Discrete Divergence, Structural Flux, and the Conservation of τ-Closure

UNNS Research Collective

Abstract

We investigate whether τ-closure, previously identified as the criterion for structural persistence under collapse, admits a conservation principle at the level of refinement evolution. Rather than treating closure as a terminal or purely classificatory property, we ask whether τ-closure constrains admissible refinement paths through a discrete continuity condition.

To this end, we introduce a discrete divergence operator defined on the refinement graph of a mechanism and formalize the notion of structural flux using admissible, non-numerical flux carriers. These carriers encode the transport of closure-relevant structure across refinement steps and are invariant under admissible renaming and canonicalization. We define τ-conservation as the requirement that all τ-closed structural states be divergence-free with respect to such a flux.

Finally, we construct a minimal structural configuration—a nontrivial recurrent refinement orbit admitting a strict closure-relevant leak—that falsifies τ-conservation if closure is attributed to that orbit. This counterexample isolates the precise condition under which τ-closure fails to function as a conservation principle.

The results clarify the status of τ-closure as either a purely terminal criterion or a law-like constraint on refinement evolution, without introducing new closure mechanisms or extending the τ-basis.

1 Introduction

In prior work, τ-closure was identified as the criterion by which recursive structures persist under refinement and survive collapse. Closure, in this sense, is not attributed to numerical convergence or stability, but to the satisfaction of irreducible structural constraints that remain invariant under admissible refinement. While this characterization establishes which structures are preserved, it does not address whether τ-closure merely classifies terminal outcomes or whether it constrains the admissible refinement evolution itself.

The present work examines this distinction by asking whether τ-closure admits a conservation principle: namely, whether closure persistence can be expressed as a discrete continuity condition governing the transport of closure-relevant structure across refinement steps. This question is structural rather than dynamical and concerns refinement evolution prior to collapse, without introducing new closure mechanisms or altering the established τ-basis.

Terminological Remark (Conservation)

Throughout this paper, the term *conservation* is used in a strictly discrete and structural sense. It does not refer to numerical invariance, physical quantities, temporal evolution, or continuous fields. Rather, conservation denotes the absence of net structural divergence in the refinement graph of a mechanism, expressed via a discrete divergence operator acting on admissible refinement steps.

Accordingly, a conserved quantity is not an evaluated magnitude but a bookkeeping condition on the transport of closure-relevant structure under refinement. All conservation statements in this work are invariant under admissible renaming and canonicalization and are formulated entirely at the mechanism level.

Non-Extension Statement

The introduction of a discrete divergence operator and structural flux formalism does not extend, modify, or enlarge the set of admissible closure mechanisms in the τ-regime. In particular, no new τ-invariants are proposed, and no additional closure classes are introduced.

The conservation principle examined here is not inferred from the numerical behavior of known constants, nor is it used to retroactively justify closure persistence. Instead, conservation is tested as an independent structural constraint on refinement evolution. Failure of conservation does not invalidate the classification of τ-closure established in previous work; it clarifies its scope by distinguishing terminal persistence from law-like constraint.

This separation prevents circularity by ensuring that closure classification is not derived from conservation assumptions, and conservation is not asserted on the basis of observed closure alone.

2 Preliminaries (Mechanisms and Refinement)

Definition 1 (Mechanism). *A* mechanism *is a tuple* $M = (\Sigma, R, C, O)$ *where* Σ *is a finite alphabet,* R *is a set of rewrite/generation rules on* Σ*,* C *is a set of admissibility constraints, and* O *specifies rule ordering/interaction.*

Definition 2 (Structural states and refinement sequence). *Let* $\mathcal{S}(M)$ *be the space of* structural states *generated by* M*, modulo admissible renaming/canonicalization. A* refinement sequence *is*

$$S_0 \to S_1 \to S_2 \to \cdots$$

where each arrow is induced by a refinement operation (e.g. deeper generation, tighter constraints, stricter rule application).

Definition 3 (τ-closure predicate (abstract)). *Fix a closure predicate*

$$\mathsf{Cl}_\tau : \mathcal{S}(M) \to \{0, 1\},$$

where $\mathsf{Cl}_\tau(S) = 1$ *means "S is τ-closed" (i.e. satisfies the relevant closure constraint) and* $\mathsf{Cl}_\tau(S) = 0$ *means "not τ-closed." This is* structural*, not numerical.*

3 The Refinement Graph and Discrete Divergence

The refinement evolution can be treated as a directed graph (or directed multigraph) whose edges are admissible refinement steps.

Definition 4 (Refinement graph). *Given a mechanism* M*, define the directed graph*

$$G(M) = (V, E),$$

where $V \subseteq \mathcal{S}(M)$ *is the set of reachable structural states, and* $(S \to T) \in E$ *iff there exists an admissible refinement step sending* S *to* T*.*

Chains and the boundary operator

We define a purely combinatorial chain complex on $G(M)$.

Definition 5 (0- and 1-chains). *Let* $\mathbb{A}$ *be an abelian group (the* flux coefficient group*; chosen below). Define the free* $\mathbb{A}$*-modules:*

$$C_0(G; \mathbb{A}) := \left\{ \sum_{v \in V} a_v \, v \ : \ a_v \in \mathbb{A}, \ \textit{finite sum} \right\}, \quad C_1(G; \mathbb{A}) := \left\{ \sum_{e \in E} b_e \, e \ : \ b_e \in \mathbb{A}, \ \textit{finite sum} \right\}.$$

Definition 6 (Boundary map)**.** *For a directed edge* $e = (u \to v)$ *define*

$$\partial(e) = v - u,$$

and extend $\mathbb{A}$*-linearly to* $\partial : C_1(G; \mathbb{A}) \to C_0(G; \mathbb{A})$*.*

Flux and divergence

Definition 7 (Flux field)**.** *A* flux field *is a function*

$$J : E \to \mathbb{A}.$$

Equivalently, it is a 1-cochain on edges with values in $\mathbb{A}$*.*

Definition 8 (Discrete divergence at a vertex)**.** *Given a flux field* J*, define the* divergence *at a vertex* $v \in V$ *by*

$$(\nabla \cdot J)(v) := \sum_{e=(v \to w)} J(e) \; - \sum_{e=(u \to v)} J(e),$$

where the first sum ranges over edges leaving v *and the second over edges entering* v*.*

Lemma 1 (Coordinate-free form)**.** *If we identify* J *with an element of* $C_1(G; \mathbb{A})$ *via* $J = \sum_{e \in E} J(e)\, e$*, then*

$$\nabla \cdot J = \partial(J) \in C_0(G; \mathbb{A}),$$

with evaluation at vertices matching the signed incidence formula above.

4 Admissible Flux Carriers (Mechanism-Level, Non-Numerical)

The key structural requirement is that "flux" must not be a numerical conserved quantity. It must be carried by *structural features that are invariant under admissible renaming/canonicalization* and are defined *at the mechanism level.*

4.1 Choosing the coefficient group $\mathbb{A}$

We now specify admissible choices for $\mathbb{A}$ and what $J(e)$ is allowed to encode.

Definition 9 (Flux carrier type)**.** *A* flux carrier type *is a functorial assignment*

$$\mathcal{F} : \mathcal{S}(M) \to \mathsf{Obj}(\mathsf{Ab}),$$

mapping each structural state to an abelian group of structural charges*, together with edge maps along refinement steps. In the present framework we use a fixed global carrier group $\mathbb{A}$ that all edges draw from (a coarse, universal charge space).*

Definition 10 (Admissible flux coefficient group)**.** *An abelian group $\mathbb{A}$ is* admissible *if its generators correspond to* structural invariants *(not evaluated numbers) that are:*

- *invariant under admissible renaming/canonicalization,*
- *definable pre-numerically (before any evaluation),*
- *stable under refinement bookkeeping (so they can be transported along edges).*

Canonical examples of admissible $\mathbb{A}$ include:

1. ***Signature group.*** *$\mathbb{A} = \mathbb{Z}[\mathcal{T}]$ where $\mathcal{T} = \{O, R, P\}$ are closure-mechanism tags (orthogonal/relaxation/projection), and $\mathbb{Z}[\mathcal{T}]$ is the free abelian group on these tags.*
2. ***Defect-type group.*** *$\mathbb{A} = \mathbb{Z}[\mathcal{D}]$ where $\mathcal{D}$ indexes* structural defect types *(e.g. constraint-violation classes, rule-incompatibility classes), again as formal symbols rather than real numbers.*
3. ***Typed multiset group.*** *$\mathbb{A}$ the Grothendieck group of finite multisets of typed motifs (rewrite motifs, constraint motifs) modulo admissible equivalence.*

4.2 Mechanism-level flux along refinement edges

A refinement step changes the mechanism's *admissible structural content.* Flux should quantify the *structural transfer* of closure-relevant content across that step.

Definition 11 (Admissible flux assignment)**.** *Let $e = (S \to T) \in E$. An* admissible flux assignment *is a map $J : E \to \mathbb{A}$ such that $J(e)$ depends only on:*

- *the pair of structural equivalence classes (S, T),*
- *the refinement action type (deeper generation, tightened constraint, stricter ordering),*
- *and the induced change in closure-relevant structural invariants,*

and is invariant under admissible renaming/canonicalization of symbols in S and T.

Interpretation. Think of $J(e)$ as the *transport* of closure-relevant structure under refinement. For instance, in the signature group case, $J(e)$ may encode that a refinement step introduces/removes a contribution of type O, R, or P to the closure bookkeeping. No numeric magnitudes are required.

5 A Formal τ-Conservation Principle

We can now state a conservation law purely structurally.

Definition 12 (τ-conservation (discrete form))**.** *A mechanism M satisfies* τ-conservation *if there exists an admissible coefficient group $\mathbb{A}$ and an admissible flux field $J : E \to \mathbb{A}$ such that:*

1. *(**Local conservation**) For every vertex $v \in V$ with $\mathsf{Cl}_\tau(v) = 1$,*
$$(\nabla \cdot J)(v) = 0.$$
2. *(**No closure without conservation**) If $(\nabla \cdot J)(v) \neq 0$ then $\mathsf{Cl}_\tau(v) = 0$. Equivalently, $\mathsf{Cl}_\tau(v) = 1 \Rightarrow (\nabla \cdot J)(v) = 0$ is taken as a hard constraint.*
3. *(**Refinement covariance**) If two refinement steps are equivalent under admissible renaming/canonicalization, they carry the same flux value in $\mathbb{A}$.*

Proposition 1 (Interpretation of the Conservation Test)**.** *The conservation test succeeds if there exists at least one admissible carrier group $\mathbb{A}$ and admissible flux J for which τ-closed vertices are exactly the divergence-free vertices:*
$$\mathsf{Cl}_\tau(v) = 1 \iff (\nabla \cdot J)(v) = 0.$$

6 Minimal Counterexample that Falsifies τ-Conservation

A falsifier must show that *closure can occur without a divergence-free law*, or that *divergence-free flow can occur without closure.* A falsifier must show that *closure can occur without a divergence-free law*, or that *divergence-free flow can occur without closure.* Your roadmap demands the stronger, cleaner falsifier:

There exists a τ-closed state v such that for every admissible flux field J,
$$(\nabla\cdot J)(v) \neq 0.$$

We now construct the smallest refinement graph pattern that forces this failure.

6.1 The minimal graph pattern

Consider the directed graph with three vertices $\{A, B, C\}$ and edges

$$A \to B, \quad B \to A, \quad A \to C.$$

This is a 2-cycle ($A \leftrightarrow B$) with a single refinement "leak" from A to a new state C.

Definition 13 (Meta-stable two-cycle with a leak)**.** *A refinement graph contains a* meta-stable two-cycle with a leak *if there exist distinct states A, B, C such that*

$$A \to B, \;\; B \to A, \;\; A \to C \;\; \textit{are admissible refinement edges,}$$

and additionally the refinement step $A \to C$ corresponds to a strict tightening that destroys the cycle at the next refinement level (so C is not equivalent to A or B).

6.2 Why this falsifies conservation (structural argument)

Assume $\mathsf{Cl}_\tau(A) = 1$ due to the existence of a nontrivial recurrent orbit $A \leftrightarrow B$ under finite refinement, while deeper refinement admits the strict leak $A \to C$

Now fix any admissible coefficient group $\mathbb{A}$ and any admissible flux field $J : E \to \mathbb{A}$. Compute divergence at A:

$$(\nabla\cdot J)(A) = \underbrace{J(A \to B) + J(A \to C)}_{\text{outgoing}} - \underbrace{J(B \to A)}_{\text{incoming}}.$$

Because $A \leftrightarrow B$ is a cycle, the only way to have local conservation on the cycle itself is to balance $J(A \to B)$ with $J(B \to A)$ (up to the chosen orientation). But the extra outgoing edge $A \to C$ forces an uncompensated term unless there exists an incoming edge into A carrying exactly the same charge as $A \to C$ with opposite sign.

In the *minimal* pattern, there is no additional incoming edge into A besides $B \to A$. Therefore, for $(\nabla\cdot J)(A) = 0$ to hold, we must have

$$J(A \to C) = 0 \quad \text{in } \mathbb{A}.$$

But $A \to C$ is, by construction, a strict refinement leak that changes closure-relevant structure. Admissibility requires $J(A \to C)$ to encode that change, i.e. $J(A \to C) \neq 0$. Hence $(\nabla\cdot J)(A) \neq 0$ for all admissible J.

Proposition 2 (Minimal falsifier)**.** *If there exists a τ-closed (or proto-closed-certified) structural state A that participates in a nontrivial recurrent orbit but admits a strict refinement leak $A \to C$ that is closure-relevant, then τ-conservation fails: there is no admissible flux field J with*

$$\mathsf{Cl}_\tau(A) = 1 \;\Rightarrow\; (\nabla\cdot J)(A) = 0.$$

6.3 A concrete mechanism-level instantiation (no numbers)

We now realize the minimal pattern using a mechanism with:

- an alternation rule producing a 2-cycle under mild constraints,
- and a stricter refinement constraint that admits a leak to a new state.

Definition 14 (Toy mechanism M_{leak})**.** *Let $\Sigma = \{a, b\}$ and rules R:*

$$r_1:\ a \Rightarrow b, \qquad r_2:\ b \Rightarrow a, \qquad r_3:\ a \Rightarrow aa.$$

Ordering O is "apply exactly one rule per refinement step."

Constraints are refined in two stages:

- *At coarse refinement, $C^{(0)}$ forbids r_3 (so only r_1, r_2 are admissible). This yields the recurrent orbit $a \leftrightarrow b$ (states $A = [a]$, $B = [b]$).*
- *At deeper refinement, $C^{(1)}$ is tightened so that r_3 becomes admissible for a whenever a particular structural predicate holds (e.g. a canonicalization tag is present). This introduces an admissible edge $A \to C$ where $C = [aa]$ is not equivalent to A or B under the equivalence rules.*

Graph realization. Under $C^{(0)}$, $A \to B$ and $B \to A$ exist and form a nontrivial cycle. Under $C^{(1)}$, the additional admissible step $A \to C$ exists, yielding the minimal "leak".

Why it is closure-relevant. The step $A \to C$ changes the generative structure class (it is not a renaming, not a canonicalization, and not equivalent by admissible tolerances), hence any admissible flux carrier that records closure-relevant change must assign $J(A \to C) \neq 0$.

Therefore A cannot be both τ-closed and divergence-free in any admissible flux model on this graph.

6.4 Empirical Validation via Structural Instrumentation

To assess whether the structural conservation principle formulated in this paper is operationally meaningful, we implemented a discrete evaluation instrument that realizes the refinement graph, admissible flux assignment, discrete divergence operator, and falsifier detection exactly as specified in Sections **??**–**??**.

The instrument operates entirely at the mechanism level. Its inputs are mechanisms $M = (\Sigma, R, C, O)$ together with admissible refinement traces; its outputs are refinement graphs, admissible structural flux assignments, per-vertex discrete divergence values, and (when present) explicit falsifying witnesses. No numerical quantities, continuous dynamics, time parameters, or probabilistic weights are introduced at any stage.

Minimal falsifier instantiation. We instantiated the minimal counterexample described in Section **??** as a concrete mechanism exhibiting (i) a nontrivial recurrent two-cycle under coarse refinement, and (ii) a strict, closure-relevant refinement step (a "leak") admitted only under tightened admissibility constraints. This mechanism satisfies the τ-closure predicate Cl_τ, yet admits an unavoidable imbalance in any admissible structural flux assignment.

Running the instrument on this mechanism yields a refinement graph with three structural equivalence classes and eleven admissible refinement edges. The computed discrete divergence is nonzero at the cycle source and at the leak sink, while remaining zero at the complementary cycle vertex. The instrument detects exactly one minimal "two-cycle with leak" witness, and the conservation test is falsified precisely at that vertex.

Control mechanisms. To verify that the instrument does not trivially falsify conservation, we evaluated several control mechanisms. A pure recurrent cycle without refinement leaks yields identically zero divergence and no witnesses, and therefore satisfies conservation. Mechanisms lacking nontrivial recurrent refinement behavior fail the closure predicate and are excluded from conservation testing. These controls confirm that the instrument discriminates conservative from non-conservative proto-closure rather than rejecting closure indiscriminately.

Result. The empirical outcome matches the theoretical prediction exactly: proto-closure does not imply a conservation principle. The falsifier identified in this paper is not hypothetical but operational, and conservation fails if and only if the structural conditions identified in Section **??** are realized.

All evaluation artifacts, including refinement traces, refinement graphs, divergence reports, and witness data, are publicly available and linked alongside this paper, providing a complete and reproducible record of the validation.

Mechanism	Cl_τ	Conservation	Witness
Pure cycle	1	satisfied	none
Balanced cycle	1	satisfied	none
Two-cycle with leak	1	falsified	detected
Non-recurrent	0	n/a	n/a

7 Operational Summary (Structural Evaluation Procedure)

To evaluate τ-conservation in practice:

1. Construct $G(M)$ from refinement traces (states as equivalence classes, edges as admissible refinements).
2. Choose an admissible carrier group $\mathbb{A}$ (e.g. signature group $\mathbb{Z}[O, R, P]$).
3. Define an admissible $J : E \to \mathbb{A}$ from rule/constraint deltas (purely structural).
4. Evaluate $(\nabla \cdot J)(v)$ at states marked τ-closed by your closure predicate.
5. **Falsification criterion:** detect a minimal "two-cycle with leak" at any state labeled closed. If such a pattern occurs with $J(A \to C) \neq 0$, conservation fails.

A Separation of Closure Classification and Conservation Testing

This appendix clarifies the logical separation between structural closure classification and conservation testing within the UNNS substrate. The distinction is not methodological but structural and is enforced by a minimal falsifier identified in the main text.

A.1 Closure Classification (Established Results)

Previous work established a binary classification of mechanisms under refinement: *PROTO-CLOSED* versus *STRUCTURAL-COLLAPSE*. This classification is determined entirely at the mechanism level and is independent of numerical evaluation.

In particular:

- Mechanisms are evaluated as tuples $M = (\Sigma, R, C, O)$ rather than as numerical realizations.
- Structural persistence is decided by the existence of nontrivial recurrent behavior under refinement, modulo admissible equivalence.
- Collapse is terminal and eliminative: once a mechanism fails to sustain closure under refinement, it cannot be reinstated.

This classification exhausts the question of survival under collapse and admits no intermediate or graded verdicts.

A.2 Independence of Conservation from Closure

The present work establishes that structural closure does not, by itself, imply a conservation principle governing refinement evolution. A mechanism may satisfy the proto-closure criterion while admitting net structural divergence at closed states.

The minimal falsifier is a purely structural configuration consisting of:

- a nontrivial recurrent refinement orbit, and
- a strict closure-relevant refinement leak that destroys recurrence at deeper refinement.

In such a configuration, every admissible assignment of structural flux yields nonzero discrete divergence at the closed state. Therefore, no divergence-free refinement law can be inferred from closure persistence alone.

This demonstrates that conservation must be tested independently and cannot be assumed as a consequence of closure classification.

A.3 Operational Conservation Test

Conservation testing proceeds as follows:

1. Construct the refinement graph $G(M) = (V, E)$ from admissible refinement traces, with vertices identified as structural equivalence classes.
2. Select an admissible flux carrier group $\mathbb{A}$ whose generators represent closure-relevant structural invariants and are invariant under renaming and canonicalization.
3. Define an admissible structural flux assignment $J : E \to \mathbb{A}$ from rule and constraint deltas.
4. Evaluate the discrete divergence $(\nabla \cdot J)(v)$ at vertices v classified as τ-closed.
5. Apply falsification by detecting minimal cycle–leak patterns at such vertices.

A mechanism satisfies conservation only if there exists at least one admissible flux assignment for which all τ-closed vertices are divergence-free.

A.4 Classification and Tagging

To preserve the binary verdict discipline of closure classification, conservation outcomes are treated as orthogonal diagnostic tags rather than new ontological classes.

- **Closure verdict:** $\mathsf{Cl}_\tau(v) = 1$ (PROTO-CLOSED) or $\mathsf{Cl}_\tau(v) = 0$ (STRUCTURAL-COLLAPSE).
- **Conservation tag (defined only when $\mathsf{Cl}_\tau(v) = 1$):**
  - *Conservative*: $(\nabla \cdot J)(v) = 0$,
  - *Non-conservative*: $(\nabla \cdot J)(v) \neq 0$.

No conservation outcome alters the closure verdict. In particular, failure of conservation does not reinstate collapse or invalidate proto-closure.

A.5 Design Constraints for Conservation Evaluation

Any implementation of conservation testing must satisfy the following constraints:

- Flux carriers must be non-numerical and defined pre-evaluation.
- Flux assignments must be invariant under admissible renaming and canonicalization.
- Refinement graphs must distinguish recurrent cycles from strict refinement leaks.
- Numerical magnitudes, time evolution, and continuous dynamics are inadmissible.

These constraints ensure compatibility with collapse monotonicity and prevent conservation from being misinterpreted as a dynamical or physical law.

A.6 Structural Implication

The separation demonstrated here establishes that closure persistence and conservation are independent structural properties. Closure classification determines which mechanisms survive collapse, while conservation testing determines whether refinement evolution is subject to a divergence-free structural constraint.

The minimal falsifier shows that this separation is fundamental and not an artifact of evaluation procedure.