

The Triadic Coherence Condition

A Predictive Framework for Coherence
Under Continuous Non-Stationary Streams

Vinicius Parreira

Independent Researcher

figueiredo@pm.me

March 2026

Abstract

We establish the **Triadic Coherence Condition** (TCC): any information system maintaining bounded-error fidelity to a continuously drifting environment must implement exactly three mutually irreducible functional roles, *Potentiality* ($\mathcal{P}$), *Actuation* ($\mathcal{A}$), and *Convergence* (Φ), and each is indispensable.

The necessity proof derives this ternary structure without presupposing it. The Belief Update Lemma shows that any coherence-restoring update function must depend on three arguments: the current belief state, the most recent action, and the incoming environmental signal. Removing any one argument destroys drift-coherence. Three independent formal frameworks ground this result: POMDP credit-assignment, Katsuno-Mendelzon belief-update semantics, and Input-to-State Stability theory. Sufficiency follows from a *Conditional Borromean Cascade Stability* theorem: under four explicit ISS hypotheses, the composite Lyapunov functional $V_P + V_A + V_\Phi$ is non-increasing, and removal of any component causes divergence.

We introduce the **Triadic Coherence Index** $\text{TCI} \in (0, 1]$, computed via the debiased Sinkhorn divergence with proven $O(n^{-1/2})$ convergence. Minimising the total estimation error over window size yields the **Coherence-Delay Uncertainty Principle**: $\mathcal{E}_{\min} = \frac{3}{2}C_K^{2/3}\zeta^{1/3}$, with optimal window $n^* = (C_K/\zeta)^{2/3}$, establishing that perfect coherence is thermodynamically impossible under continuous drift. The **Coercive Masking Corollary** formalises a qualitatively distinct failure mode: highly action-effective systems maintain an apparently healthy TCI by coercing the environment to conform to a stale world model, collapsing catastrophically once exogenous drift exceeds their coercive capacity.

Two complementary trackers provide empirical validation. The **Triadic Gaussian Tracker** confirms all three failure modes via Borromean ablation (full $\bar{V}_{\text{total}} = 0.027$ vs. FM-1 24.259), verifies the slow-drift estimation rate (slope -0.436), and demonstrates the $\zeta^{1/3}$ scaling law (slope 0.316) on synthetic data. The **Triadic Particle Tracker** extends validation to non-Gaussian, non-linear environments via Sequential Monte Carlo ablation and shows coercive masking under moderate coupling: at $\alpha = 0.3$ and $\lambda = 3.0$, the standard TCI remains high (0.949 ± 0.012) while TCI^E drops to 0.697 ± 0.046 .

On the ELEC2 and Bikes passive benchmarks, ADWIN remains the strongest generic detector on the residual streams, while the dynamic n_t^* strategy still improves the TCI family over fixed windows. On the KuaiRand logged benchmark (367 eligible users, healthy-phase calibration), TCI^E raises bubble detection from 45.8% to 69.2% and collapse detection from 54.8% to 70.0%, while ADWIN remains near zero on the raw watch-ratio stream.

Code: <https://github.com/jovemexausto/triadic-coherence>

Keywords: triadic coherence condition; drift-coherence; ISS stability; Wasserstein distance; Sinkhorn divergence; concept drift; Triadic Coherence Index; coercive masking; coherence-delay uncertainty.

MSC 2020: 93D05, 93D30, 68P99.

Contents

1	Introduction	5
1.1	The Capurro Trilemma and the Vacant Space	6
2	Mathematical Definitions	7
2.1	State Spaces and Information State	7
2.2	ISS Definitions	8
3	Necessity: Triadic Irreducibility	8
3.1	External Characterization: Drift-Coherence	8
3.2	The Belief Update Lemma	9
3.3	Triadic Irreducibility as Structural Consequence	11
3.4	Relation to Peirce’s Reduction Thesis	11
4	Sufficiency	12
4.1	Hypotheses	12
4.2	Main Stability Theorem	12
5	Self-Referential Extension	13
6	Connection to Ruffini’s Algorithmic Regulator Theorem	14
6.1	The Algorithmic Regulator Theorem	14
6.2	TCC as the Structural Account of Ruffini’s $M(W:R) > 0$	14
6.3	The Coercive Masking Corollary	16
7	The Three Failure Modes	17
8	The Triadic Coherence Index	19
8.1	The Estimable TCI	20
8.2	TCC-Compliance and Drift-Coherence Are Not Equivalent	23
9	Empirical Validation: The Triadic Gaussian Tracker	24
9.1	Formalization	24
9.2	Experiment A: Borromean Ablation	25
9.3	Experiment B: Sinkhorn TCI Sample Complexity and the Coherence-Delay U-Curve	26
9.4	Experiment C: ELEC2 Early-Warning Lead Time	28
9.5	Experiment D: Dynamic n_t^* Optimization on ELEC2	29
9.6	Experiment E: Coercive Masking in the TPT	31
9.7	Experiment F: Bikes Early-Warning Lead Time	33
9.8	Experiment G: KuaiRand Logged Benchmark	34
9.9	Discussion: Limitations of the TGT	35
10	Related Work	36
11	Structural Minimality	38
11.1	Why Exactly Three	38
11.2	Convergent Evidence	39

12 Open Questions	39
13 Conclusion	40
A Code Repository	44
B KuaiRand Logged Benchmark Protocol	44

1 Introduction

In 2021, Zillow’s automated home-purchasing system lost approximately \$500M after its price model failed to track a market-regime shift while its purchasing engine continued executing stale predictions. In 2012, Knight Capital lost \$440M in forty-five minutes when a deployment error caused its order-routing engine to become non-reproducible across replicas. Social-media recommendation systems continuously refine user models and deploy deterministic rankings, but their feedback operators fail to correct for the distributional shift that the recommendations themselves induce, producing filter-bubble phenomena documented at scale.

Each failure instantiates a distinct structural deficiency in an information system operating under a drifting environment. Each is predictable, diagnosable, and preventable given a formal account of what structural components a drift-coherent information system must possess. Providing that account is the purpose of this paper.

The formal question is the following:

What are the necessary and sufficient architectural conditions for an information system to maintain bounded-error fidelity to a continuously drifting environment, without requiring a central coordinating authority?

The theory of information has approached this question from three directions without convergence. Capurro and Hjørland [2003] identified the resulting fragmentation as a trilemma: any framework covering information syntax, semantics, and dynamics simultaneously sacrifices rigor in at least one dimension. The present paper occupies the vacant space in Capurro’s trilemma by grounding all three dimensions in a single formal framework derived from control theory and belief-update semantics.

Contributions

Principal results.

- 1. Necessity** (section 3): The Belief Update Lemma derives ternary structure from drift-coherence and action-effectiveness alone, via POMDP credit-assignment and Katsuno-Mendelzon update semantics, with no triadic presupposition. Passive observers are correctly identified as genuinely dyadic. Triadic irreducibility (theorem 3.7) follows as a corollary.
- 2. Conditional Borromean Cascade Stability** (section 4): Under four explicit ISS hypotheses (H-ISS-P/A/ Φ /Env), the composite Lyapunov functional V_{total} is non-increasing and diverges upon any single component removal (theorem 4.2).
- 3. Self-Referential Fixed Point** (section 5): The recursive TCC with meta-actuation operator Π_t admits a fixed point in **QBS** (theorem 5.1). A restricted fixed-point theorem (theorem 5.1) grounds the engineering approximation over a local endomorphism class.
- 4. Ruffini Correspondence** (section 6): Each TCC failure mode collapses Ruffini’s regulation gap Δ (proposition 6.3), establishing the TCC as the internal structural account of what sustains $M(W:R) > 0$.

Supporting contributions.

- a. **Formalization** (section 2): Metric-space definitions of $(\mathcal{P}, \mathcal{A}, \Phi)$ with derivable failure modes and falsifiable predictions.
- b. **Triadic Coherence Index** (section 8): Estimable scalar diagnostic $\text{TCI} \in (0, 1]$ via the debiased Sinkhorn divergence S_ε , with proven $O(n^{-1/2})$ estimation rate (proposition 8.7) and the **Coherence-Delay Uncertainty Principle** (proposition 8.9): minimum achievable tracking error $\mathcal{E}_{\min} = \frac{3}{2}C_K^{2/3}\zeta^{1/3}$, optimal window $n^* = (C_K/\zeta)^{2/3}$, empirically confirmed at slope 0.337.
- c. **Empirical Validation** (section 9): Two complementary trackers validate the theory across settings. The Triadic Gaussian Tracker (TGT) provides Bormean ablation (full $\bar{V}_{\text{total}} = 0.027$ vs FM-1 24.259), slow-drift estimation verification (slope -0.436), and the Coherence-Delay U-curve (slope 0.316, confirming $\mathcal{E}_{\min} \propto \zeta^{1/3}$). The Triadic Particle Tracker (TPT) extends validation to non-Gaussian, non-linear environments via Sequential Monte Carlo and isolates coercive masking under moderate coupling. On the passive ELEC2 and Bikes benchmarks, ADWIN remains the strongest generic detector on the residual streams, while dynamic n_t^* improves the TCI family over fixed windows. On the KuaiRand logged benchmark, calibrated on healthy random-exposure windows, TCI^E improves bubble detection from 45.8% to 69.2% and collapse detection from 54.8% to 70.0% on 367 eligible users, while ADWIN on raw watch ratio remains near zero. The Coercive Masking phenomenon is identified and formalised (corollary 6.5 and definition 6.7).
- d. **Structural Minimality** (section 11): Three independent arguments (epistemic Q1/Q2/Q3 ordering, tetrad-factoring, tiered tradition convergence) and six intellectual traditions with explicit Tier 1/2/3 evidential weighting.
- e. **Open Questions** (section 12): Combinatorial uniqueness of the decomposition and convergence rate of the Π_t self-application tower, with Coercive Masking detectability as a third open direction.

Scope and limits of the claims. The necessity proof (lemma 3.4) applies to *action-effective* systems (definition 3.3); passive observers are genuinely dyadic and explicitly outside the triadic scope. The ISS proof in section 4 is complete modulo four hypotheses (H-ISS-P/A/ Φ /Env) stated explicitly; verification for concrete systems is application-dependent. The Lawvere fixed-point theorem (section 5) establishes *existence* of a fixed point within a restricted endomorphism class; convergence rate remains open (conjecture 12.2). The Triadic Gaussian Tracker and Triadic Particle Tracker (section 9) provide controlled empirical validation in Gaussian-linear and non-linear non-Gaussian settings respectively; high-dimensional extensions remain open (section 12).

1.1 The Capurro Trilemma and the Vacant Space

In 2003, Capurro and Hjørland [2003] published a mapping of the concept of information across intellectual history. The survey revealed a field fragmented into three irreconcilable traditions, each with its virtues and its limits.

The **syntactic** tradition, inaugurated by Shannon [1948], treats information as the reduction of uncertainty a signal produces relative to a probability distribution. It is mathematically precise, empirically validated, and deliberately blind to semantic content.

The **semantic** tradition, developed by Floridi [2011], recovers content at the cost of applicability. For Floridi, information must be truthful: false data is *misinformation*, not low-quality information. This is philosophically defensible but inapplicable to real streaming systems, which are noisy, adversarial, and partially false.

The **pragmatic** tradition, represented by Bateson [1972] and developed by Hofkirchner [2013], attempts synthesis via emergent structures. It is conceptually the richest and the least operationally specified.

The *Capurro Trilemma* is the observation that any theory of information attempting to cover all three aspects pays an irrecoverable cost in at least one: Shannon maximises rigour at the cost of semantics; Floridi maximises semantics at the cost of applicability; Hofkirchner maximises scope at the cost of predictive specificity. What is absent from all three is a theory that is simultaneously **formal** (with derivable theorems), **semantic** (addressing what information *represents*), and **dynamic** (treating change as the central phenomenon rather than a complication to be eliminated). The Triadic Coherence Condition occupies this vacant space.

2 Mathematical Definitions

2.1 State Spaces and Information State

Definition 2.1 (State space). Let S be a non-empty measurable space (the *signal space*). The *information state space* $\mathcal{I}$ is the space of all probability measures on S equipped with a metric $d_{\mathcal{I}}$ compatible with weak convergence. Unless otherwise specified, $d_{\mathcal{I}}$ denotes the L^2 -Wasserstein distance W_2 .

Definition 2.2 (Potentiality). The *Potentiality* at time t is a measurable map

$$\mathcal{P}(S, t) \in \mathcal{I},$$

representing the system’s current probability measure over possible world states. $\mathcal{P}$ is ε -coherent at time t if $W_2(\mathcal{P}(S, t), \mathcal{P}^*(S, t)) \leq \varepsilon$, where $\mathcal{P}^*(S, t)$ denotes the true marginal distribution of the environment at time t .

Remark 2.3 (Computational realization of W_2). The L^2 -Wasserstein distance W_2 is chosen as the metric on $\mathcal{I}$ for its geometric faithfulness: unlike f -divergences such as KL divergence, W_2 inherits the geometry of the underlying signal space S and behaves well under spatial translations of the distribution support, a desirable property when concept drift shifts the support of $\mathcal{P}^*$. In high-dimensional signal spaces, exact computation of W_2 costs $O(n^3 \log n)$ via network-flow solvers. For streaming applications requiring real-time Φ_t updates, the entropic regularization of Cuturi [2013] replaces W_2 with the Sinkhorn divergence S_ε , computable in $O(n^2/\varepsilon^2)$ via parallelizable matrix-scaling iterations [Altschuler et al., 2017]. S_ε converges to W_2^2 as $\varepsilon \rightarrow 0$ and breaks the curse of dimensionality of raw W_2 at the cost of a bias proportional to ε [Genevay et al., 2019]. The theoretical results hold for W_2 ; their computational realization in streaming contexts is demonstrated in the TGT experiments (section 9; code at appendix A).

Definition 2.4 (Actuation). The *Actuation operator* is a measurable map

$$\mathcal{A}_t: \mathcal{I} \rightarrow \mathcal{O},$$

where $\mathcal{O}$ is the output space (distinct from lowercase o , which denotes a specific output element $o \in \mathcal{O}$). $\mathcal{A}_t$ is δ -reproducible if for any $p \in \mathcal{I}$ and any two executions ω, ω' , $d_{\mathcal{O}}(\mathcal{A}_t(p; \omega), \mathcal{A}_t(p; \omega')) \leq \delta$. A deterministic actuation is 0-reproducible.

Definition 2.5 (Convergence operator). The *Convergence operator* is a measurable map

$$\Phi_t: \mathcal{I} \times \mathcal{O} \times S \rightarrow \mathcal{I},$$

mapping the current information state, the most recent actuation, and an incoming signal to an updated information state. Φ_t is *coherence-restoring* if there exist $\beta \in \mathcal{KL}$ and $\gamma \in \mathcal{K}$ such that for all $t \geq 0$,

$$W_2(\mathcal{P}(S, t), \mathcal{P}^*(S, t)) \leq \beta(W_2(\mathcal{P}(S, 0), \mathcal{P}^*(S, 0)), t) + \gamma\left(\sup_{s \leq t} \|u(s)\|\right),$$

where $u(s)$ represents the environmental drift input at time s .

Definition 2.6 (TCC-compliant system). An information system $I(S, t)$ is *TCC-compliant* if it is equipped with all three operators $(\mathcal{P}, \mathcal{A}, \Phi)$ satisfying definitions 2.2, 2.4 and 2.5 respectively. The composite information state is

$$I(S, t) := (\mathcal{P}(S, t), \mathcal{A}_t, \Phi_t) \in \mathcal{I} \times \mathcal{F}_{\mathcal{O}} \times \mathcal{F}_{\mathcal{I}}.$$

2.2 ISS Definitions

Definition 2.7 (Class $\mathcal{K}$, $\mathcal{KL}$). A function $\gamma: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is class $\mathcal{K}$ if it is continuous, strictly increasing, and $\gamma(0) = 0$. A function $\beta: \mathbb{R}_{\geq 0}^2 \rightarrow \mathbb{R}_{\geq 0}$ is class $\mathcal{KL}$ if $\beta(\cdot, t) \in \mathcal{K}$ for each t and $\beta(r, \cdot) \rightarrow 0$ as $t \rightarrow \infty$ for each $r > 0$.

Definition 2.8 (Input-to-State Stability, ISS). A system $\dot{x} = f(x, u)$ on $\mathbb{R}^n$ is *ISS* if there exist $\beta \in \mathcal{KL}$ and $\gamma \in \mathcal{K}$ such that for all x_0 and all essentially bounded u ,

$$\|x(t)\| \leq \beta(\|x_0\|, t) + \gamma(\text{ess sup}_{s \leq t} \|u(s)\|).$$

Equivalently [Sontag, 1989], there exists a smooth ISS-Lyapunov function V with $\alpha_1(\|x\|) \leq V(x) \leq \alpha_2(\|x\|)$ and $\|x\| \geq \chi(\|u\|) \Rightarrow \nabla V \cdot f(x, u) \leq -\alpha_3(\|x\|)$ for $\alpha_1, \alpha_2, \alpha_3, \chi \in \mathcal{K}$.

3 Necessity: Triadic Irreducibility

3.1 External Characterization: Drift-Coherence

The necessity proof derives the three-component structure from two external properties: drift-coherence and action-effectiveness. Both are stated in terms of W_2 bounds and total variation distances, with no reference to the triadic decomposition.

Definition 3.1 (Drift-Coherence). An information system $I(S, t)$ is *drift-coherent* if there exist $\beta \in \mathcal{KL}$ and $\gamma \in \mathcal{K}$ such that for all $t \geq 0$ and all environmental inputs with $\text{ess sup}_{s \leq t} \|u(s)\| \leq M < \infty$:

$$W_2(\mathcal{P}(S, t), \mathcal{P}^*(S, t)) \leq \beta(W_2(\mathcal{P}(S, 0), \mathcal{P}^*(S, 0)), t) + \gamma(M).$$

The system maintains bounded fidelity to the true environmental distribution under bounded drift, with initial errors decaying over time. This is the ISS condition (definition 2.8) applied to the information state space $\mathcal{I}$, with no triadic structure presupposed.

Remark 3.2 (Ontological status of $\mathcal{P}^*$). The oracle distribution $\mathcal{P}^*(S, t)$ is the objective generative distribution of the environment at time t . It exists independently of the information system and is accessible only to an external observer assessing drift-coherence. The system’s internal state is mapped to a probability distribution $\mathcal{P}(S, t) \in \mathcal{I}$ by its own representational dynamics. Drift-coherence is an externally evaluated property: the question is whether this internal measure tracks $\mathcal{P}^*$ within bounded W_2 error. The system requires no explicit probabilistic awareness; it must merely maintain a state configuration that, under an appropriate homomorphic mapping, corresponds to a distribution tracking $\mathcal{P}^*$ in the ISS sense.

Definition 3.3 (Action-Effectiveness). The actuation operator $\mathcal{A}_t$ is *action-effective* in environment $\mathcal{X}$ if there exist distinct outputs $o, o' \in \mathcal{O}$ and a state $s \in S$ such that the environmental transition distributions differ:

$$D_{\text{TV}}\left(P(\cdot \mid s_t = s, o) \parallel P(\cdot \mid s_t = s, o')\right) > 0.$$

A system is *passive* if $D_{\text{TV}} = 0$ for all o, o', s : actions leave state transitions unaffected.

3.2 The Belief Update Lemma

Lemma 3.4 (Belief Update Lemma). *Let $I(S, t)$ be a drift-coherent, action-effective information system. Any function U whose repeated application maintains drift-coherence must depend on at least three distinct arguments:*

- (i) *the current information state $p = \mathcal{P}(S, t) \in \mathcal{I}$,*
- (ii) *the most recent actuation output $o \in \mathcal{O}$, and*
- (iii) *an incoming environmental signal $s_{t+1} \in S$.*

Each argument is indispensable.

Proof. We show by contradiction that omitting any single argument destroys drift-coherence.

(i) Necessity of p . Suppose U omits p , so $\mathcal{P}(S, t+1) = U(o, s_{t+1})$. The recursive belief state $p = \mathcal{P}(S, t)$ is the *minimal sufficient statistic* for the history $H_t = (s_1, o_1, \dots, s_t, o_t)$ [Kaelbling et al., 1998]: by the Fisher-Neyman factorization theorem, the distribution of future observations is conditionally independent of H_{t-1}

given p_t (the Markov property of POMDP belief states). By the Data Processing Inequality [Csiszár and Körner, 1981]: for any function g ,

$$I(\mathcal{P}^*(t+1); g(o, s_{t+1})) \leq I(\mathcal{P}^*(t+1); H_{t+1}),$$

with strict inequality when environmental drift is temporally structured, i.e., when $\mathcal{P}^*(t)$ exhibits temporal correlation. A memoryless $U(o, s_{t+1})$ discards H_{t-1} , incurring an information loss $\Delta I_t > 0$ at each step. Since ΔI_t accumulates monotonically, no $\beta \in \mathcal{KL}$ decay is achievable against adversarial temporally structured drift, and drift-coherence fails.

(ii) **Necessity of o .** Suppose U omits o , so $\mathcal{P}(S, t+1) = U(p, s_{t+1})$. By definition 3.3, the transition distribution $P(s_{t+1} \mid s_t, o)$ depends on o . Without o , the update conflates two distinct causes of the observed signal: exogenous drift and the environment’s response to the system’s own actuation. This is the *credit-assignment problem* [Sutton and Barto, 1998]: the update misattributes action-driven signals to drift, producing a biased posterior. In the POMDP formulation [Kaelbling et al., 1998], the correct belief update is

$$b'(s') = \eta \cdot P(o_{\text{obs}} \mid s', a) \sum_s P(s' \mid s, a) b(s),$$

which requires the action a explicitly via the transition kernel $P(s' \mid s, a)$. Omitting $a = o$ degenerates the kernel to $P(s' \mid s)$, which is incorrect whenever actions influence transitions (definition 3.3).

The formal belief change literature provides independent grounding. Katsuno and Mendelzon [1992] distinguish *belief revision* (accommodating new information about a *static* world) from *belief update* (tracking a *dynamic* world). The TCC operates in the update regime: $\mathcal{P}^*$ itself drifts. Φ_t is therefore a KM-update operator. Bonanno [2024]’s Kripke-Lewis characterization clarifies the indispensability of o : in the Lewis selection function, the most-similar-worlds ordering depends on what the system did; without o , worlds where s_{t+1} was caused by the system’s own action are conflated with worlds where it resulted from exogenous drift. The resulting bias in $\mathcal{P}$ accumulates, and for sufficiently action-effective environments $W_2(\mathcal{P}, \mathcal{P}^*)$ grows without bound, violating drift-coherence.

(iii) **Necessity of s_{t+1} .** Suppose U omits s_{t+1} , so $\mathcal{P}(S, t+1) = U(p, o)$. The information state evolves from the prior state and the action, but receives no environmental evidence. Under adversarial drift, $\mathcal{P}^*$ changes in ways not predictable from p and o alone, and no information flows from $\mathcal{P}^*$ into $\mathcal{P}$. The system operates in a closed loop. The gap $W_2(\mathcal{P}(t), \mathcal{P}^*(t))$ grows at least as fast as the drift rate $\zeta = \|\dot{\mathcal{P}}^*\|_{W_2}$, violating the $\gamma(M)$ bound of definition 3.1. $\square$

Remark 3.5 (Passive observer systems). Lemma 3.4 requires action-effectiveness (definition 3.3). For passive systems where $D_{\text{TV}} = 0$ (monitoring, forecasting, read-only sensors), actions leave transitions unaffected. The POMDP kernel reduces to an HMM filter, o drops from the transition kernel, and the system is genuinely *dyadic*: $(\mathcal{P}, \Phi)$ alone suffices to maintain drift-coherence. Passive observers thus represent the boundary case where the triadic necessity argument is inapplicable.

Corollary 3.6 (Ternary Irreducibility of the Update Function). *The drift-coherent update function $U: \mathcal{I} \times \mathcal{O} \times S \rightarrow \mathcal{I}$ is irreducibly ternary: every binary decomposition*

either embeds one of $\{p, o, s_{t+1}\}$ as an implicit slot, which reintroduces the ternary structure internally, or fails drift-coherence directly.

Proof. By lemma 3.4, all three arguments are necessary. Any binary map over two of the three either (a) fails drift-coherence directly, or (b) encodes the missing argument implicitly, producing an obfuscated ternary map. $\square$

3.3 Triadic Irreducibility as Structural Consequence

The step from “the update function U requires three arguments” to “the system must implement three architectural components” is the following: any partition of the system into fewer than three components either assigns all three arguments to one component (hiding the ternary structure inside it) or distributes the arguments across the boundary, forcing the boundary itself to act as a third role. Neither escape route produces a genuinely dyadic architecture. The three TCC components $(\mathcal{P}, \mathcal{A}, \Phi)$ are precisely the minimal realization of the three roles p, o, s_{t+1} that lemma 3.4 establishes as necessary.

Theorem 3.7 (TCC Triadic Irreducibility). *Let $I(S, t)$ be any drift-coherent, action-effective information system. For any partition of $I(S, t)$ into two functional components C_1, C_2 , at least one of C_1 or C_2 must internally implement the three-argument structure of lemma 3.4. Every strictly dyadic decomposition fails drift-coherence for action-effective systems.*

Proof. By corollary 3.6, drift-coherence requires $U: \mathcal{I} \times \mathcal{O} \times S \rightarrow \mathcal{I}$. Any partition into C_1, C_2 must either: (a) assign all three arguments to a single component, which then internally implements the full ternary structure, or (b) distribute the arguments across C_1 and C_2 , requiring an interface that passes one argument across the boundary, thereby constituting a third functional role. In either case, three roles are present. The TCC components $(\mathcal{P}, \mathcal{A}, \Phi)$ are precisely these roles: $\mathcal{P}$ maintains p , $\mathcal{A}_t$ produces o , and Φ_t maps $(p, o, s_{t+1}) \rightarrow \mathcal{I}$. $\square$

The derivation is external: ternary structure follows from drift-coherence and action-effectiveness, both defined in terms of W_2 bounds and total variation distances, with no reference to TCC-compliance.

Theorem 3.8 (Scope of theorem 3.7). *Theorem 3.7 establishes that three functional roles are necessary for drift-coherence in action-effective systems. Whether $(\mathcal{P}, \mathcal{A}, \Phi)$ is the unique minimal decomposition remains open: the combinatorial uniqueness question is stated as conjecture 12.1.*

3.4 Relation to Peirce’s Reduction Thesis

The Belief Update Lemma provides an independent derivation of ternary structure from ISS formalism. Peirce’s Reduction Thesis (PRT) constitutes independent corroborating evidence rather than the proof foundation. Burch [1991] proved PRT algebraically in Peircean Algebraic Logic; Hereth Correia and Pöschel [2011] strengthened it. The mapping $\Phi_t: \mathcal{I} \times \mathcal{O} \times S \rightarrow \mathcal{I}$ is triadic in Peirce’s sense, with necessity established independently by lemma 3.4.

Koshkin [2024] argues that certain PRT proofs employ constructions that may be question-begging. This debate is orthogonal to the TCC necessity argument: the

Belief Update Lemma stands on ISS formalism alone, and the Koshkin controversy leaves it unaffected.

4 Sufficiency: Conditional Borromean Cascade Stability

4.1 Hypotheses

The following hypotheses are required. Their verification for a concrete system is application-dependent. We exhibit a class of systems satisfying them in example 4.4.

- (H-ISS-P) The potentiality subsystem, driven by the convergence output Φ_t , admits an ISS-Lyapunov function $V_P: \mathcal{I} \rightarrow \mathbb{R}_{\geq 0}$ with comparison functions $(\alpha_{1P}, \alpha_{2P}, \alpha_{3P}, \chi_P)$.
- (H-ISS-A) The actuation subsystem, driven by $\mathcal{P}(S, t)$, admits an ISS-Lyapunov function V_A with gains $(\alpha_{1A}, \alpha_{2A}, \alpha_{3A}, \chi_A)$.
- (H-ISS- Φ) The convergence subsystem, driven by actuation outputs and environmental inputs, admits an ISS-Lyapunov function V_Φ with gains $(\alpha_{1\Phi}, \alpha_{2\Phi}, \alpha_{3\Phi}, \chi_\Phi)$.
- (H-ISS-Env) The environmental drift input $u(s)$ is essentially bounded: $\text{ess sup}_{s \geq 0} \|u(s)\| \leq M < \infty$, and the drift Lipschitz constant L_u satisfies $L_u < \lambda_{\min}(\Phi_t)$, where $\lambda_{\min}(\Phi_t)$ denotes the minimum contraction rate of the convergence operator. Equivalently: the *update bandwidth* of Φ_t strictly exceeds the *drift velocity* of the environment.

Remark 4.1 (Lyapunov construction for measure-valued $\mathcal{P}$). When $\mathcal{P}(S, t)$ is a probability measure on a Polish space S (rather than a vector in $\mathbb{R}^n$), the ISS hypotheses above require a Lyapunov functional on the metric space $(\mathcal{M}(S), W_2)$, which is a non-linear metric space lying outside the scope of classical finite-dimensional ISS. The natural construction follows the Wasserstein gradient flow theory of [Ambrosio et al. \[2005\]](#): take

$$V_P = \frac{1}{2} W_2(\mathcal{P}(S, t), \mathcal{P}^*(S, t))^2.$$

Under the Jordan–Kinderlehrer–Otto (JKO) scheme, Φ_t is coherence-restoring (definition 2.5) if and only if its induced velocity field on $\mathcal{M}(S)$ points toward $\mathcal{P}^*$ in the Otto–Riemannian sense, i.e., the dissipation inequality $\frac{d}{dt} V_P \leq -\alpha_{3P}(V_P) + \chi_P(\|u\|)$ holds along measure-valued trajectories. The Kalman instance in example 4.4 verifies (H-ISS-P/A/ Φ /Env) for the Gaussian case directly, where $W_2(\mathcal{P}, \mathcal{P}^*)^2 = \|\hat{\mu}_t - \mu_t\|^2 \in \mathbb{R}$. Verification for general measure-valued systems is application-dependent; the JKO framework provides the natural tool. For infinite-dimensional ISS in Banach and Hilbert spaces see [Mironchenko and Prieur \[2020\]](#).

4.2 Main Stability Theorem

Theorem 4.2 (Borromean Cascade Stability). *Under hypotheses (H-ISS-P), (H-ISS-A), (H-ISS- Φ), and (H-ISS-Env), the composite system $I(S, t) = (\mathcal{P}, \mathcal{A}, \Phi)$ with cascade structure*

$$\Phi_t \xrightarrow{\text{updates}} \mathcal{P}(S, t) \xrightarrow{\text{drives}} \mathcal{A}_t \xrightarrow{\text{feeds back}} \Phi_t$$

is globally asymptotically stable at the equilibrium $I^* = (\mathcal{P}^*, \mathcal{A}^*, \Phi^*)$: the composite Lyapunov functional $V_{\text{total}} := V_P + V_A + V_\Phi$ is monotonically non-increasing along trajectories, and $V_{\text{total}}(t) \rightarrow 0$ as $t \rightarrow \infty$ for all bounded environmental inputs. Moreover, removing any single component from the cascade disrupts the monotone decrease of V_{total} in general.

Proof. Cascade stability. Applying the ISS Cascade Theorem [Sontag, 1995, Theorem 1] twice: the sub-cascade $(\Phi) \rightarrow (\mathcal{P})$ is ISS by (H-ISS-P) and (H-ISS- Φ); the extended cascade $(\Phi \rightarrow \mathcal{P}) \rightarrow (\mathcal{A})$ is ISS by (H-ISS-A). Hence the full cascade is ISS, equivalent to global asymptotic stability for bounded inputs. By the ISS-Lyapunov characterisation [Sontag, 1989], each V_i is non-increasing along its subsystem’s trajectories, and cross-terms are absorbed by the χ_i gains [Dashkovskiy et al., 2010]. Hence $V_{\text{total}} = V_P + V_A + V_\Phi$ is non-increasing.

Divergence upon component removal.

- *Remove Φ :* By proposition 7.6, V_P grows unboundedly when the environment drifts and $\Phi_t \equiv \text{id}$; hence $\dot{V}_P > 0$ is possible.
- *Remove $\mathcal{P}$:* V_A requires bounded input from $\mathcal{P}$ (hypothesis H-ISS-A). Without $\mathcal{P}$, the input to $\mathcal{A}_t$ is undefined or unbounded; H-ISS-A loses its applicability.
- *Remove $\mathcal{A}$:* V_Φ requires actuation outputs as inputs (H-ISS- Φ). Without $\mathcal{A}_t$, the feedback loop to Φ_t is severed; the system cannot reduce V_P , which grows by proposition 7.1.

In all three cases V_{total} is not guaranteed non-increasing. $\square$

Remark 4.3 (Why “Borromean”). The Borromean rings are three topological rings linked such that removing any one releases the other two [Cromwell et al., 1998]. theorem 4.2 establishes an analogous property: no two-element subset of $\{\mathcal{P}, \mathcal{A}, \Phi\}$ suffices for stability.

Example 4.4 (Concrete instantiation satisfying H-ISS-P/A/ Φ). Let $\mathcal{P}^*(S, t) \sim \mathcal{N}(\mu_t, \Sigma)$ with

$$\mu_t = \mu_0 + \int_0^t \eta(s) ds, \quad \|\eta\|_\infty < M.$$

Let $\mathcal{P}(S, t) = \mathcal{N}(\hat{\mu}_t, \Sigma)$ be the Kalman-filtered estimate, $\mathcal{A}_t(\mathcal{P}) = \hat{\mu}_t$ (mean-actuation), and Φ_t the standard Kalman update rule.

Then $V_P = W_2(\mathcal{P}, \mathcal{P}^*)^2 = \|\hat{\mu}_t - \mu_t\|^2$, $V_A = \|\mathcal{A}_t(\mathcal{P}) - \mathcal{A}_t^*\|^2$, and $V_\Phi = \|\text{innovation covariance}\|_F^2$ are all standard ISS-Lyapunov functions for their respective subsystems [Jazwinski, 1970], so (H-ISS-P/A/ Φ) hold.

5 Self-Referential Extension

The base TCC $(\mathcal{P}, \mathcal{A}, \Phi)$ maintains drift-coherence when the environment drifts within a stable parameter envelope. When the drift dynamics themselves change (meta-drift, meaning abrupt shifts in drift rate ζ or correlation structure), the base operators’ fixed parameters become inadequate.

A meta-actuation operator Π_t reconfigures the parameters of $(\mathcal{P}, \mathcal{A}, \Phi)$ in response to detected regime shifts, applying the triadic logic recursively to the parameter space. This creates a vertical hierarchy, not a flat tetradic extension: Π_t applies the same triadic structure one organizational level above.

Theorem 5.1 (Restricted Fixed Point). *Let $\mathbf{QBS}$ be the cartesian closed category of quasi-Borel spaces. In the recursive TCC equipped with Π_t , there exists a restricted class $\mathcal{G} \subset \mathcal{F}(\mathcal{I}, \mathcal{I})$ of endomorphisms and a fixed information state $I(S, t)^*$ such that $\text{TCC}(I(S, t)^*) = I(S, t)^*$, provided Π_t is locally point-surjective over $\mathcal{G}$.*

In the TGT, $\mathcal{G}$ consists of Kalman gain recalibrations, process-noise adjustments, and observation-noise rescalings, all δ -stable by construction. The fixed point is the steady-state Kalman gain $K^* = \sigma_w^2 / (\sigma_w^2 + \sigma_{\text{obs}}^2)$.

Formal proof of theorem 5.1 and the convergence rate of the Π_t tower are stated as open questions in section 12.

6 Connection to Ruffini’s Algorithmic Regulator Theorem

6.1 The Algorithmic Regulator Theorem

Ruffini [2026] proves an independent result in algorithmic information theory (AIT), recasting the classical Good Regulator Theorem [Conant and Ashby, 1970] in a distribution-free, single-trajectory framework. The paper contains three theorems; the result cited here is Theorem 1 of Ruffini [2026], which we refer to as the *Algorithmic Regulator Theorem* (ART) for brevity.¹

Let W be a world and R a regulator. Ruffini defines the *regulation gap* $\Delta(W, R) := K(O_{W, \emptyset}) - K(O_{W, R})$, where $K(\cdot)$ denotes Kolmogorov complexity and $O_{W, R}$ is the world trajectory under regulator R . He further defines $M(W : R)$ as the algorithmic mutual information between world and regulator.

Theorem 6.1 (Ruffini 2026, paraphrased). *For any world-regulator pair (W, R) ,*

$$\Pr[(W, R) \mid x] \leq C \cdot 2^{M(W : R)} \cdot 2^{-\Delta(W, R)},$$

for a universal constant C . Consequently, any regulator sustaining $\Delta(W, R) > 0$ must maintain positive algorithmic mutual information $M(W : R) > 0$ with the world.

6.2 tcc as the Structural Account of Ruffini’s $M(W : R) > 0$

Ruffini’s framework is a *dyad*: (W, R) . The regulator R is a single object whose internal structure is not further analyzed. The TCC asks the next question: *what internal structure must R have to sustain $\Delta > 0$ under continuous, non-stationary streaming?*

¹Ruffini’s own terminology is the “Good Algorithmic Regulator” (GAR) result; the label ART is adopted here for conciseness and to distinguish it from Ruffini’s Theorems 2 and 3.

Lemma 6.2 (Drift Decay of $M(W : R)$). *Let $\mathcal{P}(S, t) = \mathcal{P}(S, t_0)$ be frozen after t_0 . Let the world drift at rate $\zeta > 0$ such that $W_2(\mathcal{P}^*(t), \mathcal{P}^*(t_0)) \geq \zeta(t - t_0)$. Under the assumption that $K(O_{W,R} | \mathcal{P}(S, t))$ is bounded above by a monotone increasing function of $W_2(\mathcal{P}, \mathcal{P}^*)$, so that the output trajectory's conditional complexity grows with the model's divergence from reality, we have:*

$$M(W : R)(t) \leq \frac{C_0}{1 + \zeta(t - t_0)} \rightarrow 0 \quad \text{as } t \rightarrow \infty,$$

for some constant $C_0 > 0$.

Proof. By proposition 7.1, static $\mathcal{P}$ implies $W_2(\mathcal{P}(t), \mathcal{P}^*(t)) \geq \zeta(t - t_0)$. The algorithmic mutual information satisfies $M(W : R) \leq K(R) - K(R | W)$. Under the assumption, $K(R | W) = K(O_{W,R} | \mathcal{P})$ grows with $W_2(\mathcal{P}, \mathcal{P}^*)$, so $M(W : R)$ decreases at least as fast as $C_0/(1 + W_2(\mathcal{P}, \mathcal{P}^*)) \leq C_0/(1 + \zeta(t - t_0))$. $\square$

The monotonicity assumption in lemma 6.2 holds generically under non-mean-reverting drift: as the frozen model diverges from reality, it loses predictive power and $K(O_{W,R} | \mathcal{P})$ grows. It *fails* when drift is mean-reverting and the frozen model remains partially accurate by reversion. In that case, $M(W : R)$ may not collapse to zero; instead it stabilizes at a level determined by the residual alignment between $\mathcal{P}(S, t_0)$ and the stationary distribution of the mean-reverting process. Case (i) of proposition 6.3 is therefore qualified: $\Delta \rightarrow 0$ under non-mean-reverting drift; under mean-reverting drift, Δ stabilizes at a value that may be positive but is bounded by the residual predictive power of the frozen model.

Proposition 6.3 (TCC $\Rightarrow$ Ruffini's $M(W : R) > 0$). *An TCC-compliant system necessarily maintains $M(W : R) > 0$. Conversely, each TCC component failure causes the regulation gap Δ to collapse or become non-measurable:*

- (i) Static- $\mathcal{P}$ failure: *Under non-mean-reverting drift, lemma 6.2 gives $M(W : R) \rightarrow 0$, hence $\Delta \rightarrow 0$. Under mean-reverting drift, Δ may stabilize at a reduced but positive value; see the qualification above.*
- (ii) Non-deterministic- $\mathcal{A}$ failure: *$K(O_{W,R})$ becomes non-computable given $\mathcal{P}$ alone; Δ is ill-defined and the regulator ceases to be well-defined in Ruffini's framework.*
- (iii) Absent- Φ failure: *Without Φ , actuation feedback leaves $\mathcal{P}$ unupdated; $K(O_{W,R}) \rightarrow K(O_{W,\emptyset})$ as $t \rightarrow \infty$, so $\Delta \rightarrow 0$.*

Proof. For (i): by lemma 6.2, static $\mathcal{P}$ under non-mean-reverting drift implies $M(W : R) \rightarrow 0$, hence $\Delta \rightarrow 0$. For (ii): non-deterministic $\mathcal{A}$ makes $K(O_{W,R})$ substrate-dependent and Δ non-reproducible across replicas. For (iii): absent Φ means no feedback loop; $\mathcal{P}$ never updates on W , so $M(W : R)$ falls to zero by the same argument as case (i) applied to the actively diverging posterior. $\square$

Authority without center. The political-philosophical consequence of the TCC is that coherence without central authority constitutes a structural theorem, not a utopian ideal. TCC-compliant systems exhibit three properties simultaneously: (i) every agent holds a local representation of $\mathcal{P}$, and TCC-compliance guarantees their

convergence without any single authoritative version; (ii) the transition function is identical for all agents by δ -reproducibility, requiring no arbiter for $\mathcal{A}$; (iii) convergence emerges from the algebraic structure, requiring no external coordinator for Φ . The authority emerges from the whole because the whole has the triadic structure.

6.3 The Coercive Masking Corollary

Proposition 6.3 characterizes three mechanisms by which a TCC failure mode collapses the regulation gap Δ . The coercive-masking experiment in the TPT (section 9.6) reveals a fourth, qualitatively distinct failure mechanism, one the standard ART analysis conflates with FM-1 rather than treating separately.

Definition 6.4 (Coercive Regime). An information system $I(S, t) = (\mathcal{P}, \mathcal{A}, \Phi)$ operates in the *coercive regime* if:

- (i) $\mathcal{A}_t$ is action-effective (definition 3.3) with coupling coefficient $\alpha > 0$: the transition kernel is $P(s_{t+1} \mid s_t, o) = \mathcal{N}(\mu^*(t) + \alpha(o - \mu^*(t)), \sigma_w^2)$;
- (ii) $\mathcal{P}$ is static (proposition 7.1): $\hat{\mu}_t \equiv \hat{\mu}_0$ and $\mathcal{A}_t(\mathcal{P}) \equiv \hat{\mu}_0$;
- (iii) the exogenous drift is bounded: $|\zeta| < \alpha \cdot |\hat{\mu}_0 - \mu_0|$.

In the coercive regime, the system's stale actions continuously pull μ^* toward the frozen estimate $\hat{\mu}_0$: $\mu^*(t+1) = \mu^*(t) + \zeta + \alpha(\hat{\mu}_0 - \mu^*(t))$, which is a mean-reverting process with equilibrium $\hat{\mu}_0 + \zeta/\alpha$. When $|\zeta/\alpha|$ is small, μ^* is held near $\hat{\mu}_0$, suppressing V_P and making the system appear healthy by the TCI.

Corollary 6.5 (Coercive Masking). *In the coercive regime (definition 6.4), the standard TCI diagnostic $\sigma_P(t) = 1/(1 + V_P(t))$ fails to detect FM-1: $\sigma_P(t) \approx 1$ for all $t \leq t_{\text{mask}}$, where $t_{\text{mask}} = O(\alpha^{-1} \log(\alpha/\zeta))$ is the coercive horizon. Beyond t_{mask} , exogenous drift overwhelms the coercive capacity and V_P diverges catastrophically.*

Proof sketch. Under condition (iii) of definition 6.4, the gap $\delta_t = \mu^*(t) - \hat{\mu}_0$ satisfies $\delta_{t+1} = (1 - \alpha)\delta_t + \zeta$. The fixed point is $\delta^* = \zeta/\alpha$, so $V_P(t) \leq \frac{1}{2}(\zeta/\alpha)^2 + \sigma_w^2/(2\alpha)$ in expectation, bounded independently of t , masking FM-1 for as long as condition (iii) holds. Once $|\zeta| \cdot t_{\text{mask}} \geq \alpha \cdot |\hat{\mu}_0 - \mu_0|$, the coercive capacity is exhausted and V_P grows unboundedly by proposition 7.1. $\square$

Remark 6.6 (Coercive masking and the Ruffini gap). A coercive agent maintains $\Delta = K(O_{W,\emptyset}) - K(O_{W,R}) > 0$ by *reducing* $K(O_{W,\emptyset})$ rather than increasing $M(W : R)$: the agent's continuous coercive actions simplify the world's output trajectory, compressing $O_{W,\emptyset}$ toward $O_{W,R}$, rather than improving the regulator's model to track a complex world. In Ruffini's framework, both mechanisms produce $\Delta > 0$, but they are structurally distinct: the first is adaptive (the regulator updates to track a complex world), the second is coercive (the regulator forces the world to match its stale model). The TCC's V_P diagnostic requires additional information about whether the world's low complexity is intrinsic or imposed in order to distinguish between them.

Definition 6.7 (Effort-Corrected TCI). The *effort signal* at time t is any monotone proxy for the corrective actuation burden required to keep the world aligned with the system’s present output. A generic observable surrogate is

$$E_t^{\text{obs}} := |\mathcal{A}_t(\mathcal{P}) - \mathcal{A}_{t-1}(\mathcal{P})|,$$

the magnitude of change in the system’s actuation. In controlled synthetic settings, a sharper proxy is the realized coercive effort

$$E_t^{\text{sim}} := \alpha |a_t - s_t^{\text{free}}|,$$

where s_t^{free} is the state that would have arisen before applying control. In a well-adapted system, effort is low because the world and the model are aligned. In a coercive system, effort grows as the actuation must bridge an increasing model–world gap.

The *effort-corrected TCI* is

$$\text{TCI}^E(I(S, t), t) := \min(\sigma_P(t) \cdot e^{-\lambda E_t/E_0}, \sigma_A(t), \sigma_\Phi(t)),$$

where $E_0 > 0$ is a reference effort scale and $\lambda > 0$ is a sensitivity parameter. TCI^E penalizes increasing actuation effort, correctly diagnosing FM-1 even in the coercive regime where $\sigma_P(t) \approx 1$.

The empirical evidence for coercive masking comes from the TPT experiment (section 9.6). In the passive FM-1 regime, the standard TCI and TCI^E coincide (0.795 ± 0.062), because no realized forcing of the world occurs. Under moderate coupling ($\alpha = 0.3$, $\lambda = 3.0$), the same frozen model reaches $\text{TCI} = 0.949 \pm 0.012$ while $\text{TCI}^E = 0.697 \pm 0.046$, showing that the apparent coherence is being purchased by intervention. The effect is already visible at $\alpha = 0.2$, where the mean masking gap $\text{TCI} - \text{TCI}^E$ is 0.194 in the parameter sweep.

The coercive masking phenomenon provides a formal account of a class of real-world failures: highly influential systems (financial trading algorithms, content recommendation engines) can *appear* healthy by standard drift metrics while actively suppressing the environmental variation that would reveal their staleness. Failure arrives catastrophically once exogenous drift exceeds the coercive capacity, rather than through the gradual degradation visible in passive environments.

7 The Three Failure Modes

The TCC yields three characteristic and identifiable failure modes, one for each component, each with a closed-form Lyapunov divergence and a documented high-consequence instantiation.

Proposition 7.1 (Failure Mode 1: Static- $\mathcal{P}$ collapse). *Suppose Φ_t ceases to update $\mathcal{P}$, so $\mathcal{P}(S, t) = \mathcal{P}(S, t_0)$ for all $t > t_0$. If the environment undergoes distributional shift $\mathcal{P}^*(S, t) \neq \mathcal{P}^*(S, t_0)$ for large t , then $W_2(\mathcal{P}(S, t), \mathcal{P}^*(S, t)) \rightarrow \infty$ and the expected actuation error $\mathbb{E}[d_{\mathcal{O}}(\mathcal{A}_t(\mathcal{P}), \mathcal{A}_t^*)]$ is unbounded.*

Proof. $\mathcal{P}(S, t) = \mathcal{P}(S, t_0)$ is fixed while $\mathcal{P}^*$ drifts, so $W_2 \rightarrow \infty$ by definition. Lipschitz continuity of $\mathcal{A}_t$ (standard regularity) then propagates the divergence to actuation error. $\square$

Remark 7.2 (The Streaming Degradation Problem). FM-1 manifests concretely in streaming data systems. Static approximate nearest neighbor indices fix their partition geometry at build time. Under continuous streaming, the data distribution drifts while the partition remains frozen, causing progressive centroid drift, cell imbalance, and monotonic recall erosion with no remedy short of a full $O(N \cdot d)$ rebuild. This is the engineering realization of $W_2(\mathcal{P}, \mathcal{P}^*) \rightarrow \infty$ for a static $\mathcal{P}$ (proposition 7.1).

Remark 7.3 (Zillow Offers, 2021). Zillow’s automated home-purchasing system maintained price-prediction models that were updated at a timescale insufficient for the 2021 real-estate market regime change. The resulting actuation error, purchasing homes above market value, cost approximately \$500M before the business unit was closed. This instantiates proposition 7.1.

Proposition 7.4 (Failure Mode 2: Non-deterministic- $\mathcal{A}$ divergence). *If $\mathcal{A}_t$ is not δ -reproducible for any $\delta < \infty$, then $H(\mathcal{A}_t \mid \mathcal{P}(S, t)) > 0$ for all t : actuation carries information not determined by the state model. Distributed replicas of the system diverge, and the composite resists auditing and replication.*

Proof. Non- δ -reproducibility implies $\exists \omega \neq \omega'$ with $\mathcal{A}_t(p; \omega) \neq \mathcal{A}_t(p; \omega')$, giving positive conditional entropy. Auditability requires reconstructing actuation from state; this fails when conditional entropy is positive. $\square$

Remark 7.5 (Knight Capital, 2012). A deployment error activated legacy code alongside new code in Knight Capital’s execution engine, producing non-deterministic order routing. With accurate market models ($\mathcal{P}$) but non-reproducible $\mathcal{A}_t$, the system issued $\approx 4\,000\,000$ erroneous trades in 45 minutes, resulting in a \$440M loss. This instantiates proposition 7.4.

Proposition 7.6 (Failure Mode 3: Absent- Φ monopolization). *Let $\Phi_t \equiv \text{id}$ (no update). In any closed-loop deployment, actuation feedback shifts $\mathcal{P}^*$; since $\mathcal{P}$ cannot track, the distributional mismatch grows at a rate lower-bounded by the rate of environmental response to actuations. In open-loop settings, if actuation load is allocated by a Hebbian rule $\dot{s}_j = \eta \nu_j s_j$, then $\lim_{t \rightarrow \infty} s_{j^*} / \sum_j s_j = 1$ where $j^* = \arg \max_j s_j(0)$: winner-take-all collapse.*

Proof. The closed-loop case follows from $\mathcal{P}^*(t+1) \leftarrow f(\mathcal{P}^*(t), \mathcal{A}_t(\mathcal{P}(t)))$ with $\mathcal{P}$ fixed, so the gap $W_2(\mathcal{P}, \mathcal{P}^*)$ grows at least as fast as the variation of f in its second argument, which is strictly positive whenever $\mathcal{A}_t(\mathcal{P}) \neq \mathcal{A}_t^*$. The Hebbian limit follows from the ODE $\dot{s}_j = \eta \nu_j s_j$ with constant load ν_j , whose solution is $s_j(t) = s_j(0)e^{\eta \nu_j t}$; normalization then yields the stated limit. $\square$

Remark 7.7 (Social-media echo chambers). Recommender systems that maintain accurate user models ($\mathcal{P}$) and act on them ($\mathcal{A}$) but apply no distribution-shift correction (Φ) produce recommendation feedback that shifts $\mathcal{P}^*$ toward the prior $\mathcal{P}$, endogenously reinforcing initial biases. This is an instance of proposition 7.6.

Remark 7.8 (Hebbian Instability as FM-3). The winner-take-all limit of proposition 7.6 can be derived explicitly in the neuromorphic mapping where partition j has access frequency $\nu_j(t)$ (firing rate) and size $s_j(t)$ (synaptic weight). Pure Hebbian routing assigns load proportionally to size: $\nu_j \propto s_j$, giving the ODE $\dot{s}_j(t) = \eta \cdot \nu_j(t) \cdot s_j(t) = \eta s_j(t)^2 / \sum_k s_k(t)$. Normalization yields $s_{j^*}(t) / \sum_j s_j(t) \rightarrow 1$ as $t \rightarrow \infty$ where $j^* = \arg \max_j s_j(0)$: the partition with the largest initial size

monopolizes all load. The coherence-restoring Φ_t prevents this by introducing non-linear inhibitory feedback: a penalty $\phi_j \propto (\nu_j/\rho^*)^\gamma$ with $\gamma > 1$ creates super-linear negative feedback above the homeostatic set-point ρ^* , bounding $\nu_j^* \leq \rho^*(1 + \epsilon_\gamma)$.

Remark 7.9 (Falsifiability). propositions 7.1, 7.4 and 7.6 are falsifiable. To falsify proposition 7.1, one must exhibit a system with static $\mathcal{P}$ that maintains bounded error under adversarial distributional drift without rebuilding. The streaming degradation literature records no such case.

8 The Triadic Coherence Index

TCC-compliance is a binary property: a system either satisfies definition 2.6 or it fails to do so. In practice, systems degrade continuously: $\mathcal{P}$ stalens gradually, $\mathcal{A}$ drifts toward non-determinism incrementally, Φ weakens before failing completely. The Lyapunov functionals V_P , V_A , V_Φ of section 4 already measure this degradation; a normalization that makes the three scores commensurate and combines them into a single computable diagnostic completes the picture.

Definition 8.1 (Component Coherence Scores). Let $I(S, t) = (\mathcal{P}, \mathcal{A}, \Phi)$ be an information system with ISS-Lyapunov functionals V_P , V_A , V_Φ as in section 4. Define the three *component coherence scores* at time t as:

$$\sigma_P(t) := \frac{1}{1 + V_P(t)} = \frac{1}{1 + \frac{1}{2}W_2(\mathcal{P}(S, t), \mathcal{P}^*(S, t))^2}, \quad (1)$$

$$\sigma_A(t) := \exp(-H(\mathcal{A}_t | \mathcal{P}(S, t))), \quad (2)$$

$$\sigma_\Phi(t) := \frac{1}{1 + V_\Phi(t)}. \quad (3)$$

Each score takes values in $(0, 1]$, equals 1 when the corresponding component is ideal, and approaches 0 as the component degrades.

Definition 8.2 (Triadic Coherence Index). The *Triadic Coherence Index* of a system $I(S, t)$ at time t is

$$\text{TCI}(I(S, t), t) := \min(\sigma_P(t), \sigma_A(t), \sigma_\Phi(t)) \in (0, 1].$$

The min aggregation is the correct choice given the Borromean property (theorem 4.2): system coherence is bounded by the weakest component, and no strength in two components compensates for failure in the third. An arithmetic mean would obscure localized failures; the min preserves the diagnostic identity of the bottleneck.

Remark 8.3 (Bottleneck sensitivity and effort-corrected TCI). The use of $\min(\sigma_P, \sigma_A, \sigma_\Phi)$ is deliberately bottleneck-sensitive: it is a diagnostic of the weakest active subsystem, not a smooth summary of overall system quality. Aggregators such as arithmetic or geometric means may be reported as descriptive summaries, but they are not valid replacements for the structural diagnostic given by definition 8.2.

In action-effective environments, however, high actuation capacity can keep σ_P artificially high by forcing the world toward a stale model. When such coercive masking is possible, a conservative extension is the *effort-corrected TCI*

$$\text{TCI}^E(I(S, t), t) := \min(\sigma_P(t)e^{-\lambda E_t/E_0}, \sigma_A(t), \sigma_\Phi(t)),$$

where E_t is any monotone proxy for corrective actuation burden, $E_0 > 0$ is a reference scale, and $\lambda > 0$ controls sensitivity. In the synthetic TPT we operationalise E_t as realized coercive effort $\alpha|a_t - s_t^{\text{free}}|$; in deployed systems, action increments, energy expenditure, or control-surface usage can play the same role. For logged recommendation benchmarks, the preferred instantiation is the per-user KL divergence between the current exposure-tag distribution and the healthy random-exposure baseline distribution, aggregated sequentially over time and then batched across users.

Remark 8.4 (Optimality discussion). The scaling law behind proposition 8.7 is not an artefact of the sliding-window estimator alone. The two terms in the bound, $n^{-1/2}$ and $\zeta\tau$, encode a structural bias–variance trade-off that any estimator with finite effective memory must confront: the first term reflects the statistical limit of estimating a distribution from finitely many samples under subgaussian noise, and the second reflects the mismatch induced by drift when the estimator necessarily looks at stale information.

Informally, one may interpret this as a minimax lower-bound phenomenon over the class of trajectories with bounded drift rate ζ : if two admissible trajectories differ only by a perturbation below the resolution induced by finite memory, they are indistinguishable at the level of observations, and any estimator must incur an error of the same order. The estimator used in proposition 8.7 achieves the resulting cube-root scaling up to constants, so the $\zeta^{1/3}$ law should be read as near-optimal under the stated regularity assumptions rather than as a peculiarity of a single windowing scheme.

More concretely, any estimator with finite effective memory must incur both a statistical error term of order $n^{-1/2}$ and a drift-induced bias of order ζn , so no choice of estimator can simultaneously eliminate both contributions. This argument applies to estimators with finite effective memory, including sliding-window and exponentially weighted schemes.

This interpretation is consistent with similar exponents in non-stationary on-line learning and adaptive filtering, although typically in cumulative or algorithm-specific formulations. Optimizing the balance between these unavoidable contributions yields the cube-root scaling, suggesting that it reflects a minimax rate rather than an artifact of a specific construction. Any improvement beyond this regime should therefore come from stronger structural assumptions on the drift, not from a different choice of finite-memory estimator.

8.1 The Estimable TCI

Definition 8.1 requires $\mathcal{P}^*(S, t)$, the true oracle distribution of the environment, to compute $\sigma_P(t)$. In any deployed system, $\mathcal{P}^*$ is unobservable by construction: a system with complete knowledge of $\mathcal{P}^*$ would have nothing to estimate. We therefore introduce an estimable version of σ_P using an empirical approximation of $\mathcal{P}^*$.

Definition 8.5 (Empirical Potentiality Estimator). Let $\{s_{t-n+1}, \dots, s_t\}$ be the n most recent environmental signals observed by the system. The *empirical potentiality estimator* is the empirical measure:

$$\hat{\mathcal{P}}^*(S, t) := \frac{1}{n} \sum_{k=t-n+1}^t \delta_{s_k},$$

where δ_{s_k} is the Dirac measure at observation s_k .

Definition 8.6 (Estimable TCI). The *debiased Sinkhorn divergence* between measures μ, ν is

$$S_\varepsilon(\mu, \nu) := T_\varepsilon(\mu, \nu) - \frac{1}{2}T_\varepsilon(\mu, \mu) - \frac{1}{2}T_\varepsilon(\nu, \nu),$$

where T_ε is the raw entropically regularized optimal transport cost. Unlike T_ε , the debiased divergence satisfies $S_\varepsilon(\mu, \mu) = 0$, restoring the metric property and enabling the dimension-independent convergence rate of proposition 8.7.

The *estimable component* $\hat{\sigma}_P$ replaces $\mathcal{P}^*$ with $\hat{\mathcal{P}}^*$ in eq. (1) using S_ε :

$$\hat{\sigma}_P(t) := \frac{1}{1 + \frac{1}{2}S_\varepsilon\left(\mathcal{P}(S, t), \hat{\mathcal{P}}^*(S, t)\right)^2}.$$

The *estimable TCI* is

$$\widehat{\text{TCI}}(I(S, t), t) := \min(\hat{\sigma}_P(t), \sigma_A(t), \sigma_\Phi(t)).$$

Note that σ_A and σ_Φ remain exactly computable without oracle access: $\sigma_A = e^{-H(\mathcal{A}_t|\mathcal{P})}$ is estimated from replica divergence logs, and $\sigma_\Phi = 1/(1 + V_\Phi)$ from the convergence subsystem's Lyapunov functional.

Proposition 8.7 (Estimation Error Bound for $\hat{\sigma}_P$). *Let $\mathcal{P}^*(S, t)$ be K -subgaussian with drift rate $\zeta = \text{ess sup}_s \|\dot{\mathcal{P}}^*(s)\|_{W_2}$. Let $\hat{\mathcal{P}}^*$ be the sliding-window empirical estimator with window n (definition 8.5), and let τ be the lag between the center of the window and t . Then under the entropic regularization of remark 2.3:*

$$|\sigma_P(t) - \hat{\sigma}_P(t)| \leq \frac{C_K \cdot (n^{-1/2} + \zeta\tau)}{(1 + V_P(t))^2},$$

where C_K depends only on the subgaussian constant K .

Proof. By the triangle inequality for S_ε :

$$|S_\varepsilon(\mathcal{P}, \mathcal{P}^*) - S_\varepsilon(\mathcal{P}, \hat{\mathcal{P}}^*)| \leq S_\varepsilon(\mathcal{P}^*, \hat{\mathcal{P}}^*).$$

Under K -subgaussian $\mathcal{P}^*$, Theorem 3 of [Genevay et al. \[2019\]](#) establishes that the debiased Sinkhorn divergence achieves a dimension-independent parametric rate: $\mathbb{E}[S_\varepsilon(\mathcal{P}^*, \hat{\mathcal{P}}^*)] \leq C_{K,\varepsilon} n^{-1/2}$, where $C_{K,\varepsilon}$ depends on the subgaussian norm K and the regularization parameter ε . (Theorem 2 of the same paper, which proves that the Sinkhorn potentials lie in a Sobolev ball independent of the measures, is the enabling lemma for this rate.) The result holds when $\mathcal{P}^*$ is supported on bounded subsets of $\mathbb{R}^d$ with C^1 cost; for unbounded subgaussian measures, see [Mena and Weed \[2019\]](#) for an improvement. This rate is specific to the debiased S_ε ; the raw cost T_ε retains an entropy bias $T_\varepsilon(\mu, \mu) > 0$ that prevents the dimension-independent convergence from holding. An additional drift error of $\zeta\tau$ arises because $\hat{\mathcal{P}}^*$ is centered τ steps in the past. The bound on $|\sigma_P - \hat{\sigma}_P|$ follows from the Lipschitz continuity of $x \mapsto 1/(1 + \frac{1}{2}x^2)$ with derivative $\leq 1/(1 + V_P)^2$. $\square$

proposition 8.7 makes $\widehat{\text{TCI}}$ operationally grounded: the estimation error decreases as the window n grows and increases proportionally to drift rate ζ . The

required window size for target error η is $n = O(C_{K,\varepsilon}^2/\eta^2)$, independent of dimension.

But this characterization reveals a deeper structure. When window lag is accounted for explicitly ($\tau = n/2$ for a centred window), the total estimation error decomposes as a sum of two opposing terms. Minimising over n yields a fundamental lower bound on achievable estimation accuracy that holds regardless of the estimator or window strategy chosen.

Definition 8.8 (Lag-Inclusive Estimation Error). For a sliding window of size n centred at lag $\tau = n/2$, the total expected estimation error is:

$$\mathcal{E}(n) = C_K n^{-1/2} + \frac{1}{2} \zeta n,$$

where the first term is the *variance* (decreasing in n) and the second is the *lag-bias* (increasing in n).

Proposition 8.9 (Coherence-Delay Uncertainty Principle). *Under definition 8.8, the estimation error $\mathcal{E}(n)$ is minimised at the optimal window*

$$n^* = \left(\frac{C_K}{\zeta} \right)^{2/3},$$

yielding the minimum achievable tracking error

$$\mathcal{E}_{\min} = \frac{3}{2} C_K^{2/3} \zeta^{1/3}.$$

Proof. Setting $d\mathcal{E}/dn = 0$: $-\frac{1}{2}C_K n^{-3/2} + \frac{1}{2}\zeta = 0$, giving $n^* = (C_K/\zeta)^{2/3}$. Substituting back: $\mathcal{E}(n^*) = C_K(C_K/\zeta)^{-1/3} + \frac{1}{2}\zeta(C_K/\zeta)^{2/3} = C_K^{2/3}\zeta^{1/3} + \frac{1}{2}C_K^{2/3}\zeta^{1/3} = \frac{3}{2}C_K^{2/3}\zeta^{1/3}$. $\square$

Remark 8.10 (Interpretation: a thermodynamic floor). Proposition 8.9 establishes that perfect coherence ($V_P = 0$) is *impossible* under continuous drift, regardless of the estimator used. Every drift-coherent system must carry a residual tracking error at least $\mathcal{E}_{\min} \propto \zeta^{1/3}$. This is the information-theoretic analogue of a thermodynamic limit: as in Heisenberg's uncertainty principle, the two sources of error (variance and lag) trade against each other with no common lower bound of zero. The cube-root scaling has a concrete engineering consequence: a fifty-fold increase in drift rate ζ raises the minimum error floor by only $50^{1/3} \approx 3.7\times$, while simultaneously forcing the optimal window n^* to contract by $50^{2/3} \approx 14\times$. The optimal window $n^* = (C_K/\zeta)^{2/3}$ provides a *self-tuning rule*: given a real-time estimate $\hat{\zeta}_t$ of the local drift rate, the system should maintain window size $n_t^* = (C_K/\hat{\zeta}_t)^{2/3}$ to operate at the minimum of the error U-curve at every step. This transitions the TCI from a passive diagnostic into an active self-optimising control mechanism.

Remark 8.11 (Practical estimation of V_Φ for black-box Φ_t). The estimable TCI requires $\sigma_\Phi = 1/(1 + V_\Phi)$, where V_Φ is the ISS-Lyapunov functional for the convergence subsystem. When Φ_t is a black box, with internal dynamics unknown and ISS comparison functions $(\alpha_{1\Phi}, \alpha_{2\Phi}, \alpha_{3\Phi}, \chi_\Phi)$ unavailable, V_Φ admits no closed-form derivation. In such cases, V_Φ is estimated empirically via Counterexample-Guided Inductive Synthesis (CEGIS): a learner proposes a candidate $\hat{V}_\Phi$ from observed state trajectories $\{(\Phi_t(p_i, o_i, s_i), p_i)\}_i$; a verifier checks the dissipation condition $\Delta\hat{V}_\Phi \leq 0$

against sampled trajectories using local Lipschitz bounds; failed verification returns counterexamples to the learner. The resulting $\hat{V}_\Phi$ yields a conservative lower-bound estimate $\hat{\sigma}_\Phi = 1/(1 + \hat{V}_\Phi) \leq \sigma_\Phi$, usable in TCI for any opaque convergence subsystem.

8.2 TCC-Compliance and Drift-Coherence Are Not Equivalent

A structural asymmetry in the framework deserves explicit statement.

Drift-coherence implies TCC-compliance. By lemma 3.4 and theorem 3.7, any drift-coherent, action-effective system must implement the three functional roles $(\mathcal{P}, \mathcal{A}, \Phi)$. Drift-coherence is the stronger behavioral condition; TCC-compliance is its necessary architectural consequence.

TCC-compliance is necessary but requires additional dynamical conditions for drift-coherence. A system may implement $(\mathcal{P}, \mathcal{A}, \Phi)$ structurally while failing to satisfy the ISS hypotheses (H-ISS-P/A/ Φ /Env) required for theorem 4.2. The components may be present but improperly tuned, have inadequate bandwidth, or lack sufficient expressiveness. TCC-compliance is the necessary architectural precondition; the ISS hypotheses are the sufficient dynamical conditions.

An analogy: an engine requires cylinders, spark plugs, and fuel injectors (TCC-compliance = architectural pattern) but will not run without fuel and proper timing (ISS hypotheses = dynamical conditions). The design target is drift-coherence. TCC-compliance is the structural precondition; the ISS hypotheses are what verify the structure is functioning correctly.

Proposition 8.12 (TCI Characterizes TCC-Compliance). *Let $I(S, t)$ be an TCC-compliant system under hypotheses (H-ISS-P/A/ Φ /Env). Then:*

- (i) $\text{TCI}(I(S, t), t) = 1$ if and only if $V_P(t) = 0$, $H(\mathcal{A}_t | \mathcal{P}) = 0$, and $V_\Phi(t) = 0$ simultaneously, i.e., the system is at the equilibrium $I^* = (\mathcal{P}^*, \mathcal{A}^*, \Phi^*)$.
- (ii) $\text{TCI}(I(S, t), t) \rightarrow 0$ if and only if at least one of the three failure modes is active: $V_P(t) \rightarrow \infty$ (FM-1), $H(\mathcal{A}_t | \mathcal{P}) \rightarrow \infty$ (FM-2), or $V_\Phi(t) \rightarrow \infty$ (FM-3).
- (iii) The identity of $\arg \min(\sigma_P, \sigma_A, \sigma_\Phi)$ names the bottleneck component: the component whose degradation is the binding constraint on system coherence at time t .

Proof. (i) Each score is a strictly decreasing function of its Lyapunov functional or conditional entropy: $\sigma_P = 1 \Leftrightarrow V_P = 0$, $\sigma_A = 1 \Leftrightarrow H(\mathcal{A}_t | \mathcal{P}) = 0$, $\sigma_\Phi = 1 \Leftrightarrow V_\Phi = 0$. The min equals 1 iff all three equal 1 simultaneously.

(ii) As $V_P(t) \rightarrow \infty$, $\sigma_P(t) \rightarrow 0$, so $\text{TCI} \rightarrow 0$; similarly for $V_\Phi(t) \rightarrow \infty$. As $H(\mathcal{A}_t | \mathcal{P}) \rightarrow \infty$, $\sigma_A(t) = e^{-H} \rightarrow 0$, so $\text{TCI} \rightarrow 0$. The converse holds because $\text{TCI} \rightarrow 0$ requires at least one $\sigma_i \rightarrow 0$, which by the above requires the corresponding functional to diverge.

(iii) $\text{TCI} = \min_i \sigma_i$; the argmin identifies which subsystem's functional is largest, i.e., which component is furthest from its ideal state. $\square$

9 Empirical Validation: The Triadic Gaussian Tracker

The preceding sections establish the TCC mathematically. This section provides self-contained empirical validation through the *Triadic Gaussian Tracker* (TGT), a minimal, analytically tractable information system that instantiates all three TCC components and admits closed-form computation of all Lyapunov functionals.

The TGT serves three functions: (i) it verifies the Borromean Cascade Stability theorem by controlled ablation of each component under pure drift (section 9.2); (ii) it validates the estimation error bound of proposition 8.7 numerically and characterizes the lag-drift tradeoff (section 9.3); and (iii) it demonstrates the TCI as a pre-failure early-warning diagnostic on a real concept-drift dataset (section 9.4). Code and reproduction instructions are provided in appendix A.

9.1 Formalization

The TGT extends example 4.4. Let the environment produce observations from a drifting Gaussian:

$$s_t \sim \mathcal{N}(\mu^*(t), \sigma_{\text{obs}}^2), \quad \mu^*(t) = \mu_0 + \zeta t + \sigma_w B_t,$$

where B_t is standard Brownian motion. The TGT components are:

- $\mathcal{P}(S, t) = \mathcal{N}(\hat{\mu}_t, \hat{\sigma}_t^2)$: Gaussian belief state.
- $\mathcal{A}_t(\mathcal{P}) = \hat{\mu}_t$: deterministic mean extraction ($H(\mathcal{A}_t | \mathcal{P}) = 0$, $\delta = 0$).
- Φ_t : Kalman measurement update (see example 4.4).

All three Lyapunov functionals are available in closed form (example 4.4), making the TGT the ideal minimal test bed for all three TCC results.

Lemma 9.1 (TGT satisfies the ISS hypotheses). *The TGT with $\sigma_w > 0$ and $\sigma_{\text{obs}} > 0$ satisfies all four ISS hypotheses (H-ISS-P/A/ Φ /Env) of theorem 4.2:*

- (i) **H-ISS-P**: $V_P(t) = \frac{1}{2}W_2(\mathcal{P}, \mathcal{P}^*)^2 = \frac{1}{2}(\hat{\mu}_t - \mu^*(t))^2$ is a valid ISS-Lyapunov functional; the Kalman update contracts V_P at rate $1 - K_t$ per step, where $K_t \in (0, 1)$ is the Kalman gain.
- (ii) **H-ISS-A**: $V_A = H(\mathcal{A}_t | \mathcal{P}) = 0$ since $\mathcal{A}_t = \hat{\mu}_t$ is deterministic; H-ISS-A is satisfied trivially.
- (iii) **H-ISS- Φ** : the Kalman recursion for $\hat{\sigma}_t^2$ converges to the steady-state variance $\hat{\sigma}_\infty^2 = \frac{1}{2}(\sigma_w^2 + \sqrt{\sigma_w^4 + 4\sigma_w^2\sigma_{\text{obs}}^2})$, a stable fixed point; $V_\Phi = |\hat{\sigma}_t^2 - \hat{\sigma}_\infty^2|$ decays geometrically with rate $1 - K_\infty$.
- (iv) **H-ISS-Env**: the environmental drift rate ζ is bounded by assumption; H-ISS-Env holds whenever $\zeta < (1 - K_\infty)\sigma_{\text{obs}}$.

Proof. Items (i)–(iii) follow directly from the Kalman filter’s convergence properties [Sontag, 1989, Mironchenko and Prieur, 2020]. Item (iv) is the H-ISS-Env condition of theorem 4.2 applied to the scalar Gaussian case; the contraction rate equals the steady-state Kalman gain $K_\infty = \hat{\sigma}_\infty^2/(\hat{\sigma}_\infty^2 + \sigma_{\text{obs}}^2)$. $\square$

Table 1: TGT experimental parameters. All experiments use `uv sync` from the repository for exact environment reproduction.

Parameter	Exp A	Exp B	Exp C/D
T (steps)	500	5,000	45,312 (ELEC2)
Seeds	42	0–49	fixed
σ_{obs}	1.0	1.0	1.0
ζ	0.02	0.001	estimated
σ_w	0.1	0.0 / 0.035	N/A
ε (Sinkhorn)	0.1	0.1	0.1
Window n	N/A	5–100	100
EMA α	N/A	N/A	0.05
PH threshold	N/A	N/A	84
PH δ	N/A	N/A	0.01
Warning level	N/A	N/A	0.295
Max lead gap	N/A	N/A	5,000 steps

9.2 Experiment A: Borromean Ablation

Protocol. We run $T = 500$ steps with $\zeta = 0.02$, $\sigma_w = 0.1$, $\sigma_{\text{obs}} = 1.0$, seed = 42, under four conditions:

C1. Full tcc: standard TGT.

C2. FM-1 (static $\mathcal{P}$): $\hat{\mu}_t \equiv 0$ for all t ; Kalman update disabled.

C3. FM-2 (non-det. $\mathcal{A}$): $\mathcal{A}_t = \hat{\mu}_t + \epsilon_t$, $\epsilon_t \sim \mathcal{N}(0, 0.25)$.

C4. FM-3 (absent Φ): Kalman update disabled; no feedback.

Environment: pure drift with no action feedback ($\alpha = 0$), so the world drifts independently of the system’s actuation. The effect of action-environment coupling is addressed in the Coercive Masking Corollary (corollary 6.5).

Results. Table 2 summarizes the tail-mean Lyapunov functionals and component coherence scores.

Table 2: Ablation results ($T = 500$, $\zeta = 0.02$, influence= 0, seed= 42). Tail means computed over the final 100 steps.

Condition	$\bar{V}_P$	$\bar{V}_A$	$\bar{V}_\Phi$	$\bar{\sigma}_P$	$\bar{\sigma}_A$	$\bar{\sigma}_\Phi$
Full TCC	0.017	0.000	0.010	0.984	1.000	0.990
FM-1 (static $\mathcal{P}$)	24.259	0.000	0.000	0.040	1.000	1.000
FM-2 (noisy $\mathcal{A}$)	0.017	1.061	0.010	0.984	0.530	0.990
FM-3 (absent Φ)	2.154	0.000	5.400	0.328	1.000	0.157

The results confirm theorem 4.2: V_{total} is lowest under the full triad. FM-1 and FM-3 both fail, but via different components: the TCI correctly identifies $\mathcal{P}$ as the bottleneck in FM-1 ($\bar{\sigma}_P = 0.040$) and Φ as the bottleneck in FM-3 ($\bar{\sigma}_\Phi = 0.157$), while FM-2 is selectively degraded through $\mathcal{A}$ ($\bar{\sigma}_A = 0.530$). The passive full-triad run remains sharply concentrated near the equilibrium ($\bar{V}_{\text{total}} = 0.027$), whereas FM-1 rises to $\bar{V}_{\text{total}} = 24.259$, providing the diagnostic precision claimed in definition 8.2.




Figure 1: Borromean ablation: V_{total} (log scale, 20-step rolling mean, top) and TCI (bottom) under the four conditions. Full TCC is the only condition where TCI remains bounded away from 0. FM-1 and FM-3 both diverge in V_{total} but are distinguished by different bottleneck components: $\sigma_P \rightarrow 0$ for FM-1 (stale $\mathcal{P}$) and $\sigma_\Phi \rightarrow 0$ for FM-3 (absent Φ).

V_{total} is not sample-path monotone, consistent with theorem 4.2 which guarantees non-increasing *expectation* under the ISS hypotheses, not sample-path monotonicity of the stochastic process.

9.3 Experiment B: Sinkhorn TCI Sample Complexity and the Coherence-Delay U-Curve

Protocol. We use the TGT at stationary regime ($T = 5000$, 50 seeds, $\sigma_{\text{obs}} = 1$) and compute $\hat{\sigma}_P(t)$ using sliding-window estimators with $n \in \{5, 10, 25, 50, 100\}$. Mean absolute error (MAE) $|\sigma_P - \hat{\sigma}_P|$ is measured in the tail (last 40% of T).

The lag-drift tradeoff and the Coherence-Delay Principle. The estimation error has two competing terms (definition 8.8):

$$\mathcal{E}(n) \leq \underbrace{C_K n^{-1/2}}_{\text{variance term}} + \underbrace{L \cdot \zeta \cdot \frac{n}{2}}_{\text{lag-bias term}}. \quad (4)$$

At large n or fast drift, the lag-bias term dominates and the total error increases with n , hiding the $O(n^{-1/2})$ variance rate.

To validate proposition 8.9 empirically, we additionally run a U-curve sweep: for each of four drift rates $\zeta \in \{0.001, 0.005, 0.01, 0.05\}$, we measure the raw tracking

error $|\mu^*(t) - \bar{s}_{t-n:t}|$ across $n \in \{5, 10, 20, 50, 75, 100, 150, 200, 300, 500\}$ in a passive environment (influence = 0) over $T = 6,000$ steps and 20 seeds.




Figure 2: Coherence-Delay U-curve validation (proposition 8.9). *Left*: empirical tracking error vs. window size n for four drift rates (log scale; dotted lines are scaled theoretical curves; stars mark empirical n^*). Each curve shows the predicted U-shape: variance-dominated decay for small n , lag-dominated growth for large n . *Right*: $\mathcal{E}_{\min}$ vs. ζ on log-log scale. Fitted slope = 0.316; theoretical prediction = $1/3$.

Results. Figure 2 confirms the two predictions of proposition 8.9. The U-curves show the characteristic variance-lag tradeoff for each ζ , with empirical n^* shrinking as drift increases (from $n^* \approx 100$ at $\zeta = 0.001$ to $n^* \approx 10$ at $\zeta = 0.05$). The log-log slope of $\mathcal{E}_{\min}$ vs. ζ is **0.316**, matching the theoretical $1/3$ exponent to within measurement precision. Table 3 reports the extracted minima.

Table 3: Empirical validation of the Coherence-Delay bound ($T = 6,000$, 20 seeds, passive environment).

ζ	n_{obs}^*	$\mathcal{E}_{\min}$	$\mathcal{E}_{\min}/\zeta^{1/3}$
0.001	100	0.092	0.924
0.005	50	0.157	0.916
0.010	20	0.197	0.916
0.050	10	0.318	0.862

The ratio $\mathcal{E}_{\min}/\zeta^{1/3}$ remains close to constant (≈ 0.90 across the four drift levels), confirming the expected $\zeta^{1/3}$ scaling despite finite-window effects.

For the TCI estimation rate, we use slow drift ($\zeta = 0.001$, $\sigma_w = 0.035$) to stay near the variance-dominated regime while retaining a visible lag term at $n = 100$. The observed log-log slope is -0.436 (target -0.5); the deviation reflects the onset of lag-bias at the largest window and residual autocorrelation induced by the slow-drift regime (table 4 and fig. 3).




Figure 3: Sinkhorn TCI sample complexity in the variance-dominated regime ($\zeta = 0.001$, $n \leq 100$, 50 seeds, $T = 5000$). Empirical log-log slope = -0.436 ; theoretical target -0.5 . This is the variance-only regime of the U-curve in fig. 2.

Table 4: Sinkhorn TCI estimation error vs window size ($\zeta = 0.001$, 50 seeds, $T = 5000$).

n	MAE $ \sigma_P - \hat{\sigma}_P $
5	0.0801
10	0.0457
25	0.0233
50	0.0188
100	0.0245
Log-log slope: -0.436 (target: -0.5)	

9.4 Experiment C: ELEC2 Early-Warning Lead Time

Dataset. The ELEC2 dataset [Harries, 1999] records 45,312 hourly observations of electricity demand in New South Wales, Australia, exhibiting documented concept drift due to market deregulation. We apply a 1D Gaussian TGT to the normalized demand stream.

Protocol. Failures are detected by the Page-Hinkley test [Harries, 1999, threshold = 84, $\delta = 0.01$], yielding 113 regime-shift events on the normalized NSW demand stream. TCI warnings are recorded on each falling-edge crossing of $\hat{\sigma}_P < 0.295$. Because ELEC2 exposes no direct actuation or effort channel, this benchmark uses the estimable TCI rather than TCI^E ; the effort-corrected score is reserved for the active/coercive setting of section 9.6. Lead time is defined as the gap between a warning and the nearest subsequent Page-Hinkley event (nearest-future matching; see appendix A for implementation).




Figure 4: ELEC2 early-warning diagnostic. *Top*: TCI $\hat{\sigma}_P(t)$ (sliding window $n = 100$) with warning threshold 0.295 (dashed red). Vertical lines mark TCI warnings (red), ADWIN warnings (purple), and Page-Hinkley failures (grey). *Bottom*: residual signal fed to ADWIN (50-step rolling mean).

Results. Across the 45,312-step stream, the fixed- $n = 100$ TCI issues 58 warnings (falling-edge crossings of $\hat{\sigma}_P < 0.295$), of which **51 are matched to a subsequent Page-Hinkley event** within a maximum gap of 5,000 steps. The median lead time is **1,324 steps**; mean is 1,376; minimum 45; maximum 3,830. For comparison, ADWIN with $\delta = 0.03$ on the same residual stream issues 114 warnings, of which 101 are matched to a subsequent Page-Hinkley event, at 88.6% precision and median lead 1,476. At lower thresholds the detector becomes more conservative; at higher thresholds it produces more warnings but the same one-to-one matching rule reduces effective precision. The result should be read narrowly: on this passive benchmark ADWIN is the stronger generic change detector, but the TCI remains a tuneable structural diagnostic whose score semantics carry over to the active setting where TCI^E becomes the correct observable.

9.5 Experiment D: Dynamic n_t^* Optimization on ELEC2

Proposition 8.9 provides the self-tuning rule $n_t^* = (C_K/\hat{\zeta}_t)^{2/3}$. This experiment tests whether deploying that rule dynamically on ELEC2 outperforms fixed-window baselines in matched lead count while remaining in the same precision band.

Protocol. We estimate the local drift rate at each step using an exponentially weighted moving average (EMA, decay $\alpha = 0.05$) of the absolute change in mean between consecutive windows of $d = 50$ steps: $\hat{\zeta}_t = \text{EMA}_\alpha(|\bar{s}_{t-d:t} - \bar{s}_{t-2d:t-d}|/d)$.

The constant C_K is calibrated once on the first 2,000 steps (near-stationary): $C_K \approx 1.25 \overline{|s_t - \bar{s}|} \cdot \sqrt{100}$. The resulting $n_t^* = (C_K / \hat{\zeta}_t)^{2/3}$ is clipped to $[30, 300]$, with $\hat{\zeta}_t$ estimated from consecutive 24-step means and EMA decay $\alpha = 0.03$. Three TCI strategies are compared: fixed $n = 50$, fixed $n = 300$, and dynamic n_t^* . ADWIN on the fixed- $n = 100$ residual stream is included as an external passive baseline. TCI warnings and lead times are computed identically to section 9.4, using the same warning threshold 0.295 and one-to-one future matching.




Figure 5: Dynamic n_t^* optimization on ELEC2 ($C_K \approx 1.25 \overline{|s_t - \bar{s}|} \sqrt{100}$, EMA $\alpha = 0.03$, drift window $d = 24$). *Top*: TCI $\hat{\sigma}_P$ (100-step mean) for all three strategies; grey lines are PH regime shifts. *Middle*: dynamic window size n_t^* , contracting sharply near regime shifts and expanding during stable periods. *Bottom*: lead-time distributions, including the ADWIN baseline.

Results. The results confirm proposition 8.9’s prediction in a qualified sense. Among the TCI-based strategies, the dynamic rule detects the most failures (82 leads vs. 71, 51, and 24), consistent with the theoretical result that n_t^* minimises instantaneous tracking error at each step. Unlike the earlier coarse-window baseline, it does so while remaining in the same precision band as fixed $n = 300$ (86% vs. 86%), indicating that the extra warnings are not merely threshold chatter.

ADWIN is a stronger generic passive-stream detector on this benchmark, reaching 101 matched leads at 89% precision on the same residual signal. That does not undercut the role of the TCI: ADWIN yields a changepoint alarm, whereas the TCI

Table 5: ELEC2 early-warning comparison: fixed vs. dynamic windows and ADWIN. 113 PH failures; warning level = 0.295; max gap = 5,000.

Strategy	Warnings	Leads	Precision	Median lead	Min lead
Fixed $n = 50$	24	24	100%	490	51
Fixed $n = 100$	58	51	88%	1,324	45
Fixed $n = 300$	83	71	86%	1,681	0
Dynamic n_t^*	95	82	86%	1,604	189
ADWIN ($\delta = 0.03$)	114	101	89%	1,476	38

remains tied to the potentiality error term and therefore extends naturally to the effort-corrected TCI^E used in the active coercive-masking experiment. The Bikes benchmark below shows the same passive-stream pattern on a longer real-world demand sequence.

The trade-off is structural and expected: the fixed $n = 50$ strategy achieves perfect precision by being highly conservative, but detects only 24 matched events. Fixed $n = 300$ suppresses variance while accumulating long lead times, but still misses 42 of the 113 PH events under the matching rule. The dynamic n_t^* navigates the Coherence-Delay U-curve in real time, contracting on high- $\hat{\zeta}_t$ periods and expanding on stable ones (visible in the middle panel of fig. 5). The primary limitation is still $\hat{\zeta}_t$ estimation noise: a sharper local drift estimator would likely reduce the residual gap between the dynamic and fixed- $n = 300$ median lead times while preserving the dynamic strategy’s lead-count advantage.

9.6 Experiment E: Coercive Masking in the TPT

This experiment instantiates the Coercive Masking Corollary (corollary 6.5 and definition 6.7) in the non-Gaussian, non-linear Triadic Particle Tracker. The environment evolves according to the scalar particle-filter simulator implemented in `tpt.py`, but now with an explicit action-coupling parameter $\alpha \in [0, 1]$ so that the actuation can pull the latent state toward the stale controller output. In the synthetic setting we operationalise the effort proxy as the *realized coercive effort*

$$E_t^{\text{sim}} = \alpha |a_t - s_t^{\text{free}}|,$$

where s_t^{free} is the latent state that would have arisen at time t before applying control. This lets us distinguish mere action-world mismatch from actual world-forcing.

Protocol. We freeze $\mathcal{P}$ (FM-1) and compare two regimes: passive ($\alpha = 0$) and coercive ($\alpha > 0$). The main single-run diagnostic uses $\alpha = 0.3$ and effort penalty $\lambda = 3.0$. We then sweep $\alpha \in \{0.1, 0.2, 0.3, 0.5, 0.7, 0.9\}$ and $\lambda \in \{0.5, 1.0, 2.0, 3.0, 4.0\}$ over 8 seeds, exporting all raw and aggregated metrics from the experiment harness.

Results. The key effect is immediate. Under coercion, the stale FM-1 controller appears nearly healthy under the standard TCI (0.949 ± 0.012) because the environment is being dragged toward the frozen action, shrinking the visible tracking error to 0.276 ± 0.046 . The effort-corrected index, however, falls to 0.697 ± 0.046 , correctly diagnosing that this "coherence" is being purchased by intervention rather



Figure 6: Synthetic coercive masking in the Triadic Particle Tracker. *Top*: passive FM-1 retains high tracking error, while coercive FM-1 suppresses the visible error by pulling the world toward the frozen action. *Middle*: the action gap remains nonzero in both cases, but realized coercive effort appears only when influence is positive. *Bottom*: standard TCI stays artificially high in the coercive regime, whereas TCI^E drops once effort is penalized.

than adaptation. In the passive regime, where no forcing occurs, TCI and TCI^E coincide.

The sweep confirms that the correction is not an artifact of one parameter choice. At fixed $\lambda = 3.0$, the masking gap increases from 0.115 at $\alpha = 0.1$ to 0.481 at $\alpha = 0.9$, with a clear intermediate effect already at $\alpha = 0.2$ (gap = 0.194). At fixed $\alpha = 0.3$, increasing λ from 0.5 to 4.0 expands the gap from 0.051 to 0.313. Thus, stronger coercive capacity and stronger effort penalisation separate apparent from genuine coherence in the expected direction. The TPT therefore closes the empirical loop on corollary 6.5: standard TCI can be masked by control, whereas TCI^E remains sensitive to the intervention cost that creates the mask.

To test whether this advantage survives comparison with a generic change detector, we build an active benchmark with two known regime transitions: healthy tracking for $t < 200$, coercive masking for $200 \leq t < 400$, and coercive collapse for $t \geq 400$. The masking phase freezes $\mathcal{P}$ while maintaining moderate coupling ($\alpha = 0.3$); the collapse phase sharply increases exogenous drift. We compare thresholded TCI, thresholded TCI^E , ADWIN, PageHinkley, KSWIN, and NoDrift on the

Table 6: Multi-seed coercive masking summary for FM-1 in the TPT ($\alpha = 0.3$, $\lambda = 3.0$, 8 seeds).

Regime	$ e $	action gap	effort	TCI	TCI^E
Passive	0.686 ± 0.158	0.686 ± 0.158	0.000 ± 0.000	0.795 ± 0.062	0.795 ± 0.062
Coercive	0.276 ± 0.046	0.394 ± 0.066	0.118 ± 0.020	0.949 ± 0.012	0.697 ± 0.046




Figure 7: Multi-seed parameter sweep for coercive masking in the TPT (8 seeds). Left: mean coercive TCI. Middle: mean coercive TCI^E . Right: masking gap $\text{TCI} - \text{TCI}^E$. The standard TCI monotonically rises with influence, while the masking gap widens with both influence and effort penalty λ .

raw observation stream, giving the generic detectors a fully observable signal rather than a model-derived residual.

The benchmark isolates the point that passive ELEC2 cannot address. ADWIN is stronger on passive streams because it is a generic changepoint detector. PageHinkley is reliable but slower, KSWIN is too conservative on this signal, and NoDrift is the sanity baseline. In the active setting, however, the relevant anomaly is not merely distributional change but coercive maintenance of an apparently healthy world-model fit. Here TCI^E detects masking onset in all 12 runs with median delay 5.0, whereas ADWIN detects only 7 of 12 masking onsets and does so with median delay 279.0. For the collapse transition the same ordering persists: ADWIN detects only 5 of 12 runs, while TCI and TCI^E detect all of them, with TCI^E firing earliest. This is the regime in which the effort-corrected index provides information that a generic passive detector does not encode.

9.7 Experiment F: Bikes Early-Warning Lead Time

Dataset. The River Bikes stream records a long real-world bike-demand sequence with seasonal and weather-driven changes. We normalize the target stream to unit variance and use it as a second passive benchmark beyond ELEC2.

Protocol. As in ELEC2, Page-Hinkley events on the normalized stream define the warning reference. The same fixed-window TCI strategies ($n = 50, 100, 300$), the dynamic n_t^* rule, and ADWIN on the residual stream are evaluated under the same one-to-one future matching rule.




Figure 8: Active benchmark in the TPT with two known phase changes. Orange line: masking onset at $t = 200$. Red line: collapse onset at $t = 400$. ADWIN receives the raw observation stream. TCI and TCI^E are thresholded at 0.8.

Results. On Bikes, ADWIN again outperforms the TCI family as a generic passive detector, while the dynamic window rule improves on fixed $n = 100$ but does not close the precision gap. This confirms that the passive-stream caveat is not ELEC2-specific.

9.8 Experiment G: KuaiRand Logged Benchmark

Dataset and protocol. We next test the effort-corrected diagnostic on the KuaiRand-Pure sequential recommendation logs. The random-exposure slice is treated as healthy, the first half of the later standard-recommendation slice as bubble formation, and the second half as collapse. The benchmark is evaluated on 367 users with at least 20 interactions in each relevant phase. To keep the comparison honest, the TCI and TCI^E thresholds are calibrated only on healthy random-exposure windows, using the 20th percentile of the healthy score distribution.

Table 7: Active benchmark summary on the TPT (12 seeds). Detection is measured against the known onset of coercive masking and the known onset of coercive collapse.

Phase	Detector	Detections	Rate	Median delay
Masking onset	TCI	12/12	1.00	25.5
Masking onset	TCI ^E	12/12	1.00	5.0
Masking onset	ADWIN	7/12	0.58	279.0
Masking onset	PageHinkley	12/12	1.00	68.5
Masking onset	KSWIN	1/12	0.08	118.0
Masking onset	NoDrift	0/12	0.00	NA
Collapse onset	TCI	12/12	1.00	2.0
Collapse onset	TCI ^E	12/12	1.00	0.0
Collapse onset	ADWIN	5/12	0.42	111.0
Collapse onset	PageHinkley	12/12	1.00	37.5
Collapse onset	KSWIN	0/12	0.00	NA
Collapse onset	NoDrift	0/12	0.00	NA




Figure 9: Calibration grid for TCI^E over effort penalty λ and threshold. The masking detection rate saturates at 1.0 across most of the grid, so the useful calibration axes are delay and false positives; the paper setting ($\lambda = 3.0$, threshold 0.80) remains a low-delay compromise.

Results. The passive-baseline detectors remain informative but weak on the raw watch-ratio stream. The effort-corrected score is not uniformly faster than plain TCI, but it is consistently more sensitive: bubble detection rises from 45.8% to 69.2%, and collapse detection rises from 54.8% to 70.0%. This is the strongest empirical evidence that the effort term exposes a form of control-induced coherence that a generic passive detector misses.

9.9 Discussion: Limitations of the TGT

Four limitations bound the scope of these results. First, the Gaussian-linear setting satisfies H-ISS-P/A/ Φ /Env by construction (example 4.4); non-Gaussian, non-linear systems require case-by-case hypothesis verification. Second, the Sinkhorn rate is validated only in the variance-dominated regime ($\zeta \leq 0.001$, $n \leq 100$); high-dimensional or fast-drift systems require the debiased S_ϵ with lag correction. Third,




Figure 10: Bikes early-warning diagnostic. *Top*: TCI $\hat{\sigma}_P(t)$ (sliding window $n = 100$) with warning threshold 0.295 (dashed red). Vertical lines mark TCI warnings (red), ADWIN warnings (purple), and Page-Hinkley failures (grey). *Bottom*: residual signal fed to ADWIN (50-step rolling mean).

the ELEC2 failure definition (Page-Hinkley events) is one operational choice; different failure definitions yield different lead-time statistics. Fourth, the dynamic n_t^* strategy’s precision depends critically on $\hat{\zeta}_t$ estimation quality; a more accurate real-time drift estimator would sharpen the precision-recall frontier.

The TGT provides the minimal, reproducible, analytically tractable validation of the TCC results. High-dimensional, non-Gaussian, and non-linear extensions constitute open directions for future validation (section 12).

10 Related Work

Information theory. Shannon’s framework [Shannon, 1948] treats information as uncertainty reduction, syntactically precise but semantically neutral. Floridi’s philosophy of information [Floridi, 2011] recovers content at the cost of applicability to noisy, adversarial streams. Bateson’s pragmatic tradition [Bateson, 1972] attempts synthesis but lacks operational specification. The TCC occupies the space identified by Capurro and Hjørland [2003] as vacant: simultaneously formal, semantic, and dynamic.

Concept drift detection. A mature literature addresses distributional shift in streaming data [Gama et al., 2014]. ADWIN [Bifet and Gavalda, 2007] and DDM [Gama et al., 2004] signal *when* drift has occurred. WATCH [Faber et al., 2022]

Table 8: Bikes early-warning comparison: fixed vs. dynamic windows and ADWIN. The same one-to-one future matching rule is used as in ELEC2.

Strategy	Warnings	Leads	Precision	Median lead	Min lead
Fixed $n = 50$	572	323	56%	3,040	5
Fixed $n = 100$	989	362	37%	4,001	45
Fixed $n = 300$	3,363	389	12%	4,529	64
Dynamic n_t^*	717	335	47%	3,604	6
ADWIN ($\delta = 0.03$)	477	394	83%	4,502.5	395




Figure 11: KuaiRand logged benchmark. The random-exposure phase is treated as healthy, the first half of the later standard-recommendation slice as bubble formation, and the second half as collapse. Thresholds for TCI and TCI^E are calibrated only on healthy windows.

uses W_2 to detect shift in feature space. The TCC operates at a different level: it prescribes *what structural components* a system must have to remain coherent under drift, and predicts *which component* is failing when coherence degrades. Detectors such as ADWIN are natural candidates for implementing Φ_t ; they are instantiations, not competitors.

ISS control theory. Sontag [1989] introduced ISS and established the Lyapunov characterization. Dashkovskiy et al. [2010] proved cascade stability for ISS systems. Mironchenko and Prieur [2020] extended ISS to infinite-dimensional spaces. The TCC applies ISS to probability-measure spaces metrized by W_2 , which, to our knowledge, has no direct precedent.

Belief revision and update. Alchourrón et al. [1985] established the AGM postulates for revision. Katsuno and Mendelzon [1992] introduced the update/revision distinction critical for dynamic worlds. Bonanno [2024] provides the Kripke-Lewis characterization exploited in the Belief Update Lemma.

Algorithmic regulation. Ruffini [2026] establishes the ART: a good regulator must maintain positive algorithmic mutual information with the world. The TCC

Table 9: KuaiRand logged benchmark summary on 367 users.

Phase	Detector	Detections	Rate	Median delay
Bubble onset	TCI	168/367	0.458	27.0
Bubble onset	TCI ^E	254/367	0.692	36.0
Bubble onset	ADWIN	4/367	0.011	54.5
Bubble onset	PageHinkley	170/367	0.463	28.0
Bubble onset	KSWIN	5/367	0.014	30.0
Bubble onset	NoDrift	0/367	0.00	NA
Collapse onset	TCI	201/367	0.548	42.0
Collapse onset	TCI ^E	257/367	0.700	50.0
Collapse onset	ADWIN	1/367	0.003	52.0
Collapse onset	PageHinkley	260/367	0.708	34.0
Collapse onset	KSWIN	8/367	0.022	19.0
Collapse onset	NoDrift	0/367	0.00	NA

provides the internal structural account of how $M(W : R) > 0$ is sustained (see section 6).

Neural architectures. Kosowski et al. [2025] show that the BDH architecture instantiates the triadic structure at scale: synaptic state ($\mathcal{P}$), forward pass ($\mathcal{A}$), Hebbian update (Φ). Training without feedback collapses into catastrophic forgetting, reproducing FM-3.

11 Structural Minimality

11.1 Why Exactly Three

Any adaptive information system must answer three questions in a fixed dependency order:

Q1. What is the world? $\Rightarrow \mathcal{P}(S, t)$

Q2. What should I do? $\Rightarrow \mathcal{A}_t(\mathcal{P})$

Q3. Was I right? $\Rightarrow \Phi_t(\mathcal{P}, \mathcal{A}_t, s_{t+1})$

Q2 presupposes Q1 but cannot answer it. Q3 presupposes Q2 but cannot answer it. Q1 cannot incorporate Q3’s answer without Q2 having produced a committed action, because without a committed action there is no causal structure to attribute. This is a directed acyclic graph with exactly three nodes, matching the three arguments of the Belief Update Lemma (lemma 3.4).

Any tetradic decomposition either introduces redundancy or posits a fourth question independent of Q1–Q3. Candidate fourth questions reduce: confidence is a property of $\mathcal{P}$ ’s width (encoded in V_P); drift rate is an environmental observable used by Π_t at the meta-level; policy selection is a sub-structure of $\mathcal{A}_t$. Three formal traditions independently derived the same ternary structure by different means: ISS cascade stability (section 4), POMDP credit-assignment (lemma 3.4), and KM-update/Kripke-Lewis semantics (section 10).

11.2 Convergent Evidence

Table 10: Four traditions and their TCC correspondences.

Tradition	$\mathcal{P}$	$\mathcal{A}$	Φ
Peirce semiotics [Hartshorne and Weiss, 1931–1935]	Firstness	Secondness	Thirdness
Practopoiesis [Nikolić, 2015]	genetic	anatomical	cognitive
ISS control theory	state measure	transition	convergence
Graceful extensibility [Woods, 2015]	world model	adaptive action	coherence feedback

These traditions operated under distinct motivations and vocabularies. Their convergence on the same decomposition is evidential not because the traditions are orthogonal (Peirce influenced systems theory via cybernetics) but because the structural conclusion persists across such different conceptual routes. We do not claim this convergence proves the TCC; Tier 1 convergence (ISS, Peirce, practopoiesis) constitutes qualitative evidence that the triadic structure tracks a genuine feature of the distributed coherence problem.

12 Open Questions

O1: Combinatorial uniqueness. Theorem 3.7 proves irreducibility of $(\mathcal{P}, \mathcal{A}, \Phi)$. Whether an alternative triadic decomposition could provide an equally valid architecture remains open. A proof requires a categorical notion of isomorphism between triadic systems (conjecture 12.1).

Conjecture 12.1 (Combinatorial minimality). *Up to relabelling, $(\mathcal{P}, \mathcal{A}, \Phi)$ is the unique minimal triadic decomposition of a drift-coherent, action-effective system.*

O2: Domain generalization. The TGT provides Gaussian-linear validation; the TPT extends to non-linear non-Gaussian environments. Remaining targets: distributed ledgers (failure modes precisely documented via consensus violations) and adaptive immunity (FM-3 cytokine storm mechanism well-characterized). The TCI in its current form detects both virtual and real drift but requires label supervision to distinguish them.

O3: Convergence rate of the Π_t tower. Theorem 5.1 establishes existence of a fixed point without bounding the tower depth required. Nikolić [2015] reports empirically $d \approx 4$ for mammalian complexity, providing the only empirical estimate.

Conjecture 12.2 (Convergence rate). *The tower depth $d \leq f(K(\mathcal{X}))$ for some computable function f , where $K(\mathcal{X})$ is the Kolmogorov complexity of the environment.*

O4: Coercive masking: formal proof and validation. Three sub-problems remain open. O4a: a tight bound $t_{\text{mask}} = \Theta(\alpha^{-1} \log(\alpha/\zeta))$ via a Lyapunov argument on $\delta_t = \mu^*(t) - \hat{\mu}_0$. O4b: calibration of the effort-corrected TCI^E beyond the current KuaiRand benchmark. The real-data wrapper confirms the qualitative effect, but the precision-recall frontier still depends on the choice of threshold quantile and effort penalty λ . O4c: a more accurate $\hat{\zeta}_t$ estimator would sharpen the precision-recall frontier further. O4d: external detectability: whether monotonically increasing E_t with bounded V_P is a sufficient external signature of coercive masking without knowledge of α .

O5: Geometry-dependence of the uncertainty bound. The $\zeta^{1/3}$ floor of proposition 8.9 relies on W_2 -metrized tracking and the transport-geometric lag interpretation. Extending the bound to f -divergences and kernel MMD is open.

13 Conclusion

Summary. The Triadic Coherence Condition establishes that any drift-coherent, action-effective information system must implement exactly three mutually irreducible functional roles: Potentiality ($\mathcal{P}$), Actuation ($\mathcal{A}$), and Convergence (Φ). The necessity proof derives this ternary structure externally, without presupposing it, via the Belief Update Lemma grounded in POMDP credit-assignment and KM-update semantics. The Borromean Cascade Stability theorem establishes sufficiency: under four explicit ISS hypotheses, the composite Lyapunov functional V_{total} is non-increasing and diverges upon removal of any single component. The Coherence-Delay Uncertainty Principle establishes a thermodynamic floor: $\mathcal{E}_{\min} = \frac{3}{2} C_K^{2/3} \zeta^{1/3}$, with optimal window $n^* = (C_K/\zeta)^{2/3}$, proving that perfect coherence is impossible under continuous drift.

Empirical grounding. The Triadic Gaussian Tracker confirms all three failure modes with correct bottleneck identification; the measured estimation slope -0.436 and U-curve slope 0.316 validate the expected finite-window bounds in the slow-drift regime. The Triadic Particle Tracker extends validation to non-Gaussian non-linear environments via Sequential Monte Carlo ablation and exhibits the coercive masking effect directly: under moderate coupling ($\alpha = 0.3$), the stale FM-1 controller reaches $\text{TCI} = 0.949 \pm 0.012$ while the effort-corrected score falls to $\text{TCI}^E = 0.697 \pm 0.046$. On ELEC2 and Bikes, the passive story is the same: ADWIN is the stronger generic detector on the residual streams, while the TCI family remains an interpretable structural diagnostic rather than a best-in-class passive change detector. On the KuaiRand logged benchmark, however, the order reverses in recall: after calibrating on the healthy random-exposure windows, TCI^E improves bubble detection from 45.8% to 69.2% and collapse detection from 54.8% to 70.0% on 367 eligible users, while ADWIN on raw watch ratio remains near zero. Page-Hinkley stays competitive on delay but is less sensitive than the effort-corrected score. The Coercive Masking Corollary formalises a qualitatively distinct failure mechanism: systems with high action-effectiveness can maintain an apparently healthy TCI by coercing the environment to conform to a stale world model, with catastrophic collapse once exogenous drift exceeds the coercive capacity.

Limitations. The necessity proof requires action-effectiveness; passive observers are genuinely dyadic. The Borromean stability proof is conditional on four ISS hypotheses whose verification for non-linear systems is application-dependent. The combinatorial uniqueness of $(\mathcal{P}, \mathcal{A}, \Phi)$ and the convergence rate of the Π_t tower remain open (section 12).

Take-home. Any drift-coherent information system requires all three of $\mathcal{P}$, $\mathcal{A}$, and Φ . Their joint absence is a structural failure waiting to be triggered. Their joint presence, under the ISS hypotheses, is sufficient for bounded-error coherence under bounded drift. Systems that centralise $\mathcal{P}$ (control the world model), monopolise $\mathcal{A}$ (control actuation), and suppress external Φ (block corrective feedback) operate in FM-1, FM-2, and FM-3 simultaneously. The architecture is the theorem.

Acknowledgements

The author thanks the communities of control theory, category theory, and semiotics for the foundational work on which this paper depends.

References

- Carlos E. Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50(2):510–530, 1985. doi: 10.2307/2274239. The AGM axioms for belief revision (static world). Contrasted with belief update (dynamic world) by [Katsuno and Mendelzon \[1992\]](#).
- Jason Altschuler, Jonathan Weed, and Philippe Rigollet. Near-linear time approximation algorithms for optimal transport via Sinkhorn iteration. In *Advances in Neural Information Processing Systems 30 (NeurIPS)*, pages 1964–1974, 2017.
- Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. *Gradient Flows in Metric Spaces and in the Space of Probability Measures*. Lectures in Mathematics ETH Zürich. Birkhäuser, Basel, 2005. doi: 10.1007/978-3-7643-8722-8.
- Gregory Bateson. *Steps to an Ecology of Mind*. Ballantine Books, New York, 1972.
- Albert Bifet and Ricard Gavaldà. Learning from time-changing data with adaptive windowing. In *Proceedings of the 2007 SIAM International Conference on Data Mining (SDM)*, pages 443–448, 2007. doi: 10.1137/1.9781611972771.42.
- Giacomo Bonanno. A Kripke-Lewis semantics for belief update and belief revision. *Artificial Intelligence*, 328:104065, 2024. doi: 10.1016/j.artint.2023.104065. Provides a unified Kripke-Lewis conditional-belief characterization of both KM-update and AGM-revision. Without the action argument o , the Lewis selection function cannot distinguish action-caused from drift-caused signals, grounding the necessity of o in the TCC’s Belief Update Lemma (ii).
- Robert W. Burch. *A Peircean Reduction Thesis: The Foundations of Topological Logic*. Texas Tech University Press, Lubbock, TX, 1991.

- Rafael Capurro and Birger Hjørland. The concept of information. *Annual Review of Information Science and Technology*, 37(1):343–411, 2003. doi: 10.1002/aris.1440370109.
- Roger C. Conant and W. Ross Ashby. Every good regulator of a system must be a model of that system. *International Journal of Systems Science*, 1(2):89–97, 1970. doi: 10.1080/00207727008920220.
- Peter R. Cromwell, Cary Malkiewich, and Dale Rolfsen. The Borromean rings. *The Mathematical Intelligencer*, 20(1):53–62, 1998. doi: 10.1007/BF03024401.
- Imre Csiszár and János Körner. *Information Theory: Coding Theorems for Discrete Memoryless Systems*. Academic Press, New York, 1981. Standard reference for the Data Processing Inequality: $I(X; f(Y)) \leq I(X; Y)$ for any function f .
- Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in Neural Information Processing Systems 26 (NeurIPS)*, pages 2292–2300, 2013.
- Sergey N. Dashkovskiy, Björn S. Rüffer, and Fabian R. Wirth. Small gain theorems for large scale systems and construction of ISS lyapunov functions. *SIAM Journal on Control and Optimization*, 48(6):4089–4118, 2010. doi: 10.1137/090746483.
- Malte Faber, Magda Sakr, and Annette Bauer. WATCH: Wasserstein change point detection for high-dimensional time series data. In *2022 IEEE International Conference on Big Data*, pages 3268–3277, 2022. doi: 10.1109/BigData55660.2022.10020525.
- Luciano Floridi. *The Philosophy of Information*. Oxford University Press, Oxford, 2011. doi: 10.1093/acprof:oso/9780199232383.001.0001.
- João Gama, Pedro Medas, Gladys Castillo, and Pedro Rodrigues. Learning with drift detection. In *Advances in Artificial Intelligence — SBIA 2004*, volume 3171 of *Lecture Notes in Computer Science*, pages 286–295. Springer, 2004. doi: 10.1007/978-3-540-28645-5_29.
- João Gama, Indrė Žliobaitė, Albert Bifet, Mykola Pechenizkiy, and Abdelhamid Bouchachia. A survey on concept drift adaptation. *ACM Computing Surveys*, 46(4):44:1–44:37, 2014. doi: 10.1145/2523813.
- Aude Genevay, Lénaïc Chizat, Francis Bach, Marco Cuturi, and Gabriel Peyré. Sample complexity of Sinkhorn divergences. In *Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 89 of *Proceedings of Machine Learning Research*, pages 1574–1583, 2019.
- Michael Harries. SPLICE-2 comparative evaluation: Electricity pricing. Technical report, University of New South Wales, 1999. Source of the ELEC2 concept-drift benchmark dataset, 45,312 hourly observations of electricity prices in New South Wales with documented distributional shift. Available at <https://www.openml.org/d/151>.
- Charles Hartshorne and Paul Weiss, editors. *Collected Papers of Charles Sanders Peirce*, volume 1–6. Harvard University Press, Cambridge, MA, 1931–1935.

- Joachim Hereth Correia and Reinhard Pöschel. Peircean algebraic logic and Peirce’s reduction thesis. *Semiotica*, 186(1–4):141–167, 2011. doi: 10.1515/semi.2011.062.
- Wolfgang Hofkirchner. *Emergent Information: A Unified Theory of Information Framework*, volume 3 of *World Scientific Series in Information Studies*. World Scientific, Singapore, 2013. doi: 10.1142/8347.
- Andrew H. Jazwinski. *Stochastic Processes and Filtering Theory*, volume 64 of *Mathematics in Science and Engineering*. Academic Press, New York, 1970.
- Leslie Pack Kaelbling, Michael L. Littman, and Anthony R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101(1–2):99–134, 1998. doi: 10.1016/S0004-3702(98)00023-X.
- Hirofumi Katsuno and Alberto O. Mendelzon. On the difference between updating a knowledge base and revising it. In Peter Gärdenfors, editor, *Belief Revision*, volume 29 of *Cambridge Tracts in Theoretical Computer Science*, pages 183–203. Cambridge University Press, 1992. Distinguishes *belief update* (world changes) from *belief revision* (new information about a static world); establishes the KM postulates for update. The TCC’s Φ_t is a KM-update operator.
- Sergiy Koshkin. Peirce’s reduction thesis and its proofs. *Transactions of the Charles S. Peirce Society*, 60(1):1–28, 2024. doi: 10.2979/trancharpeirsoc.60.1.01.
- Adrian Kosowski, Przemysław Uznański, Jan Chorowski, Zuzanna Stamirowska, and Michał Bartoszkiewicz. The dragon hatchling: The missing link between the Transformer and models of the brain. arXiv preprint arXiv:2509.26507, 2025. Pathway, Palo Alto. Code: <https://github.com/pathwaycom/bdh>.
- Gonzalo Mena and Jonathan Weed. Statistical bounds for entropic optimal transport: Sample complexity and the central limit theorem. *Advances in Neural Information Processing Systems*, 32, 2019. arXiv:1905.11882. Extends Genevay et al. (2019) to unbounded subgaussian measures with an exponentially improved bound.
- Andrii Mironchenko and Christophe Prieur. Input-to-state stability of infinite-dimensional systems: Recent results and open questions. *SIAM Review*, 62(3): 529–614, 2020. doi: 10.1137/19M1291248.
- Danko Nikolić. Practopoiesis: Or how life fosters a mind. *Journal of Theoretical Biology*, 373:40–61, 2015. doi: 10.1016/j.jtbi.2015.03.003.
- Giulio Ruffini. The algorithmic regulator. *Entropy*, 28(3):257, 2026. doi: 10.3390/e28030257. Published 26 February 2026. The “Good Algorithmic Regulator” (GAR) result cited here is Theorem 1 of this paper. The paper contains three theorems; the label “Algorithmic Regulator Theorem (ART)” used in the present work refers specifically to Theorem 1.
- Claude E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27(3):379–423, 1948. doi: 10.1002/j.1538-7305.1948.tb01338.x.
- Eduardo D. Sontag. Smooth stabilization implies coprime factorization. *IEEE Transactions on Automatic Control*, 34(4):435–443, 1989. doi: 10.1109/9.28018.

Eduardo D. Sontag. On the input-to-state stability property. *European Journal of Control*, 1(1):24–36, 1995. doi: 10.1016/S0947-3580(95)70005-X.

Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA, 1998.

David D. Woods. Four concepts for resilience and the implications for the future of resilience engineering. *Reliability Engineering & System Safety*, 141:5–9, 2015. doi: 10.1016/j.res.2015.03.018.

A Code Repository

All implementations — the Triadic Gaussian Tracker (TGT), the Triadic Particle Tracker (TPT), and the ELEC2 benchmark experiments — are publicly available at:

<https://github.com/jovemexausto/triadic-coherence>

The repository contains:

- `experiments/tcc/tgt.py` — canonical Triadic Gaussian Tracker experiments for ablation, sample-complexity, U-curve, and ELEC2 diagnostics;
- `experiments/tcc/tpt.py` — Triadic Particle Tracker, coercive-masking experiments, and effort-corrected TCI^E;
- `experiments/run_tgt.py` and `experiments/run_experiment.py` — reproducible entry points for all current figures and CSV artifacts;
- `experiments/artifacts/` — generated summary tables used to update the manuscript.

All experiments use `uv` for dependency management (`uv sync`) and are reproducible with fixed random seeds as specified in each script.

B KuaiRand Logged Benchmark Protocol

This protocol instantiates an ethically motivated logged benchmark on top of the KuaiRand sequential recommendation logs. The goal is to test whether an effort-corrected coherence score can surface exposure concentration and filter-bubble-like drift.

Data and signals

The primary state signal is `watch_ratio`. Auxiliary feedback signals include `click` and `like`; `dwel`, `follow`, and `share` are reserved for secondary validation when available. For user u at time t , let $P_{u,t}$ denote the inferred preference state, $A_{u,t}$ the exposure policy, and $\Phi_{u,t}$ the composite feedback signal.

Effort proxy

The effort proxy is defined as the divergence between the current exposure distribution and the healthy random-exposure baseline distribution:

$$E_{u,t} := D_{\text{KL}} \left(\pi_{u,t}^{\text{current}} \parallel \pi_{u,t}^{\text{healthy}} \right).$$

This choice is directional, decomposable, and directly interpretable as the cost of keeping the observed trajectory close to a healthy baseline. In settings where only aggregate exposure counts are available, total variation distance may be used as a fallback, but KL is the preferred default.

Protocol

The benchmark is executed in three phases.

- 1. Healthy.** Exposure follows the random baseline, diversity is unconstrained, and effort remains low.
- 2. Bubble.** Exposure becomes progressively concentrated, so that $E_{u,t}$ rises while standard coherence may remain artificially high.
- 3. Collapse.** The concentrated exposure regime stops tracking the healthy baseline, generic detectors begin to fire, and the effort-corrected index should lead the baselines.

Evaluation

Evaluation is carried out sequentially within each user session and then batched across users with strict temporal splits. Reported metrics are:

- detection rate of the bubble/collapse phases,
- median detection delay,
- false positives in the healthy phase,
- sensitivity to λ in TCI^E ,
- comparison against ADWIN, PageHinkley, KSWIN, and NoDrift.

Controls

Required controls include shuffled labels, random phase boundaries, no- coercion ablations, and removal of auxiliary feedback signals. These controls ensure that the benchmark does not reduce to a hindsight-tuned logged-bandit exercise.