

# Generative Collapse Dynamics (UMCP/GCD): Enabling Cross-Domain Comparability via Contract-Frozen Kernel Invariants and Typed Return

Whitepaper — release line v2.1.5 (March 2026)

Clement Paulus

ORCID: 0009-0000-6069-8234

TZ: America/Chicago

Paper DOI: 10.5281/zenodo.18819238

Canon (context): 10.5281/zenodo.17756705 (PRE) · 10.5281/zenodo.18072852 (POST)

Weld-ID (context): W-2025-12-31-PHYS-COHERENCE

Repo: GCD (release line: v2.1.5) (commit: ) (Zenodo software DOI: 10.5281/zenodo.18616616)

**Jurisdiction / Tier discipline.** This whitepaper operates under the UMCP contract-first rule: the Tier-1 ledger  $\{\omega, F, S, C, \tau_R, \kappa, \text{IC}\}$  is computed *only* by the kernel on the bounded trace  $\Psi(t) \in [0, 1]^n$  under the active frozen contract, and these symbols are reserved (no redefinition by prose, domain semantics, or overlays). Tier-0 specifies the measurable interface—observables  $x(t)$ , adapter  $N_K$ , bounds/OOR and missingness policy, and return settings  $(\|\cdot\|, \eta, D_\theta)$ —and therefore defines the object being measured. Tier-2 objects (scores, labels, narratives, auxiliary diagnostics) may be introduced only as diagnostics and must not be used as regime or weld gates unless promoted through an explicit seam, a new frozen run, and a recorded weld decision. Segments with  $\tau_R = \infty_{\text{rec}}$  are *typed* as censored for return credit under the default rule-set (no return observed under the contract), and must be treated as such in any continuity accounting.

## Abstract

Generative Collapse Dynamics (GCD) is a contract-first measurement discipline governed by a single admissibility axiom: *collapse is generative; only what returns through collapse is real*. In UMCP terms, “real” is not metaphysical; it is an audit rule: a claim receives epistemic credit only if (i) it is evaluated on a declared bounded trace  $\Psi(t) \in [0, 1]^n$  produced by a declared adapter  $N_K$  from observables  $x(t)$ , and (ii) it exhibits finite re-entry time  $\tau_R(t)$  under the same frozen evaluation rules across the collapse/return boundary. We formalize the Tier-1 kernel ledger  $\{\omega, F, S, C, \kappa, \text{IC}\}$  and the typed return output  $\tau_R$ , state the Tier-1 identities that hold for every admissible trace, and introduce two diagnostics that recur across domains without altering Tier-1 meanings: the heterogeneity gap  $\Delta := F - \text{IC}$  and coherence efficiency  $\rho := \text{IC}/F$  (defined when  $F > 0$ ). We present the frozen parameter set  $(\varepsilon, p, \alpha, \text{tol}_{\text{seam}})$  and the seam-derived structural constants  $(c^*, \omega_{\text{trap}})$  that close the seam consistently across 16 domain closures. The conjunctive four-gate regime classification (Stable/Watch/Collapse with Critical overlay) and the seam budget identity  $\Delta\kappa = R \cdot \tau_R - (D_\omega + D_C)$  are stated and connected to the kernel. Finally, we record the scale-ladder summary (406 objects across 11 rungs) as a cross-scale comparability stress test under a single kernel, and summarize the 28 structural identities and 46 lemmas verified to machine precision.

## Scope and non-negotiables

1. **Tier-1 reservation.** The symbols  $\{\omega, F, S, C, \tau_R, \kappa, \text{IC}\}$  denote kernel outputs on  $\Psi_\varepsilon(t)$  under the frozen contract. They are not redefined by overlays, narrative interpretation, or domain-specific seman-

tics.

2. **Structural-change rule.** Any change to the embedding  $N_K$ ; the return neighborhood generator  $D_\theta$ ; the metric  $\|\cdot\|$ ; tolerance  $\eta$ ; weights  $\mathbf{w}$ ; clipping level  $\varepsilon$ ; or any closure entering seam accounting constitutes a structural change. Continuity claims across structural change require an explicit seam specification and (if adopted as “the same object”) a weld.

## 0.1 Run identity and freeze statement

A UMCP evaluation is identified by a run key that binds *what was measured to how it was evaluated*. We denote this key by RunID and treat it as immutable for the purpose of comparison:

$$\text{RunID} \equiv (\Pi, N_K, \varepsilon, \mathbf{w}, \|\cdot\|, \eta, D_\theta, \mathcal{C}), \quad (1)$$

where  $\Pi$  is the acquisition/pipeline declaration,  $N_K$  the adapter,  $\varepsilon$  the log-safety clip,  $\mathbf{w}$  the frozen weights,  $(\|\cdot\|, \eta)$  the return metric and tolerance,  $D_\theta$  the return neighborhood generator, and  $\mathcal{C}$  the active frozen contract/closure registry. Two results are audit-comparable only if they share the same RunID (or else the comparison is explicitly staged across a seam). Any modification to any component of RunID constitutes a structural change and must be declared before inference.

**Axiom** (Axiom-0 — The Return Axiom). *Collapse is generative; only what returns through collapse is real.*

Every definition, identity, regime label, and seam decision in this paper is derived from this single admissibility constraint. “Real” is operational: a claim receives epistemic credit if and only if it exhibits finite re-entry ( $\tau_R \neq \infty_{\text{rec}}$ ) under the frozen contract. No additional axiom is imported; classical results emerge as degenerate limits when degrees of freedom are removed from the kernel (§10).

# 1 Measured object: observables, adapter, bounded trace

## 1.1 Observables and time base

Let  $x(t)$  denote the observed data stream at discrete index  $t \in \{0, 1, 2, \dots\}$ . A UMCP measurement is the pair  $(x(t), \Pi)$ , where  $\Pi$  is the declared acquisition and preprocessing pipeline. Audit-equivalence requires that the following be declared:

1. **Units and frame.** Units, reference frame, and any coordinate conventions for each observable channel.
2. **Time base.** Sampling cadence, windowing/segmentation rule, and any resampling or interpolation.
3. **Timestamps.** Absolute timestamps (preferred) or an explicit statement that  $t$  is an abstract index.
4. **Pipeline provenance.** Instrument + preprocessing steps (software versions, calibration steps, filters, transforms).

If any component is unspecified, the induced trace  $\Psi(t)$  is not reproducible and the run is not audit-equivalent.

## 1.2 Adapter (embedding)

An adapter (embedding)  $N_K$  maps observables into a bounded trace:

$$N_K : x(t) \mapsto \Psi(t) \in [0, 1]^n, \quad (2)$$

where  $n$  is fixed by the adapter schema and does not vary within a run. The adapter is part of the measured object: changing  $N_K$  changes what it means to measure  $x(t)$ .

**Minimum adapter declaration.** The adapter specification must include, at minimum:

1. **Channel map.** The ordered channel list  $\Psi(t) = (c_1(t), \dots, c_n(t))$  and the meaning of each channel (including any physical units prior to normalization).
2. **Normalization.** The explicit map from raw values to  $[0, 1]$  for each channel, including constants, priors, baselines, or calibration tables, and their provenance.
3. **Bounds policy.** The out-of-range (OOR) rule: clip, saturate, rescale, drop, or flag; and how OOR affects subsequent computation and receipts.
4. **Missingness policy.** The rule for missing/undefined values: impute, mask, skip, or fail-fast; and the audit log field that records the event.

Absent an explicit declaration, the default assumption is *no assumption*: the run is non-auditable.

**Weights and evaluation interface.** While  $\mathbf{w}$  is a kernel input, the choice of weights is part of the evaluation interface and must be frozen together with the adapter. The statement “ $\Psi(t) \in [0, 1]^n$ ” is incomplete without both the channel map and  $\mathbf{w}$ , because the kernel is a weighted functional on  $\Psi_\varepsilon(t)$ .

**Log-safety and boundary honesty.** Fix  $\varepsilon \in (0, \frac{1}{2})$  by the frozen contract and define componentwise  $\varepsilon$ -clipping:

$$\Psi_\varepsilon(t) := \text{clip}(\Psi(t), \varepsilon, 1 - \varepsilon). \tag{3}$$

All log-domain quantities—specifically  $\kappa(t) = \sum_i w_i \ln(c_{i,\varepsilon}(t))$  and  $\text{IC}(t) = \exp(\kappa(t))$ —are computed on  $\Psi_\varepsilon(t)$ . Boundary operations are not silent conveniences: every OOR correction and every  $\varepsilon$ -induced substitution must be recorded in the run log. Silent boundary substitution is nonconformant under the contract because it changes IC and  $\kappa$  while denying provenance.

### 1.3 Run identity and freeze statement

The run key RunID is defined in §0.1 and is treated as immutable for audit-comparison. All results in this paper are interpreted relative to that frozen RunID; comparisons across changes to any component of RunID are descriptive only unless explicitly staged across a declared seam and weld.

## 2 Kernel (Tier-1) and notation

This section fixes notation and states the Tier-1 kernel objects computed from the  $\varepsilon$ -clipped bounded trace  $\Psi_\varepsilon(t) \in [\varepsilon, 1 - \varepsilon]^n$ . The kernel is evaluated pointwise in  $t$ ; all quantities below are deterministic functions of  $(\Psi_\varepsilon(t), \mathbf{w})$  under the active frozen contract.

### 2.1 Weights

Let  $\mathbf{w} = (w_1, \dots, w_n)$  satisfy

$$w_i \geq 0 \quad (i = 1, \dots, n), \quad \sum_{i=1}^n w_i = 1. \tag{4}$$

Weights are part of the evaluation interface and are frozen within a run (see RunID). They are not fitted post hoc.

## 2.2 Trace components and clipping

Write the clipped trace in components

$$\Psi_\varepsilon(t) = (c_{1,\varepsilon}(t), \dots, c_{n,\varepsilon}(t)), \quad c_{i,\varepsilon}(t) \in [\varepsilon, 1 - \varepsilon]. \quad (5)$$

When unambiguous, we suppress the  $\varepsilon$  subscript and write  $c_i(t)$  for  $c_{i,\varepsilon}(t)$ ; log-domain expressions always use the clipped values.

## 2.3 Fidelity and drift

Define the weighted fidelity and drift

$$F(t) := \sum_{i=1}^n w_i c_{i,\varepsilon}(t), \quad \omega(t) := 1 - F(t). \quad (6)$$

**Identity 1** (Complementarity). *For all  $t$ ,*

$$F(t) + \omega(t) = 1. \quad (7)$$

*Proof.* Immediate from the definition  $\omega(t) := 1 - F(t)$ .  $\square$

## 2.4 Log-integrity and integrity composite

Define the log-integrity and integrity composite

$$\kappa(t) := \sum_{i=1}^n w_i \ln(c_{i,\varepsilon}(t)), \quad \text{IC}(t) := \exp(\kappa(t)) = \prod_{i=1}^n c_{i,\varepsilon}(t)^{w_i}. \quad (8)$$

The  $\varepsilon$ -clipping guarantees  $c_{i,\varepsilon}(t) \in (0, 1)$ , hence  $\kappa(t) \in (-\infty, 0)$  and  $\text{IC}(t) \in (0, 1)$ .

**Identity 2** (Exponentiation). *For all  $t$ ,*

$$\text{IC}(t) = \exp(\kappa(t)) \quad \text{and} \quad \kappa(t) = \ln(\text{IC}(t)). \quad (9)$$

*Proof.* By definition,  $\text{IC}(t) = \exp(\kappa(t))$ . Since  $\text{IC}(t) > 0$ , applying  $\ln(\cdot)$  yields  $\ln(\text{IC}(t)) = \kappa(t)$ .  $\square$

## 2.5 Entropy functional

Define the Bernoulli-field entropy functional

$$S(t) := - \sum_{i=1}^n w_i [c_{i,\varepsilon}(t) \ln c_{i,\varepsilon}(t) + (1 - c_{i,\varepsilon}(t)) \ln(1 - c_{i,\varepsilon}(t))]. \quad (10)$$

*Remark 1* (Role of  $S$ ).  $S(t)$  is a dispersion functional on bounded traces. Any thermodynamic interpretation is an overlay and must be declared as Tier-2;  $S$  itself is a Tier-1 kernel output.

## 2.6 Curvature proxy

Let  $\text{std}(\mathbf{c}(t))$  denote the population standard deviation of the multiset  $\{c_{1,\varepsilon}(t), \dots, c_{n,\varepsilon}(t)\}$ . The canon default proxy is

$$C(t) := \frac{\text{std}(\mathbf{c}(t))}{0.5}. \quad (11)$$

Because  $c_{i,\varepsilon}(t) \in [\varepsilon, 1 - \varepsilon] \subset [0, 1]$ , we have  $\text{std}(\mathbf{c}(t)) \leq 0.5$  and therefore  $C(t) \in [0, 1]$  under the trace bounds.

## 2.7 Integrity bound

**Identity 3** (Integrity bound). For  $c_{i,\varepsilon}(t) \in (0, 1)$  and weights  $\mathbf{w}$  with  $w_i \geq 0$  and  $\sum_i w_i = 1$ ,

$$\text{IC}(t) \leq F(t). \quad (12)$$

*Proof.* Let  $a_i = c_{i,\varepsilon}(t) \in (0, 1)$ . By Jensen’s inequality applied to the concave function  $\ln(\cdot)$ ,  $\sum_i w_i \ln a_i \leq \ln(\sum_i w_i a_i)$  for nonnegative weights summing to one. Exponentiating both sides yields  $\prod_i a_i^{w_i} \leq \sum_i w_i a_i$ . Substituting the definitions  $\text{IC}(t) = \prod_i a_i^{w_i}$  and  $F(t) = \sum_i w_i a_i$  gives  $\text{IC}(t) \leq F(t)$ . Equality holds if and only if all  $a_i$  are equal (homogeneous trace).  $\square$

*Remark 2* (Solvability interpretation). For  $n = 2$  equal-weight channels, the system  $c_1 + c_2 = 2F$ ,  $c_1 c_2 = \text{IC}^2$  has real solutions  $c_{1,2} = F \pm \sqrt{F^2 - \text{IC}^2}$  if and only if  $\text{IC} \leq F$ . The integrity bound is therefore the *solvability condition* for recovering individual channel values from  $(F, \text{IC})$ .

*Remark 3* (Relationship to classical inequalities). The proof above uses Jensen’s inequality on concave  $\ln(\cdot)$ , which has been known since Cauchy (1821). The mathematical content of  $\text{IC} \leq F$ —that the weighted geometric mean does not exceed the weighted arithmetic mean—is classical. We make this explicit because the contribution of the integrity bound is not the inequality itself but the following four consequences of the architecture:

1. **Axiom-produced domain.** The bounded trace  $\Psi_\varepsilon(t) \in [\varepsilon, 1 - \varepsilon]^n$  and normalized weights  $\sum w_i = 1$  are not modelling choices—they are consequences of the admissibility axiom (Axiom 0.1) and the frozen contract. The axiom produces the conditions under which Jensen applies; Jensen does not produce the axiom.
2. **Solvability interpretation.** The reinterpretation of  $\text{IC} \leq F$  as a *solvability condition*—the requirement for individual channel recovery from aggregate invariants—is not part of the classical inequality and gives the bound operational content beyond “the geometric mean is small.”
3. **Composition laws.** The structural consequences  $\text{IC}_{12} = \sqrt{\text{IC}_1 \cdot \text{IC}_2}$  (geometric composition) and  $F_{12} = (F_1 + F_2)/2$  (arithmetic composition) are properties of the architecture, not of Jensen’s inequality. The heterogeneity gap’s variance decomposition  $\Delta \approx \text{Var}(\mathbf{c})/(2\bar{c})$  (Proposition 2) connects the bound to Fisher information geometry.
4. **Cross-domain universality.** That  $\text{IC} \leq F$  holds with zero violations across 16 domains, 5,413 tests, and 406 scale-ladder objects is a validation of the architecture—the inequality is guaranteed by definition, but its *relevance* to particle physics, finance, evolution, and cosmology simultaneously is not. The classical inequality is the degenerate limit (§10); the contribution is the architecture that makes it universally diagnostic.

## 3 Frozen parameters and seam-derived constants

The kernel requires a small set of constants that are *frozen*—consistent across the seam, the same rules on both sides of every collapse-return boundary. These constants are not prescribed by convention; they are the unique values where seams close consistently across all domain closures.

### 3.1 Frozen contract parameters

Parameter	Value	Symbol	Role
Guard band	$10^{-8}$	$\varepsilon$	Pole guard at $\omega = 1$ ; log-stability
Drift cost exponent	3	$p$	Unique integer: $\omega_{\text{trap}}$ is Cardano root of $x^3 + x - 1 = 0$
Curvature coefficient	1.0	$\alpha$	Unit coupling: $D_C = \alpha \cdot C$
Seam tolerance	0.005	$\text{tol}_{\text{seam}}$	Width where $\text{IC} \leq F$ holds at 100% across all 16 domains
Normalization domain	$[0, 1]$	$[a, b]$	Bounded trace domain

### 3.2 Seam-derived structural constants

Two constants emerge from the frozen parameters by solving structural equations:

**Logistic self-dual fixed point  $c^*$ .** The channel value that maximizes  $S + \kappa$  per channel, solving  $\ln((1 - c)/c) + 1/c = 0$ :

$$c^* \approx 0.7822. \quad (13)$$

**Trapping threshold  $\omega_{\text{trap}}$ .** The drift value where the drift cost function  $\Gamma(\omega) = \omega^p/(1 - \omega + \varepsilon)$  equals  $\alpha$ , giving the boundary between returnable and trapped drift. With  $p = 3$ ,  $\omega_{\text{trap}}$  is the unique real root of the depressed cubic  $x^3 + x - 1 = 0$  (Cardano root):

$$\omega_{\text{trap}} \approx 0.6823, \quad c_{\text{trap}} = 1 - \omega_{\text{trap}} \approx 0.3177. \quad (14)$$

*Remark 4 (Why  $p = 3$ ).*  $p = 3$  is the unique integer exponent for which the trapping threshold satisfies a solvable depressed cubic with a single real root. No other integer value of  $p$  yields a Cardano-solvable algebraic form for  $\omega_{\text{trap}}$ . This is discovered by the seam, not chosen by convention.

## 4 Regime classification

The kernel maps continuous invariants to discrete regime labels via a conjunctive four-gate criterion. These gates are frozen per run and sourced from the active contract.

### 4.1 Four-gate criterion

$$\begin{aligned} \text{STABLE} : & \quad \omega < 0.038 \wedge F > 0.90 \wedge S < 0.15 \wedge C < 0.14, \\ \text{WATCH} : & \quad 0.038 \leq \omega < 0.30 \text{ (or Stable gates not all satisfied)}, \\ \text{COLLAPSE} : & \quad \omega \geq 0.30. \end{aligned} \quad (15)$$

Stability is *conjunctive*: all four gates must be satisfied simultaneously. Relaxing any single gate (e.g.,  $S$  exceeds 0.15 while  $\omega$  remains small) moves the system to Watch.

### 4.2 Critical severity overlay

A Critical overlay accompanies any regime when integrity is dangerously low:

$$\text{CRITICAL} : \quad \text{IC} < 0.30. \quad (16)$$

Critical is not a regime; it is a severity flag. A system can be Watch + Critical (moderate drift, but multiplicative coherence has collapsed) or even Stable + Critical (all drift/entropy/curvature gates pass, but a single near- $\varepsilon$  channel suppresses IC).

*Remark 5 (Geometric slaughter).* A single channel  $c_k \rightarrow \varepsilon$  drives  $IC \rightarrow 0$  regardless of the remaining channels, because the weighted geometric mean is multiplicatively fragile:  $IC = \prod_i c_i^{w_i}$ , so one near-zero factor annihilates the product. Meanwhile  $F$  (the weighted arithmetic mean) can remain healthy. This decoupling of  $F$  from  $IC$  is the mechanism behind Stable + Critical configurations and explains why the heterogeneity gap  $\Delta = F - IC$  is the primary diagnostic of structural vulnerability.

### 4.3 Fisher space partition

Under the four-gate criterion with the frozen thresholds, the regime fractions of the Bernoulli manifold are:

Regime	Condition	Fisher space fraction
Stable	All four gates satisfied	12.5%
Watch	$0.038 \leq \omega < 0.30$ (or gates broken)	24.4%
Collapse	$\omega \geq 0.30$	63.1%

Stability is rare—87.5% of the manifold lies outside it. This is structural, not accidental: the conjunctive gate makes stability demanding, which is precisely what gives return from collapse its meaning.

## 5 Seam budget and cost closures

The seam budget is the accounting identity that determines whether a return from collapse receives epistemic credit. It is computed from the frozen contract and the Tier-1 kernel outputs.

### 5.1 Drift cost closure

The drift cost function  $\Gamma$  maps drift  $\omega$  to a debit in the seam ledger:

$$\Gamma(\omega) := \frac{\omega^p}{1 - \omega + \varepsilon}, \quad (17)$$

where  $p = 3$  and  $\varepsilon = 10^{-8}$  are frozen. The drift debit is  $D_\omega = \Gamma(\omega)$ .

### 5.2 Curvature cost closure

The curvature cost couples channel dispersion into the budget:

$$D_C := \alpha \cdot C, \quad (18)$$

where  $\alpha = 1.0$  is frozen and  $C$  is the curvature proxy from the kernel.

### 5.3 Budget identity

The seam budget balances return credit against drift and curvature costs:

$$\Delta \kappa_{\text{budget}} = R \cdot \tau_R - (D_\omega + D_C), \quad (19)$$

where  $R$  is the return credit rate and  $\tau_R$  is the re-entry delay. If  $\tau_R = \infty_{\text{rec}}$ , the credit term is zero: no return, no credit.

## 5.4 Seam residual and weld decision

The seam residual compares the budget against the actual ledger change in log-integrity:

$$s := \Delta \kappa_{\text{budget}} - \Delta \kappa_{\text{ledger}}, \quad (20)$$

where  $\Delta \kappa_{\text{ledger}} = \kappa(t_1) - \kappa(t_0)$ . A weld PASS requires:

1. **Finite return.**  $\tau_R \neq \infty_{\text{rec}}$  (typed return observed).
2. **Residual closure.**  $|s| \leq \text{tol}_{\text{seam}} = 0.005$ .
3. **Identity check.** The ratio  $i_r = \text{IC}(t_1)/\text{IC}(t_0)$  satisfies  $i_r \approx \exp(\Delta \kappa_{\text{ledger}})$ .

If  $\tau_R = \infty_{\text{rec}}$ , return credit is censored by typing (default rule), and no weld may claim continuity credit.

## 6 Return typing and censoring

Return is a contract object. The kernel does not infer return; it receives a return neighborhood generator and a metric as frozen evaluation settings. This section defines the admissible candidate set, the re-entry delay  $\tau_R$ , and the typed censoring rule used when no return is observed under the active contract.

### 6.1 Return neighborhood and admissible candidates

Fix a closure-defined return domain (candidate set)  $D_\theta(t)$  at time  $t$ , parameterized by  $\theta$ , together with a frozen metric  $\|\cdot\|$  on  $\Psi$ -space and a frozen tolerance  $\eta > 0$ . The admissible candidate set is

$$U_\theta(t) := \{ u \in D_\theta(t) : \|\Psi(t) - \Psi(u)\| \leq \eta \}. \quad (21)$$

Here  $u$  ranges over the evaluation indices available to the run (e.g.,  $u \in \{0, 1, \dots, t\}$  for a retrospective search, or a specified horizon window if the contract restricts search). The horizon rule is part of the closure specification: changing it changes  $D_\theta$  and therefore changes the measured object.

*Remark 6* (Contract dependence). The triple  $(D_\theta, \|\cdot\|, \eta)$  is part of RunID. Any modification to the return domain generator, metric, tolerance, or horizon constitutes a structural change and must be declared before inference.

### 6.2 Re-entry delay

Define the (typed) re-entry delay:

$$\tau_R(t) := \begin{cases} \min\{t - u : u \in U_\theta(t)\}, & \text{if } U_\theta(t) \neq \emptyset, \\ \infty_{\text{rec}}, & \text{if } U_\theta(t) = \emptyset. \end{cases} \quad (22)$$

Thus  $\tau_R(t)$  is either a nonnegative integer (finite return observed under the contract) or the typed sentinel  $\infty_{\text{rec}}$  (no return observed under the contract).

**Note.** Unless the active contract specifies return evaluation on  $\Psi_\varepsilon$ , the return metric is evaluated on  $\Psi$  (unclipped) while log-domain quantities are evaluated on  $\Psi_\varepsilon$ .

*Remark 7* (Typed censoring and return credit).  $\infty_{\text{rec}}$  is a typed censoring value, not a large number. It asserts: *under the active contract and within the active horizon, no admissible candidate  $u$  was found.* Unless a different censoring rule is explicitly declared and frozen, segments with  $\tau_R = \infty_{\text{rec}}$  receive zero return credit in any seam or continuity accounting, and must be treated as censored rather than as finite-but-large delays.

## 7 Diagnostics from the invariant ledger (Tier-2; non-gating)

This section defines Tier-2 diagnostics constructed from the Tier-1 ledger. They are permitted as *descriptive* quantities (ranking, visualization, clustering, annotation), but they are *not* regime or weld gates unless explicitly promoted through a declared seam and a new frozen run. In particular, no Tier-2 score may override (i) the active regime gates or (ii) typed return censoring.

### 7.1 Heterogeneity gap

**Definition 1** (Heterogeneity gap). For each  $t$ , define

$$\Delta(t) := F(t) - \text{IC}(t). \quad (23)$$

**Proposition 1** (Nonnegativity). For all  $t$ ,  $\Delta(t) \geq 0$ .

*Proof.* Immediate from the Tier-1 integrity bound  $\text{IC}(t) \leq F(t)$  (Identity 3).  $\square$

**Operational meaning.**  $F(t)$  is an arithmetic aggregate of channel confidence, while  $\text{IC}(t)$  is multiplicative and therefore bottlenecked by weak channels. Large  $\Delta(t)$  indicates imbalance: the mean can remain moderate while multiplicative coherence is driven toward the guard band. Equivalently,  $\Delta(t)$  is the gap between “average adequacy” and “weakest-link survivability” under the frozen weights.

**Proposition 2** (Variance decomposition). For equal weights  $w_i = 1/n$  and small heterogeneity, the heterogeneity gap admits the leading-order expansion

$$\Delta(t) \approx \frac{\text{Var}(\mathbf{c}(t))}{2\bar{c}(t)}, \quad (24)$$

where  $\bar{c}(t) = F(t)$  is the channel mean and  $\text{Var}(\mathbf{c}(t))$  is the channel variance. Thus  $\Delta$  is proportional to the Fisher information contribution from channel heterogeneity.

**Diagnostic-only class labels (conventional).** The following thresholds are optional labels for reporting and visualization only; they are not gates.

Condition on $\Delta(t)$	Diagnostic label
$\Delta(t) < 10^{-6}$	homogeneous
$10^{-6} \leq \Delta(t) < 10^{-2}$	coherent
$10^{-2} \leq \Delta(t) < 5 \times 10^{-2}$	heterogeneous
$\Delta(t) \geq 5 \times 10^{-2}$	fragmented

### 7.2 Coherence efficiency

When  $F(t) > 0$ , define the coherence-efficiency ratio

$$\rho(t) := \frac{\text{IC}(t)}{F(t)} \in (0, 1). \quad (25)$$

**Operational meaning.**  $\rho(t)$  measures how much of the arithmetic mean survives multiplicatively under the frozen weights. At fixed  $F(t)$ , smaller  $\rho(t)$  indicates stronger bottlenecking by one (or several) weak channels. When reporting  $\rho$ , always report alongside  $\Delta$  (or  $\text{IC}$ ) so that “efficiency loss” is not confused with low fidelity.

### 7.3 Regimes from the four-gate criterion

The v2.1.5 release line uses the conjunctive four-gate regime classification defined in §4. The active gates are:

$$\text{Stable: } \omega < 0.038 \wedge F > 0.90 \wedge S < 0.15 \wedge C < 0.14; \quad \text{Watch: } 0.038 \leq \omega < 0.30; \quad \text{Collapse: } \omega \geq 0.30. \quad (26)$$

These gates are frozen per run and sourced from the active contract. Any modification to the gate thresholds constitutes a structural change requiring a seam declaration.

## 8 Cross-domain patterns (16 domains; expressed in kernel terms)

The kernel exists to make comparison possible without importing domain-native units into the invariant ledger. The v2.1.5 release line validates 16 domain closures, each supplying a declared adapter  $N_K$  that produces a bounded trace  $\Psi(t) \in [0, 1]^n$ ; after that point, the quantities  $(F, \omega, \kappa, \text{IC}, S, C, \tau_R)$  are the *same mathematical objects* in every domain. The 16 domains are: GCD, RCFT, kinematics, security, Weyl cosmology, astronomy, nuclear physics, quantum mechanics, finance, atomic physics (118 elements), materials science, Standard Model (31 particles, 10 proven theorems), everyday physics, evolution, dynamic semiotics, and continuity theory. Accordingly, the statements in this section are expressed in kernel terms: they refer to the invariant ledger and its Tier-2 diagnostics  $(\Delta, \rho)$ , not to the native physical units of the source measurements.

### 8.1 Particle and atomic physics (reported)

#### Reported patterns.

- **Identity conformance at scale.** A set of 31 Standard Model particles and 118 chemical elements is reported to satisfy the Tier-1 identities (Complementarity, Exponentiation, Integrity bound) with zero violations under the active adapters.
- **Generation monotonicity (integrity ordering).** “Generation monotonicity” is reported as a strict ordering in integrity-derived quantities across fermion generations, expressed as a decrease in IC (and typically  $\rho$ ) from first- to third-generation particles.
- **Integrity cliff at the quark→hadron boundary.** A sharp discontinuity is reported at the quark→hadron transition, expressed as an abrupt drop in IC and/or  $\rho$ , often accompanied by a spike in  $\Delta$  (weak-channel bottlenecking).

**Kernel meaning.** These are statements about  $(F, \text{IC}, \Delta, \rho, \omega)$  computed from domain traces. They do not require (and do not imply) any claim about domain-native units beyond what is already encoded in the adapter.

### 8.2 Quantum mechanics and materials (reported)

#### Reported patterns.

- **Complementarity cliff (channel co-failure).** A “complementarity cliff” is reported when wave- and particle-proxy channels simultaneously approach the guard band. The kernel signature is: (i)  $F$  remains nonzero (the mean does not vanish), (ii) IC collapses toward the guard band (multiplicative failure), hence  $\Delta$  spikes, and (iii) return may fail in the affected window, yielding typed censoring  $\tau_R = \infty_{\text{rec}}$ .
- **Phase-transition signatures in materials.** Phase transitions are reported as joint excursions in dispersion diagnostics: spikes in  $\Delta$  (imbalance) and  $C$  (dispersion), with corresponding drops in coherence efficiency  $\rho$ .

**Kernel meaning.** The “cliff” language denotes a regime where a small subset of channels becomes bottlenecking. In kernel terms, the arithmetic aggregate can remain stable while the multiplicative aggregate collapses.

### 8.3 Nuclear, RCFT, cosmology (reported)

**Reported patterns.**

- **Critical points as dispersion maxima (nuclear / RCFT).** In nuclear closures and RCFT closures, “critical points” are reported as local maxima in  $\Delta$  (imbalance) and/or  $C$  (dispersion), often coincident with increased entropy  $S$  and elevated drift  $\omega$ .
- **Cosmological heterogeneity (near-zero channel suppression).** Cosmological closures are reported to exhibit persistent heterogeneity: one or more near-zero channels suppress IC toward the guard band while leaving  $F$  moderate, producing large  $\Delta$  and small  $\rho$ .

**Kernel meaning.** These reports describe the same mechanism: multiplicative suppression by weak channels in an otherwise moderate mean. In cosmology, the claim is not “the universe is low-fidelity,” but “coherence is bottlenecked by one or more channels.”

### 8.4 Finance, evolution, everyday systems (reported)

**Reported patterns.**

- **Finance: stress as weak-channel dominance.** Systemic stress is reported as spikes in  $\Delta$  (weak-channel suppression of IC) and drops in  $\rho$ , with  $\omega$  typically increasing as coherence degrades.
- **Evolution: anti-proof configurations and human heterogeneity.** Cancer is reported as an “anti-proof” configuration: local channel fidelity can remain high while organism-level IC collapses (large  $\Delta$ , low  $\rho$ ). Humans are reported with  $\Delta \approx 0.34$  under the relevant adapter.
- **Everyday physics: identity preservation under macroscopic adapters.** Macroscopic systems (thermodynamics, optics, electromagnetism) are reported to satisfy the Tier-1 identities under their declared adapters, supporting single-kernel comparability across scale.

**Kernel meaning.** In each case, the diagnostic content is the same:  $\Delta$  isolates imbalance,  $\rho$  quantifies multiplicative survival relative to the mean, and typed return censoring ( $\tau_R = \infty_{\text{rec}}$ ) forbids continuity credit where return is not observed.

## 9 Scale-ladder summary (406 objects; 11 rungs; reported)

The scale-ladder run applies a single kernel to a heterogeneous catalog: 406 objects spanning 11 scale rungs (from Planck length through the cosmic horizon). The run is reported to use a frozen  $\varepsilon$  and equal weights, so that cross-object comparison is driven by the trace values rather than by tuned weighting. The items below are reported outcomes expressed strictly in kernel terms.

1. **Exact complementarity.** For every object evaluated, the Tier-1 identity holds:

$$F + \omega = 1. \tag{27}$$

This is not a trend; it is an identity enforced by the kernel definition.

2. **Fidelity is not monotone in scale.** Mean fidelity does not increase or decrease monotonically with physical size. The highest reported mean  $F$  occurs at the galactic rung (characteristic scale  $\sim 10^{21}$  m), while Planck-scale objects are reported to have near-zero  $F$  under the scale-ladder adapter.

3. **A single incoherence mechanism (weak-channel suppression).** Across rungs,  $\Delta$  is reported to diagnose incoherence by the same mechanism: one (or more) near-zero channels suppress IC toward the guard band even when the arithmetic mean remains moderate. In kernel terms, incoherence appears as large  $\Delta = F - \text{IC}$  and small  $\rho = \text{IC}/F$ .
4. **Geological objects maximize low-integrity incidence.** The geological rung is reported to contain the largest fraction of objects with extremely low integrity, quantified as  $\text{IC} < 0.01$ . This is a statement about multiplicative collapse in the trace (and therefore large  $\Delta$ ), not about low mean fidelity.
5. **Watch regime is rare under the four-gate criterion.** Under the conjunctive four-gate regime classification (§4), only 10/406 objects fall in WATCH. This highlights that intermediate drift with all Stable gates otherwise satisfied is uncommon; objects cluster toward Stable or Collapse under the four-gate criterion.
6. **Coherence efficiency separates nuclear from cosmological.** Coherence efficiency  $\rho = \text{IC}/F$  is reported to be maximized in nuclear objects (approximately 72.6%) and minimized in cosmological objects (approximately 17.6%). Interpreted purely in kernel terms: for comparable mean fidelity, the nuclear rung preserves multiplicative coherence, while the cosmological rung loses it to weak-channel suppression.
7. **Specialization as imbalance (biological rungs).** Biological specialization is reported to be quantified by the ratio  $\Delta/F$ . Highly specialized cells are reported to exhibit large  $\Delta/F$ , with near-zero IC despite moderate  $F$ , while yeast is reported to remain comparatively balanced (small  $\Delta$ ), hence higher  $\rho$ .

**Reporting rule.** These scale-ladder claims are recorded as reported summaries under the v2.1.5 release line. Any attempt to promote a scale-ladder pattern into a gate, a closure, or a continuity claim must bind the claim to a frozen RunID and (when structural changes are involved) to an explicit seam and weld decision.

## 10 Degenerate limits and the arrow of derivation

Classical results emerge as degenerate limits when degrees of freedom are removed from the GCD kernel. The arrow of derivation runs *from* Axiom-0 *to* the classical result, never the reverse.

GCD structure	Degenerate operation	Classical limit
$\text{IC} \leq F$	Strip channel semantics, weights, $\varepsilon$	Weighted AM–GM inequality
$S$ (Bernoulli field)	Remove collapse field ( $c_i \in \{0, 1\}$ )	Shannon entropy
$F + \omega = 1$	Remove cost function and trace	Unitarity
$\text{IC} = \exp(\kappa)$	Remove channel structure	Exponential map

The integrity bound  $\text{IC} \leq F$  is not “the AM–GM inequality applied to GCD”; it is the solvability condition for recovering individual channels from aggregate invariants, independently derived from Axiom-0 on the bounded trace. The classical inequality is what remains when the channel semantics are discarded.

## 11 Structural identities and lemma inventory

The Tier-1 kernel and its Tier-0 protocol machinery yield 28 structural identities and 46 lemmas, verified computationally to machine precision ( $< 10^{-16}$  residual) across all domain closures. The full derivation chain is available via three diagnostic scripts in the repository; the key results are summarized below.

### 11.1 Three core identities (Tier-1; always true by construction)

1. **Duality identity:**  $F + \omega = 1$ . The complementary partition of the bounded trace.
2. **Integrity bound:**  $IC \leq F$ . Solvability condition; equality iff homogeneous.
3. **Log-integrity relation:**  $IC = \exp(\kappa)$ . Link between multiplicative and additive coherence.

### 11.2 Key derived results (from the 28 identity set)

- **Flat manifold.** The Fisher metric on the Bernoulli manifold is  $g_F(\theta) = 1$ —all structure comes from the embedding, not intrinsic curvature.
- **One function.**  $S$  and  $\kappa$  are projections of one function:  $f(\theta) = 2 \cos^2 \theta \cdot \ln(\tan \theta)$ , verified to residual  $< 10^{-16}$ .
- $p = 3$  **uniqueness.**  $\omega_{\text{trap}}$  is the Cardano root of  $x^3 + x - 1 = 0$ ; no other integer  $p$  yields a solvable algebraic form.
- **Solvability.** For  $n = 2$  channels,  $c_{1,2} = F \pm \sqrt{F^2 - IC^2}$  requires  $IC \leq F$  for real solutions.
- **Low-rank closures.** 5 closure diagnostics span only 4 effective dimensions (PCA)—the closure algebra is low-rank.
- **Composition laws.**  $IC$  composes geometrically ( $IC_{12} = \sqrt{IC_1 \cdot IC_2}$ );  $F$  composes arithmetically ( $F_{12} = (F_1 + F_2)/2$ ).
- **Regime partition.** Under the four-gate criterion: Collapse 63.1%, Watch 24.4%, Stable 12.5% of Fisher space.

### 11.3 Lemma coverage

The 46 lemmas (Lemmas 1–42 plus extensions) cover: range bounds (L1), monotonicity (L2, L12), sensitivity control (L3, L7), well-posedness of  $\tau_R$  (L8), curvature bounds (L10), entropy envelopes (L5, L15, L16), Lipschitz continuity (L23), seam composition (L20), return probability (L29), and empirical discoveries including return-collapse duality (L35), super-exponential convergence (L39), and the entropy-integrity anti-correlation at  $c^*$  (L41). All lemmas are verified computationally in the test suite (5,413 tests across 16 domain closures).

## 12 Falsifiable predictions

A framework that only describes and organizes does not yet constitute a theory. The following predictions are derived from the kernel architecture and have not yet been independently tested. Each is stated in falsifiable form: a specific empirical outcome that, if contradicted, would require revision of the framework or its domain closures.

### 12.1 Confinement-scale inversion universality

The Standard Model closure documents an “IC cliff”—a 98% drop in  $IC$  at the quark→hadron boundary—followed by  $IC$  recovery at the atomic scale when new measurable channels (electronic, bulk) are added to the adapter. GCD predicts this cliff-and-recovery pattern is *universal for confinement boundaries*:

1. At any boundary where constituent degrees of freedom become confined (inaccessible to the adapter),  $IC$  will collapse while  $F$  may remain moderate, producing a spike in  $\Delta$ .
2. When the composite system acquires new measurable channels not present in the constituent adapter,  $IC$  will recover.

3. The recovery magnitude scales as  $(n_{\text{new}}/n_{\text{total}})^{1/n_{\text{total}}}$  to leading order in the geometric mean.

**Test.** Build adapters for transitions not yet in the closure set—atom→molecule, molecule→crystal, neuron→circuit—and verify whether the cliff-recovery pattern and its scaling hold. Failure to observe the cliff at a genuine confinement boundary would falsify the prediction.

## 12.2 $c^*$ clustering in persistent systems

The logistic self-dual fixed point  $c^* \approx 0.7822$  maximizes  $S + \kappa$  per channel. GCD predicts that systems which persist over evolutionary, geological, or operational timescales—i.e., systems that *return* ( $\tau_R \neq \infty_{\text{rec}}$ )—will exhibit mean channel values that cluster near  $c^*$ . Systems far from  $c^*$  should either drift toward it over time or fail to return.

**Test.** Across all 16 domain closures, compute the mean channel value  $\bar{c}$  for objects classified as Stable (all four gates satisfied) versus Collapse ( $\omega \geq 0.30$ ). The prediction is:  $\bar{c}_{\text{Stable}}$  clusters near  $c^*$  while  $\bar{c}_{\text{Collapse}}$  does not. A domain where Stable objects systematically avoid  $c^*$  would falsify the prediction.

## 12.3 Heterogeneity gap as a leading indicator

The variance decomposition  $\Delta \approx \text{Var}(\mathbf{c}) / (2\bar{c})$  (Proposition 2) implies that  $\Delta$  responds to channel divergence *before* the arithmetic mean  $F$  deteriorates—because variance increases before the mean shifts. GCD predicts that in time-series data,  $\Delta$  will spike before regime transitions:

$$t_{\Delta\text{-spike}} < t_{\text{regime change}}. \quad (28)$$

**Test.** In domains with temporal evolution (finance, materials under temperature ramp, biological development), compute  $\Delta(t)$  and verify whether  $\Delta$  peaks before the system crosses a regime gate. If  $\Delta$  consistently lags rather than leads, the prediction is falsified.

## 12.4 Trapping threshold as return boundary

The trapping threshold  $\omega_{\text{trap}} \approx 0.6823$  is seam-derived (the Cardano root of  $x^3 + x - 1 = 0$ ). GCD predicts that for any domain with empirically measured  $\tau_R$ , the distribution of return delays will show a transition from finite  $\tau_R$  to  $\tau_R = \infty_{\text{rec}}$  near  $\omega \approx \omega_{\text{trap}}$ : systems with  $\omega > \omega_{\text{trap}}$  should exhibit  $\tau_R = \infty_{\text{rec}}$  with overwhelming frequency, while systems with  $\omega < \omega_{\text{trap}}$  should predominantly return.

**Test.** Across all domain closures, bin objects by  $\omega$  and compute the fraction with  $\tau_R = \infty_{\text{rec}}$  in each bin. The prediction is a sharp transition near  $\omega_{\text{trap}}$ . A smooth, gradual relationship between  $\omega$  and return frequency, with no inflection near 0.6823, would falsify the prediction.

**Status.** None of these predictions has been independently verified as of the v2.1.5 release. They are stated here to make the framework falsifiable and to distinguish it from a purely descriptive system. Independent testing—especially by researchers outside the GCD development lineage—is the necessary next step for any claim beyond internal consistency.

# 13 Continuity across structural change

Continuity is not assumed; it is evaluated. If continuity across a structural change is claimed, a seam must be explicitly declared and a weld decision recorded. The seam budget formulas are stated in §5; this section states the structural-change rule and seam declaration requirements.

### 13.1 Structural change and seam declaration

A *seam* names the boundary between two evaluations  $\text{RunID}_{\text{pre}}$  and  $\text{RunID}_{\text{post}}$ . It is defined by (i) the comparison points  $(t_0, t_1)$  (or windows), (ii) the identity-check rule in force, and (iii) the frozen closure registry used to compute seam costs and return credit. Absent an explicit seam declaration, comparisons across structural changes are descriptive only and do not carry continuity credit.

## 14 Citation and reproducibility

**Verification.** The v2.1.5 release line contains 5,413 tests across 16 domain closures, 122 closure modules, and 46 lemmas. The three core Tier-1 identities ( $F + \omega = 1$ ,  $\text{IC} \leq F$ ,  $\text{IC} = \exp(\kappa)$ ) are verified with zero violations across the full test suite. The 28 structural identities are re-derivable by running the diagnostic scripts in the repository.

### Whitepaper citation (Zenodo; fixed artifact)

**Preferred.** C. Paulus, *Generative Collapse Dynamics (UMCP/GCD): Kernel invariants, return typing, and cross-domain diagnostics*, Zenodo, March 2026. DOI: [10.5281/zenodo.18819238](https://doi.org/10.5281/zenodo.18819238).

### Repository citation (GitHub; release line primary, commit optional)

**Primary (release line).** Generative Collapse Dynamics (UMCP/GCD). Validator line v2.1.5 (March 2026). Repository: <https://github.com/calebpruett927/GENERATIVE-COLLAPSE-DYNAMICS>. MIT License.

**Preferred (release line + immutable pointer).** Generative Collapse Dynamics (UMCP/GCD). Validator line v2.1.5 (March 2026). Repository: <https://github.com/calebpruett927/GENERATIVE-COLLAPSE-DYNAMICS>. MIT License. Commit: .

### Canon anchors (context)

PRE: [10.5281/zenodo.17756705](https://doi.org/10.5281/zenodo.17756705). POST: [10.5281/zenodo.18072852](https://doi.org/10.5281/zenodo.18072852). These anchors define the contract-first invariant skeleton and continuity-law context for the present release line.

### Endnote (admissibility stance; non-gating)

Axiom-0 is enforced as an admissibility rule: if return is not observed under the active contract, the system is typed as no-return ( $\tau_R = \infty_{\text{rec}}$ ) and receives zero return credit under the default censoring rule. No-return is recorded and censored (typed); any change that converts no-return to return must be declared as a structural seam and evaluated under a weld decision.

## Acknowledgements

This whitepaper is a release-line report of the *Generative Collapse Dynamics (UMCP/GCD)* reference implementation and its validator suite: every identity claim, invariant definition, regime label, and reported cross-domain summary in the present document is intended to be reproducible as a literal proof or computed artifact from the repository under the pinned release line/commit ([1]; see also the runnable CasePack, [2]).

Archival support and fixed-artifact citation are provided via Zenodo, with canon anchoring through the PRE/POST works that define the invariant skeleton and weld/continuity context ([3, 4]). Additional

practitioner-facing consolidation and audit discipline are carried by the broader toolkit/protocol manuscripts in the corpus ([5, 6, 7, 8]).

## Lineage and provenance

**Canon anchors (PRE → POST).** The UMCP/GCD invariant skeleton and the continuity-law stance used here are canon-anchored by *The Episteme of Return* (PRE) and *The Physics of Coherence: Recursive Collapse & Continuity Laws* (POST) ([3, 4]). These anchors fix the contract-first posture (bounded trace, frozen closures, typed censoring of no-return) and constrain continuity claims across structural change to explicit seam/weld bookkeeping.

**Protocol lineage (contract-first measurement).** The measurement discipline underlying the present release line proceeds from the contract-first method texts in the corpus: the Universal Measurement Contract Protocol, the Collapse Metric Protocol, and the canon/provenance bookkeeping note ([6, 7, 8]). Seam admissibility and continuity geometry are treated as explicit objects (not implicit assumptions) in the seam taxonomy work ([9]).

**Runnable proofs and reproducibility.** All results reported in this whitepaper are claims about what the reference implementation computes and what its validator suite proves under the declared adapter and frozen RunID: in particular, (i) Tier-1 quantities are computed as explicit functions on the bounded trace under the active contract, (ii) Tier-2 diagnostics are explicitly non-gating and cannot overwrite Tier-1 symbols, and (iii) any change converting typed no-return to return is treated as a structural seam requiring an explicit weld decision. The repository is therefore the primary proof substrate for the present document ([1]), and the UMCP CasePack provides runnable audit examples in a fixed publication form ([2]).

**Cross-domain audit lineage (context, not a substitution for the kernel).** Where the paper reports cross-domain patterns, those statements are summaries of computed outputs expressed in kernel terms, not imported domain-native theory. The broader audit series and calibration notes provide additional context for how UMCP/GCD is applied across domains while preserving contract discipline ([10, 11, 12, 13, 5]).

## References

- [1] Clement Paulus. Generative-collapse-dynamics, 2026. Reference implementation repository (GitHub). URL: <https://github.com/calebpruett927/GENERATIVE-COLLAPSE-DYNAMICS>.
- [2] Clement Paulus. Umcp casepack publication, 2026. Runnable CasePack publication anchor. doi: [10.5281/zenodo.18226878](https://doi.org/10.5281/zenodo.18226878).
- [3] Clement Paulus. The episteme of return, November 2025. PRE canon anchor for UMCP/GCD. doi: [10.5281/zenodo.17756705](https://doi.org/10.5281/zenodo.17756705).
- [4] Clement Paulus. The physics of coherence: Recursive collapse & continuity laws, December 2025. POST canon anchor; Weld-ID: W-2025-12-31-PHYS-COHERENCE. doi: [10.5281/zenodo.18072852](https://doi.org/10.5281/zenodo.18072852).
- [5] Clement Paulus. The universal regime diagnostic toolkit: A practitioner’s guide to collapse-aware signal analysis, 2025. Canonical invariant definitions ( $\omega$ ,  $F$ ,  $S$ ,  $C$ ,  $\tau_R$ , IC); three-agent epistemic field model (Present  $A_1$ , Archive  $A_2$ , Unknown  $A_3$ ); regime algorithm with empirical thresholds. doi: [10.5281/zenodo.16734906](https://doi.org/10.5281/zenodo.16734906).
- [6] Clement Paulus. Universal measurement contract protocol. Manuscript, 2025. Contract-first measurement discipline; Tier-1 invariants reserved.
- [7] Clement Paulus. Collapse metric protocol. Manuscript, 2025. Kernel metric computation protocol and audit scaffolding.
- [8] Clement Paulus. Provenance & canon note (v1.0.1). Manuscript, 2025. Canon provenance, anchoring, and continuity bookkeeping.
- [9] Clement Paulus. Geometry of admissible seams. Manuscript, 2025. Seam taxonomy and admissibility geometry for continuity claims.
- [10] Clement Paulus. Regime-aware procurement systems via rcft/ulrc: Formalizing budget integrity through collapse dynamics, 2025. 37 numbered equations; GLRC cross-domain operator; 4th regime Critical Loss (IC < 0.30); field synchrony operator  $S_{\text{sync}}$ . doi: [10.5281/zenodo.16660740](https://doi.org/10.5281/zenodo.16660740).
- [11] Clement Paulus. Empirical calibration of the universal measurement contract protocol via hawking-analog polariton fluids, 2025. Calibrates UMCP invariants against Falque et al. (PRL, DOI 10.1103/t5dh-rx6w); sonic horizon  $v(x_h) = c_s(x_h)$  as collapse boundary; post-collapse  $\exists \omega < 0$ . doi: [10.5281/zenodo.16623285](https://doi.org/10.5281/zenodo.16623285).
- [12] Clement Paulus. Ising anyons, tqft, and the collapse regime: A gcd/umcp audit of topological quantum computation, 2025. Super-exponential leakage suppression  $\omega_{n+1} = (\omega_n)^5$ ; source: Iulianelli et al., Nat. Commun. 16, 4558 (2025);  $d_\alpha \approx -2.47$  flags Collapse. doi: [10.5281/zenodo.16745545](https://doi.org/10.5281/zenodo.16745545).
- [13] Clement Paulus. Self-sustained frictional cooling in active matter: A gcd/umcp regime audit, 2025. 5-regime refinement (adds transitional states); source: Antonov et al., Nat. Commun. 16, 7235 (2025); frictional cooling operator  $\mathcal{F}_c = \sum [v_i - v_j]^2 \delta(t - t_{\text{coll}})$ ; cluster stability  $\Gamma = F e^{-S} e^{-\lambda C / (1 + \tau_R)}$ . doi: [10.5281/zenodo.16757373](https://doi.org/10.5281/zenodo.16757373).

## License

© 2026 Clement Paulus.

This paper is shared under the *Creative Commons Attribution 4.0 International* license (CC BY 4.0).

You may copy, share, and modify this paper (including for commercial use), as long as you: (i) give credit to the author, (ii) link to the license, and (iii) say if you made changes.

License text: <https://creativecommons.org/licenses/by/4.0/>