The Dynamic Completion of the UNNS Substrate

From Structural Existence to Structural Dynamics

UNNS Research Collective

Abstract

The Unbounded Nested Number Sequences (UNNS) Substrate has previously established a framework for structural existence based on recursive generability, consistency, and τ-closure. These results identify which structures *can survive* infinite refinement. However, survival alone does not explain which structures are physically realized, nor how structures move, interact, or stabilize within the substrate.

This paper introduces the *Dynamic Completion of UNNS*, extending the theory from a taxonomy of admissible structures to a dynamical framework governed by variational principles. We formalize the Principle of Least Divergence as a substrate-level selection law, introduce a minimal refinement metric defining structural trajectories, and establish a correspondence between refinement symmetries and conserved quantities. This culminates in the *UNNS–Noether Correspondence*, a discrete analogue of Noether's theorem for recursive substrates.

1 Introduction

The UNNS Substrate was developed to answer a foundational question:

> What does it mean for a mathematical structure to *exist* independently of representation, construction, or observation?

Earlier work established that existence in the UNNS Substrate is not determined by provability, syntactic definability, or numerical convergence, but by *structural survivability under refinement*. This led to the identification of three irreducible structural stages:

- Generability (Φ)
- Consistency (Ψ)
- τ-closure

A structure is admissible in the UNNS Substrate if and only if it survives these filters. However, this framework remains fundamentally *static*.

2 Structural Flux and Divergence

To construct a dynamical theory, we require a quantity that measures how structure is redistributed during refinement. This role is played by *structural flux*.

2.1 Structural Mass

We introduce a scalar functional

$$m : \mathcal{S} \to \mathbb{R}_{\geq 0},$$

called *structural mass*, which measures the amount of coherent structure present in a state.

Interpretation. Structural mass is not numerical magnitude, probability, or energy. It is a substrate-internal measure of organizational integrity. Informally:

- High $m(S)$ indicates strong closure, consistency, and low defect density
- Low $m(S)$ indicates fragmentation, leakage, or instability

Minimal requirements. The functional m is required to satisfy only:

1. **Monotonicity under collapse**: Operator XII does not increase m
2. **Boundedness**: $m(S)$ is finite for all admissible S
3. **Path sensitivity**: m may vary along refinement paths even if τ-closure holds

No further assumptions are imposed at this stage.

2.2 Structural Flux

Consider a refinement step

$$S \xrightarrow{\mathcal{R}} S'.$$

[Structural Flux] The structural flux along a refinement edge is defined as

$$J(S \to S') := m(S') - m(S).$$

Flux measures the local gain or loss of structural mass during refinement. Positive flux corresponds to consolidation; negative flux corresponds to leakage.

2.3 Divergence along Refinement Paths

Let

$$P = (S_0 \to S_1 \to \cdots \to S_n)$$

be a finite refinement path.

[Path Divergence] The divergence accumulated along P is defined as

$$\mathrm{Div}(P) := \sum_{k=0}^{n-1} |m(S_{k+1}) - m(S_k)| - |m(S_n) - m(S_0)|\,.$$

Meaning.

- The first term measures total structural activity
- The second term measures net structural change
- Their difference captures internal cancellation, oscillation, and leakage

2.4 Candidate Structural Mass Functions

The structural mass functional $m(S)$ was introduced axiomatically in Section 2.1. For computational exploration and empirical validation, we now propose several *candidate* definitions. These are not asserted as canonical, but as testable instantiations satisfying the required axioms.

Candidate 1: Consistency Fraction. Let $\mathcal{R}_{\text{syn}}(S)$ denote the set of all syntactically admissible refinement edits, and $\mathcal{R}_{\text{adm}}(S) \subset \mathcal{R}_{\text{syn}}(S)$ those preserving consistency. Define

$$m_1(S) := \frac{|\mathcal{R}_{\text{adm}}(S)|}{|\mathcal{R}_{\text{syn}}(S)|}.$$

This measures the fraction of refinement directions that preserve structural consistency.

Candidate 2: Closure Stability. Let $\delta_\tau(S, n)$ denote the cumulative defect growth under τ-refinement up to depth n. Define

$$m_2(S) := \exp\left(- \limsup_{n \to N} \frac{\delta_\tau(S, n)}{n} \right),$$

for fixed finite cutoff N in computational settings.

This captures resistance to closure defect accumulation.

Candidate 3: Operator Stability Eigenvalue. Let $\mathcal{T}_\tau$ denote the τ-closure operator restricted to admissible states reachable from S. Define

$$m_3(S) := \lambda_{\max}(\mathcal{T}_\tau),$$

the maximal eigenvalue governing persistence under refinement.

Remark. All three candidates are bounded, nonnegative, and monotone under collapse. Agreement or disagreement among these definitions under refinement is itself an empirical signal and does not invalidate the theory.

A perfectly conservative refinement has $\text{Div}(P) = 0$.

3 The Least Divergence Variational Principle

We now state the central dynamical principle.

[Least Divergence] Among all admissible refinement trajectories connecting structural states, physically stable structures are those that minimize expected divergence across their refinement path ensembles.

This is a variational principle: stability is not defined pointwise, but as a global property of refinement behavior.

3.1 Ensemble Formulation

Let $\mathcal{P}(S)$ denote the set of admissible refinement paths originating from state S under allowed perturbations.

[Expected Divergence] The expected divergence of S is

$$\langle \mathrm{Div} \rangle(S) := \mathbb{E}_{P \in \mathcal{P}(S)} \left[\mathrm{Div}(P) \right].$$

[Stable Structure] An admissible structure S is *stable* if

$$\langle \mathrm{Div} \rangle(S)$$

is locally minimal with respect to perturbations of refinement rules, ordering, or coupling.

Stability is therefore a property of *behavior under refinement*, not of static form.

3.2 Probability Measure on Refinement Paths

The definition of expected divergence requires a probability measure on the space of admissible refinement paths. At the foundational level, we adopt a minimal and explicit assumption.

[Admissible Path Measure] At each refinement step, admissible primitive edits are selected according to a uniform probability distribution conditioned on admissibility, bounded refinement depth, and coupling constraints.

This induces a product measure on finite refinement paths.

Scope. The uniform measure is not claimed to be unique or physically privileged. It serves as a neutral baseline sufficient to define expected divergence and stability.

Alternative measures (e.g. weighted by edit cost, historical frequency, or environmental constraints) may be explored in computational validation and do not alter the structural results established here.

4 Compatibility with τ-closure

We now show that the Least Divergence Principle is consistent with, and strictly stronger than, τ-closure.

[Admissibility Precondition] If a structure S is stable under the Least Divergence Principle, then S is τ-closed.

Proof. Suppose S is not τ-closed. Then there exists a refinement sequence along which structural defects grow without bound. By boundedness of m, this implies persistent negative flux and therefore unbounded divergence along admissible paths.

Hence $\langle \mathrm{Div} \rangle(S)$ cannot be minimal. Contradiction. □

[Non-Equivalence] There exist τ-closed structures that are not stable under the Least Divergence Principle.

Proof. τ-closure guarantees survival but does not constrain oscillatory or leaking behavior during refinement. Structures may return to an invariant form only asymptotically while exhibiting high intermediate flux variation.

Such structures have nonzero expected divergence and therefore fail the variational criterion. □

These results establish the logical hierarchy:

$$\text{Stability} \;\Rightarrow\; \tau\text{-closure} \;\Rightarrow\; \text{admissibility}.$$

5 Perfectly Conservative Invariants

[Perfectly Conservative Invariant] A structure S is called a perfectly conservative invariant if

$$\mathrm{Div}(P) = 0 \quad \text{for all admissible refinement paths } P.$$

Such structures exhibit:

- Zero structural leakage
- Path-independent refinement behavior
- Maximal robustness under perturbation

These invariants represent extrema of the Least Divergence Principle and are natural candidates for universally observed constants.

The identification and classification of such invariants is deferred to later sections and computational validation.

6 A Minimal Refinement Metric

To move from admissibility and stability to dynamics, refinement must be endowed with a notion of distance. Without such a metric, refinement remains a sequence of discrete steps without geometry, and concepts such as velocity, curvature, or optimality cannot be defined.

The goal of this section is to introduce a *minimal* refinement metric: sufficient to define trajectories and geodesics, yet weak enough to avoid overfitting or embedding physical assumptions prematurely.

6.1 Structural State Space

Recall that each structural state is defined by a mechanism

$$M = (\Sigma, R, C, O),$$

where Σ is a symbolic alphabet, R a rule set, C a constraint system, and O a set of admissible operators.

Let $\mathcal{S}$ denote the set of admissible structural states.

We emphasize that $\mathcal{S}$ is not a numerical space. Distance must therefore be defined in terms of *structural effort*, not numerical proximity.

6.2 Primitive Refinement Edits

We define a finite set of *primitive refinement edits* $\mathcal{E}$, each representing an elementary transformation of a mechanism. Examples include:

- modification of a rule in R
- introduction or removal of a constraint in C
- extension or restriction of the alphabet Σ
- application or removal of an operator in O

Each edit $e \in \mathcal{E}$ is assigned a nonnegative cost $c(e) > 0$.

Minimality requirement. Costs are chosen to reflect *structural effort*, not expressive power. No edit is allowed zero cost, preventing trivialization of distance.

6.3 Definition of the Refinement Metric

[Refinement Distance] Let $S, T \in \mathcal{S}$. The refinement distance between S and T is defined as

$$d(S,T) := \inf \left\{ \sum_{k=1}^{n} c(e_k) \;\middle|\; e_1 \circ e_2 \circ \cdots \circ e_n(S) = T \right\},$$

where the infimum is taken over all finite sequences of primitive edits that transform S into T.

This definition parallels edit-distance metrics, but operates on mechanisms rather than syntactic strings.

6.4 Metric Properties

The function $d : \mathcal{S} \times \mathcal{S} \to \mathbb{R}_{\geq 0}$ satisfies the axioms of a metric on $\mathcal{S}$.

Proof.

- *Non-negativity*: follows from $c(e) > 0$
- *Identity*: $d(S,T) = 0$ implies no edits are required, hence $S = T$
- *Symmetry*: edits can be inverted with equal cost by definition
- *Triangle inequality*: concatenation of edit sequences yields $d(S,U) \leq d(S,T) + d(T,U)$

□

6.5 Refinement Trajectories

A refinement sequence

$$S_0 \to S_1 \to \cdots \to S_n$$

now defines a discrete trajectory in $(\mathcal{S}, d)$.

[Refinement Velocity] The refinement velocity at step k is defined as

$$v_k := d(S_k, S_{k+1}).$$

[Refinement Acceleration] The refinement acceleration at step k is defined as

$$a_k := v_{k+1} - v_k.$$

These quantities characterize how rapidly a structure moves through refinement space and how that motion changes under perturbation.

6.6 Default Edit Cost Scheme

The refinement quasi-metric depends on the cost function $c(e)$ assigned to primitive edits. For computational work, we adopt a minimal default scheme.

- Symbol addition or removal:

$$c(\text{add\_symbol}) = 1, \quad c(\text{remove\_symbol}) = 1 + \nu,$$

where ν counts induced constraint violations.

- Rule modification:
$$c(\text{modify\_rule}) = d_{\text{struct}}(R, R'),$$
where d_{struct} measures rule-complexity difference.

- Operator application:
$$c(\text{apply\_operator}) = w(O),$$
with $w(O)$ a fixed operator complexity weight.

Scope. This scheme is not unique. Its role is to provide a neutral baseline for geodesic and divergence computations. Structural conclusions are required to be robust across reasonable variations of $c(e)$.

7 Structural Entanglement

Thus far, refinement dynamics have been defined for isolated mechanisms. However, physically relevant structures rarely exist in isolation. Interaction arises when multiple mechanisms share symbolic resources, rules, or operators.

This section introduces *structural entanglement*: the coupling of refinement dynamics between mechanisms within the UNNS Substrate.

7.1 Coupled Mechanisms

Let
$$M_A = (\Sigma_A, R_A, C_A, O_A), \quad M_B = (\Sigma_B, R_B, C_B, O_B)$$
be two admissible mechanisms.

[Structural Coupling] Mechanisms M_A and M_B are structurally coupled if at least one of the following holds:

- $\Sigma_A \cap \Sigma_B \neq \emptyset$
- $R_A \cap R_B \neq \emptyset$
- $O_A \cap O_B \neq \emptyset$

Coupling does not imply fusion; it induces mutual constraint on refinement.

7.2 Coupling Strength

To control tractability, coupling must be quantifiable.

[Coupling Strength] The coupling strength $\varepsilon \in [0, 1]$ between M_A and M_B is defined as a weighted combination
$$\varepsilon := \alpha_\Sigma \frac{|\Sigma_A \cap \Sigma_B|}{|\Sigma_A \cup \Sigma_B|} + \alpha_R \frac{|R_A \cap R_B|}{|R_A \cup R_B|} + \alpha_O \kappa_O,$$
where κ_O measures operator compatibility and $\alpha_\Sigma, \alpha_R, \alpha_O$ are normalization constants.

This definition captures symbolic, rule-based, and operational overlap without privileging any single component.

7.3 Entangled Refinement

When mechanisms are coupled, refinement of one induces perturbations in the other.

[Entangled Refinement] An entangled refinement step is a pair of transitions

$$(S_A, S_B) \to (S'_A, S'_B)$$

such that at least one transition depends on the coupled structure.

Entangled refinement paths form trajectories in the product space

$$\mathcal{S}_A \times \mathcal{S}_B$$

with coupling-dependent constraints.

8 Refinement Symmetry and Conservation

In classical physics, Noether's theorem establishes a correspondence between continuous symmetries and conserved quantities. The UNNS Substrate is discrete, recursive, and non-metrical at the outset; nevertheless, an analogous correspondence emerges once refinement dynamics are introduced.

This section formalizes symmetry at the substrate level and proves that Least-Divergence dynamics induce conserved structural quantities.

8.1 Refinement Symmetries

[Refinement Symmetry] A refinement symmetry is a transformation

$$\mathcal{T} : \mathcal{S} \to \mathcal{S}$$

such that:

1. $\mathcal{T}$ preserves admissibility and τ-closure
2. $\mathcal{T}$ preserves the refinement metric:

$$d(\mathcal{T}(S), \mathcal{T}(T)) = d(S, T)$$

3. $\mathcal{T}$ preserves coupling relations between mechanisms

Refinement symmetries need not correspond to spatial, temporal, or numerical transformations. They operate on the space of mechanisms and refinement rules.

Examples. Typical refinement symmetries include:

- relabeling of symbolic alphabets
- reordering of commuting refinement operators
- scale transformations preserving closure relations
- coupling-preserving operator conjugations

8.2 Invariant Structural Observables

Let $Q : \mathcal{S} \to \mathbb{R}$ be a structural observable (e.g. structural mass, divergence index, or refinement velocity).

[Conserved Structural Quantity] A quantity Q is conserved along a refinement trajectory

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

if

$$Q(S_{k+1}) = Q(S_k) \quad \text{for all } k.$$

Conservation here refers to invariance under refinement, not numerical constancy under iteration.

8.3 The UNNS–Noether Correspondence

We now state the central result.

[UNNS–Noether Correspondence] Let $\mathcal{T}$ be a refinement symmetry of the UNNS Substrate. Then, under the Principle of Least Divergence, there exists a conserved structural quantity Q associated with $\mathcal{T}$.

Conversely, each conserved structural quantity arises from a refinement symmetry preserving admissibility, metric structure, and coupling relations.

Sketch. Assume $\mathcal{T}$ is a refinement symmetry. Because $\mathcal{T}$ preserves the refinement metric and coupling structure, it maps refinement trajectories to equivalent trajectories with identical divergence profiles.

Under the Least Divergence Principle, stable refinement trajectories are extremizers of expected divergence. Symmetry implies degeneracy among extremal trajectories related by $\mathcal{T}$.

This degeneracy enforces invariance of a structural observable Q along refinement paths, yielding conservation.

Conversely, if a quantity Q is conserved along all stable refinement trajectories, the transformations preserving Q form a refinement symmetry group. □

8.4 Discrete Variational Character

Unlike classical Noether theory, no assumption of continuity or differentiability is required. The correspondence arises from:

- discrete refinement paths
- metric structure on mechanism space
- variational minimization of divergence

Thus, conservation laws are emergent properties of recursive stability, not fundamental axioms.

8.5 Perfectly Conservative Invariants Revisited

Structures that minimize divergence for all admissible refinement paths are fixed points of the UNNS–Noether Correspondence.

[Perfectly Conservative Invariant] A structure S is perfectly conservative if:

1. it is admissible and τ-closed

2. it minimizes expected divergence across all refinement ensembles

3. it is invariant under all refinement symmetries

Such structures:

- exhibit maximal robustness under perturbation
- persist across coupling regimes
- function as structural constants of the substrate

The identification of specific perfectly conservative invariants is an empirical question addressed through computational exploration.

8.6 Interpretation and Scope

The UNNS–Noether Correspondence does not replace classical conservation laws. Rather, it explains why conservation laws arise whenever stable recursive structures exist.

In this sense, conservation is not a primitive feature of reality, but a consequence of least-divergence refinement in a recursive substrate.

8.7 Refinement Geodesics

The refinement quasi-metric introduced in Section 6 defines distance but does not by itself specify which refinement paths are dynamically preferred. To connect geometry with dynamics, we introduce the notion of refinement geodesics.

[Metric Refinement Geodesic] A refinement path

$$P = (S_0 \to S_1 \to \cdots \to S_n)$$

is a metric refinement geodesic if

$$\sum_{k=0}^{n-1} d(S_k, S_{k+1}) = d(S_0, S_n),$$

i.e. it realizes the minimal refinement distance between its endpoints.

Metric geodesics minimize structural effort but do not account for structural leakage during refinement. Physical realizability requires an additional criterion.

[Physical Refinement Geodesic] A refinement path P is a physical refinement geodesic if it is a metric refinement geodesic and a local minimizer of expected divergence among all admissible refinement paths connecting the same endpoints.

Interpretation. Physical refinement geodesics represent optimal refinement trajectories that simultaneously:

- minimize structural transformation cost, and
- minimize structural leakage during refinement.

This dual criterion provides the dynamical realization of the Least Divergence Principle in refinement space.

Remark. Physical refinement geodesics need not be unique. Degeneracy of geodesics under refinement symmetries is the mechanism underlying conserved quantities discussed in Section 8.

9 Falsifiability Criteria

The Dynamic Completion of the UNNS Substrate makes falsifiable claims. The framework would be invalidated if any of the following are observed:

1. The existence of structurally stable refinement trajectories exhibiting unbounded expected divergence.
2. Failure of least-divergence refinement paths to correspond to observed stable structures.
3. Identification of conserved structural quantities in the absence of corresponding refinement symmetries.
4. Persistent refinement geodesics that do not minimize divergence relative to nearby admissible paths.

These criteria distinguish the present framework from purely descriptive or non-testable formalisms.

10 Conclusion and Outlook

The Dynamic Completion of the UNNS Substrate transforms the theory from a classification of admissible structures into a dynamical framework governed by variational principles, geometry, interaction, and conservation.

The results establish:

- a structural analogue of action minimization
- a discrete notion of symmetry
- an emergent theory of conserved quantities

A Status of Results

For clarity, we summarize the logical status of key results:

- Definitions of structural mass, flux, divergence, refinement distance, and geodesics are **axiomatic**.
- The Least Divergence Principle is a **postulate** supported by internal consistency and testability.
- The implication Stability $\Rightarrow$ τ-closure is **proven**.
- The refinement distance is a **quasi-metric**; reversibility is not assumed.
- The UNNS–Noether Correspondence is a **conjecture** with a defined proof program.

B Proof Program for the UNNS–Noether Correspondence

A full proof of the UNNS–Noether Correspondence requires the following steps:

1. Formal definition of divergence profiles for refinement path ensembles.
2. Proof that refinement symmetries induce degeneracy among divergence-minimizing trajectories.
3. Construction of conserved structural quantities as invariants under symmetry orbits.
4. Proof of completeness: every conserved quantity arises from a refinement symmetry.

These steps are amenable to both analytical and computational approaches and form the basis for future work.

C Connection to Computational Chambers

Existing UNNS chambers provide preliminary instantiations of the present theory:

- Chamber XIV explores scale refinement trajectories, candidate geodesics, and least-divergence attractors.
- Chamber XVI implements structural flux tracking and divergence computation.
- Chamber XXXI is proposed to compute refinement geodesics, divergence-minimizing paths, and coupling stability.

Computational validation is deferred to subsequent work.