}(E)$. A necessary condition is existence of parameters θ such that:

$$\Delta\overline{V} \text{ reduces } |\overline{V} - V_{\text{target}}| \quad \text{and} \quad \|\Delta\overline{I}\|_\infty \text{ remains within the guardrail.}$$

Operationally, if empirical sweeps over θ show that any parameter value producing the required V-shift necessarily violates at least one protected component bound, then no admissible θ exists for that τ family.

10.7 NC5: Nondegeneracy is required before τ is even evaluated

Admissibility post-Ω_{4b} presupposes that the Ω band is coherent:

$$\frac{|E_\Omega|}{|E|} \in [a,b] \subset (0,1).$$

Thus, if a run yields $\frac{|E_\Omega|}{|E|} \notin [a,b]$, then *no* τ operator can be declared admissible on that run under the fixed admissibility protocol, because the upstream condition for meaningful post-selection stabilization fails.

10.8 NC6: Robustness prerequisite (multi-seed necessity)

Let $\mathcal{R}$ be a prescribed set of seeds producing ensembles $\{E^{(r)}\}_{r\in\mathcal{R}}$ under the same configuration. A necessary (but not sufficient) robustness condition is that there exists at least one parameter value θ such that:

$$\exists\,\theta \in \Theta : \quad \text{(NC1–NC5) hold for a nontrivial fraction of } r \in \mathcal{R}.$$

If every seed violates the improvement inequality (NC1) or violates protection bounds (NC3), the τ family is ruled out as admissible under that configuration class.

10.9 Summary of Necessary Conditions

For a τ family to be admissible post-Ω_{4b}, it is necessary that:

1. (NC0) τ_θ is well-defined on E_Ω and -fixed (no target/backreaction).
2. (NC1) τ must reduce $\left|\overline{V} - V_{\text{target}}\right|$ by at least $\Delta \max(|V_{\text{target}}|, \varepsilon)$.
3. (NC2) τ cannot leave $\overline{V}$ invariant if $\Delta > 0$.
4. (NC3) Protected macro-invariant drifts must lie inside the componentwise envelope.
5. (NC4) There must exist a compatibility window where required V-shift does not force invariant violation.
6. (NC5) The upstream Ω acceptance rate must be nondegenerate ($[a,b]$).
7. (NC6) The above must occur robustly across seeds (at least occasionally), otherwise the family is excluded.

These conditions are necessary; passing them does not guarantee admissibility because admissibility additionally requires a coherent lock window, cross-family sanity, and sustained performance across runs.

11 Summary

The $\to$ coupling hypothesis asserts the existence of a parameter regime where acts as a *stabilizing completion* on the $_{4b}$-selected band: it contracts the stationarity residual while preserving -protected macro invariants. The admissible families are those that cannot “cheat” by moving targets or destabilizing invariants, and that exhibit a robust lock window across seeds and generator families.