

The Structuralization of Empiricism

Formalizing the Structural Conditions Under Which Empiricism Stabilizes Knowledge

Jeremy C. Jones (ORCID 0009-0007-2515-3774)—HoldingLight LLC

contact@universalcollapse.com

© 2026 | CC BY 4.0

Version: v1.0—Prepared 2026–04

Abstract

Empiricism has proven remarkably durable—shared evidence, reproducibility, and revision under disconfirmation have produced cumulative progress across centuries and disciplines. What has not been formalized is the architecture that makes those practices converge rather than drift into pseudo-consensus, frozen paradigm, or institutional inertia. This paper specifies that architecture. A minimal recursive kernel $(\Omega, K, C^K, x^*, R, U)$ formalizes what empirical practice already does; three portable stabilization signatures—redundancy-driven consensus (S_1), neutrality-induced resolution delay (S_2), and constraint-sweep hysteresis (S_3)—serve as cross-domain stress tests of empirical integrity. Update integrity names the structural condition under which records are permitted to modify constraints. Neutrality and commitment are identified not merely as epistemic virtues or personal dispositions but as structurally grounded moves with signal-state preconditions. Five Level 3 divergence claims define local failure conditions and a global portability challenge. The contribution is architectural: the framework does not replace empirical practice but specifies the conditions under which it stabilizes—conditions now explicit, pathologies now named, claims now exposed to operational challenge.

Keywords: empiricism; reproducibility; records; constraint propagation; update integrity; coherence; falsifiability; philosophy of science

1. The Problem Empiricism Solves Without Naming

A finding has been replicated across a dozen studies. The results converge. The effect appears robust. By every visible indicator, this is empirical consensus.

Now look closer. The studies share an instrument calibrated to the same standard. They draw from overlapping populations. They cite a common methodological lineage. The “independent replications” are correlated trials of the same underlying assumption. The convergence is measurable, but structurally weaker than it appears.

Empirical practice has developed tools for handling this: independence requirements, calibration audits, multi-site replication, preregistration, transparent reporting, and adversarial review. These safeguards do not always work, and they are not always applied. But when empirical inquiry converges rather than merely appears to converge, some version of this architecture is usually present. Reproducibility concerns whether the same data, methods, code, and analysis conditions yield consistent results, while replicability concerns whether separate studies collecting their own data obtain consistent results for the same scientific question (National Academies of Sciences, Engineering, and Medicine [NASEM], 2019).

What has not been made explicit is the architecture they instantiate together.

Many procedures of empirical practice—observe, measure, test, replicate—are well described in local methodological literatures: measurement, uncertainty, replication, statistical inference, and reporting. What is less often stated in one place is the shared architecture by which those procedures stabilize knowledge across observers and over time. Why does redundancy produce convergence under some conditions and only the appearance of convergence under others? When does revision genuinely refine a constraint set, and when does it merely simulate refinement? What distinguishes a stable empirical regime from a self-sealing one?

These are not metaphysical questions. They are architectural ones, and empirical practice already contains the answers in distributed form. This paper makes that architecture explicit.

We call the project the **Structuralization of Empiricism**: the articulation of the minimal recursive structure under which empirical systems generate durable, corrigible knowledge.

The thesis:

Empirical knowledge stabilizes when constraint-governed selection, durable records, and recursive update remain aligned across observers and over time.

This claim does not challenge empirical practice. It does not introduce new physical equations, override established measurements, or supersede experimental protocols. It identifies the structural conditions that empirical practice has always relied on when it genuinely converges—and, by the same architecture, what fails when convergence becomes illusory.

The contribution is architectural. What follows develops this claim: a minimal recursive kernel (§4), the role of records in grounding objectivity (§5), the operational signals by which stabilization can be detected (§6), the integrity conditions the architecture depends on (§7), and the falsifiable divergence claims that make it testable (§11).

Triad placement. This paper occupies the stabilization position in the Universal Collapse Theory (UCT) standards layer. *Records Across Nature, Life, and Mind* defines the persistence layer: what records are and

why they make collapse cumulative. The *Update Integrity Standard* (UIS v1.0) defines the operational governance layer: how update loops remain corrigible through record discipline, independence audit, and symmetric update rules. The present paper defines the middle architecture: the structural conditions under which empirical inquiry stabilizes rather than drifting into pseudo-consensus, frozen paradigm, or institutional inertia. For the broader framework, see end matter.

2. Relation to Existing Traditions

The structural question this paper asks—under what conditions does empirical inquiry stabilize rather than fragment?—has been approached from several directions. The traditions below identify pieces of the stabilization problem; SoE routes those pieces through a single record-constraint-update architecture without replacing their local accounts.

Popper and Kuhn. Popper (1959) treated falsifiability as a demarcation criterion for scientific claims; here, falsifiability is treated not as a static property of propositions but as a structural feature of the update loop—a system preserves falsifiability when U remains sensitive to records that could revise K (Section 7), a condition the divergence claims in Section 11 operationalize. Kuhn (1962) described how paradigms stabilize normal science and how anomaly accumulation produces crisis; in the terms introduced here, paradigm stability corresponds to constraint-alignment across recursive cycles, and paradigm crisis corresponds to a failure of update integrity—records accumulate that the active constraint set can no longer absorb without structural revision. Lakatos (1970) refined this picture by distinguishing progressive research programs—whose protective-belt modifications predict novel facts—from degenerating ones, whose modifications absorb anomalies ad hoc; in SoE terms, the latter signals a localized update-integrity failure in which records accumulate but only the protective periphery is revised.

Bayesian epistemology and Peirce. Bayesian epistemology provides one of the most developed formal accounts of belief updating; the recursive kernel is compatible with Bayesian updating, and the update rule $U(K, x^*, R)$ can be read as a generalized posterior revision. AGM belief-revision theory likewise formalized contraction and revision as operations on theories (Alchourrón, Gärdenfors, & Makinson, 1985); SoE does not replace that formal tradition, but relocates the update problem inside record-bearing empirical systems where record durability, independence, and update integrity determine whether revision remains corrigible. What the present framework adds is the observation that priors and background assumptions are accumulated constraint sets that can freeze, producing the self-sealing dynamics Section 7 treats as update-integrity failures. Peirce's (1878) pragmatic convergence thesis finds an operational echo in Signal S₁: redundant independent records should drive inter-observer agreement in measurable trends—not assumed, but tested.

Social epistemology and neighboring positions. Longino (1990) argued that objectivity is a function of critical interaction within communities rather than individual cognition; the present framework is compatible with this view and adds a structural mechanism—shared, independently accessible records are the interface through which critical interaction produces convergence. Goldman's veritistic social epistemology and Kitcher's account of the organization of cognitive labor reinforce the same point from the social side: epistemic success is shaped by the practices and institutions through which inquiry is organized, not by isolated individual inference alone (Goldman, 1999; Kitcher, 1990, 1993). The paper's agnosticism about what ultimately exists places it in the neighborhood of structural realism (Worrall,

1989; Ladyman & Ross, 2007) and is compatible with constructive empiricism (van Fraassen, 1980), though it commits to neither. The contribution is architectural, not metaphysical.

Thus, the framework's contribution is not to adjudicate among falsification, Bayesianism, belief revision, social epistemology, or structural realism, but to specify the stabilization architecture that empirical systems instantiate when records remain durable, constraints remain revisable, and update remains corrigible.

3. Claim Levels and Scope

The paper operates at three claim levels and one bounded scope. Keeping them distinct prevents the framework from reading as more than it claims to be.

3.1 Architectural Claims (Level 1)

These define the minimal recursive kernel: Ω , K , C^K , x^* , R , U . They describe a structural pattern present in record-bearing update systems. They do not assert ontological primacy, metaphysical finality, or universal necessity. If future work replaces this kernel with a more precise structural description, the framework updates accordingly.

3.2 Interpretive Claims (Level 2)

These reparameterize empirical practice in structural terms: records as constraint-carriers, redundancy as convergence driver, neutrality as latency source, hysteresis as memory signature. They do not overturn existing scientific definitions; they provide a unifying lens under which those definitions can be related. They are compatible with current statistical, physical, and methodological models.

3.3 Divergence Claims (Level 3)

These assert measurable stabilization signatures (S_1 – S_3). They are the empirically testable portion of the framework; Level 1 and Level 2 claims can still be rejected on grounds of incoherence, redundancy, lack of explanatory gain, or poor fit with established practice. If redundancy fails to predict convergence, if neutrality fails to alter resolution latency, if hysteresis fails under record accumulation, or if integrity interventions do not alter stabilization, the framework must be revised or constrained.

3.4 Scope Limits

This paper does not claim that coherence is a moral property, that stabilization implies truth, that all reality reduces to recursive collapse, that the framework replaces domain-specific science, or that metaphysical disputes are resolved by structuralization. It claims only that empiricism stabilizes under identifiable structural conditions. No additional primitives, theological commitments, ethical prescriptions, or cosmological extensions are introduced here; those belong to separate works and must not be inferred from this one.

3.5 Review Target

This paper does not ask the reader to accept Universal Collapse Theory as a metaphysical system. It asks whether empirical stabilization can be usefully modeled as a record-bearing constraint-update process: Ω is the space of live hypotheses or interpretations; K is the active constraint set; x^* is the selected result; R is the durable record; U is the update rule that modifies K in response to R . The framework should be

accepted provisionally only if it clarifies known empirical practices, distinguishes real convergence from pseudo-convergence, and yields operational tests beyond metaphor. It should be revised or rejected if S_1 , S_2 , and S_3 do not add predictive or diagnostic value beyond existing statistical, methodological, and social-epistemic models. Companion documents carry their own claims and should be evaluated separately: *Records Across Nature, Life, and Mind* for the cross-domain theory of records, and the *Update Integrity Standard* for the operational governance specification.

4. The Minimal Structural Kernel of Empirical Stabilization

Empiricism is a recursive constraint-refinement process. Let:

- Ω denote a structured possibility space (hypotheses, parameter values, interpretations).
- K denote the active constraint set (measurement rules, background assumptions, calibration standards, logical consistency requirements).
- C^K denote constraint-conditioned selection or realization from Ω under K .
- x^* denote the realized outcome of selection.
- R denote the record generated by x^* (data traces, instrument logs, publications, memory traces).
- U denote the update rule that modifies K in light of R .

Notation note. C^K denotes constraint-conditioned selection in general—the structural pattern by which a possibility space narrows to a realized outcome under active constraints. It is a method-level formalism, not a claim about a new microphysical collapse mechanism. The same notation appears across UCT’s physics, biology, and mind domains; here it is applied to empirical inquiry as a record-bearing update system.

Triad routing. The record term R is treated only minimally here; its full operational definition and cross-domain role are developed in *Records Across Nature, Life, and Mind*. The formal apparatus more generally is distributed across the standards triad: Records carries the persistence formalism, the *Update Integrity Standard* carries the measurement and threshold formalism, and the present paper develops the stabilization kernel that connects them.

The recursive cycle can be expressed schematically:

$$\Omega + K \rightarrow C^K \rightarrow x^* \rightarrow R \rightarrow U(K, x^*, R) \rightarrow K'$$

Empirical inquiry unfolds as repeated iteration of this loop. A possibility space Ω is constrained by K . Selection C^K yields a concrete outcome x^* . The outcome produces durable records R . Records feed back through U to refine K . The refined constraint set K' governs subsequent selection.

The formulation is intentionally minimal: it does not specify the content of Ω or K , and imposes no metaphysical stance on what ultimately exists. It formalizes what empirical practice already does—hypotheses are constrained, outcomes realized, records preserved, assumptions updated.

The structural question becomes: **under what conditions does this loop converge toward stability rather than oscillate, fragment, or self-seal?** Stability here does not mean immobility. It means constraint-alignment across iterations—later updates reduce inconsistency rather than amplify it.

Several features of empirical practice can now be recognized as structural safeguards rather than merely cultural conventions. Independence requirements prevent correlated pseudo-redundancy. Calibration ensures that records correspond to external constraints. Replication increases redundancy. Falsifiability preserves revisability in U . Preregistration constrains hypothesis flexibility within Ω . Simmons, Nelson, and Simonsohn (2011) show why this last constraint matters: undisclosed flexibility in data collection, analysis, and reporting can inflate false-positive rates, so *ex ante* rules for stopping, variables, exclusions, conditions, and covariates protect empirical update from post hoc search. Chambers (2013) developed the Registered Reports format as one operationalization of this constraint: peer review of methods and predictions occurs before data collection, so that publication decisions track the rigor of the design rather than the direction of the result—a structural mechanism for keeping U sensitive to disconfirming records.

The loop preserves stability under three conditions:

- **Constraint discipline** (K is explicit and revisable).
- **Record integrity** (R is durable and independently accessible).
- **Update corrigibility** (U allows revision rather than freezing).

When these hold, empirical systems tend toward dynamic stabilization. When they fail, instability emerges—through drift, echo chambers, pseudo-consensus, or ideological lock-in.

This structuralization does not replace empirical science. It names the conditions under which empirical science works.

5. Records and the Emergence of Objectivity

Empiricism depends on an underexamined assumption: that public records stabilize shared knowledge. The Structuralization of Empiricism makes this explicit.

In the recursive kernel, R (record) is not incidental. It is the stabilizing interface between observers.

In this account, objectivity does not require observers to be unbiased; it emerges when shared records constrain observers across recursive update cycles.

5.1 Records as Constraint-Carriers

A record is any durable trace that persists beyond the moment of selection, is independently accessible, and constrains future updates. Examples include instrument readouts, archived datasets, published results, logged experimental conditions, and replication outcomes.

Records function as constraint-carriers. Once produced, they limit the admissible interpretations in Ω for subsequent cycles. If a measurement yields x^* , and that result is redundantly recorded and accessible, the next iteration of K cannot ignore it without violating constraint discipline.

Objectivity is therefore not metaphysical neutrality. It is structural constraint propagation through records.

5.2 Redundancy and Convergence

A central empirical regularity follows: when records are independently redundant, disagreement tends to decay. This is formalized as Signal S_1 (Section 6). The opening vignette is exactly this problem in reverse: when ostensibly independent replications share a record lineage, convergence can appear stronger than its true independence warrants.

S_1 depends on independence. Correlated records—shared bias, shared instrumentation error, copied data—do not increase constraint strength; they simulate redundancy. Record redundancy must therefore satisfy two structural conditions: durability (records persist) and independence (records are not trivially correlated). When both hold, convergence is expected. When either fails, apparent consensus may be illusory.

5.3 Failure Modes: Pseudo-Objectivity

Several well-known empirical pathologies are structural degradations of R, U, or independence conditions:

- **Pseudo-redundancy:** Correlated datasets mistaken for independent confirmation.
- **Record suppression:** Selective publication or deletion of disconfirming traces.
- **Constraint freezing:** Updates U that no longer revise K despite new records.
- **Echo-chamber amplification:** High redundancy within a closed network but low external independence.

The last failure mode admits a finer-grained distinction. Nguyen (2020) separates epistemic bubbles, in which independent records are simply absent from a network, from echo chambers, in which external records are present but actively discredited. The structural intervention differs: bubbles require expanded record access, while echo chambers require restoring the sensitivity of U to disconfirming records. Both are integrity failures, but they fail at different points in the loop.

These failure modes do not invalidate empiricism. They show what happens when record integrity conditions are violated—and they are the specific pathologies the architecture predicts, not surprises.

5.4 Objectivity as Emergent Stabilization

Under the recursive kernel, objectivity is multi-agent constraint-alignment under shared record access. Observers do not need identical priors. They need access to shared, independent records that iteratively constrain divergence.

This interpretation does not alter existing empirical definitions of measurement or statistics. It explains why they function.

6. Stabilization Signals: The Coherence-State Family (S_1 – S_3)

The recursive kernel describes how empirical systems update. The question now becomes: how can stabilization be detected operationally?

The Structuralization of Empiricism introduces a small family of portable signals that function as stress-tests of constraint-alignment across domains. These are not new physical laws. They are observable stabilization signatures in record-bearing, update-capable systems.

Coherence, in this framework, is defined minimally as constraint-alignment across recursive update steps. It is not moral, teleological, or normative. It is dynamic stabilization under constraint propagation.

Three canonical signals operationalize this stabilization.

6.1 S_1 — Redundancy → Consensus

Definition. When independently generated records increase under preserved independence conditions, inter-observer disagreement decreases.

Operational proxies. Inter-rater reliability trends (e.g., Cohen’s κ), meta-analytic variance reduction, agreement rates across independent instruments, cross-laboratory reproducibility rates.

Expected directionality. As independent redundancy increases, variance between observers decreases.

Control conditions. Independence must be explicitly tested (e.g., different labs, different sampling frames); shared calibration artifacts must be controlled; correlated biases must be measured.

Falsifier. If redundancy increases but disagreement does not decrease under verified independence, S_1 fails locally.

The Open Science Collaboration’s 100-study replication project provides a concrete S_1 caution: replication effects averaged about half the magnitude of original effects, and only 36% of replications reached statistical significance despite 97% of original studies doing so, so convergence must be audited as a measured property rather than assumed from the existence of prior positive records (Open Science Collaboration, 2015).

6.2 S_2 — Neutrality → Resolution Delay

Definition. Under reduced prior constraint—when priors are weak or symmetry between competing hypotheses is high—resolution latency increases and selection C^K becomes slower or less stable.

Operational proxies. Increased decision time under balanced priors, extended convergence time in statistical estimation under uninformative priors, greater instability in ambiguous perceptual tasks, increased iteration counts before convergence in algorithmic systems.

Expected directionality. As prior constraint strength decreases, time-to-stable-selection increases.

Control conditions. Task difficulty must be controlled independently from neutrality; noise levels must be measured; stopping rules must be pre-declared.

Falsifier. If resolution speed does not change under systematically varied prior neutrality, S_2 fails locally.

S_2 reflects not weakness, but structural symmetry: when constraints do not privilege one region of Ω , stabilization takes longer.

6.3 S_3 — Constraint Sweeps → Hysteresis

Definition. When constraints are varied bidirectionally (swept up and down), system responses may exhibit path dependence. This reflects the persistence of records in K .

Operational proxies. Loop area in parameter sweeps ($\oint y \, dx$), return-path discrepancy in decision thresholds, persistence of belief states after stimulus removal, persistence of model or classification states after constraint reversal.

Expected directionality. Systems with accumulated records exhibit path-dependent stabilization.

Control conditions. External lag must be separated from intrinsic hysteresis; instrumental inertia must be measured; sweep rates must be controlled.

Falsifier. If no path dependence appears where record accumulation is expected, S_3 fails locally.

6.4 Interpreting the Coherence-State Family

The three signals describe distinct but related stabilization behaviors. S_1 captures cross-agent alignment under redundancy. S_2 captures symmetry sensitivity under weak constraints. S_3 captures memory effects under constraint variation.

Together, they form a minimal operational family for detecting stabilization in recursive empirical systems. They do not assume a particular ontology. They are intended to be tested across physical, biological, cognitive, and institutional domains, subject to the scope conditions in §9. They do not define what is true; they detect when systems converge under constraint.

The schema below compresses the operational structure of the four primary test families that support the Level 3 divergence claims: S_1 – S_3 and update integrity. Cross-domain portability (Claim 5) is assessed by comparing the performance of these tests across domains satisfying the scope conditions in §9, rather than by a separate test family. Each row gives a generic manipulation, a generic metric, a null/failure condition, and the location at which full implementation is specified. Per the triad division, SoE defines the test schema; the *Update Integrity Standard* specifies the empirical ledger and reporting standards; domain demonstration papers instantiate each test in context.

Claim	Generic manipulation	Generic metric	Null / failure condition	Full implementation
S_1: redundancy → consensus	Increase verified-independent records	Inter-rater reliability (κ), meta-analytic variance reduction, cross-laboratory reproducibility rates	Redundancy increases but disagreement does not fall under verified independence	UIS empirical ledger + domain demonstration paper
S_2: neutrality → resolution delay	Reduce prior constraint or increase symmetry among alternatives	Time-to-resolution, convergence latency, instability duration in ambiguous tasks	Prior-neutrality varies but resolution latency does not	UIS empirical ledger + perceptual/cognitive demonstration paper
S_3: sweeps → hysteresis	Sweep a constraint bidirectionally (up and down)	Loop area ($\oint y \, dx$), return-path discrepancy	Bidirectional sweep produces simple reversibility where record accumulation predicts path dependence	UIS empirical ledger + constraint-sweep / domain demonstration paper
Update integrity	Independently improve or degrade R, U, or independence conditions	Specified in UIS empirical ledger	Verified integrity change produces no predicted change in stabilization metrics	UIS empirical ledger and reporting standards

Table 1. Schema for Level 3 testable claims. Each row states a generic manipulation, a generic metric, the local null/failure condition for that claim, and the location at which full operational implementation is specified. The schema is intentionally domain-neutral; instantiated metrics, control conditions, and threshold values are supplied in the Update Integrity Standard and in domain-specific demonstration papers.

6.5 Coherence as a State Family, Not a Scalar

Coherence is not a single scalar quantity. It is a family of stabilization behaviors detectable through agreement trends (S_1), resolution latency patterns (S_2), and hysteresis signatures (S_3). Empirical systems rarely display all three signals uniformly. Section 9 states the structural preconditions under which these signals apply. This avoids overreach and preserves falsifiability.

7. Update Integrity and Corrigibility

The recursive kernel assumes that constraint sets (K) are revisable under new records (R). This revisability is not guaranteed. It depends on structural conditions.

Empiricism stabilizes not only because records exist, but because records are permitted to modify constraints. We call this condition **update integrity**.

7.1 What Is Update Integrity?

Update integrity holds when records are preserved without falsification, redundancy is genuine rather than manufactured, evidence standards are symmetric, and constraints remain revisable under legitimate disconfirmation. Formally, update integrity is the condition under which $U(K, x^*, R)$ remains sensitive to R . If U becomes insensitive to R , stabilization degenerates into rigidity rather than convergence.

This distinction is structural, not moral. A system may appear coherent internally while failing update integrity. Such systems display path-dependent persistence (S_3) without redundancy-driven convergence (S_1)—a diagnostic signature the architecture predicts.

7.2 Failure Modes Viewed From the Update Side

The failure modes introduced in Section 5.3 are each integrity failures viewed from the update side rather than the record side. If pseudo-redundancy is the opening failure mode, constraint freezing is its mirror failure: records exist, but U stops allowing them to revise K . Record corruption distorts R . Pseudo-redundancy manufactures S_1 . Constraint freezing paralyzes U . Asymmetric standards bias U 's response to R .

The point is not merely that such pathologies exist—they are well-known. The point is that they can be read as recurrent structural failures of the same recursive architecture rather than as a miscellaneous list of sociological defects.

7.3 Corrigibility as a Structural Property

Corrigibility is the capacity of a system to reduce constraint misalignment over time. A system is corrigible when disconfirming records propagate, updates modify K proportionally, independence conditions are preserved, and sweep tests reveal appropriate reversibility (bounded hysteresis).

Cross-reference. The operational specification of update integrity—audit procedures, corruption detection rules, measurement metrics, and institutional implementation guidance—is formalized in the *Update Integrity Standard* (UIS v1.0). This section provides the conceptual framing; UIS provides the operational specification.

7.4 Structuralization, Not Moralization

Update integrity is not a moral property. A system may act in good faith yet violate update integrity structurally. Conversely, structural safeguards can improve stabilization even when actors are imperfect. The Structuralization of Empiricism does not claim that science is corrupt or broken. It claims that empirical stability depends on preserving update integrity across recursive cycles.

8. Neutrality and Commitment as Structural Moves

Empiricism contains implicit decision rules about when to remain neutral and when to commit. The Structuralization of Empiricism locates both in the signal architecture.

Neutrality, structurally, is the appropriate response to symmetric or weakly-resolved constraints. When constraints do not privilege one region of Ω over another, premature commitment risks false stabilization—and the architecture predicts exactly this through S_2 : under weak or symmetrical constraints, resolution latency increases. Delayed commitment preserves update integrity. Constraint-aware neutrality is therefore not indecision; it is a structural safeguard grounded in the kernel itself.

Commitment, structurally, is the appropriate response to convergent records with bounded memory and preserved integrity. When redundant independent records converge (S_1), when neutrality-induced delay no longer changes outcome (S_2 plateau), when sweep tests show bounded hysteresis (S_3 controlled), and when update integrity remains satisfied, commitment becomes structurally warranted. Commitment here does not mean certainty. It means provisionally accepting constraint refinements into K , using K' for predictive and practical action, and maintaining revisability under future disconfirmation. It is bounded stabilization under constraint.

The deeper point is that these are not merely epistemic virtues or personal dispositions. They are also structural moves with specifiable signal-state preconditions. A system that remains neutral when S_1 has converged is as structurally mismatched as one that commits when S_2 has not plateaued. Which move is appropriate is therefore not merely dispositional. It depends on the recursive state the system is actually in.

Cross-reference. The operational decision rules—stopping criteria, pre-declared thresholds, multi-level alignment procedures, stabilization-mode classifications—are formalized in the *Update Integrity Standard* (UIS v1.0). This section establishes the structural grounds; UIS provides the protocols.

9. Scope Conditions and Domain of Validity

The framework applies to systems that satisfy three structural preconditions:

- **Structured possibility (Ω):** The system contains distinguishable alternative states or hypotheses.
- **Durable records (R):** Outcomes leave persistent traces accessible beyond the moment of realization.
- **Recursive update (U):** Constraint sets (K) are capable of modification in response to records.

Where these conditions fail, the coherence-state family (S_1 – S_3) may not apply or may require reinterpretation.

The framework may not apply, or may apply only weakly, when records are inaccessible or non-durable; when update rules are opaque or non-correctable; when independence cannot be audited; when stability is imposed by authority rather than record-sensitive update; or when the system has no meaningful possibility space beyond retrospective interpretation. These exclusions are not edge cases for the framework to absorb. They mark the structural boundary of the architecture, and the divergence claims of §11 become inapplicable rather than merely weakened in such regimes.

10. What This Is Not

Five scope limits are worth stating explicitly.

Not a replacement or an expansion. The framework does not introduce new physical equations, override established measurements, or supersede experimental protocols, and it does not propose a cosmology, metaphysics, or theory of everything. It addresses one thing: the architecture of empirical stabilization.

Not a metaphysical or moral doctrine. Coherence is treated descriptively as dynamic stabilization under constraint, not as a fundamental substance, a teleological tendency, or an ethical commandment. Ethical or metaphysical extensions may be derived elsewhere, but they are not entailed here.

Not anti-relativist or anti-constructivist. The framework does not deny that observers have priors, that social structures influence interpretation, or that context shapes constraint formation. It states only that durable records and revisable constraints enable convergence under defined conditions.

Not self-sealing. The falsification conditions are stated in Section 11. Failure of any divergence claim constrains the framework’s scope.

Not the full record theory or update-integrity standard. The cross-domain definition of records and the operational governance specification for update integrity are supplied by companion documents (*Records Across Nature, Life, and Mind* and the *Update Integrity Standard* respectively). This paper carries only the stabilization architecture between them.

11. Level 3 Divergence Claims and Falsifiers

The framework makes five Level 3 claims: statements about observable differences expected if the structural model is correct. Each is narrow and testable. Failure under controlled conditions requires revision or rejection.

11.1 Divergence Claim 1 — Redundancy as a Convergence Driver

Claim. Verified independent redundancy predicts measurable reduction in inter-observer disagreement.

This is stronger than “replication is good.” It asserts a directional effect: as independent redundancy increases, disagreement should decrease in a predictable manner.

Falsifier. If disagreement remains unchanged under verified independence and increasing redundancy, the S_1 structural claim fails locally.

11.2 Divergence Claim 2 — Neutrality-Induced Latency

Claim. Under systematically reduced prior constraint, resolution latency increases.

Falsifier. If resolution speed is invariant under controlled prior-neutrality manipulation, the S_2 claim fails locally.

11.3 Divergence Claim 3 — Record-Dependent Hysteresis

Claim. Systems that accumulate records exhibit measurable path dependence under bidirectional constraint sweeps.

Falsifier. If systems with accumulated records show no path-dependent discrepancy under controlled sweeps, the S_3 structural claim fails locally.

11.4 Divergence Claim 4 — Integrity Interventions Alter Stabilization

Claim. Interventions that independently improve or degrade R, U, or independence conditions should measurably alter stabilization outcomes in the predicted direction. Integrity-preserving interventions increase S_1 reliability, regulate S_2 appropriately, and reduce pathological S_3 lock-in; integrity-degrading interventions show the inverse pattern.

Falsifier. If independently verified integrity improvements or degradations produce no corresponding change in stabilization metrics across repeated controlled comparisons, the update-integrity claim fails locally.

11.5 Divergence Claim 5 — Cross-Domain Portability

Claim. The S_1 – S_3 stabilization family applies across domains satisfying the recursive preconditions of Section 9.

Falsifier. If the signals hold in one domain but systematically fail in another under equivalent structural conditions, portability weakens.

11.6 Global Failure Condition

The five claims above describe local falsifiers: the failure of a single signal in a single domain weakens that signal's claim, not the architecture as a whole. A global failure condition is therefore also stated. The framework's portability claim weakens if, across multiple domains satisfying the preconditions of §9 and the reporting requirements of the *Update Integrity Standard*, independently specified S_1 – S_3 tests repeatedly fail to discriminate stable convergence from pseudo-convergence, latency from noise, or hysteresis from ordinary lag better than domain-local statistical, methodological, and social-epistemic models alone. That is the acceptance/rejection target for the architecture as a portable framework.

These claims do not assert that coherence is universal, that stabilization guarantees truth, or that convergence implies metaphysical correctness. They assert that record-bearing recursive systems exhibit measurable stabilization behaviors under constraint propagation. Failure of any divergence claim constrains the framework's scope.

12. Conclusion

Empiricism has endured not because it is philosophically complete, but because it stabilizes knowledge under real-world conditions. This paper has proposed that such stabilization is not accidental. It arises from a specifiable recursive architecture: structured possibility (Ω), constraint-governed selection (C^K), durable records (R), revisable update (U), and measurable stabilization behaviors (S_1 – S_3).

When redundant independent records accumulate (S_1), neutrality delays premature resolution (S_2), constraint sweeps reveal bounded hysteresis (S_3), and update integrity is preserved, empirical inquiry stabilizes. When these conditions fail, instability emerges—as drift, echo-chamber consensus, frozen paradigm, or pseudo-redundant agreement. These pathologies are not surprises. They are the predicted failure modes of the architecture.

The contribution is architectural. The framework identifies the conditions under which empirical systems converge rather than fragment, names the pathologies that result when those conditions fail, and provides portable signals that can stress-test stabilization across domains satisfying the preconditions of Section 9. The falsification conditions are stated in Section 11.

The architecture supports one claim:

Empirical knowledge stabilizes when constraint propagation, record durability, and corrigible update remain aligned across recursive cycles.

This is not a replacement for empiricism. It is an articulation of what empirical practice has always depended on when it converges. The conditions are now specifiable. The failure modes are now named. The claims are now exposed to operational challenge.

Notation Key

The following symbols and technical terms appear throughout the paper. Definitions are given in compact form; full development is in the indicated sections.

Symbol / Term	Meaning
Ω	Structured possibility space (hypotheses, parameter values, interpretations). See §4.
K	Active constraint set (measurement rules, background assumptions, calibration standards, logical consistency requirements). See §4.
C^K	Constraint-conditioned selection from Ω under K . A method-level formalism, not a microphysical mechanism. See §4.
x^*	Realized outcome of selection. See §4.
R	Durable record generated by x^* (data traces, instrument logs, publications). See §4–§5; full cross-domain treatment in <i>Records Across Nature, Life, and Mind</i> .
U	Update rule modifying K in light of R . Read schematically as $U(K, x^*, R) \rightarrow K'$. See §4, §7.
S_1	Stabilization signature: redundancy $\rightarrow$ consensus. Verified independent redundancy predicts measurable reduction in inter-observer disagreement. See §6.1.
S_2	Stabilization signature: neutrality $\rightarrow$ resolution delay. Reduced prior constraint or symmetry among alternatives extends time-to-stable-selection. See §6.2.
S_3	Stabilization signature: constraint sweeps $\rightarrow$ hysteresis. Bidirectional constraint sweeps produce path-dependent responses where records have accumulated. See §6.3.
Update integrity	Condition under which $U(K, x^*, R)$ remains sensitive to R . Operationalized in the <i>Update Integrity Standard</i> . See §7.
Coherence	Dynamic stabilization under constraint propagation. Used descriptively, not as a moral, teleological, or normative property. See §6.
Pseudo-redundancy	Correlated records mistaken for independent confirmation. Simulates S_1 convergence without strengthening it. See §5.3.
Constraint freezing	U becomes insensitive to R despite the accumulation of new records. Stabilization degenerates into rigidity rather than convergence. See §7.

References

- Alchourrón, C. E., P. Gärdenfors, and D. Makinson. 1985. “On the logic of theory change: Partial meet contraction and revision functions.” *Journal of Symbolic Logic* 50 (2): 510–530. <https://doi.org/10.2307/2274239>
- Chambers, C. D. 2013. “Registered Reports: A new publishing initiative at Cortex.” *Cortex* 49 (3): 609–610. <https://doi.org/10.1016/j.cortex.2012.12.016>
- Goldman, A. I. 1999. *Knowledge in a Social World*. Oxford University Press.
- Jones, J. C. 2025. *Universal Collapse Theory—Foundations of Collapse* (WP01 v2.0). HoldingLight LLC. PhilArchive.
- Jones, J. C. 2026a. *Records Across Nature, Life, and Mind: The Persistence Layer of Collapse* (v2.0). HoldingLight LLC. <https://doi.org/10.17605/OSF.IO/7H6DY>
- Jones, J. C. 2026b. *Update Integrity Standard (UIS v1.0): A Structural Ethic for Preserving Corrigibility in Record-Based Systems*. HoldingLight LLC. <https://doi.org/10.17605/OSF.IO/DWM29>
- Kitcher, P. 1990. “The division of cognitive labor.” *The Journal of Philosophy* 87 (1): 5–22. <https://doi.org/10.2307/2026796>
- Kitcher, P. 1993. *The Advancement of Science: Science without Legend, Objectivity without Illusions*. Oxford University Press.
- Kuhn, T. S. 1962. *The Structure of Scientific Revolutions*. University of Chicago Press.
- Ladyman, J., and D. Ross. 2007. *Every Thing Must Go: Metaphysics Naturalized*. Oxford University Press.
- Lakatos, I. 1970. “Falsification and the methodology of scientific research programmes.” In *Criticism and the Growth of Knowledge*, edited by I. Lakatos and A. Musgrave, 91–196. Cambridge University Press.
- Longino, H. E. 1990. *Science as Social Knowledge: Values and Objectivity in Scientific Inquiry*. Princeton University Press.
- National Academies of Sciences, Engineering, and Medicine (NASEM). 2019. *Reproducibility and Replicability in Science*. Washington, DC: The National Academies Press. <https://doi.org/10.17226/25303>
- Nguyen, C. T. 2020. “Echo chambers and epistemic bubbles.” *Episteme* 17 (2): 141–161. <https://doi.org/10.1017/epi.2018.32>
- Open Science Collaboration. 2015. “Estimating the reproducibility of psychological science.” *Science* 349 (6251): aac4716. <https://doi.org/10.1126/science.aac4716>
- Peirce, C. S. 1878. “How to make our ideas clear.” *Popular Science Monthly* 12: 286–302.
- Popper, K. R. 1959. *The Logic of Scientific Discovery*. Hutchinson.
- Simmons, J. P., L. D. Nelson, and U. Simonsohn. 2011. “False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant.” *Psychological Science* 22 (11): 1359–1366. <https://doi.org/10.1177/0956797611417632>

van Fraassen, B. C. 1980. *The Scientific Image*. Oxford University Press.

Worrall, J. 1989. “Structural realism: The best of both worlds?” *Dialectica* 43 (1–2): 99–124.

AI Disclosure. AI tools were used to assist with manuscript preparation. The underlying theory, arguments, and interpretive claims are the author’s own, and the author takes full responsibility for the content.

Citation: Jones, J. C. (2026). The Structuralization of Empiricism: Formalizing the Structural Conditions Under Which Empiricism Stabilizes Knowledge. HoldingLight LLC.

Series: Universal Collapse Theory — Standards Layer

Companions: Records Across Nature, Life, and Mind; Update Integrity Standard (UIS v1.0)

© 2026 Jeremy C. Jones — HoldingLight LLC