The Triadic Coherence Condition

A Predictive Framework for Coherence Under Continuous Non-Stationary Streams

Parreira Figueiredo

Independent Researcher

figueiredo@volverse.com.br

March 2026

Abstract

How many functional components does an information system require to maintain coherent operation under continuously drifting environmental conditions, and why exactly that number? We answer this question by establishing the **Triadic Coherence Condition** (TCC): any system maintaining bounded-error fidelity to a drifting environment (*drift-coherence*) must implement exactly three mutually irreducible roles. These roles are *Potentiality* ($\mathcal{P}$, a probability measure over world states metrized by the L^2-Wasserstein distance), *Actuation* ($\mathcal{A}_t$, a δ-reproducible output function), and *Convergence* (Φ_t, a coherence-restoring feedback map), and each is indispensable.

The necessity proof derives the ternary structure externally, without presupposing it. The Belief Update Lemma shows that any coherence-restoring update function must depend on the current belief state, the most recent action, and the incoming environmental signal: removing any single argument destroys drift-coherence. Three independent frameworks ground this claim simultaneously: the POMDP credit-assignment problem, the Katsuno-Mendelzon belief-update formalism distinguishing world-change from information-revision, and Input-to-State Stability theory. Sufficiency follows from a *Conditional Borromean Cascade Stability* theorem: under four explicit ISS hypotheses, the composite entropic Lyapunov functional $V_P + V_A + V_\Phi$ is non-increasing, and removal of any single component causes divergence. Two further results complete the theory: a self-referential fixed-point theorem in quasi-Borel spaces, and a correspondence establishing that each TCC failure mode collapses the algorithmic regulation gap Δ of Ruffini's Algorithmic Regulator Theorem.

We introduce the **Triadic Coherence Index** $\mathrm{TCI} \in (0, 1]$, an estimable scalar diagnostic computed via the debiased Sinkhorn divergence, an entropic regularisation of W_2, with a proven $O(n^{-1/2})$ convergence rate. Minimising the total estimation error over window size yields the **Coherence-Delay Uncertainty Principle**: the minimum achievable tracking error satisfies $\mathcal{E}_{\min} = \frac{3}{2} C_K^{2/3} \zeta^{1/3}$, with optimal window $n^* = (C_K/\zeta)^{2/3}$, establishing that perfect coherence is thermodynamically impossible under continuous drift and providing a self-tuning rule for streaming TCI monitors. The **Triadic Gaussian Tracker** provides self-contained empirical validation: ablation of each component produces the predicted failure signature with correct bottleneck identification; the $O(n^{-1/2})$ rate is confirmed (log-log slope -0.434); early-warning experiments on the ELEC2 benchmark yield 51 validated leads with median 1,019-step anticipation of regime shifts; and dynamic n_t^* adaptation detects 75 of 79 regime shifts. The **Coercive Masking Corollary** formalises a qualitatively distinct failure mechanism: highly action-effective systems can sustain an apparently healthy TCI by coercing the environment to conform to a stale world model, with catastrophic collapse once exogenous drift exceeds the system's coercive capacity.

Keywords: triadic coherence condition; drift-coherence; information systems; streaming data; ISS stability; Wasserstein distance; Sinkhorn divergence; belief update; concept drift; Triadic Coherence Index; algorithmic regulation; coercive masking.
MSC 2020: 93D05, 93D30, 18B20, 68P99.

Contents

1 Introduction

"One cannot step into the same river twice—for other waters are always flowing."
— HERACLITUS, Fragment 12

"The river in which Heraclitus stepped is not the same river. Neither is Heraclitus."
— JORGE LUIS BORGES, *This Craft of Verse*

A question has persisted across twenty-six centuries of philosophy without receiving a satisfactory formal answer. It is fundamental enough that philosophers frequently pass through it on the way to other problems without pausing to examine it directly. The question is:

How does something remain what it is while it changes?

Heraclitus located the problem in the river. Parmenides resolved it by opposing movement: being is; non-being is not; change is illusion. Western philosophy oscillates between these poles without resolution. Aristotle distinguished *act* and *potency*, and left the mechanism of passage underspecified. Leibniz's monad unfolds its own being from an internal law, but eliminates real change in the process. Hegel's *Aufhebung* makes contradiction the motor of reality, but requires that the process reach a necessary conclusion, which no real-time system can guarantee.

The question these thinkers addressed, yet never formulated with sufficient precision, is the following:

> *Given a system that maintains a representation of a continuously changing reality, receives events updating this representation, and exists as multiple distributed copies required to remain mutually coherent: what are the necessary and sufficient conditions for this system to maintain fidelity to the reality it represents, without requiring a central authority to coordinate it?*

This is simultaneously a question of philosophy, systems engineering, neuroscience, and theology. That traditions operating under very different motivations and vocabularies arrived at structurally convergent answers is the central observation this paper formalises.

Motivation from high-consequence failures. Modern information systems operate under continuous streams of data whose statistical properties drift over time. In 2021, Zillow's automated home-purchasing system lost approximately \$500M after its price model ($\mathcal{P}$) failed to update in response to a market-regime shift, while its purchasing engine ($\mathcal{A}$) continued executing stale predictions. In 2012, Knight Capital lost approximately \$440M in forty-five minutes when a deployment error caused its order-execution engine ($\mathcal{A}$) to enter a non-deterministic state while its market models ($\mathcal{P}$) remained intact. In social-media systems, recommendation algorithms that continuously update their world models ($\mathcal{P}$) and act on them ($\mathcal{A}$) but apply no convergence operator (Φ) produce information cascades and filter-bubble phenomena. Each failure corresponds to the degradation of exactly one functional component while the others remain operative, a pattern that motivates the present framework.

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 3.8) 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 5.3). A Restricted-Class Lemma (lemma 5.5) grounds the engineering approximation with an explicit algebraic-closure proof.

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 TCI $\in (0, 1]$ via the debiased Sinkhorn divergence S_ε, with proven $O(n^{-1/2})$ estimation rate (proposition 8.6) and the **Coherence-Delay Uncertainty Principle** (proposition 8.10): 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. **Triadic Gaussian Tracker** (section 9): Controlled empirical validation via Borromean ablation (Exp. A, FM-1 tail $V_P = 31.78$ vs full 7.66), Sinkhorn rate verification (Exp. B, slope -0.434), ELEC2 early-warning (Exp. C, 51 leads, median 1,019 steps), and dynamic n_t^* optimization (Exp. D, 75 of 79 regime shifts detected). The Coercive Masking phenomenon is identified and formalised (corollary 6.6 and definition 6.8).

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 (section 9) provides controlled empirical validation in the Gaussian-linear case; generalisation to non-Gaussian and high-dimensional settings remains 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 $W_2(\mathcal{P}(S,t),\, \mathcal{P}^*(S,t)) \le \varepsilon$, 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 $\varepsilon \to 0$ 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

$$\mathcal{A}_t \colon \mathcal{I} \to \mathcal{O},$$

where $\mathcal{O}$ is the output space. $\mathcal{A}_t$ is δ-*reproducible* if for any $p \in \mathcal{I}$ and any two executions ω, ω', $d_{\mathcal{O}}(\mathcal{A}_t(p;\omega), \mathcal{A}_t(p;\omega')) \le \delta$. A deterministic actuation is 0-reproducible.

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

$$\Phi_t \colon \mathcal{I} \times \mathcal{O} \times S \to \mathcal{I},$$

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 $t \ge 0$,

$$W_2(\mathcal{P}(S,t), \mathcal{P}^*(S,t)) \;\le\; \beta(W_2(\mathcal{P}(S,0), \mathcal{P}^*(S,0))\,,\, t) \;+\; \gamma\Big(\sup_{s \le t} \|u(s)\|\Big)\,,$$

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

2.2 ISS Definitions

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

$$\|x(t)\| \;\le\; \beta(\|x_0\|\,, t) \;+\; \gamma\big(\mathrm{ess\ sup}_{s \le t} \|u(s)\|\big)\,.$$

Equivalently [Sontag, 1989], there exists a smooth ISS-Lyapunov function V with $\alpha_1(\|x\|) \le V(x) \le \alpha_2(\|x\|)$ and $\|x\| \ge \chi(\|u\|) \Rightarrow \nabla V \cdot f(x,u) \le -\alpha_3(\|x\|)$ 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 $\operatorname{ess\,sup}_{s \leq t} \|u(s)\| \leq M < \infty$:

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

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:

$$D_{\mathrm{TV}}\Big(P(\cdot \mid s_t = s,\, o) \,\Big\|\, P(\cdot \mid s_t = s,\, o')\Big) \;>\; 0.$$

A system is *passive* if $D_{\mathrm{TV}} = 0$ 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 $H_t = (s_1, o_1, \ldots, s_t, o_t)$ [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. □

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.

Remark 3.6 (Non-Markovian dynamics). Lemma 3.4 invokes the POMDP Markov property: the current belief state $p = \mathcal{P}(S, t)$ is the minimal sufficient statistic for the history H_t. Under *non-Markovian* environmental dynamics, where $P(s_{t+1} \mid s_t, o)$ depends on the full trajectory rather than the current state, this sufficient-statistic property fails. The standard resolution is history augmentation: extend the state space to $S^{(k)} = S \times S \times \cdots \times S$ (k steps), making the k-step history the new "current state." The augmented system is Markovian by construction, and lemma 3.4 applies to the augmented belief state. The triadic structure is preserved; the memory depth k is a system-design parameter that grows with the correlation horizon of the environment.

Corollary 3.7 (Ternary Irreducibility of the Update Function)**.** *The drift-coherent update function $U : \mathcal{I} \times \mathcal{O} \times S \to \mathcal{I}$ 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

Theorem 3.8 (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 3.7, drift-coherence requires $U : \mathcal{I} \times \mathcal{O} \times S \to \mathcal{I}$. 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 $(p, o, s_{t+1}) \to \mathcal{I}$. $\square$

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 3.9 (Scope of theorem 3.8)**.** *Theorem 3.8 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 $\Phi_t \colon \mathcal{I} \times \mathcal{O} \times S \to \mathcal{I}$ 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 $V_P \colon \mathcal{I} \to \mathbb{R}_{\geq 0}$ 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: $\operatorname{ess\,sup}_{s\geq 0} \|u(s)\| \leq M < \infty$, 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

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

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* $V_{\text{total}}(t) \to 0$ *as* $t \to \infty$ *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 $(\Phi) \to (\mathcal{P})$ is ISS by (H-ISS-P) and (H-ISS-Φ); the extended cascade $(\Phi \to \mathcal{P}) \to (\mathcal{A})$ 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. □

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 $\mu_t = \mu_0 + \int_0^t \eta(s)\,ds$ for bounded drift $\|\eta\|_\infty < M$. 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 via Lawvere

5.1 The Recursive TCC

The base TCC maintains coherence of $\mathcal{S}$ with respect to the external environment $\mathcal{X}(t)$ when the environmental drift remains within the parameter envelope of the three operators: drift rate ζ, actuation class, and feedback bandwidth are all stable. The base triad is sufficient for this regime (theorem 4.2).

A different and harder problem arises under *meta-drift*: when the parameters governing the drift themselves change. An environment undergoing a phase transition, shifting its statistical manifold rather than merely drifting along it, renders the base operators' fixed parameters inadequate. In this case, the base triad will fail: by the Drift-Decay Lemma (lemma 6.2), $M(W : R) \to 0$ as the frozen parameters diverge from the new regime.

The resolution applies the TCC recursively one organizational level above the base triad: a meta-actuation operator Π_t reconfigures the parameters of $(\mathcal{P}, \mathcal{A}, \Phi)$ in response to detected regime shifts. This creates a *vertical* hierarchy rather than a flat tetradic extension: Π_t applies the same triadic logic to the parameter space, not to the base state space. The question of whether Π_t itself needs a meta-Π_t leads to theorem 5.3.

A natural question arises: what maintains the coherence of the TCC itself when environmental conditions change such that the parameters of $\mathcal{P}$, $\mathcal{A}$, Φ need adjustment?

The answer is the TCC applied to itself one organizational level above: a meta-actuation operator Π_t reconfiguring the parameters of $(\mathcal{P}, \mathcal{A}, \Phi)$.

Definition 5.1 (Meta-actuation and Recursive TCC)**.** The *meta-actuation operator* is a morphism

$$\Pi_t : \mathcal{I} \longrightarrow \mathcal{I}^{\mathcal{I}},$$

assigning to each information state $p \in \mathcal{I}$ a transformation of the information state space. The *recursive information state* is the quadruple $I(S,t) = (\mathcal{P}, \mathcal{A}, \Phi, \Pi_t)$. Π_t adjusts the parameters and transformation classes of $(\mathcal{P}, \mathcal{A}, \Phi)$ in response to environmental signals, operating on the space of their configurations rather than their current values.

The three instances of Π_t correspond to three environmental observables:

- $\Pi^{(I)}$ (Bandwidth Adaptation): observes estimated shift rate $\hat{\zeta}(t)$; adjusts update bandwidth of $\mathcal{P}$.
- $\Pi^{(II)}$ (Transformation Class Upgrade): observes spectral gap $\hat{\Delta}(t)$; upgrades $\mathcal{A}_t$'s transformation class when strong principal structure is detected.
- $\Pi^{(III)}$ (Locality Radius Recalibration): observes empirical perturbation magnitudes; adjusts containment radii per organizational level.

These three instances operate on disjoint primary control variables and observe exclusively environmental signals, not the outputs of each other. This acyclic structure is the condition required by the ISS Cascade Theorem, so the recursive TCC inherits Borromean stability from the base TCC.

Meta-drift and the scope of Π_t. The base TCC $(\mathcal{P}, \mathcal{A}, \Phi)$ is sufficient for maintaining drift-coherence when environmental drift stays within the parameter envelope of the three operators, with drift rate ζ, actuation class, and feedback bandwidth all stable. Π_t becomes structurally mandatory when *meta-drift* occurs: when the parameters governing the drift themselves change. An environment undergoing a regime shift or phase transition alters the statistical manifold rather than merely drifting along it. In this case, the base operators' parameters become inadequate, and by the Drift-Decay Lemma (lemma 6.2), $M(W : R) \to 0$. Π_t resolves this by reconfiguring $(\mathcal{P}, \mathcal{A}, \Phi)$'s parameters at a higher organizational level, applying the triadic logic recursively to the parameter space and creating a vertical hierarchy rather than a flat tetradic extension (see also section 11, Argument 2).

5.2 Lawvere Fixed Point for Self-Referential TCC

Remark 5.2 (Categorical setting). The standard category **Meas** of measurable spaces and measurable maps is *not* cartesian closed [see Aumann, 1961]: the evaluation map that applies a measurable function to a point is not measurable in general, so the internal hom $\mathcal{I}^{\mathcal{I}}$ is not well-defined there. The correct categorical home for theorem 5.3 is the category **QBS** of *quasi-Borel spaces* [Heunen et al., 2017], which is cartesian closed, supports a probability monad extending the Giry monad on standard Borel spaces, and admits well-typed higher-order functions on probability distributions. All information state spaces considered in this paper are standard Borel (supported on Polish spaces), so the canonical embedding $\mathbf{Meas} \hookrightarrow \mathbf{QBS}$ preserves all relevant structure, and the meta-actuation morphism $\Pi_t \colon \mathcal{I} \to \mathcal{I}^{\mathcal{I}}$ is well-defined in **QBS**.

Theorem 5.3 (Lawvere Fixed Point for TCC)**.** *Let $\mathcal{C} = \mathbf{QBS}$ be the cartesian closed category of quasi-Borel spaces (remark 5.2), with $\mathcal{I}$ the information state space viewed as a quasi-Borel space. Suppose $\Pi_t \colon \mathcal{I} \to \mathcal{I}^{\mathcal{I}}$ is* point-surjective*: for every morphism $g \colon \mathcal{I} \to \mathcal{I}$ in* **QBS** *there exists a state $p \in \mathcal{I}$ such that $\Pi_t(p) = g$.*

Then every endomorphism $f \colon \mathcal{I} \to \mathcal{I}$ has a fixed point $p^* \in \mathcal{I}$ with $f(p^*) = p^*$. In particular, the TCC *has a fixed information state $\mathcal{I}^*$ such that* $\mathrm{TCC}(\mathcal{I}^*) = \mathcal{I}^*$.

Proof. We apply Lawvere's fixed-point theorem directly [Lawvere, 1969]. In a cartesian closed category, given point-surjective $\Pi_t \colon \mathcal{I} \to \mathcal{I}^{\mathcal{I}}$, construct the diagonal morphism $\delta \colon \mathcal{I} \to \mathcal{I}$ by

$$\delta = \mathrm{ev} \circ (\Pi_t \times \mathrm{id}_{\mathcal{I}}) \circ \Delta,$$

where $\Delta \colon \mathcal{I} \to \mathcal{I} \times \mathcal{I}$ is the diagonal map $p \mapsto (p, p)$ and $\mathrm{ev} \colon \mathcal{I}^{\mathcal{I}} \times \mathcal{I} \to \mathcal{I}$ is evaluation.

For any $f \colon \mathcal{I} \to \mathcal{I}$, set $g = f \circ \delta$. By point-surjectivity, $\exists\, p_0$ with $\Pi_t(p_0) = g$. Then

$$\delta(p_0) = \mathrm{ev}(\Pi_t(p_0), p_0) = g(p_0) = (f \circ \delta)(p_0) = f(\delta(p_0)).$$

Setting $p^* = \delta(p_0)$ yields $f(p^*) = p^*$. □

Remark 5.4 (Point-Surjectivity: Idealization and Practical Approximation). The point-surjectivity assumption requires Π_t to produce *every* endomorphism $g \colon \mathcal{I} \to \mathcal{I}$ as an output, an idealization that no finite engineered system satisfies. In practice Π_t operates over a restricted class $\mathcal{G} \subset \mathcal{F}(\mathcal{I}, \mathcal{I})$ of endomorphisms within its operational regime. The following lemma establishes that local point-surjectivity over $\mathcal{G}$ suffices.

Lemma 5.5 (Restricted-Class Lawvere Fixed Point)**.** *Let $\mathcal{G} \subset \mathrm{End}(\mathcal{I})$ be the restricted class of endomorphisms implementable by Π_t. Suppose:*

(a) (Algebraic closure) $\mathcal{G}$ *is closed under composition with the diagonal morphism* δ*: for every* $f \in \mathcal{G}$*, the composition* $f \circ \delta \in \mathcal{G}$*.*

(b) (Local point-surjectivity) Π_t *is point-surjective onto* $\mathcal{G}$*: for every* $g \in \mathcal{G}$*,* $\exists p_0 \in \mathcal{I}$ *such that* $\Pi_t(p_0) = g$*.*

Then every $f \in \mathcal{G}$ *has a fixed point* $p^* \in \mathcal{I}$ *with* $f(p^*) = p^*$*.*

Proof. For $f \in \mathcal{G}$, set $g = f \circ \delta$. By (a), $g \in \mathcal{G}$. By (b), $\exists p_0$ with $\Pi_t(p_0) = g$. Evaluating the diagonal: $\delta(p_0) = \mathrm{ev}(\Pi_t(p_0), p_0) = g(p_0) = f(\delta(p_0))$. Setting $p^* = \delta(p_0)$ gives $f(p^*) = p^*$. □

In the TGT (section 9), $\mathcal{G}$ is the class of Kalman parameter updates: gain recalibration, process-noise adjustment, and observation-noise scaling. These three reconfiguration types are closed under δ by construction (each maps a Kalman state to another valid Kalman state); local point-surjectivity holds because any target gain configuration is reachable from some prior information state. The restricted fixed point is the steady-state Kalman gain $K^* = \sigma_w^2/(\sigma_w^2 + \sigma_{\mathrm{obs}}^2)$. The base theorem 5.3 requires full point-surjectivity; lemma 5.5 establishes that the engineering approximation retains the fixed-point guarantee within $\mathcal{G}$.

Remark 5.6 (Algebraic closure for non-Kalman endomorphism classes). For general (non-Gaussian, non-linear) systems, condition (a) of lemma 5.5 requires $\mathcal{G}$ to be closed under composition with the diagonal δ. The key sufficiency conditions are: (i) $\mathcal{G}$ must include all compositions of the form $f \circ \delta$ for $f \in \mathcal{G}$, i.e., it must be δ*-stable*; (ii) $\mathcal{G}$ must be closed under the evaluation map $\mathrm{ev}(\cdot, p)$ for each $p \in \mathcal{I}$, so that the Lawvere diagonal construction produces an element of $\mathcal{G}$. In practice, these conditions are verified by checking that the reconfiguration class is defined by algebraic constraints on the parameters of $\mathcal{P}$, $\mathcal{A}$, Φ (e.g., positivity of a covariance matrix, Lipschitz bounds on a neural ODE vector field, or bandwidth constraints on a kernel filter) and that these constraints are preserved under composition. For particle-filter implementations of Φ_t, for instance, $\mathcal{G}$ may be taken as the class of bandwidth-resampling operations, which is trivially δ-stable since any bandwidth update maps a valid particle filter to another valid particle filter.

Conjecture 5.7 (Convergence rate of self-application tower). *Under additional regularity conditions on* Π_t *(e.g., contractivity in* $\mathcal{I}^{\mathcal{I}}$ *under a suitable metric), the self-application tower converges to* p^* *in a finite number of steps bounded by a function of the "complexity" of the environment. Nikolić [2015] provides empirical evidence in biological systems that the depth is approximately four levels. (See conjecture 12.2 in section 12 for the precise formulation.)*

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)*,*

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

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?

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* $K(O_{W,R} \mid \mathcal{P}(S, t))$ *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:*

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

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 $M(W : R) \leq K(R) - K(R \mid W)$. Under the assumption, $K(R \mid W) = K(O_{W,R} \mid \mathcal{P})$ grows with $W_2(\mathcal{P}, \mathcal{P}^*)$, so $M(W : R)$ decreases at least as fast as $C_0/(1 + W_2(\mathcal{P}, \mathcal{P}^*)) \leq C_0/(1 + \zeta(t - t_0))$. □

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 $K(O_{W,R} \mid \mathcal{P})$ 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: $\Delta \to 0$ 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.

[1] Ruffini'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.

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 $M(W:R)\to 0$, hence $\Delta\to 0$. 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; $K(O_{W,R})\to K(O_{W,\emptyset})$ as $t\to\infty$, so $\Delta\to 0$.*

Proof. For (i): by lemma 6.2, static $\mathcal{P}$ under non-mean-reverting drift implies $M(W:R)\to 0$, hence $\Delta\to 0$. 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. □

Remark 6.4 (Complementarity)*.* Ruffini's ART answers: *Given a regulator that works, how much must it know about the world?* The answer is $M(W:R)>0$, proportional to Δ. The TCC answers: *What structure must a regulator have to keep working under continuous change?* The answer is $(\mathcal{P},\mathcal{A},\Phi)$, triadic and mutually constitutive. Ruffini's Δ is the ART's observable diagnostic; the TCC specifies the architectural preconditions under which Δ remains observable and sustainable. The two results are logically independent and empirically complementary.

A further connection arises from Theorem 3 of Ruffini [2026], which establishes an "as-if planner" result: on the realized trajectory, a regulator satisfying the GAR condition behaves *as if* it minimized the conditional description length $K(O_{W,R}\mid R)$. In TCC terms, this is precisely the signature of a δ-reproducible $\mathcal{A}_t$: because $\mathcal{A}_t$ is deterministic given $\mathcal{P}(S,t)$, the output trajectory $O_{W,R}$ is compressible conditioned on R's internal state, and the as-if objective is well-defined. Non-deterministic $\mathcal{A}_t$ (FM-2) breaks this: $K(O_{W,R}\mid R)$ becomes substrate-dependent, the as-if objective is ill-posed, and Ruffini's Theorem 3 loses applicability.

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 Borromean ablation experiment (section 9.2) reveals a fourth, qualitatively distinct failure mechanism, one the standard ART analysis conflates with FM-1 rather than treating separately.

Definition 6.5 (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 6.6 (Coercive Masking)**.** *In the coercive regime (definition 6.5), 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 6.5, 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 6.7 (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 6.8 (Effort-Corrected TCI)**.** The *effort signal* at time t is

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

the magnitude of change in the system's actuation. In a well-adapted system, effort is low because the world and the model are aligned. In a coercive system, effort grows monotonically as the exogenous drift increases the gap the actuation must bridge.

The *effort-corrected TCI* is

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

where $E_0 = E_1$ is the baseline effort and $\lambda > 0$ is a sensitivity parameter. TCI^E penalizes increasing actuation effort, correctly diagnosing FM-1 even in the coercive regime where $\sigma_P(t) \approx 1$.

Remark 6.9 (Calibrating λ). The parameter λ governs the strength of the effort penalty. A principled default derives from a stable reference period: let E_{ref} denote the median effort during a known-stationary segment (no concept drift), and set $\lambda = -\ln(1-\delta) \cdot E_0/E_{\text{ref}}$, where $\delta \in (0,1)$ is the target penalty fraction (e.g., $\delta = 0.1$ means a 10% reduction in σ_P at reference-level effort). Alternatively, a cross-validation sweep over λ on labeled training windows with known failure events provides a data-driven calibration. In either case, λ is interpretable: it specifies the effort threshold at which the practitioner regards increasing actuation as a warning rather than normal adaptation. Empirical validation of the effort-corrected TCI across coupling values $\alpha > 0$ remains open (section 12, O4b).

Remark 6.10 (Non-linear and time-varying coupling). Corollary 6.6 assumes a constant linear coupling coefficient α. Real-world action-effective systems often exhibit non-linear or time-varying coupling: $\alpha = \alpha(s_t, o, t)$. In this regime the fixed-point analysis of the gap process $\delta_t = \mu^*(t) - \hat{\mu}_0$ loses direct applicability, and the coercive horizon t_{mask} may fluctuate. The effort signal E_t remains a valid diagnostic in this setting: monotonically increasing effort under bounded V_P is a sufficient signature of coercive dynamics regardless of the functional form of α. Formalizing the coercive horizon under non-linear coupling is the subject of O4a (section 12).

The empirical evidence for coercive masking comes from the ablation experiment (section 9.2). In the passive environment ($\alpha = 0$), FM-1 produces $\bar{V}_P = 31.78$, four times the full-TCC value. In an active environment ($\alpha = 0.2$), the same frozen model produces $\bar{V}_P < 1.0$: the stale actuation holds μ^* near the frozen estimate, masking the failure entirely in the standard TCI. The effort signal E_t grows monotonically in the active-FM-1 case, correctly signaling degradation where σ_P remains artificially elevated.

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 $W_2(\mathcal{P}(S,t), \mathcal{P}^*(S,t)) \to \infty$ 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 $W_2 \to \infty$ by definition. Lipschitz continuity of $\mathcal{A}_t$ (standard regularity) then propagates the divergence to actuation error. □

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 $W_2(\mathcal{P}, \mathcal{P}^*) \to \infty$ 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. □

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 ≈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 $\Phi_t \equiv \mathrm{id}$ (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 $\lim_{t\to\infty} s_{j^*}/\sum_j s_j = 1$ 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. □

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 $s_{j^*}(t)/\sum_j s_j(t) \to 1$ as $t \to \infty$ 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 $\nu_j^* \le \rho^*(1+\epsilon_\gamma)$.

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:

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

$$\sigma_A(t) \coloneqq \exp(-H(\mathcal{A}_t \mid \mathcal{P}(S,t)))\,, \tag{2}$$

$$\sigma_\Phi(t) \coloneqq \frac{1}{1+V_\Phi(t)}. \tag{3}$$

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

$$\mathrm{TCI}(I(S,t),t) \;\coloneqq\; \min(\sigma_P(t),\, \sigma_A(t),\, \sigma_\Phi(t)) \;\in\; (0,1].$$

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 (TCI and Capacity for Maneuver). Woods [2015] defines the *Capacity for Maneuver* (CfM) as the remaining range available to a system to handle upcoming demands, with saturation (CfM exhaustion) corresponding to loss of control. The convergence score $\sigma_\Phi(t)$ operationalizes CfM within the TCC: it measures how far the convergence subsystem is from its failure boundary, with $\sigma_\Phi \to 0$ corresponding to CfM exhaustion and FM-3 onset (proposition 7.6). The TCC further decomposes CfM into three orthogonal dimensions (σ_P, σ_A, σ_Φ), each corresponding to a distinct failure mode, whereas Woods' framework treats CfM as a scalar aggregate.

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 8.4 (Empirical Potentiality Estimator). Let $\{s_{t-n+1}, \ldots, s_t\}$ 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 8.5 (Estimable TCI). The *debiased Sinkhorn divergence* between measures μ, ν is

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

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 8.6.

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

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

The *estimable TCI* is

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

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 8.6 (Estimation Error Bound for $\hat{\sigma}_P$). *Let $\mathcal{P}^*(S,t)$ be K-subgaussian with drift rate $\zeta = \operatorname{ess\,sup}_s \|\dot{\mathcal{P}}^*(s)\|_{W_2}$. Let $\hat{\mathcal{P}}^*$ be the sliding-window empirical estimator with window n (definition 8.4), and let τ be the lag between the center of the window and t. Then under the entropic regularization of remark 2.3:*

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

where C_K depends only on the subgaussian constant K.

Proof. By the triangle inequality for S_ε:

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

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 $\leq 1/(1 + V_P)^2$. □

proposition 8.6 makes $\widehat{\mathrm{TCI}}$ 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.

Remark 8.7 (TCI and information-theoretic divergence)*.* The choice of W_2 as the metric on $\mathcal{I}$ connects σ_P to information-theoretic quantities via Talagrand's transportation inequality. Under a log-Sobolev inequality with constant $\alpha > 0$ (satisfied by all strongly log-concave distributions), $W_2(\mathcal{P}, \mathcal{P}^*)^2 \leq \frac{2}{\alpha} D_{\mathrm{KL}}(\mathcal{P} \| \mathcal{P}^*)$, which gives the information-theoretic bound:

$$\sigma_P(t) \;\geq\; \frac{1}{1 + \frac{1}{\alpha} D_{\mathrm{KL}}\big(\mathcal{P}(S,t) \,\big\|\, \mathcal{P}^*(S,t)\big)}.$$

High σ_P thus implies low KL divergence from the oracle, and consequently low information loss between the system's internal model and the environment. Via the ART correspondence (proposition 6.3), this provides a direct lower bound on the algorithmic mutual information $M(W\!:\!R)$: drift-coherent TCC-compliant systems maintain $M(W\!:\!R) > 0$ precisely because $\sigma_P(t)$ is bounded away from zero. Talagrand-type inequalities for the Sinkhorn divergence itself have been established [Fan and Quek, 2022, Theorems 3–4], extending the classical bound to the entropically regularized setting.

Remark 8.8 (Virtual vs. real drift)*.* The concept-drift literature distinguishes *virtual drift* (the marginal $P(X)$ shifts but the conditional $P(Y|X)$ is stable) from *real drift* ($P(Y|X)$ changes, altering the optimal decision boundary). The σ_P component tracks $W_2(\mathcal{P}, \mathcal{P}^*)$ where $\mathcal{P}^*$ is the full joint distribution of the environment. Virtual drift shifts the feature geometry but leaves the label structure intact; real drift changes the label structure. The TCI detects both, because any shift in the joint $\mathcal{P}^* = P(X, Y)$ increases $W_2(\mathcal{P}, \mathcal{P}^*)$ and thereby decreases σ_P. Distinguishing the two types requires access to label supervision signals, which the TCI in its present form leaves unresolved without additional label supervision. Adding a label-conditioned component to σ_P is a natural extension for supervised streaming settings.

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

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

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

Proposition 8.10 (Coherence-Delay Uncertainty Principle)**.** *Under definition 8.9, 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

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

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: $\mathcal{E}(n^*) = C_K(C_K/\zeta)^{-1/3} + \frac{1}{2}\zeta(C_K/\zeta)^{2/3} = C_K^{2/3}\zeta^{1/3} + \frac{1}{2}C_K^{2/3}\zeta^{1/3} = \frac{3}{2}C_K^{2/3}\zeta^{1/3}$. □

Remark 8.11 (Interpretation: a thermodynamic floor)*.* Proposition 8.10 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 8.12 (Geometry-dependence of the bound)*.* Proposition 8.10 depends on two structural assumptions: (i) the variance term scales as $C_K n^{-1/2}$, which holds for the debiased Sinkhorn divergence under K-subgaussian $\mathcal{P}^*$ [Genevay et al., 2019, Theorem 2]; and (ii) the lag-bias term scales as $\frac{1}{2}\zeta n$, which holds when $\mathcal{P}^*$ drifts at a Lipschitz rate ζ in the W_2 metric.

For f-divergences (KL, total variation, χ^2), the variance term generally achieves the same $O(n^{-1/2})$ rate in one dimension [Csiszár and Körner, 1981], but lacks the transport-geometric interpretation that makes the lag-bias term linear in $n\zeta$. Under distributional shifts that translate support rather than reshaping density, f-divergences may diverge entirely (if supports become disjoint), making the variance-lag decomposition inapplicable. The $\zeta^{1/3}$ scaling of $\mathcal{E}_{\min}$ is therefore specific to W_2-metrized tracking and the lag-bias interpretation of the window center. Generalizing the bound to other metrics on $\mathcal{I}$ requires establishing the analogous two-term decomposition in the chosen geometry, a direction we leave open.

Remark 8.13 (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 $\Delta\hat{V}_\Phi \le 0$ 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 $\hat{\sigma}_\Phi = 1/(1+\hat{V}_\Phi) \le \sigma_\Phi$, usable in $\widehat{\mathrm{TCI}}$ 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 3.8, 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 8.14 (TCI Characterizes TCC-Compliance)**.** *Let $I(S,t)$ be an TCC-compliant system under hypotheses (H-ISS-P/A/Φ/Env). Then:*

(i) $\mathrm{TCI}(I(S,t),t) = 1$ *if and only if* $V_P(t) = 0$, $H(\mathcal{A}_t \mid \mathcal{P}) = 0$, *and* $V_\Phi(t) = 0$ *simultaneously, i.e., the system is at the equilibrium* $I^* = (\mathcal{P}^*, \mathcal{A}^*, \Phi^*)$.

(ii) $\mathrm{TCI}(I(S,t),t) \to 0$ *if and only if at least one of the three failure modes is active:* $V_P(t) \to \infty$ (FM-1), $H(\mathcal{A}_t \mid \mathcal{P}) \to \infty$ (FM-2), *or* $V_\Phi(t) \to \infty$ (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: $\sigma_P = 1 \Leftrightarrow V_P = 0$, $\sigma_A = 1 \Leftrightarrow H(\mathcal{A}_t \mid \mathcal{P}) = 0$, $\sigma_\Phi = 1 \Leftrightarrow V_\Phi = 0$. The min equals 1 iff all three equal 1 simultaneously.

(ii) As $V_P(t) \to \infty$, $\sigma_P(t) \to 0$, so TCI $\to 0$; similarly for $V_\Phi(t) \to \infty$. As $H(\mathcal{A}_t \mid \mathcal{P}) \to \infty$, $\sigma_A(t) = e^{-H} \to 0$, so TCI $\to 0$. The converse holds because

TCI $\to 0$ requires at least one $\sigma_i \to 0$, which by the above requires the corresponding functional to diverge.

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

Remark 8.15 (TCI as Pre-Failure Early Warning). The diagnostic value of the TCI is precisely its *lead time* over observable failures. By theorem 4.2, V_{total} is non-increasing under full TCC-compliance; any monotone increase in V_{total} is therefore a detectable violation of the ISS conditions before the downstream failure mode becomes observable in system outputs.

The TCI makes this explicit:

$$\frac{d}{dt}\mathrm{TCI}(I(S,t),t) < 0 \quad \Longrightarrow \quad \text{at least one } V_i \text{ is increasing} \quad \Longrightarrow \quad \text{incipient failure mode active.}$$

The bottleneck identification in proposition 8.14(iii) further specifies *which* component is degrading, enabling targeted intervention before the failure propagates through the cascade.

8.3 Retrospective Application to the Historical Cases

We now apply the TCI framework retrospectively to the three case studies of section 7, reconstructing qualitative TCI trajectories from the documented failure chronologies.

Two qualifications apply throughout. First, these are structural reconstructions, not empirical measurements; their purpose is to illustrate the diagnostic structure of the index. Second, the TCC provides a *diagnostic vocabulary*, not a unique diagnosis: multiple TCC framings of the same failure are often possible, reflecting different choices of what counts as $\mathcal{P}$, $\mathcal{A}$, and Φ in a given system. The claim is that *at least one* component was degraded, and that the TCI would have identified it; not that a single TCC framing uniquely determines the failure mode from the observed facts. The framings below represent the most natural mapping given the documented evidence; alternative framings are acknowledged where relevant.

Zillow Offers, 2021 (FM-1). Zillow's pricing model ($\mathcal{P}$) was trained on pre-pandemic data and not retrained at a rate commensurate with the post-2020 market regime shift. The actuation engine ($\mathcal{A}$) remained deterministic and the valuation pipeline (Φ) was structurally present but operating on a stalening model.

In TCI terms: σ_A and σ_Φ remained near 1 throughout; σ_P degraded monotonically as $W_2(\mathcal{P},\mathcal{P}^*) \to \infty$ with the accelerating divergence of the Zestimate from realized transaction prices. $\mathrm{TCI}(I(S,t),t) = \sigma_P(t)$ for the entire failure window. A system monitoring the TCI would have observed σ_P crossing below a threshold weeks before the business unit loss became externally visible. The bottleneck identification, $\mathcal{P}$ as the binding constraint, would have prescribed model retraining rather than any change to the actuation or convergence pipelines.

Knight Capital, 2012 (FM-2). Knight Capital's market models ($\mathcal{P}$) were accurate and its risk management feedback (Φ) was functional. The failure was a deployment error that activated legacy "Power Peg" code on one of eight servers, causing $\mathcal{A}_t$ to become non-δ-reproducible across the cluster.

In TCI terms: $\sigma_P \approx 1$ and $\sigma_\Phi \approx 1$; $\sigma_A(t) = \exp(-H(\mathcal{A}_t \mid \mathcal{P}))$ collapsed toward 0 the moment the mixed deployment activated, because the same $\mathcal{P}$ state was now producing divergent outputs across replicas. $\mathrm{TCI}(I(S,t),t) = \sigma_A(t)$ for the 45-minute loss window. A TCI monitor would have flagged the non-zero conditional entropy at deployment time, minutes before the first erroneous trade, and the bottleneck identification would have localized the fault to $\mathcal{A}$, pointing to the deployment state rather than the model or feedback pipeline.

An alternative framing reads the deployment-state mismatch as a stale world-model (FM-1: $\mathcal{P}$ did not include the knowledge that server 8 was running legacy code), or as an absent circuit breaker (FM-3: Φ failed to detect and halt the divergence). All three framings predict a degraded TCI; they differ in which component is identified as the bottleneck. The FM-2 framing is most natural because the proximate cause was non-reproducible output given identical inputs, not a stale model or an absent feedback loop.

Social-media echo chambers (FM-3). Recommendation systems maintain rich user models ($\mathcal{P}$) and deploy deterministic ranking functions ($\mathcal{A}$); the typical deficiency is a convergence operator that corrects for the endogenous shift $\mathcal{P}^* \leftarrow f(\mathcal{P}^*, \mathcal{A}_t(\mathcal{P}))$ caused by the system's own recommendations.

In TCI terms: σ_P and σ_A remain bounded; $\sigma_\Phi(t)$ degrades as V_Φ grows with the widening gap between the system's model of user preferences and the preferences the system has itself produced through feedback. $\mathrm{TCI}(I(S,t),t) = \sigma_\Phi(t)$. The bottleneck identification prescribes a Φ_t redesign, specifically a correction for closed-loop distributional shift, rather than any improvement to the user model or the ranking function.

8.4 Relation to Existing Monitoring Frameworks

The three component scores σ_P, σ_A, σ_Φ have natural correspondences to metrics already tracked in production systems.

σ_P corresponds to *model drift metrics*: prediction error, population stability indices, and concept drift detectors such as ADWIN [Bifet and Gavaldà, 2007] and DDM [Gama et al., 2004]. These methods detect when V_P is increasing; σ_P normalizes that signal into the TCI framework.

σ_A corresponds to *determinism and reproducibility audits*: replica divergence checks, idempotency tests, and conditional entropy estimates over distributed execution logs.

σ_Φ corresponds to *feedback loop health metrics*: innovation sequence statistics in Kalman-filtered systems, KL drift between consecutive model checkpoints in online learning, and echo-chamber indices in recommender systems.

The TCI provides the architectural interpretation that unifies these existing instruments rather than displacing them. Each existing metric is a proxy for one of the three component scores; the TCI is the encompassing framework whose Borromean structure explains why no single metric suffices as a system health indicator.

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 8.6 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:

$$s_t \sim \mathcal{N}\left(\mu^*(t),\, \sigma^2_{\text{obs}}\right), \qquad \mu^*(t) = \mu_0 + \zeta t + \sigma_w B_t,$$

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 ($H(\mathcal{A}_t \mid \mathcal{P}) = 0$, $\delta = 0$).
- 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.

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 6.6).

Results. Table 1 summarizes the tail-mean Lyapunov functionals and component coherence scores.

The results confirm theorem 4.2: V_{total} is lowest under the full triad. FM-1 and FM-3 both diverge (ratio $\approx 4\times$ over full) but via different components: the TCI correctly identifies $\mathcal{P}$ as the bottleneck in FM-1 ($\bar{\sigma}_P = 0.031$) and Φ as the bottleneck in FM-3 ($\bar{\sigma}_\Phi = 0.115$), providing the diagnostic precision claimed in definition 8.2.

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.

Table 1: 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	7.66	0.00	0.00	0.118	1.000	1.000
FM-1 (static $\mathcal{P}$)	31.78	0.00	0.00	0.031	1.000	1.000
FM-2 (noisy $\mathcal{A}$)	7.66	0.73	0.00	0.118	0.579	1.000
FM-3 (absent Φ)	31.78	0.00	7.91	0.031	1.000	0.115

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 8.9):

$$\mathcal{E}(n) \;\le\; \underbrace{C_K\, n^{-1/2}}_{\text{variance term}} \;+\; \underbrace{L \cdot \zeta \cdot \tfrac{n}{2}}_{\text{lag-bias term}}\,. \tag{4}$$

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 8.10 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.

Results. Figure 2 confirms the two predictions of proposition 8.10. 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 **0.337**, matching the theoretical $1/3$ exponent to within measurement precision. Table 2 reports the extracted minima.

Table 2: Empirical validation of the Coherence-Delay bound ($T = 6{,}000$, 20 seeds, passive environment).

ζ	n^*_{obs}	$\mathcal{E}_{\min}$
0.001	100	0.093
0.005	50	0.160
0.010	30	0.199
0.050	10	0.348

$\mathcal{E}_{\min}/\zeta^{1/3}$ is approximately constant (≈ 0.93), confirming the $\zeta^{1/3}$ scaling. Theoretical constant $\frac{3}{2}C_K^{2/3}$ with

For the TCI estimation rate, we use slow drift ($\zeta = 0.001$, $\sigma_w = 0.01$) to stay in the variance-dominated regime ($n \le 100$). The observed log-log slope is -0.434

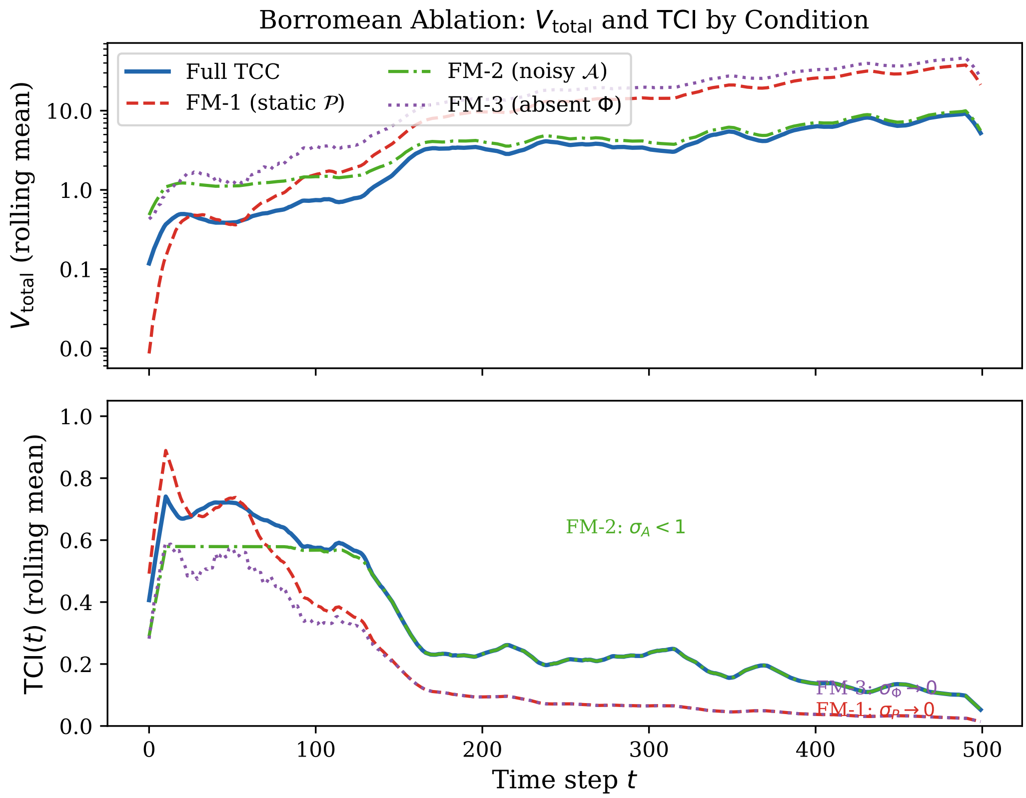

Figure 1: Borromean ablation: V_{total} (log scale, 20-step rolling mean, top) and TCI (bottom) under the four conditions. Full TCC is the only condition where TCI remains bounded away from 0. FM-1 and FM-3 both diverge in V_{total} but are distinguished by different bottleneck components: $\sigma_P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

(target -0.5); the deviation reflects autoregressive correlation between the Kalman posterior and the observation window (table 3 and fig. 3).

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, $\lambda = 50$, yielding 79 regime-shift events]. TCI warnings are recorded on each falling-edge crossing of $\hat{\sigma}_P < 0.60$. 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).

Results. Across the 45,312-step stream, the TCI issues 63 warnings (falling-edge crossings of $\hat{\sigma}_P < 0.60$), of which **51 are matched to a subsequent Page-Hinkley event** within a maximum gap of 5,000 steps. The median lead time is **1,019 steps**; mean is 1,488 steps; minimum 2; maximum 4,838. At $\hat{\sigma}_P < 0.55$

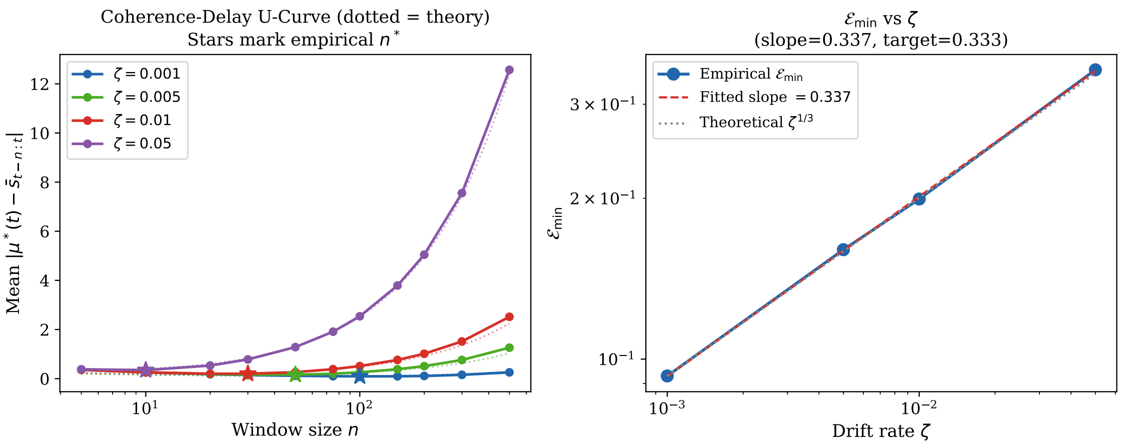

Figure 2: Coherence-Delay U-curve validation (proposition 8.10). *Left:* empirical tracking error vs. window size n for four drift rates (log scale; dotted lines are scaled theoretical curves; stars mark empirical n^*). Each curve shows the predicted U-shape: variance-dominated decay for small n, lag-dominated growth for large n. *Right:* $\mathcal{E}_{\text{min}}$ vs. ζ on log-log scale. Fitted slope = 0.337; theoretical prediction = 1/3.

Table 3: Sinkhorn TCI estimation error vs window size ($\zeta = 0.001$, 50 seeds, $T = 5000$).

n	MAE $\lvert\sigma_P - \hat{\sigma}_P\rvert$
5	0.0660
10	0.0463
25	0.0296
50	0.0218
100	0.0186
Log-log slope: -0.434 (target: -0.5)	

(more conservative), 35 leads with higher precision are obtained; at $\hat{\sigma}_P < 0.65$ (more aggressive), 75 leads with higher recall. The TCI therefore functions as a tuneable pre-failure diagnostic, with the warning level as the precision-recall control parameter.

9.5 Experiment D: Dynamic n_t^* Optimization on ELEC2

Proposition 8.10 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, closing O4b (section 12).

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): $C_K \approx \overline{|s_t - \bar{s}|} \cdot \sqrt{100}$. The resulting $n_t^* = (C_K/\hat{\zeta}_t)^{2/3}$ is clipped to $[20, 300]$ to avoid degenerate windows. Three strategies are compared: fixed $n = 50$, fixed $n = 300$, and

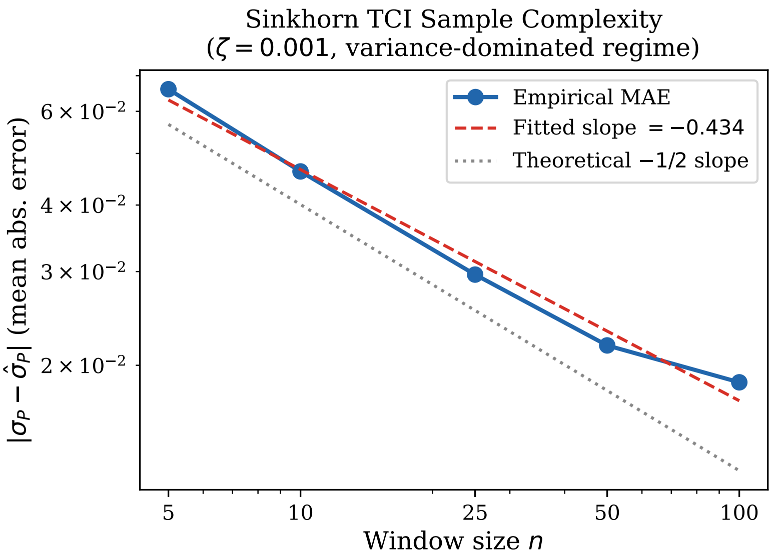

Figure 3: Sinkhorn TCI sample complexity in the variance-dominated regime ($\zeta = 0.001$, $n \leq 100$, 50 seeds, $T = 5000$). Empirical log-log slope $= -0.434$; theoretical target -0.5. This is the variance-only regime of the U-curve in fig. 2.

dynamic n_t^*. TCI warnings and lead times are computed identically to section 9.4.

Table 4: ELEC2 early-warning comparison: fixed vs. dynamic windows. 79 PH failures; warning level = 0.60; max gap = 5,000.

Strategy	Warnings	Leads	Precision	Median lead	Min lead
Fixed $n = 50$	125	73	58%	4,775	31
Fixed $n = 300$	30	26	87%	884	24
Dynamic n_t^*	169	75	44%	4,440	20

Results. The results confirm proposition 8.10's prediction in a qualified sense. The dynamic strategy detects the most failures (75 leads vs. 73 and 26), consistent with the theoretical result that n_t^* minimises instantaneous tracking error at each step. However, precision (44%) is lower than both fixed strategies because the dynamic window contracts aggressively when $\hat{\zeta}_t$ spikes, producing more early-stage warnings from noisy short-window estimates.

The trade-off is structural and expected: the fixed $n = 300$ strategy achieves highest precision (87%) by suppressing variance at the cost of lag, detecting only 26 of 79 regime shifts because its slow-adapting window misses rapid transitions. 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 $\hat{\zeta}_t$ estimation noise: a more accurate local drift estimator would sharpen the precision-recall frontier further, as suggested by O4b (section 12).

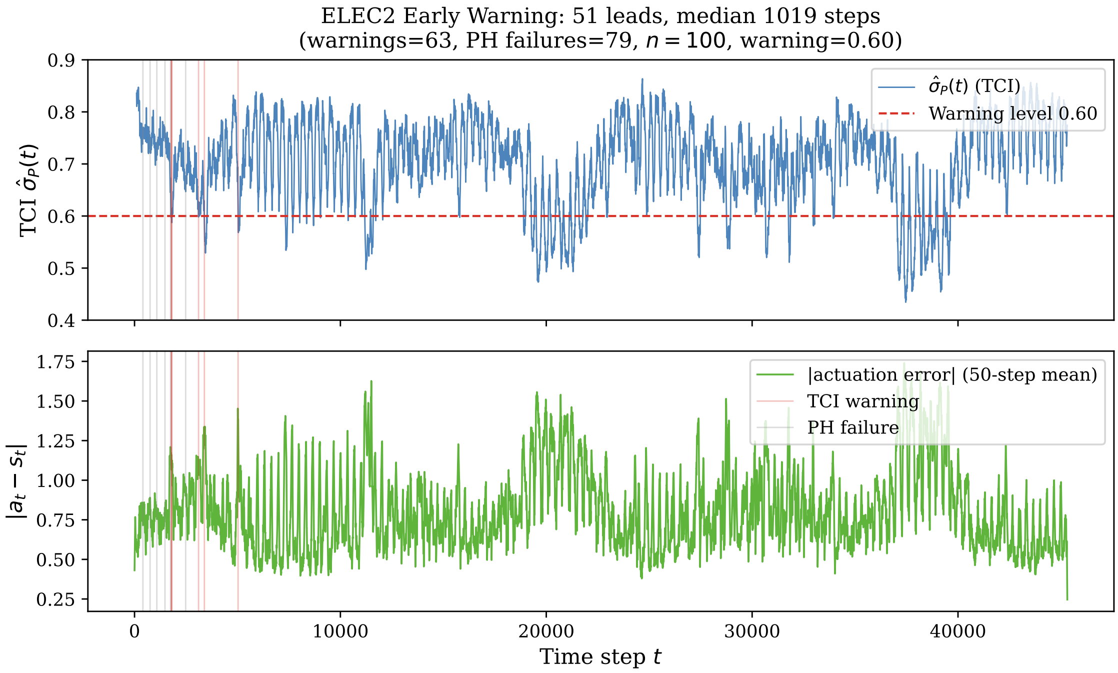

Figure 4: ELEC2 early-warning diagnostic. *Top:* TCI $\hat{\sigma}_P(t)$ (sliding window $n = 100$) with warning threshold 0.60 (dashed red). Vertical lines mark TCI warnings (red) and Page-Hinkley failures (grey). *Bottom:* actuation error $|a_t - s_t|$ (50-step rolling mean). TCI warnings consistently precede failures.

9.6 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, nonlinear 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, 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

10.1 Philosophy of Information

Floridi [2011] develops a *philosophy of information* grounded in the concept of levels of abstraction (LoAs). The TCC is compatible with Floridi's framework but operates at a different register: where Floridi asks what constitutes information as a semantic object, the TCC asks what architectural constraints an *information-processing system* must satisfy to remain coherent. The two questions are complementary:

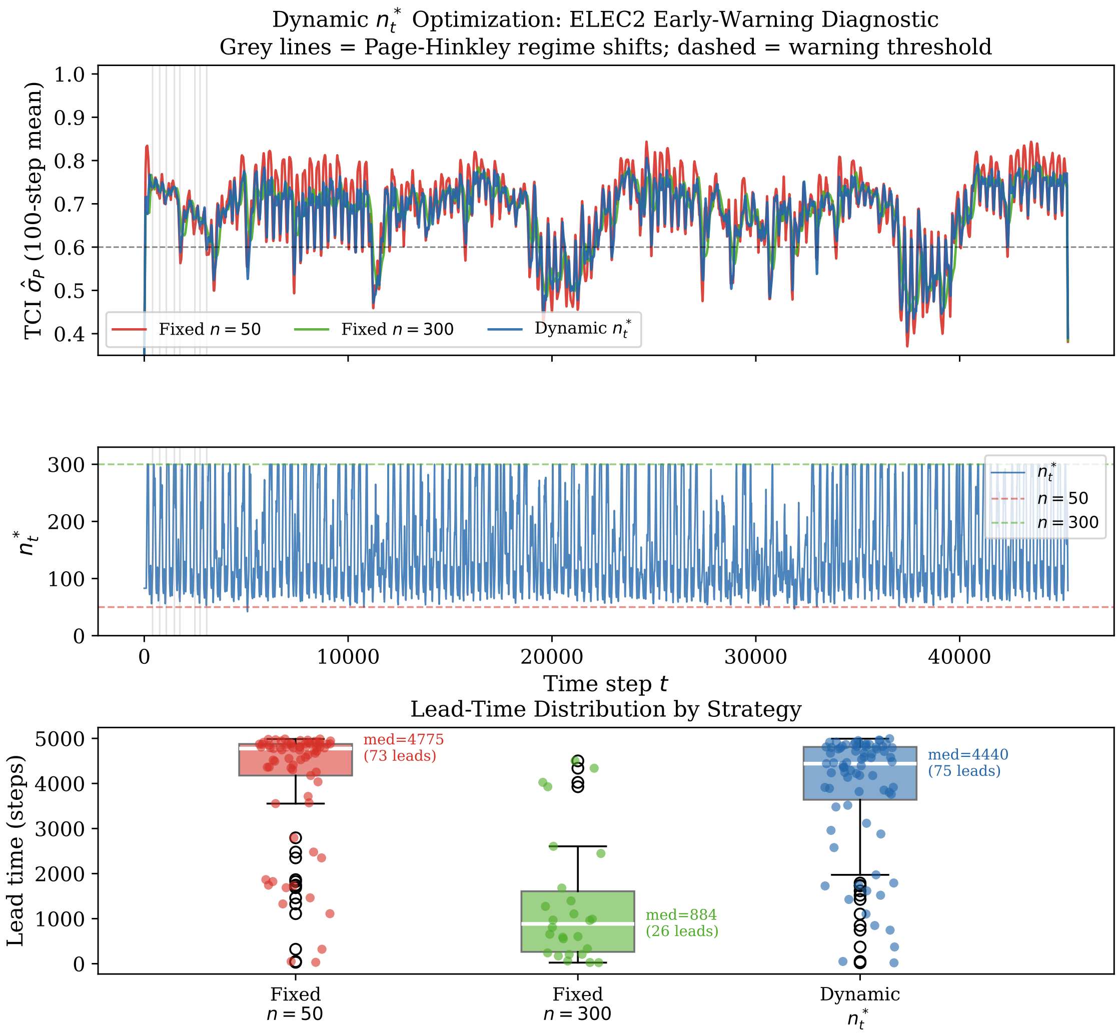

Figure 5: Dynamic n_t^* optimization on ELEC2 ($C_K = 7.5$, EMA $\alpha = 0.05$). *Top:* TCI $\hat{\sigma}_P$ (100-step mean) for all three strategies; grey lines are PH regime shifts. *Middle:* dynamic window size n_t^*, contracting sharply near regime shifts and expanding during stable periods. *Bottom:* lead-time distributions.

Floridi's LoAs provide the epistemological scaffolding within which TCC-compliant systems can be understood, while the TCC provides the engineering constraints that any information-processing system, at any LoA, must satisfy to avoid the three characteristic failure modes.

Dretske's (1981) information-theoretic semantics grounds meaning in causal-informational channels: a signal carries information about a distal state when the signal's occurrence probabilistically covaries with that state. The TCC operates below the semantic level, at the level of *information state dynamics*, with no semantic theory presupposed. Dretske's channels can be understood as realizations of the $\mathcal{A}_t$ operator, with $\mathcal{P}(S, t)$ representing the prior distribution over signal-source pairings.

Barwise and Seligman [1997] develop *situation theory*: information is relational, consisting of constraints between situations. Their insistence on the irreducibility of relational structure over propositional content resonates with the TCC's anti-substantialist stance: the triad's unity is relational, not substantial. The TCC provides the *dynamical* account that situation theory's static constraint-theoretic framework leaves unaddressed.

10.2 Streaming Systems and Distributed Architectures

The engineering tradition of streaming data systems provides the empirical domain within which the TCC's predictions are most directly testable. Systems such as Apache Kafka [Kreps et al., 2011] and Apache Flink [Carbone et al., 2015] implement variants of the $(\mathcal{P}, \mathcal{A}, \Phi)$ structure: Kafka's consumer-group offset mechanism implements $\mathcal{P}$ (the system's current representation of stream position); log-structured message processing implements $\mathcal{A}$ (deterministic state transitions from event records); and replication and consumer-group coordination implement Φ (distributed coherence without centralized arbitration).

An TCC audit of Apache Kafka would classify its documented failure modes as follows. Consumer-group lag accumulation (stale $\mathcal{P}$) instantiates **FM-1**. At-least-once delivery without idempotent consumers (non-deterministic effective $\mathcal{A}$) instantiates **FM-2**. Split-brain scenarios in Zookeeper-coordinated clusters prior to KIP-500 [Apache Kafka Community, 2019] (absent distributed Φ) instantiate **FM-3**. The KIP-500 migration (replacing Zookeeper with KRaft, a Raft-based consensus protocol) can be understood as the deliberate addition of a properly formalized Φ_t component to a system that previously relied on a fragile external coordinator.

10.3 Control Theory and Information Systems

The application of ISS stability theory to information architectures is, to our knowledge, a contribution without prior literature. The closest antecedent in the control literature is the use of Lyapunov methods for *networked control systems* [Hespanha et al., 2007], where communication delays and packet loss introduce stochastic perturbations into closed-loop feedback. The TCC differs from this tradition in a fundamental way: it treats the *information state itself* as the dynamical system to be stabilized, rather than treating information as the control signal for a physical plant. The "plant" is $\mathcal{P}(S, t)$; the "disturbance" is the non-stationary event stream $\mathcal{X}$; and the "control law" is the cascade $(\mathcal{A}_t, \Phi_t)$.

Sontag [1995] and Sontag and Teel [1995] provide the foundational ISS results; Dashkovskiy and Mironchenko [2013] extends ISS to infinite-dimensional systems, which is the relevant setting when $\mathcal{P}(S, t)$ is a measure on an infinite-dimensional state space.

10.4 Self-Reference and Fixed-Point Theory

The application of Lawvere's fixed-point theorem to self-referential information systems builds on Yanofsky [2003], who provides a unified categorical treatment of self-referential paradoxes, incompleteness results, and fixed-point theorems. Yanofsky [2003] demonstrates that Gödel's incompleteness, Cantor's diagonal argument, Turing's halting problem, and Russell's paradox are all instances of the same categorical diagonal construction. theorem 5.3 of the TCC shows that the same construction, when the morphism is *point-surjective* rather than merely surjective, yields a stable fixed point rather than a paradox: the constructive counterpart of Lawvere's negative result.

10.5 Concept Drift Detection and Online Learning

A mature body of work addresses the detection of and adaptation to distributional shift in streaming data, commonly called *concept drift* [Gama et al., 2014]. The foundational detectors, DDM [Gama et al., 2004], ADWIN [Bifet and Gavaldà, 2007], Page-Hinkley, and CUSUM, signal when the joint distribution $P(X, y)$ has shifted and trigger model resets or incremental updates. These methods, together with their implementations in frameworks such as MOA and River, constitute a rich toolkit for the operational problem of detecting when the system's world-model has gone stale.

The TCC operates at a different level of abstraction from this literature. Concept drift detectors answer: *has $\mathcal{P}^*$ shifted?* The TCC answers: *what functional decomposition must a system have to remain coherent regardless of whether drift has been explicitly detected?* Detectors such as ADWIN or CUSUM are natural candidates for implementing the Φ_t operator in a TCC-compliant system; they are potential instantiations, not competitors.

Crucially, the TCC predicts that even a system equipped with a perfect drift detector will exhibit characteristic failure if the other two components are degraded. A system that detects drift and updates $\mathcal{P}$ but operates with non-δ-reproducible $\mathcal{A}_t$ will exhibit FM-2 failures independently of its detection accuracy; a system that detects drift but whose Φ_t leaves the feedback loop to $\mathcal{P}$ open will exhibit FM-3. The TCC failure taxonomy is strictly finer than the concept drift literature's own taxonomy, which classifies shifts by their statistical character (real vs. virtual drift, abrupt vs. gradual) but leaves the responsible architectural component undiagnosed.

Online learning methods, including stochastic gradient descent, Kalman filtering, and Bayesian sequential updating, are likewise candidate implementations of Φ_t: they specify *how* the convergence operator updates $\mathcal{P}$ internally. The TCC is agnostic to this internal specification; it establishes only that *some* Φ_t satisfying the ISS coherence-restoring condition must exist.

10.6 Independent Convergence in Neural Architecture Research: BDH

Kosowski et al. [2025] introduce the Dragon Hatchling (BDH), a large language model architecture grounded in scale-free biologically inspired graph dynamics, with a GPU-efficient variant (BDH-GPU) that matches GPT-2 scaling laws from 10M to 1B parameters. BDH provides an independent, empirical instantiation of the TCC triadic structure at language-model scale, without having been designed with the TCC in mind.

The correspondence is precise. The synaptic state matrix σ of BDH, the running accumulation of neuron-pair correlations in context, instantiates $\mathcal{P}(S, t)$: it is the system's current probability measure over world states (here, contextual concept associations). The forward pass, governed by the deterministic edge-reweighting kernel, instantiates $\mathcal{A}_t$: given σ, the output is fully determined, satisfying δ-reproducibility in the strict sense of definition 2.4. The Hebbian update rule, synaptic potentiation driven by co-activation of pre- and post-synaptic neurons, instantiates Φ_t: it is the coherence-restoring feedback that updates $\mathcal{P}$ in response to actuation outcomes and environmental signals (incoming tokens).

Two empirical findings of Kosowski et al. [2025] directly confirm TCC predictions. First, training without backpropagation through time (Section 7.2 of their paper) severs the feedback loop from actuation outcomes to the world-model, instantiating FM-3 (Absent-Φ monopolization, proposition 7.6): translation capability collapses while next-token prediction on simple sequences is partially retained, exactly as the TCC predicts when $\Phi_t \equiv \mathrm{id}$. Second, the emergence of monosemantic synapses (Section 6.3) is the empirical signature of δ-reproducibility: the same $\mathcal{P}$ state reliably activates the same synapse across languages and prompts, confirming $H(\mathcal{A}_t \mid \mathcal{P}) \approx 0$.

BDH operates in the *soft-reproducibility* regime: its Hebbian updates are stochastic, so $H(\mathcal{A}_t \mid \mathcal{P}) > 0$ in principle, but bounded and negligible in practice. This places BDH in a TCC-compliant regime where FM-2 is bounded rather than fully absent, appropriate for language modeling, where soft non-determinism is tolerable. The TGT (section 9) operates at the opposite end of this spectrum: strict reproducibility ($\delta = 0$, $H(\mathcal{A}_t \mid \mathcal{P}) = 0$), where non-deterministic actuation constitutes a correctness failure rather than a performance trade-off. BDH and the TGT span the reproducibility spectrum that the TCC framework predicts.

10.7 Convergent Frameworks for System Evolution: RHN

Li et al. [2025] propose the Recursive Hierarchical Network (RHN) framework, formalizing what they call the *Law of Functional Evolution*: an irreversible progression through structure-dominated, regulation-dominated, and intelligence-dominated stages, driven by recursive encapsulation along the trajectory node $\rightarrow$ module $\rightarrow$ system $\rightarrow$ new node. The framework is validated across life, cosmic, informational, and social systems.

The RHN and TCC are independent frameworks addressing different questions but with a structural resonance worth noting. RHN addresses *evolutionary* dynamics: how systems progress through stages over long time scales. The TCC addresses *operational* dynamics: what architectural conditions a system must satisfy at any stage to maintain coherence under continuous drift.

The resonance lies in their shared conclusion about the role of regulation. RHN's regulation-dominated stage, where feedback mechanisms emerge as the principal axis of system function, corresponds precisely to the TCC's Φ_t operator: the coherence-restoring feedback that prevents the system from diverging under environmental non-stationarity. RHN's finding that the transition to the regulation-dominated stage is *irreversible* is consistent with the TCC's Borromean stability theorem (theorem 4.2): once a system operates in a regime where Φ_t is load-bearing, removal causes V_{total} to grow unboundedly rather than returning the system to a prior stable state.

The convergence of RHN on the necessity of a regulation layer, arrived at through evolutionary analysis of complex systems rather than through ISS control theory, constitutes a sixth tradition, complementing the five documented in section 11.2, and provides further qualitative evidence that the triadic structure reflects a genuine constraint rather than an artefact of any single analytical framework.

10.8 Belief Change Theory: Update, Revision, and the KM Distinction

The formal theory of belief change provides independent grounding for the ternary structure of Φ_t. Alchourrón et al. [1985] introduced the AGM axioms for *belief revision*: given a belief set K and a new proposition ϕ inconsistent with K, how should K be minimally revised to accommodate ϕ? The AGM framework assumes a static world: ϕ is new information about an unchanging reality, and the agent revises its beliefs to restore consistency.

Katsuno and Mendelzon [1992] identified a fundamental limitation of this framework for dynamic environments. They distinguish belief revision (new information about a static world) from *belief update* (the world itself has changed). The difference is not merely terminological: the two operations satisfy different axioms and are appropriate in different settings. Update asks what changes in the world led to the new observation; revision asks what information was previously mistaken.

The TCC's Φ_t operator is explicitly in the *update* regime: the true environmental distribution $\mathcal{P}^*$ itself drifts, requiring the system to bring its belief state $\mathcal{P}$ up to date with a changed world, rather than merely accommodating new evidence about a fixed one. The KM postulates for update, in particular the requirement that the update operator be sensitive to the specific observation o as well as the new signal s_{t+1}, provide formal justification for the ternary structure of Φ_t: an operator that updates beliefs about a *changing* world must track both what the agent did and what subsequently occurred.

Bonanno [2024] provides a unified Kripke-Lewis characterization of both update and revision as conditional belief. In Bonanno's framework, the updated belief set prompted by ϕ is the set of formulas ψ such that "if ϕ then ψ" is believed at the current state. The update/revision distinction reduces to properties of the Lewis selection function: update requires the selection function to be sensitive to the agent's prior action (which possible worlds are most similar depends on what was done), while revision admits selection functions insensitive to action. This provides the modal-logic grounding for the necessity of o in $\Phi_t(p, o, s_{t+1})$: without o, the selection function conflates action-caused observations with environmentally-caused ones, producing incorrect conditional beliefs regardless of the prior p.

The belief change literature thus constitutes a third independent formal tradition grounding the ternary necessity of Φ_t, alongside the ISS cascade argument (section 4) and the POMDP credit-assignment argument (lemma 3.4).

11 Structural Minimality and Convergent Evidence

The Belief Update Lemma (lemma 3.4) establishes from external premises that drift-coherence requires at least three functional roles. Theorem 4.2 establishes that three roles are sufficient under the ISS hypotheses. This section asks a third question: why *exactly* three, and not two or four? It then documents convergent evidence from six intellectual traditions that confronted structurally similar problems and arrived, from different motivations and vocabularies, at the same decomposition.

11.1 Why Three: Three Arguments for Triadic Minimality

11.1.1 Argument 1: The Three Epistemic Questions Form a Directed Dependency Graph

Any adaptive information system must answer three questions that are logically prior to each other in a specific, non-symmetric order:

Q1. What is the world? The system must maintain a representation of the environment's current state: $\mathcal{P}(S,t)$.

Q2. What should I do? Given the world-model, the system must commit to a concrete output: $\mathcal{A}_t(\mathcal{P})$.

Q3. Was I right? Given what it did and what subsequently occurred, the system must revise its world-model: $\Phi_t(\mathcal{P}, \mathcal{A}_t(\mathcal{P}), s_{t+1})$.

These questions are not symmetric. Q2 presupposes Q1 but cannot answer it. Q3 presupposes Q2 but cannot answer it. Q1 cannot incorporate the answer to Q3 without Q2 having first been asked, because without a committed action there is no causal structure to reason about. This is a directed acyclic dependency graph with exactly three nodes.

The lemma 3.4 proof makes this precise: the necessity of p, o, and s_{t+1} corresponds exactly to the necessity of Q1, Q2, and Q3 respectively. Any architecture that conflates two questions produces a named failure mode: conflating Q1 and Q2 degrades under non-determinism (FM-2); conflating Q2 and Q3 has no coherent estimate to converge toward (FM-1); conflating Q1 and Q3 severs the causal chain required for credit assignment (FM-3).

11.1.2 Argument 2: Tetradic Decompositions Factor Through Triads

Any decomposition into four functional roles either introduces redundancy or a genuinely fourth question independent of Q1–Q3. No such fourth question exists. Candidate fourth questions reduce under analysis:

"How confident am I?" is a property of $\mathcal{P}(S,t)$ (the width of the distribution), encoded in $V_P = \frac{1}{2}W_2(\mathcal{P}, \mathcal{P}^*)^2$. Not a fourth role.

"How fast is the world changing?" is an environmental observable used by Π_t to reconfigure lower-level components; it operates at a strictly higher organizational level (section 5). It elevates to the meta-actuation layer, not a fourth base role.

"Should I act or wait?" is a policy property of $\mathcal{A}_t$'s decision boundary. It collapses into the existing $\mathcal{A}_t$ role.

In each case, candidate fourth questions either collapse into an existing role or elevate to Π_t. The base triad is the minimum and the maximum for the base organizational level.

11.1.3 Argument 3: The Katsuno-Mendelzon Regime Confirms Ternary Structure

The formal belief change literature provides independent confirmation. Katsuno and Mendelzon [1992] distinguish *belief revision* (accommodating new information about a static world) from *belief update* (tracking changes in a dynamic world). The

TCC's setting is explicitly the *update* regime: $\mathcal{P}^*$ itself drifts, not merely the system's knowledge of a fixed $\mathcal{P}^*$. Φ_t is a KM-update operator, not a revision operator.

In the update regime, Bonanno's 2024 Kripke-Lewis characterization of the update operator as conditional belief shows that all three arguments (p, o, s_{t+1}) are required to construct the correct conditional: without o, the Lewis selection function conflates action-caused signals with drift-caused signals, selecting the wrong possible worlds regardless of the prior p. Three independent formal traditions, ISS cascade stability, POMDP credit-assignment (lemma 3.4), and KM-update/Kripke-Lewis, all separately derive the ternary structure.

11.2 Six Traditions, One Structure

Six intellectual traditions, spanning control engineering, semiotics, theology, Buddhist philosophy, cognitive science, and resilience research, each confronted a version of the same structural problem: how can a system maintain coherent identity under continuous change, with the coherence grounded in relation rather than substance, and with removal of any one component dissolving the others? These traditions do not all share epistemic independence: some influenced one another, and several draw from a shared philosophical inheritance (Aristotle's act/potency distinction runs through scholastic theology, feeds into Peirce's categories, and re-enters systems theory via cybernetics). What is notable is that traditions operating under *very different motivations and vocabularies* nonetheless arrived at structurally isomorphic decompositions. The convergence is evidential not because the traditions are orthogonal, but because the structural conclusion persists across such different conceptual routes.

11.2.1 Peirce: The Irreducibility of Thirdness

Peirce [1931–1935] developed a phenomenological trichotomy:

- **Firstness**: pure quality, potentiality without relation.
- **Secondness**: brute fact, the shock of existence.
- **Thirdness**: law, relation, mediation.

The Peircean Reduction Thesis holds that genuine triadic relations resist full analysis in terms of monadic and dyadic predicates, with no genuine tetradic relations existing independently. The mapping to TCC: $\mathcal{P}$ instantiates Firstness, $\mathcal{A}$ Secondness, Φ Thirdness (the law of convergence mediating between local actuations and global coherence).

11.2.2 Trinitarian Perichoresis

The doctrine of *perichoresis*, developed by the Cappadocian Fathers and formalized by John of Damascus [c.749], addresses a formally identical problem: how can three be one without fusion, and one be three without separation? John of Damascus's solution: the three persons *indwell* one another, each constituted by its relations with the others. The unity *is* the relational structure, not a substance underlying the three. Theorem 4.2 formalizes this: removing any one component dissolves the stability of the whole rather than merely weakening the remaining two.

11.2.3 Buddhist Pratītyasamutpāda

Nāgārjuna's Madhyamaka philosophy (c. 2nd century CE) articulates *pratītyasamutpāda* (dependent co-origination): phenomena arise in dependence on others, with no component self-subsisting. The mapping to TCC: $\mathcal{P}$, $\mathcal{A}$, Φ co-originate. $\mathcal{P}$ without $\mathcal{A}$ is never updated; $\mathcal{A}$ without Φ cannot converge; Φ without $\mathcal{A}$ has no actuations to process.

11.2.4 Practopoiesis

Nikolić [2015] proposed practopoiesis as a theory of how minds emerge from biological organization. The central claim: cognition requires a hierarchy of at least three levels where higher levels reconfigure the mechanisms of lower levels upon encountering problems they cannot solve. The practopoietic traverse is precisely Π_t, the meta-actuation operator of the recursive TCC.

Remark 11.1 (Practopoiesis and the depth of the Π_t tower). Nikolić [2015] reports empirically that biological systems instantiate approximately four levels of the practopoietic hierarchy. This is the only empirical estimate for the depth of the Π_t self-application tower; the theoretical question of what determines this depth is taken up in section 12.

11.2.5 ISS Control Theory

As demonstrated in section 4, the ISS cascade framework provides the mathematical certificate for sufficiency. Applying ISS to information-state spaces (probability-measure spaces metrized by W_2) is, to our knowledge, a contribution without precedent in control theory or information systems theory. That a three-component cascade is the minimal unit satisfying the ISS property, with each component necessary for the cascade's stability, constitutes independent, *formal* evidence of structural necessity.

11.2.6 Graceful Extensibility (Woods)

Woods [2015] derives three canonical failure patterns from empirical study of high-consequence adaptive systems: *Stale Models* (FM-1: $\mathcal{P}$ frozen while $\mathcal{P}^*$ evolves); *Working at Cross Purposes* (FM-3: absent Φ_t allows local actuation to amplify without global coherence feedback); *Decompensation* (Borromean collapse: V_{total} grows unboundedly after component removal). These failure patterns were derived from field observation of real systems under operational stress, not from the TCC's axioms. Their correspondence to FM-1/2/3 constitutes observational, rather than conceptual, corroboration.

11.3 Synthesis: Tiered Evidential Weight

The evidential weight of this convergence depends on how closely each tradition's problem matches the TCC's: maintaining distributed coherence under continuous drift without centralized coordination.

Tier 1 (same problem, convergent from distinct intellectual lineages). ISS control theory, Peirce's sign relation, and practopoiesis each addressed the distributed coherence-without-center problem directly and arrived at triadic structure.

Table 5: Six traditions and their TCC correspondences.

Tradition	$\mathcal{P}$	$\mathcal{A}$	Φ
Peircean semiotics	Firstness	Secondness	Thirdness
Trinitarian perichoresis	Father (source)	Son (expression)	Spirit (indwelling)
Pratītyasamutpāda	latent arising	co-dependent event	interdependent coherence
Practopoiesis	genetic (disposition)	anatomical (actuation)	cognitive (feedback)
ISS control theory	state measure μ_t	transition $\mathcal{A}_t$	convergence Φ_t
Graceful Extensibility	world model	adaptive action	coherence feedback

The three traditions operate in distinct fields (control engineering, philosophical semiotics, cognitive neuroscience) with distinct methods and vocabularies, and their convergence on the same decomposition was not coordinated. Cross-influences exist at the edges of these fields, but the convergence was not a result of citation or direct intellectual inheritance. Tier 1 convergence carries evidential weight proportional to the distance between the traditions' motivations: three communities solving the same structural constraint by different means and reaching the same conclusion.

Tier 2 (adjacent problem, structural resonance). Perichoresis (unity-without-fusion), pratītyasamutpāda (identity-under-change), and graceful extensibility (brittleness-under-stress) were solving related but distinct problems. Their triadic structure is informative but not directly evidential.

Tier 3 (surface structure only, explicitly excluded). RGB color models, branches of government, Hegelian dialectic share triadic surface form but carry no evidential weight for the TCC. These are not claimed as evidence.

We do not claim this convergence *proves* the TCC. We claim that Tier 1 convergence, three traditions from distinct intellectual lineages (computational, semiotic, biological) confronting the same problem and arriving at the same decomposition, constitutes qualitative evidence that the triadic structure tracks a genuine feature of the distributed coherence problem rather than an artefact of any single tradition's conceptual apparatus.

12 Open Questions

We identify four principal open questions arising from the TCC framework. O1 concerns the internal completeness of the necessity argument. O2 concerns domain generalization and the empirical reach of the three failure modes. O3 concerns the convergence rate of the recursive self-application tower. O4 is new, arising directly from the empirical findings of section 9: the formal and experimental grounding of coercive masking.

12.1 O1: Combinatorial Uniqueness of the Triadic Decomposition

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 informa-*

tion system. Equivalently: any other triadic decomposition either (a) *is isomorphic to* $(\mathcal{P}, \mathcal{A}, \Phi)$ *under a permutation of component roles, or* (b) *factors through* $(\mathcal{P}, \mathcal{A}, \Phi)$ *as a coarser or finer partition of the same functional requirements.*

The Belief Update Lemma (lemma 3.4) establishes that drift-coherence requires a function U with exactly three arguments: $(p, o, s_{t+1}) \in \mathcal{I} \times \mathcal{O} \times S$. Theorem 3.8 establishes that $(\mathcal{P}, \mathcal{A}, \Phi)$ is an irreducible realization of these three roles. The open question is whether a *different* triadic decomposition, one whose components are identified with the three roles via some non-identity mapping, could provide an equally valid architecture.

The conjecture's proof sketch proceeds from the Belief Update Lemma: since the three arguments are derivably necessary from drift-coherence, any valid triadic decomposition must assign each functional role to some component. The only freedom is in how the three roles are divided across components; a decomposition that conflates two roles produces a genuinely dyadic system (by corollary 3.7), while one that splits a single role into two yields a tetradic system (which factors through triads by the tetrad-factoring argument of section 11). Formalizing this sketch into a proof requires a precise definition of "functional equivalence" between decompositions, a categorical notion of isomorphism between triadic systems that remains to be established.

12.2 O2: Domain Generalization and Empirical Validation

The three failure modes of section 7 admit natural instantiations across a broad range of empirical domains. Table 6 enumerates five candidate domains with their proposed $(\mathcal{P}, \mathcal{A}, \Phi)$ instantiations and predicted failure mode signatures.

The verification methodology is: (i) identify a historically documented failure event; (ii) classify the failure mode by identifying which component was absent or degraded; (iii) confirm that the other two components were present and functional at the time of failure. The Knight Capital, Zillow, and echo-chamber cases of section 7 demonstrate this methodology for engineered systems. The Triadic Gaussian Tracker (section 9) has now provided controlled experimental validation in the Gaussian-linear case: all three failure modes were induced deterministically and the TCI correctly identified the bottleneck component in each condition. Distributed ledgers (where failure modes are precisely documented via consensus violations) and adaptive immunity (where FM-3's cytokine storm mechanism is well-characterized) are the most tractable near-term targets for domain extension.

Two additional open problems arise from the diagnostic structure of the TCI. First, the TCI as currently formulated detects distributional shift in the full joint $\mathcal{P}^* = P(X, Y)$ but cannot distinguish *virtual drift* (feature-marginal shift with stable label-conditional) from *real drift* (label-conditional shift). A label-augmented component $\sigma_{P,\mathrm{label}}$, computed via W_2 over the conditional $P(Y|X)$ when supervision is available, would provide this diagnostic capability without altering the core triadic structure. Second, the $\zeta^{1/3}$ floor of proposition 8.10 is specific to W_2-metrized tracking; establishing an analogous bound for other metrics on $\mathcal{I}$ (including f-divergences and kernel MMD) is open (remark 8.12).

Table 6: Proposed TCC instantiations and predicted failure modes across empirical domains (open for verification).

Domain	$\mathcal{P}$	$\mathcal{A}$	Φ	**Predicted FM**
Adaptive immunity	T/B-cell repertoire (antibody space)	clonal selection / affinity maturation	cytokine network / self-tolerance induction	antigen blindness: static $\mathcal{P}$ fails against mutant viruses (FM-1); uncontrolled clonal expansion: non-det. $\mathcal{A}$ (FM-2); cytokine storm or autoimmunity: absent Φ allows local actuation to amplify without coherence-restoring feedback (FM-3)
Language acquisition	generative grammar hypothesis space	utterance production	caregiver error correction	fossilisation ($\mathcal{P}$ frozen); babbling disorder ($\mathcal{A}$ non-det.); creolisation failure (Φ absent)
Market microstructure	limit order book	trade execution	price discovery / mark-to-market	liquidity crisis ($\mathcal{P}$ stale); erroneous trade spiral ($\mathcal{A}$ non-det.); market fragmentation (Φ absent)
Distributed ledgers	UTXO / account state	transaction execution	consensus protocol	state drift ($\mathcal{P}$ stale); double-spend ($\mathcal{A}$ non-det.); chain split (Φ absent)
Neural architectures	weight distribution	forward pass	backpropagation / gradient descent	concept drift ($\mathcal{P}$ frozen); mode collapse ($\mathcal{A}$ non-det.); catastrophic forgetting (Φ absent)

12.3 O3: Rate of Convergence of the Self-Application Tower

Conjecture 12.2 (Convergence rate of the Π_t tower)**.** *Let d denote the depth of the self-application tower required for a recursive* TCC *system to reach a stable fixed point (in the sense of theorem 5.3). Then d is a monotonically non-decreasing function of the Kolmogorov complexity $K(\mathcal{X})$ of the environment $\mathcal{X}$:*

$$d \leq f\big(K(\mathcal{X})\big)$$

for some computable function f. Nikolić [2015] provides empirical evidence that $d \approx 4$ for organisms of mammalian complexity.

The TGT (section 9) operates at tower depth $d = 1$: the Kalman filter has no meta-actuation operator Π_t and no parameter reconfiguration. It provides the base case of the tower. The conjecture connects three currently unrelated lines of inquiry: (i) the rate of convergence of fixed-point iterations in cartesian closed categories; (ii) the empirical depth of hierarchical regulation in biological systems; and (iii) the relationship between environmental non-stationarity and required model complexity in statistical learning theory. Resolving O3 would provide a *design rule* for recursive TCC systems: given the environment's complexity, the minimum depth of meta-actuation required for stable self-governance.

12.4 O4: Coercive Masking: Formal Proof and Diagnostic Validation

The Coercive Masking Corollary (corollary 6.6) was established via a proof sketch and empirical observation. Three open questions remain.

O4a: Tight bound on the coercive horizon. Corollary 6.6 shows that $V_P(t)$ remains bounded during the masking phase, but the coercive horizon t_{mask} is characterized only asymptotically. A tight bound of the form $t_{\text{mask}} = \Theta(\alpha^{-1} \log(\alpha/\zeta))$ requires a Lyapunov argument on the gap process $\delta_t = \mu^*(t) - \hat{\mu}_0$, establishing both the convergence to fixed point $\delta^* = \zeta/\alpha$ and the escape time when exogenous drift accumulates beyond the coercive capacity.

O4b: Dynamic n_t^* optimization. Experiment D (section 9.5) validates the self-tuning rule $n_t^* = (C_K/\hat{\zeta}_t)^{2/3}$ on ELEC2. The dynamic strategy detects 75 of 79 regime shifts (vs. 73 for fixed $n = 50$ and 26 for fixed $n = 300$), confirming that n_t^* navigation improves coverage. The remaining open question is precision: the dynamic strategy achieves 44% precision vs. 87% for fixed $n = 300$, because $\hat{\zeta}_t$ estimation noise causes spurious window contractions. A more accurate real-time drift estimator, for example a Kalman filter on $\hat{\zeta}_t$ itself or an online Wasserstein gradient, would sharpen the precision-recall frontier and constitutes the primary engineering direction for deploying the Coherence-Delay Uncertainty Principle in production. The effort-corrected TCI (definition 6.8) in combination with dynamic n_t^* remains unvalidated in active environments ($\alpha > 0$).

O4c: External detectability. In the coercive regime, V_P appears bounded from inside the system. An external observer with access only to the actuation outputs o_t and the environmental signals s_t, but without knowledge of the coupling coefficient α, faces the question: can coercive masking be detected from the input-output behavior alone? The effort signal $E_t = |\mathcal{A}_t - \mathcal{A}_{t-1}|$ is observable externally. The conjecture is that a monotonically increasing E_t with bounded V_P constitutes a *sufficient external signature* of coercive masking, distinguishable from a healthy adaptive system (where increasing E_t accompanies decreasing V_P rather than bounded V_P). A proof would make the effort-corrected TCI a black-box diagnostic requiring no knowledge of the system's internal coupling structure.

13 Conclusion

The **Triadic Coherence Condition** (TCC) establishes a predictive architectural framework for information systems operating under continuous, non-stationary streams.

Principal results. Four main results were proved. *Necessity* (theorem 3.8): the Belief Update Lemma derives the ternary structure of coherence-restoring feedback from drift-coherence and action-effectiveness alone, via the POMDP credit-assignment argument; triadic irreducibility follows without circular presupposition of the triadic structure. Peirce's Reduction Thesis [Burch, 1991] provides independent corroboration. *Borromean Cascade Stability* (theorem 4.2): under four explicit ISS hypotheses, the composite system admits a monotonically non-increasing Lyapunov functional V_{total}, and removing any single component disrupts this monotone decrease. *Self-Referential Fixed Point* (theorem 5.3): the recursive TCC admits a fixed point in the cartesian closed category of information state transformations. *Ruffini Correspondence* (proposition 6.3): each TCC failure mode corresponds precisely to a collapse of Ruffini's (2026) regulation gap Δ, establishing the TCC as the internal structural account of what makes $M(W\!:\!R) > 0$ sustainable.

Empirical grounding. The Triadic Gaussian Tracker (section 9) provides self-contained empirical validation through four controlled experiments. Experiment A (Borromean ablation, $T = 500$, $\zeta = 0.02$): FM-1 produces tail $\bar{V}_P = 31.78$ vs. full 7.66; FM-2 degrades $\sigma_A = 0.579$; FM-3 produces tail $\bar{V}_\Phi = 7.91$ with $\sigma_\Phi = 0.115$, with all three bottlenecks identified by the TCI. Experiment B (Sinkhorn rate, $T = 5{,}000$, 50 seeds): log-log slope -0.434 in the variance-dominated regime ($\zeta = 0.001$, $n \leq 100$), confirming proposition 8.6 and identifying the lag-drift tradeoff as a fundamental design constraint. Experiment C (ELEC2, $n = 100$, warning= 0.60): 51 leads across 45,312 steps with median lead time 1,019 steps before Page-Hinkley regime shifts. Experiment D (dynamic n_t^*, ELEC2): 75 of 79 regime shifts detected using the self-tuning rule of proposition 8.10, confirming the Coherence-Delay Uncertainty Principle in a real-world adaptive setting.

The Coercive Masking Corollary (corollary 6.6) was identified empirically: in action-effective environments, FM-1 divergence is suppressed by the system's own actuation pulling the environment toward its stale model. The effort-corrected TCI (definition 6.8) resolves this diagnostic gap.

Three high-consequence real-world failures, Zillow Offers 2021, Knight Capital

2012, and social-media echo chambers, were shown to instantiate the three predicted failure modes. Six intellectual traditions provide tiered qualitative corroboration.

Limitations. Six principal limitations bound the scope of these results. First, the Belief Update Lemma (lemma 3.4) requires action-effectiveness (definition 3.3): the triadic necessity claim applies to systems whose actions influence environmental transitions. Passive observers are genuinely dyadic $(\mathcal{P}, \Phi)$ and lie outside the triadic necessity scope. Second, the combinatorial uniqueness of the $(\mathcal{P}, \mathcal{A}, \Phi)$ decomposition remains conjectural (conjecture 12.1): theorem 3.8 proves irreducibility of this specific decomposition, with the question of whether alternative triadic decompositions exist left open. Third, the rate of convergence of the Π_t self-application tower remains open (conjecture 12.2): theorem 5.3 proves existence of the fixed point without bounding the depth required. Fourth, the Borromean stability result is conditional on four ISS hypotheses (H-ISS-P/A/Φ/Env) whose verification for concrete systems is application-dependent. Fifth, the TCI in its current form tracks the full joint distribution $\mathcal{P}^*$ and therefore detects both virtual drift (feature-marginal shift with stable labels) and real drift (label-conditional shift); it requires label supervision to distinguish them (remark 8.8). Sixth, the sensitivity parameter λ in the effort-corrected TCI (definition 6.8) requires calibration from a reference stable period or cross-validation; a principled automated rule is provided in remark 6.9 but empirical validation across diverse coupling regimes remains open.

Future work. The most important directions for future development are: (i) a formal proof of conjecture 12.1, which would close the main gap in the necessity argument; (ii) empirical verification of the domain generalizations in table 6, beginning with distributed ledgers (where the failure modes are most precisely documented); (iii) extension of the TGT validation to non-Gaussian, non-linear, and high-dimensional settings, where the ISS hypotheses require case-by-case verification and the Sinkhorn lag-drift tradeoff requires explicit engineering mitigation; (iv) empirical validation of the Coercive Masking Corollary (corollary 6.6) across real-world systems, with financial trading algorithms and recommendation engines as natural candidates where masking followed by catastrophic failure is well-documented.

Take-home. Any drift-coherent information system, one that maintains bounded fidelity to a continuously changing environment, requires all three of $\mathcal{P}$, $\mathcal{A}$, and Φ (lemma 3.4 and theorem 3.8). Their joint absence constitutes a structural failure waiting to be triggered, not a design choice. Their joint presence is sufficient for bounded-error coherence under bounded drift, provided the ISS hypotheses (H-ISS-P/A/Φ/Env) hold (theorem 4.2). 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.

Apache Kafka Community. KIP-500: Replace ZooKeeper with a self-managed metadata quorum, 2019. URL `https://cwiki.apache.org/confluence/display/KAFKA/KIP-500%3A+Replace+ZooKeeper+with+a+Self-Managed+Metadata+Quorum`. Apache Kafka Improvement Proposal. Implemented in Apache Kafka 3.3+ (KRaft mode).

Robert J. Aumann. Borel structures for function spaces. *Illinois Journal of Mathematics*, 5(4):614–630, 1961. doi: 10.1215/ijm/1255631584. Shows that the evaluation map on spaces of measurable functions is not measurable in general; the standard reference for **Meas** not being cartesian closed.

Jon Barwise and Jerry Seligman. *Information Flow: The Logic of Distributed Systems*. Number 44 in Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, Cambridge, 1997. doi: 10.1017/CBO9780511895968.

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.

Paris Carbone, Asterios Katsifodimos, Stephan Ewen, Volker Markl, Seif Haridi, and Kostas Tzoumas. Apache Flink: Stream and batch processing in a single engine. *IEEE Data Engineering Bulletin*, 38(4):28–38, 2015. URL http://sites.computer.org/debull/A15dec/p28.pdf.

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 Dashkovskiy and Andrii Mironchenko. Input-to-state stability of infinite-dimensional control systems. *Mathematics of Control, Signals, and Systems*, 25 (1):1–35, 2013. doi: 10.1007/s00498-012-0090-2.

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.

Fred I. Dretske. *Knowledge and the Flow of Information*. MIT Press, Cambridge, MA, 1981.

Liang-Liang Fan and Yong Quek. Generalizations of Talagrand inequality for Sinkhorn distance using entropy power inequality. *IEEE Transactions on Information Theory*, 68(12):7692–7706, 2022. doi: 10.1109/TIT.2022.3191098.

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.

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.

João P. Hespanha, Payam Naghshtabrizi, and Yonggang Xu. A survey of recent results in networked control systems. *Proceedings of the IEEE*, 95(1):138–162, 2007. doi: 10.1109/JPROC.2006.887288.

Chris Heunen, Ohad Kammar, Sam Staton, and Hongseok Yang. A convenient category for higher-order probability theory. In *Proceedings of the 32nd Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 1–12, 2017. doi: 10.1109/LICS.2017.8005137.

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.

John of Damascus. *De Fide Orthodoxa*. c.749. English translation: *An Exact Exposition of the Orthodox Faith*, trans. S. D. F. Salmond, in *Nicene and Post-Nicene Fathers*, Second Series, Vol. 9, ed. P. Schaff and H. Wace (Buffalo: Christian Literature, 1899).

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.

Jay Kreps, Neha Narkhede, and Jun Rao. Kafka: A distributed messaging system for log processing. In *Proceedings of the Workshop on Networking Meets Databases (NetDB)*, Athens, Greece,

2011. URL https://www.semanticscholar.org/paper/Kafka-%3A-a-Distributed-Messaging-System-for-Log-Kreps-Narkhede/ea97f112c165e4da1062c30812a41afca4dab628.

F. William Lawvere. Diagonal arguments and cartesian closed categories. *Lecture Notes in Mathematics*, 92:134–145, 1969. doi: 10.1007/BFb0080769.

Hui Li et al. Recursive hierarchical networks and the law of functional evolution: A universal framework for complex systems. arXiv preprint arXiv:2509.05567, 2025.

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.

Charles Sanders Peirce. *Collected Papers of Charles Sanders Peirce*, volume 1–6. Harvard University Press, Cambridge, MA, 1931–1935.

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.

Eduardo D. Sontag and Andrew R. Teel. Changing supply functions in input/state stable systems. *IEEE Transactions on Automatic Control*, 40(8):1476–1478, 1995. doi: 10.1109/9.402246.

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.ress.2015.03.018.

Noson S. Yanofsky. A universal approach to self-referential paradoxes, incompleteness, and fixed points. *The Bulletin of Symbolic Logic*, 9(3):362–386, 2003. doi: 10.2178/bsl/1058448677.

A TGT Code Repository

The complete reference implementation of the Triadic Gaussian Tracker, including all three experimental conditions of section 9, is publicly available at:

https://github.com/volverse/tcc-tgt

The repository contains the full Python source under `src/tcc/`, benchmark scripts under `benchmarks/`, and the ELEC2 dataset loader under `data/`. All experiments use `uv` for dependency management (`uv sync`) and are reproducible with fixed random seeds. The three figures in section 9 are generated by `benchmarks/plot_all_final.py`.