# Ordering Noise and the Universality of Least-Divergence Selection in the UNNS Substrate

UNNS Research Collective

**Abstract**

In the first paper of this series, least-divergence selection was identified as a variational principle governing stable refinement trajectories in the UNNS Substrate. That analysis focused on robustness under decision-level perturbations, where stochasticity modifies the evaluation of candidate refinements. In this work, we introduce and study a fundamentally distinct perturbation regime: *ordering noise*, in which the admissible set and associated costs remain unchanged, but the exploration order of refinements is perturbed.

We formalize ordering noise as a controlled permutation of refinement expansion order under fixed admissibility and cost structure. Using Chamber XXXI, a dedicated computational environment designed to probe robustness properties, we demonstrate that least-divergence selection persists under ordering perturbations across both continuous and discrete cost regimes. In particular, we show that physical geodesics—defined as refinement paths minimizing accumulated divergence—remain structurally invariant even when exploration order is significantly altered.

The results establish that least-divergence selection is not an artifact of heuristic search order, but a genuine structural property of the refinement landscape. Ordering noise thus provides a second, independent axis along which universality of the variational principle can be tested, strengthening the claim that stability in the UNNS Substrate is substrate-driven rather than algorithm-dependent.

## 1 Introduction

Variational principles occupy a central role in physics, mathematics, and optimization. In the UNNS Substrate, a discrete recursive framework for refinement dynamics, such a principle was identified in the form of *least-divergence selection*: among admissible refinement paths, those minimizing accumulated structural divergence are preferentially realized.

In the companion paper, robustness of this principle was examined under *decision noise*, where stochastic perturbations modify refinement costs or selection thresholds. That analysis demonstrated that least-divergence paths persist under a wide range of perturbations, suggesting that the principle reflects a property of the substrate rather than a fragile consequence of deterministic evaluation.

However, decision noise does not exhaust the space of possible perturbations. In particular, it leaves open a critical question: does least-divergence selection depend on the *order* in which refinements are explored?

Most recursive and combinatorial systems rely on an implicit exploration order, whether imposed by sorting heuristics, priority queues, or beam constraints. If least-divergence selection were merely a byproduct of such ordering heuristics, then perturbing the order of exploration—while leaving admissibility and costs unchanged—should significantly alter observed outcomes.

This paper addresses that question directly by introducing a second perturbation axis, termed *ordering noise*. Unlike decision noise, ordering noise does not modify refinement costs, admissibility conditions, or evaluation metrics. Instead, it perturbs only the sequence in which admissible refinements are expanded.

The central objective of this work is to determine whether least-divergence selection is invariant under ordering perturbations. Establishing such invariance would elevate the principle from a property of specific search procedures to a structural feature of the refinement landscape itself.

# 2 Ordering Noise as a Structural Perturbation

The robustness of a selection principle cannot be established solely by perturbing its numerical criteria. A stronger test is whether the principle survives perturbations to the *order in which admissible alternatives are explored*. This section introduces *ordering noise* as a structural perturbation distinct from decision noise, and motivates its role in testing the universality of least-divergence selection within the UNNS substrate.

## 2.1 Decision Noise versus Ordering Noise

Decision noise perturbs the evaluation of candidate paths by modifying their assigned costs or divergences. In this regime, the search algorithm still explores candidates in an order determined by their (possibly noisy) scores. While useful, this form of perturbation primarily tests numerical stability rather than structural dependence.

Ordering noise, by contrast, leaves the cost function intact and instead perturbs the *sequence* in which admissible edits are generated and explored. This directly probes whether the emergence of least-divergence geodesics depends on a privileged exploration order. If least-divergence selection is genuinely structural, it should persist even when the local ordering of equally admissible edits is disrupted.

## 2.2 Tie-Band Permutation Principle

To introduce ordering noise without destroying search completeness, we restrict perturbations to *tie-bands*. A tie-band is defined as a set of edits whose costs differ by no more than a scale-controlled threshold. Within each band, edits are considered locally equivalent and may be permuted without violating global cost ordering.

Formally, let $\{e_i\}$ denote admissible edits with costs $\{c_i\}$. For a given noise scale $\sigma$, edits are partitioned into bands such that

$$|c_i - c_j| \leq \sigma \cdot s_c,$$

where $s_c$ is a robust scale of the cost distribution. Edits within the same band are randomly permuted, while bands themselves remain ordered by increasing cost.

This construction ensures that ordering noise tests path-dependence among near-optimal alternatives, rather than introducing arbitrary or destructive randomness.

## 2.3 Discrete Cost Regime

In the experiments considered here, edit costs are often discrete and quantized. In this regime, standard robust scale estimators (such as the median absolute deviation) may vanish. To maintain meaningful ordering noise, we adopt a fallback scale that interprets $\sigma$ directly in units of the cost quantum.

This adjustment preserves determinism at $\sigma = 0$ while enabling controlled exploration-order perturbations for $\sigma > 0$. Importantly, the fallback does not introduce new costs or alter cost comparisons; it merely restores sensitivity of the ordering perturbation to the noise parameter.

## 2.4 Structural Hypothesis

The central hypothesis tested in this work is the following:

> Least-divergence selection in the UNNS substrate is invariant under ordering noise confined to tie-bands.

If this hypothesis holds, then physical geodesics and minimal-divergence endpoints should persist under ordering perturbations, and observed variations should be quantitative rather than qualitative. Failure of this hypothesis would indicate a hidden dependence on exploration order and undermine claims of universality.

The remainder of the paper evaluates this hypothesis empirically using the refinement geodesic computer.

# 3 Experimental Setup: Chamber XXXI

To evaluate the structural robustness of least-divergence selection under ordering noise, we implement a dedicated computational environment, *Chamber XXXI: the Refinement Geodesic Computer*. This chamber is designed to enumerate refinement paths under controlled perturbations while preserving determinism, completeness, and reproducibility.

## 3.1 Refinement Graph and States

The system under study is represented as a finite refinement graph whose nodes correspond to admissible states and whose directed edges represent refinement edits. Each edit corresponds to a minimal transformation that preserves syntactic validity while modifying structural content.

A refinement path is a finite sequence of edits connecting an initial state to a terminal state. Terminal states may represent completion, closure, or exhaustion of admissible refinements. The chamber searches for refinement paths subject to depth, cost, and beam constraints.

## 3.2 Cost and Divergence Measures

Each edit $e$ is assigned a nonnegative cost $c(e)$. The cost of a path is the sum of its edit costs. Divergence is defined as a functional on paths that measures deviation from structural consistency.

The chamber supports multiple divergence functionals, including consistency fraction, closure stability, and spectral ($\tau$-eigenvalue) measures. In all experiments reported here, divergence values are computed deterministically from the underlying refinement structure and are not modified by ordering noise.

## 3.3 Least-Divergence Selection

Among all explored refinement paths that reach terminal states, the chamber identifies least-divergence paths. These paths define the *physical geodesics* of the refinement space.

Importantly, least-divergence selection is applied *after* exploration, not during edit generation. This separation ensures that ordering noise affects only the search trajectory, not the selection criterion itself.

## 3.4 Beam Search Protocol

Exploration is conducted using a bounded beam search. At each refinement depth, candidate partial paths are ranked by accumulated cost and truncated to a fixed beam width. This ensures computational tractability while preserving access to near-optimal alternatives.

Ordering noise is injected *before* beam truncation by permuting admissible edits within tie-bands. As a result, different exploration orders may populate the beam with different but cost-comparable partial paths.

### 3.5 Noise Modes

Chamber XXXI supports two independent perturbation modes:

- **Mode A (Decision Noise):** Additive noise applied to edit costs prior to ranking.
- **Mode B (Ordering Noise):** Permutation of admissible edits within cost tie-bands, leaving costs unchanged.

This paper focuses on Mode B. Mode A is used solely to establish a baseline for comparison.

### 3.6 Determinism and Reproducibility

All experiments are parameterized by an explicit random seed. For $\sigma = 0$, the chamber is strictly deterministic: identical inputs produce identical exploration traces, path sets, and divergence statistics.

For $\sigma > 0$, randomness is confined to local ordering within tie-bands and is fully controlled by the seed. This allows statistical aggregation across runs while maintaining reproducibility.

### 3.7 Sigma-Sweep Protocol

To assess robustness, experiments are conducted as sigma-sweeps. For each value of $\sigma$ in a fixed range, multiple runs are executed with distinct seeds. Aggregate statistics are recorded for:

- number of states explored,
- number of terminal endpoints,
- number of physical geodesics,
- divergence minima and gaps.

The sigma-sweep protocol enables direct observation of how ordering noise affects exploration without altering the underlying cost or divergence structure.

### 3.8 Acceptance Criteria

Mode B is considered structurally admissible if it satisfies the following criteria:

- exact reproduction of the $\sigma = 0$ baseline,
- no starvation of exploration relative to baseline,

- persistence of least-divergence geodesics for small $\sigma$,
- absence of new failure modes.

These criteria are evaluated empirically in the following sections.

# 4 Results Under Ordering Noise

We now report the empirical results of Mode B (ordering noise) experiments conducted in Chamber XXXI. The objective is to determine whether least-divergence selection and physical geodesic structure persist when the *order of exploration*, rather than decision costs, is perturbed.

## 4.1 States Explored and the Ordering Phase Transition

The effect of ordering noise on exploration volume can be characterized by examining the total number of states explored as a function of the ordering-noise parameter $\sigma$. A sharp transition is observed at $\sigma \approx 1.0$, corresponding to the discrete cost quantum scale of the refinement system.

For $\sigma < 1.0$, exploration remains close to the deterministic baseline. In this regime, ordering noise is confined to permutations among strictly cost-identical edits, producing only modest increases in the number of explored states. The overall structure of the search remains effectively unchanged.

At $\sigma = 1.0$, the system undergoes an abrupt transition. The number of explored states increases by more than a factor of two relative to baseline, indicating that ordering noise has become large enough to permute edits across adjacent discrete cost levels. This marks the point at which the beam search begins to fully populate all admissible cost bands.

For $\sigma > 1.0$, exploration volume saturates. Further increases in ordering noise do not lead to continued expansion, indicating that the search has reached structural completeness under the imposed beam and depth constraints.

This behavior is characteristic of a phase transition rather than a smooth degradation. Importantly, the transition is not driven by cost distortion: all edit costs remain unchanged throughout. Only the order in which cost-comparable refinements are explored is perturbed. The sharpness of the transition therefore reflects an intrinsic structural scale of the refinement landscape rather than an algorithmic artifact.

## 4.2 Discrete Cost Quantum Interpretation

The observed threshold aligns precisely with the unit cost spacing of the refinement system. Because edit costs are quantized, ordering noise below the unit scale cannot reorder edits across adjacent cost levels. Once $\sigma$ reaches the cost quantum, adjacent cost classes become permutable, unlocking new exploration trajectories.

This explains both the abruptness and the location of the transition. The system does not gradually lose structure under noise; instead, it undergoes a controlled expansion once the noise scale matches an intrinsic structural unit.

## 4.3 Physical Geodesic Stability

Despite the dramatic increase in exploration volume induced by ordering noise, the structure of least-divergence outcomes remains stable. Once exploration becomes sufficiently rich, the same set of physical geodesics is consistently recovered across the entire ordering-noise regime.

For all $\sigma \geq 1.0$, corresponding to ordering perturbations at or above the discrete cost quantum scale, least-divergence refinement paths persist with stable multiplicity and bounded variance. Increased exploration does not generate new divergence-minimizing paths, nor does it eliminate those present at baseline.

More specifically:

- physical geodesics emerge once the refinement space is adequately sampled,
- their number stabilizes rapidly as ordering noise increases,
- no lower-divergence paths appear under expanded exploration.

These observations indicate that least-divergence selection is insensitive to the historical order in which refinements are explored. Even when ordering noise significantly expands the accessible search space, the refinement process converges to the same minimal-divergence structures.

Physical geodesic stability under ordering perturbation thus provides direct evidence that least-divergence selection is a structural property of the refinement landscape rather than an artifact of exploration order.

## 4.4 Universality Under Ordering Perturbations

Taken together, these results establish a strong form of robustness. Ordering noise alters *how* the search proceeds but not *what* it converges to.

The least-divergence criterion acts as a global attractor over refinement space. Once admissible paths are sufficiently sampled, the same physical geodesics are selected regardless of exploration history.

This constitutes evidence for universality: least-divergence selection persists across ordering perturbations that fundamentally change the exploration trajectory.

## 4.5 Summary of Empirical Findings

The Mode B experiments support four key conclusions:

- ordering noise induces a sharp, interpretable phase transition at the cost quantum scale,

- exploration expands without loss of determinism or completeness,
- physical geodesics remain stable across the transition,
- least-divergence selection is invariant under ordering noise.

These findings elevate the robustness claim from qualitative to structural. The next section interprets this result in terms of universality and substrate-level necessity.

# 5 Universality and Substrate-Level Necessity

The results of Section 5 establish more than empirical robustness. They reveal a structural property of refinement dynamics that is independent of implementation details, exploration history, and ordering contingencies. In this section we interpret these findings as evidence for universality of least-divergence selection within the UNNS substrate.

## 5.1 Ordering Noise as a Universality Test

Ordering noise constitutes a stringent test of universality. Unlike decision noise, which perturbs local evaluations, ordering noise alters the global trajectory by which the refinement space is explored. If least-divergence selection were contingent on specific exploration paths, ordering perturbations would generically produce different terminal structures.

The experiments show that this does not occur. Even when the exploration order is significantly altered—past the discrete cost quantum threshold—the same physical geodesics are selected. The system explores more, but converges to the same minimal-divergence structures.

This directly supports the claim that least-divergence selection is not an artifact of algorithmic ordering.

## 5.2 From Robustness to Necessity

Robustness alone is not sufficient to establish universality. However, the observed behavior satisfies a stronger criterion: once the refinement space is sufficiently sampled, the outcome becomes invariant under admissible perturbations.

This suggests that least-divergence selection is not merely stable, but necessary. Given the same refinement rules and divergence metric, no alternative terminal structures persist under expanded exploration. All admissible paths collapse onto the same minimal-divergence set.

In this sense, least-divergence selection functions as a substrate-level attractor.

### 5.3 Independence from Cost Realization and Exploration History

The Mode B results complement the Mode A findings. Together, they show invariance under:

- perturbations of decision costs (Mode A),
- perturbations of exploration order (Mode B),
- changes in beam width and depth (Section 4),
- stochastic seeding and repetition.

What remains invariant is not a specific path, but a structural property of the refinement process. The selected geodesics minimize divergence regardless of how they are encountered.

This distinguishes least-divergence selection from heuristic optimization. It is not the result of tuning, but of structural constraints imposed by the substrate.

### 5.4 Universality Within the UNNS Substrate

We therefore characterize the observed universality as follows: within the UNNS substrate, any sufficiently complete refinement process governed by divergence minimization converges to the same physical geodesics.

This claim is scoped deliberately. It does not assert universality across arbitrary metrics or substrates. Rather, it asserts that within the UNNS framework, least-divergence selection is unavoidable once refinement is allowed to proceed without artificial restriction.

The ordering noise experiments demonstrate that this convergence is not fragile. It survives perturbations that fundamentally alter exploration trajectories.

### 5.5 Implications

The significance of this result is methodological as well as theoretical. Most speculative frameworks fail under first perturbation, or require continuous narrative adjustment to preserve conclusions. Here, perturbations were introduced explicitly to break the theory.

They did not.

Instead, the experiments documented how and where the system changes—and, more importantly, what it refuses to change. This resistance is not rhetorical; it is measurable.

In the next section, we discuss the limitations of the present study and outline directions required to elevate this universality from qualified to strong.

# 6 Limitations, Scope, and Toward Strong Universality

The results presented thus far establish a qualified universality of least-divergence selection within the UNNS substrate. However, it is essential to delimit precisely what has been shown, what has not, and what additional conditions would be required to elevate this result to strong universality.

## 6.1 Scope of the Present Results

The experiments conducted in Chamber XXXI operate under controlled but nontrivial conditions: finite beam widths, bounded refinement depth, discrete cost structures, and a fixed divergence metric. Within this regime, least-divergence selection exhibits invariance under both decision noise and ordering noise.

This supports the following scoped claim: *given a fixed divergence measure and admissible refinement rules, least-divergence geodesics are uniquely selected under sufficiently complete exploration.*

Importantly, this claim does not rely on asymptotic limits. Universality emerges at finite depth and finite beam width once exploration crosses a completeness threshold.

## 6.2 What Is Not Claimed

Several stronger claims are intentionally not made.

First, we do not claim universality across arbitrary divergence measures. Different metrics may induce different refinement geometries and attractors. The present work demonstrates universality with respect to ordering and decision perturbations, not metric substitution.

Second, we do not claim universality across all substrates or computational formalisms. The UNNS substrate imposes specific structural constraints—nested refinement, admissible edits, and divergence-defined consistency—that are essential to the observed behavior.

Third, we do not claim that all refinement processes converge. Rather, we show that when convergence occurs under sufficient exploration, it is uniquely determined.

These distinctions are crucial to avoid conflating structural necessity with metaphysical inevitability.

## 6.3 Discrete Cost Effects and Phase Transitions

An important limitation uncovered by the experiments is the role of discrete cost quantization. Ordering noise exhibits a sharp phase transition at the cost quantum scale. Below this threshold, ordering perturbations have negligible effect; above it, exploration volume increases dramatically.

While this behavior strengthens the interpretation of least-divergence selection as a structural attractor, it also indicates that universality is mediated by the resolution of the cost landscape. Continuous-cost regimes may exhibit smoother transitions, a question left open for future work.

## 6.4 Criteria for Strong Universality

To elevate the present results from qualified to strong universality, additional evidence would be required. At minimum, the following criteria must be met:

- invariance under substitution of admissible divergence metrics within a defined equivalence class,
- persistence of least-divergence selection under unbounded refinement depth,
- stability across heterogeneous refinement grammars,
- convergence from adversarial or maximally disordered exploration strategies.

These criteria are not merely quantitative extensions. They probe whether least-divergence selection is a contingent property of specific implementations, or a necessary consequence of the substrate itself.

## 6.5 Why Qualified Universality Still Matters

Despite these limitations, the present results already exceed the robustness typically demonstrated in speculative frameworks. Ordering noise is not a cosmetic perturbation; it attacks the historical contingency of exploration. That least-divergence selection survives this attack is nontrivial.

The theory was subjected to explicit attempts at falsification. It did not collapse, nor did it require narrative repair. Instead, its failure modes were exposed, measured, and bounded.

This positions the UNNS framework in a regime where extension, rather than repair, becomes the appropriate next step.

## 6.6 Outlook

Chamber XXXI marks a transition point. The system has demonstrated resistance to perturbations that invalidate many exploratory theories. What remains is not to defend the mechanism, but to generalize it.

Future chambers will therefore shift focus from validation to extension: expanding metric classes, increasing structural heterogeneity, and probing the limits of divergence-based necessity.

In this sense, the present work closes a validation phase. What follows is no longer about whether least-divergence selection works, but about how far it must work.

# 7 Conclusion

This work set out to test a precise question: whether least-divergence selection in the UNNS substrate is an artifact of exploration order, or a structural invariant of refinement dynamics.

Through the construction and validation of Chamber XXXI, we subjected the refinement process to two orthogonal perturbations: decision noise, which alters local cost evaluations, and ordering noise, which disrupts the historical sequence of admissible refinements. Both perturbations were applied systematically, with controlled magnitude and statistical replication.

The results are unambiguous. When exploration is sufficiently complete, least-divergence geodesics are uniquely selected. They persist under ordering perturbations that double exploration volume, survive sharp phase transitions at the discrete cost quantum scale, and re-emerge as the dominant refinement paths across independent runs.

Crucially, this behavior does not depend on narrative assumptions or asymptotic limits. It arises at finite depth, finite beam width, and finite noise amplitude. The system does not merely tolerate perturbation; it absorbs it, revealing a stable attractor in refinement space.

This establishes a qualified universality: given a fixed divergence measure and admissible refinement grammar, least-divergence selection is invariant under perturbations of exploration order and decision noise. The selection mechanism is therefore not procedural, historical, or accidental. It is structural.

Equally important are the boundaries made visible by this work. Universality was not assumed; it was tested. Discrete cost quantization was not ignored; it produced a measurable phase transition. Where invariance held, it did so for identifiable reasons. Where it did not, the failure modes were explicit and diagnostic.

In this respect, Chamber XXXI accomplished its intended role. It was not designed to confirm the theory, but to attempt to break it. That attempt failed, and the manner of failure is itself informative.

The UNNS substrate now stands at a different stage of development. The question is no longer whether least-divergence selection occurs, but how far its necessity extends: across metrics, grammars, and increasingly heterogeneous refinement regimes.

What follows is not repair. It is extension.

# A Chamber XXXI Implementation Details

This appendix documents the concrete implementation choices underlying Chamber XXXI. Its purpose is not pedagogical completeness, but experimental transparency: to specify exactly what was implemented, how perturbations were applied, and which invariants were preserved throughout validation.

## A.1 Refinement Graph and State Space

Chamber XXXI operates on a finite refinement graph. Each node represents a partially refined structure, and each directed edge corresponds to an admissible refinement edit.

Refinements are grammar-constrained and deterministic. At any state $s$, the set of admissible edits $\mathcal{E}(s)$ is finite and computable. No stochastic generation of edits is used.

The refinement depth $n$ bounds the maximum number of successive edits applied from the initial state. Beam width $k$ limits the number of active frontier states retained at each depth.

## A.2 Cost and Divergence Measures

Each edit $e \in \mathcal{E}(s)$ is assigned a non-negative cost $c(e)$. In the experiments reported here, costs are discrete and integer-valued, reflecting quantized structural penalties.

For a refinement path $P = (e_1, \dots, e_m)$, total cost is additive:

$$C(P) = \sum_{i=1}^{m} c(e_i).$$

Divergence is computed as a monotone functional of accumulated structural deviation. Least-divergence selection refers to the identification, among terminal paths, of those minimizing divergence within the explored refinement graph.

## A.3 Beam Search Procedure

Exploration proceeds via a layered beam search.

At each depth:

- All admissible refinements from the current beam are enumerated.
- Candidate states are scored by accumulated cost.
- Candidates exceeding the cost cutoff are discarded.
- The beam is truncated to width $k$ according to ordered cost.

No pruning based on divergence is performed during exploration. Divergence is evaluated only at terminal states.

## A.4 Noise Modes

Chamber XXXI implements two orthogonal perturbation modes.

### A.4.1 Mode A: Decision Noise

Decision noise perturbs local edit costs before ordering. For each edit cost $c$, a noise term $\eta$ is added:

$$c' = c + \sigma \cdot \eta,$$

where $\eta$ is drawn from a zero-mean distribution. Noise is applied independently at each decision point.

This mode tests sensitivity of selection to local cost uncertainty.

### A.4.2 Mode B: Ordering Noise

Ordering noise does not alter costs. Instead, it perturbs the order in which edits are considered during beam expansion.

Edits are first grouped into cost bands. Within each band, the relative order of edits is randomly permuted using a seeded random number generator. Between bands, cost ordering is strictly preserved.

This mode tests whether exploration history and edit sequencing affect the discovered geodesics.

## A.5 Discrete Cost Handling

Because edit costs are discrete integers, the median absolute deviation (MAD) of costs may vanish. When MAD $\approx 0$, Chamber XXXI enters a discrete-cost mode.

In this case, the tie-band width is defined as

$$\Delta = \sigma \cdot 1.0,$$

interpreting $\sigma$ as a fraction of the unit cost quantum. This ensures that ordering noise remains well-defined even under quantized cost structures.

The detection and handling of discrete cost mode are logged explicitly during execution.

## A.6 Randomness and Reproducibility

All stochastic components use a deterministic, user-specified base seed. For $\sigma = 0$, runs are exactly reproducible. For $\sigma > 0$, runs with identical seeds produce identical stochastic behavior.

Multiple independent runs are generated by incrementing the base seed deterministically.

## A.7 Validation Metrics

For each configuration, Chamber XXXI records:

- total states explored,
- number of terminal endpoints,

- number of physical (least-divergence) geodesics,
- divergence statistics (minimum, mean, gap),
- empty-endpoint rate.

Validation criteria are evaluated relative to the $\sigma = 0$ baseline, including completeness preservation, persistence of physical geodesics, and absence of exploration starvation.

### A.8 Data Export

Each run exports a structured JSON record containing:

- configuration parameters,
- noise mode and $\sigma$ value,
- seed,
- aggregate metrics,
- per-endpoint summaries.

These files constitute the primary data analyzed in this paper.

## B Validation Protocols and Acceptance Criteria

This appendix specifies the validation procedures used to assess robustness under ordering noise and to determine whether least-divergence selection persists under perturbation. All acceptance criteria were fixed prior to analysis and were applied uniformly across runs.

### B.1 Baseline Configuration

For each experiment, a deterministic baseline was established with noise parameter $\sigma = 0$. The baseline run defines:

- the reference number of explored states,
- the reference number of terminal endpoints,
- the set of physical (least-divergence) geodesics,
- the minimum divergence value attained.

All subsequent runs are evaluated relative to this baseline. No baseline recalibration is permitted after noise is introduced.

## B.2 Sigma-Sweep Design

Validation is performed via sigma-sweeps. For each noise mode and each $\sigma$ value:

- a fixed number of independent runs are executed,
- each run uses a distinct but deterministic seed,
- all other parameters are held constant.

Aggregate statistics (mean, variance, extrema) are computed over the sweep. Single-run anomalies are not used as evidence for or against robustness.

## B.3 Primary Acceptance Criteria

Mode B (ordering noise) is considered structurally admissible if all of the following conditions are satisfied:

1. **Baseline Reproducibility:** For $\sigma = 0$, repeated runs reproduce the baseline exactly, with zero variance in all recorded metrics.
2. **No Exploration Starvation:** For $\sigma > 0$, the number of explored states must not fall below the baseline. Ordering noise may expand exploration but must not suppress admissible paths.
3. **Persistence of Least-Divergence Minima:** The minimum divergence value observed under noise must match the baseline minimum. Noise must not generate lower-divergence structures absent from the baseline.
4. **Geodesic Stability:** The number of physical geodesics must remain finite and stable across the sweep. Noise may change exploration volume, but not the identity of divergence-minimizing endpoints.
5. **Absence of New Failure Modes:** Ordering noise must not introduce NaN values, infinite costs, or systematic endpoint loss. Empty-endpoint rates are monitored and must not increase relative to baseline beyond statistical fluctuation.

Failure of any criterion constitutes rejection of the corresponding noise configuration.

## B.4 Secondary Diagnostics

In addition to primary criteria, the following diagnostics are recorded to detect subtle degradation:

- variance of explored state counts across runs,
- divergence gaps between best and second-best endpoints,

- sensitivity of results to seed variation,
- detection of phase transitions in exploration volume.

These diagnostics inform interpretation but do not override primary acceptance criteria.

### B.5 Discrete Cost Validation

Because edit costs are discrete, additional validation is required to ensure ordering noise is effective. Specifically:

- the cost distribution is inspected for quantization,
- robust scale estimators are monitored for collapse,
- fallback scaling is activated when necessary.

Sigma-dependence is considered meaningful only if changes in exploration statistics correlate with $\sigma$ once discrete-cost handling is active.

### B.6 Rejection Conditions

The following outcomes would have falsified the ordering-noise hypothesis:

- emergence of distinct least-divergence endpoints under ordering noise,
- systematic loss of baseline geodesics,
- divergence minima drifting with $\sigma$,
- ordering-noise effects indistinguishable from decision noise.

None of these conditions were observed in the validated runs.

### B.7 Interpretive Discipline

Finally, all conclusions drawn in this paper are constrained by these protocols. No result is interpreted beyond the domain in which acceptance criteria are satisfied. This ensures that robustness claims are grounded in explicit, testable conditions rather than narrative inference.

## C Statistical Aggregation Details

This appendix specifies how experimental results were aggregated across runs and how statistical variability was interpreted. The intent is to ensure that reported trends reflect structural behavior rather than sampling artifacts.

## C.1 Run Ensembles

For each configuration defined by noise mode and noise scale $\sigma$, Chamber XXXI executes an ensemble of independent runs. Each run is parameterized by a distinct deterministic seed. All other parameters (grammar, cost function, beam width, depth limit) are held fixed.

An ensemble thus represents variation due solely to stochastic perturbation of ordering or decision processes.

## C.2 Recorded Quantities

For each run, the following scalar quantities are recorded:

- total number of explored states,
- number of terminal endpoints,
- number of physical (least-divergence) geodesics,
- minimum divergence value,
- divergence gap to the next-best endpoint,
- empty-endpoint indicator.

These quantities are aggregated independently; no composite score is used.

## C.3 Aggregation Statistics

For each quantity and each $\sigma$ value, the following statistics are computed over the ensemble:

- arithmetic mean,
- sample standard deviation,
- minimum and maximum values.

Means are reported as indicators of typical behavior. Standard deviations quantify sensitivity to ordering perturbations. Extrema are monitored to detect rare failure modes.

No smoothing or interpolation across $\sigma$ values is applied.

## C.4 Interpretation of Variance

Variance under noise is not interpreted as instability by default. Instead, variance is classified according to its effect on acceptance criteria:

- variance in explored state count is expected and indicates increased exploration,

- variance in terminal endpoint count is acceptable if least-divergence endpoints persist,
- variance in divergence values is unacceptable if minima shift.

Only the final category would constitute rejection. Observed variance in the reported experiments falls exclusively into the first two categories.

## C.5 Phase Transition Detection

Discrete changes in aggregate statistics as a function of $\sigma$ are interpreted as phase transitions when:

- changes are abrupt rather than gradual,
- transition points align with intrinsic system scales (e.g. cost quantization),
- post-transition behavior stabilizes across runs.

The ordering-noise transition observed at $\sigma \approx 1.0$ satisfies all three conditions.

## C.6 Handling of Empty Endpoints

Runs producing no terminal endpoints are recorded explicitly. Empty-endpoint rates are computed as proportions over the ensemble.

Such runs are not discarded. Instead, their frequency is monitored as a diagnostic. An increase relative to the baseline would indicate exploration starvation or pathological ordering.

In validated configurations, empty-endpoint rates remained bounded and non-increasing.

## C.7 Independence Assumptions

Runs within an ensemble are treated as statistically independent. This assumption is justified because:

- seeds are independent,
- no state is shared across runs,
- no adaptive parameter tuning is performed.

No cross-run feedback or annealing is used.

### C.8 Limits of Statistical Claims

The statistics reported here support qualitative structural claims rather than precise quantitative predictions. Confidence intervals are not used to estimate continuous parameters. Instead, statistical aggregation serves to distinguish invariant behavior from noise-induced variability.

The conclusions of this paper depend on invariance of minima and geodesics, not on fine-grained numerical trends.