

Population-Level Methods for Structured Ensembles: From Chaotic DNA Inversions to Ensemble Leg Drafting

Population-Level Methods for Structured Ensembles: From Chaotic DNA Inversions to Ensemble Leg Drafting

Abstract

Many scientific settings require coordinating or extracting information from a large population of structurally identical units — dynamical systems, draft models, or mutating sequences — when only a population-level signal is available. This paper develops a single analytical framework for three such settings. First, we analyze a toy model of random DNA inversions, whose underlying process has been argued to be chaotic in the sense of Devaney's theory, and derive the expected divergence time of a binary sequence of length ten: the average segment length is 4.0 and, under a labeled independence approximation (projection P0), full Hamming divergence is reached after approximately 3 steps, with the exact figure expected to be larger (at least 3 steps) because re-flipping of already-inverted sites slows the per-step Hamming increment. Second, we build a moment-based baseline for broadcast control of a continuum of first-order units with distributed gains, showing that the first moment moves at rate $1/3$ per unit input, that the mean advance is immobile for sign-symmetric gain distributions, and that the minimum steering time to a target first moment of 0.1 is 0.3 time units. Third, we propose *ensemble leg drafting* for speculative decoding, in which m draft models are fused under an OR rule; with acceptance probabilities 0.5, 0.6, 0.7 and draft length 4, the ensemble advances 4.4349 tokens per cycle versus 2.7731 for the best single leg, a computed speedup of 1.5993. All quantitative claims are derived with shown arithmetic or explicitly labeled projections.

1. Introduction

A recurring structural problem across science and engineering is the following: a large population of units must be steered, queried, or made to produce useful output, but the only available actuator or signal acts on the whole population at once. In ensemble control, this arises because control can only be implemented at the population level, by broadcasting an input signal to all systems in the population [8]; the same literature emphasizes that such systems are severely underactuated and that comprehensive state feedback on each member is unavailable [5]. In large-language-model inference, the analogous problem appears in speculative decoding, where a cheap draft model proposes tokens that an expensive target model verifies in parallel [3]. In molecular biology, DNA inversion processes have been argued to be chaotic under Devaney's definition [1], raising the question of what population-level statistics of such processes are predictable at all.

This paper treats these three settings as case studies of one question: **what can be achieved at the population level when individual-level access is unavailable?** The contributions are:

1. A toy model of random DNA inversions with closed-form expectations for Hamming divergence (Sections 3.1 and 4.1).
2. A moment-based baseline for broadcast control of a continuum of first-order units, including an exact negative result (immobility of the mean for sign-symmetric gains) and a closure analysis of the moment hierarchy (Sections 3.2, 4.2).
3. A formal model of *ensemble leg drafting* — fusing m speculative-decoding draft models under an OR rule — with exact computations of effective acceptance probability, tokens per cycle, throughput, speedup, and a break-even condition, plus a labeled projection for correlated legs (Sections 3.3, 4.3).

Every quantitative claim is either derived with explicit arithmetic in Section 4 or labeled as a projection with stated assumptions. We report no empirical measurements; the paper is an analytical framework with worked numerical examples.

2. Background and Related Work

Ensemble control of populations. Ensemble Control on Lie Groups treats problems involving large ensembles of structurally identical dynamical systems arising in numerous scientific areas from quantum control and robotics to brain medicine, and emphasizes that in many such applications control can only be implemented at the population level, i.e., by broadcasting an input signal to all systems in the population [8]. The supplied abstract is truncated, so we draw only this framing from it; our broadcast-input assumption in Section 3.2 is taken directly from this setting. Moment-Based Ensemble Control frames the control of a large population — in the limit a continuum — of structurally identical dynamical systems with parametric variations as a pervasive task in science and engineering, and identifies the severely underactuated nature of such systems and the inability to avail comprehensive state feedback as the central challenges of ensemble analysis and design [5]. Our moment hierarchy in Section 3.2 and its closure failure in Section 4.2 are a direct, elementary instantiation of that challenge.

Speculative decoding for vision-language models. TABED (Test-Time Adaptive Ensemble Drafting) addresses speculative decoding for Large Vision-Language Models, noting that speculative decoding "remains largely unexplored" for LVLMs, which extend LLMs to process both image and text prompts, and benchmarking existing inference methods with small draft models on 11 datasets across diverse inputs [3]. The supplied abstract is truncated before its methodological details, so we relate to it only through what it states: it establishes the gap (draft-model acceleration for multimodal models is underexplored) and the practice of benchmarking draft-model inference across many datasets. Our Section 3.3 provides an analytical model of a specific multi-draft fusion rule that such benchmarks could evaluate; we attribute no mechanism, result, or number to [3] beyond its stated scope.

Chaos in biological sequence dynamics. Chaos in DNA inversions is a draft paper proving that the inversion process occurring in DNA mutations is chaotic in the sense of Devaney's theory [1]. The supplied abstract states only this claim. We use it as the motivation for Section 3.1: if the inversion process is chaotic, individual long-run trajectories are unpredictable, and only statistical quantities — such as the expected Hamming distance we compute — are the accessible objects of study. We make no use of, and claim no transfer of, the specific Devaney-chaos result beyond this motivation.

Decoupled multi-mode analysis. Master Functions and Equations for Perturbations of Vacuum Spherically-Symmetric Spacetimes shows how spherical symmetry permits an expansion of perturbations in scalar, vector, and tensor harmonics, with the resulting perturbative equations decoupling for modes of different parity and harmonic number, making the theory a crucial tool for understanding black-hole perturbation dynamics [4]. The supplied abstract is truncated mid-sentence, so we use only the decoupling structure it states. The methodological lesson — exploit structure to decouple a coupled problem into independent modes — governs both our moment expansion (modes = moment orders, which decouple only partially; Section 4.2) and our leg-fusion analysis (modes = draft legs, which decouple exactly when their errors are independent; Section 4.3).

Spanning sets and isolation of contributions. Prescriptive Master Integrals of Maximal Weight at Two Loops constructs a spanning set of individually pure, planar master integrals involving massless particles in four dimensions which include all maximal-weight contributions at two loops, such that every independent infrared-divergent region is individually matched by specific masters, with all other masters manifestly finite [6]. The stated design principle — choose a basis in which each problematic region is matched by exactly one element — informs our choice in Section 4.2 to isolate the first moment by restricting the input class so that higher-order effects are deferred to a separate closure analysis.

Semi-classical asymptotics. Asymptotic of Bergman Kernel (Master Thesis) gives a new proof of the pointwise asymptotic expansion for the Bergman kernel of a hermitian holomorphic line bundle where the curvature is positive and satisfies a local spectral gap condition, introducing a suitable semi-classical symbol space and symbolic calculus inspired by work of Hsiao and Savale [7]. The supplied summary gives no quantitative content transferable to ensemble control; we cite it as an instance of the pattern that asymptotic problems become tractable once the right calculus is introduced, mirrored in our use of a geometric acceptance model as the calculus in which throughput becomes closed-form (Section 3.3).

Formal verification and well-foundedness. DRAFT: A Formally Verified Constructive Proof of the Consistency of Peano Arithmetic Using Ordinal Assignments provides a modified version of Gentzen's 1936 consistency proof, based on Gödel's reformulation, with additional details and minor corrections necessary to definitively prove the well-foundedness of the cut-elimination argument in a constructive environment; the abstract truncates at this point [2]. We cite it for the methodological standard it represents: an argument resting on layered structure should be checkable layer by layer. Our Section 4 arithmetic is deliberately structured as short, independently checkable computations, though we have not machine-verified it.

QNFO corpus materials. The QNFO Consilient Synthesis v2.0 document describes a shared adelic kernel (valuation, tree boundary, Ostrowski, adeles), applies seven corrections (C1–C7) from a red-team audit dated 2026-07-25, and is restructured kernel-first with an honest kernel-membership table and mandatory symmetry sections [9]. We cite it as a model of audit-driven revision: maintaining an explicit record of which claims are derived versus projected, which we emulate in Section 5 and Appendix B. NUMERATA extends an 8-axis framework for evaluating numeral systems with a Distinction-Based Primality axis and reports a meta-analysis validation with $\text{MCS} = 0.875$ and 5 of 5 predictions confirmed, using sunburst notation [10]. The supplied entry gives no further methodological detail, so we relate to it only as a multi-axis evaluation framework whose axis-decomposition style parallels our decomposition of throughput into acceptance, length, and cost factors (Section 3.3). Connecting Geometrogenesis and Biogenesis is supplied with an empty abstract [11]; no substantive content is available from the supplied entry and we draw nothing from it. Finally, the Kappa/SIIT Cross-Validation reports a direct mathematical cross-validation of the Kappa / Scale-Invariant Information Thermodynamics framework against the Standard Model's known 1-loop renormalization-group β -functions, through the single definition $g_{\text{eff}} = g_0 \kappa$ (abstract truncated) [12]. We cite it as an instance of validating a new framework by exact reproduction of known reference quantities — the strategy we follow in Section 4.3, where the single-draft baseline is first reproduced exactly before the ensemble extension is computed.

3. Methods

3.1 Case study I: a toy model of random DNA inversions

We represent a DNA segment as a binary sequence of fixed length $L = 10$, where each site is 0 (original) or 1 (inverted). At each discrete step, a contiguous segment is selected uniformly among all $\frac{L(L+1)}{2}$ possible segments and all bits within it are flipped. This stochastic rule captures the essential feature of inversion without biochemical detail; the chaotic character of the underlying process [1] motivates a purely statistical treatment.

3.2 Case study II: moment-based broadcast control baseline

Consider a continuum of first-order units indexed by a gain parameter $\omega \in [-1, 1]$, each obeying

$$\dot{x}_\omega(t) = \omega u(t), \quad x_\omega(0) = 0,$$

with a single broadcast input $u(t)$, identical for all ω , bounded by $|u(t)| \leq u_{\text{max}} = 1$. The gain density is uniform, $\rho(\omega) = \frac{1}{2}$ on $[-1, 1]$. Define the gain-weighted moments

$$M_k(t) = \int_{-1}^1 \omega^k x_\omega(t) \rho(\omega) d\omega,$$

the mean advance $\bar{x}(t) = \int_{-1}^1 x_\omega(t) \rho(\omega) d\omega$, and the population variance $\sigma^2(t) = \int_{-1}^1 (x_\omega(t) - \bar{x}(t))^2 \rho(\omega) d\omega$. We restrict to piecewise-constant inputs $u(t) \in \{-1, +1\}$ (bang-bang), the weakest nontrivial broadcast class consistent with the bound. "Reachability" is defined only for M_1 : a value of $M_1(T)$ obtainable from $M_1(0) = 0$ by some admissible input.

3.3 Case study III: ensemble leg drafting for speculative decoding

A target model M_T generates tokens autoregressively. A set $\mathcal{L} = \{L_1, \dots, L_m\}$ of m draft legs each proposes a continuation of length γ . At each draft position j , leg L_i matches the target's sampled token with probability α_i (the leg's *acceptance profile*). Define c_i as the relative compute cost of leg L_i per draft position, $C_{\text{draft}} = \gamma \sum_{i=1}^m c_i$, and $C_{\text{verify}} = 1$ (target verification per cycle, normalized).

OR-fusion rule. The fused draft at position j is admissible if at least one leg's proposal is admissible; the fused sequence takes the token of the first admissible leg in fixed priority order. Under independent leg failures, the effective acceptance probability is

$$\alpha_{\text{eff}} = 1 - \prod_{i=1}^m (1 - \alpha_i).$$

Note that α_{eff} depends on the legs only through the product of their failure probabilities — a population-level statistic, consistent with the population-level viewpoint of [5].

Tokens per cycle. With per-position acceptance probability α i.i.d. across positions, the expected number of tokens advanced per cycle is

$$\mathbb{E}[N] = \sum_{k=0}^{\gamma-1} k \alpha^k (1 - \alpha) + \gamma \alpha^\gamma + 1,$$

where the sum counts cycles ending in rejection after k accepted drafts, the second term counts full-acceptance cycles, and the $+1$ is the bonus token supplied by the target at the rejection point. This sum equals the closed form

$$\mathbb{E}[N] = \frac{1 - \alpha^{\gamma+1}}{1 - \alpha},$$

which we verify numerically for specific values in Section 4.3 rather than relying on symbolic cancellation alone.

Throughput and speedup. Throughput is $T = \mathbb{E}[N]/(C_{\text{draft}} + C_{\text{verify}})$; the speedup factor against the best single leg is $S = T_{\text{ens}}/T_{\text{base}}$.

Correlated legs. If legs are correlated, the independence product overestimates reliability. For exchangeable legs with common pairwise error correlation ρ , the probability that all legs fail simultaneously lies between the independent value $\prod_i (1 - \alpha_i)$ and the comonotone value $\min_i (1 - \alpha_i)$; correlated regimes are treated as labeled projections, not exact results.

4. Analysis

All input numbers are stated with their sources. Model parameters in Sections 3.1–3.3 are definitions of this paper, not empirical measurements; the design-point values $\alpha_1 = 0.5$, $\alpha_2 = 0.6$, $\alpha_3 = 0.7$, $\gamma = 4$, $(c_1, c_2, c_3) = (0.2, 0.3, 0.5)$, and $M_1^{\text{target}} = 0.1$ are illustrative assumptions declared here.

4.1 Case study I: divergence time of random inversions

Input: $L = 10$ (Section 3.1). The number of contiguous segments is $\frac{10 \cdot 11}{2} = 55$.

Average segment length. The total length over all segments is $S = \sum_{k=1}^{10} k(10 - k + 1)$:

$$\begin{aligned} S &= 1 \cdot 10 + 2 \cdot 9 + 3 \cdot 8 + 4 \cdot 7 + 5 \cdot 6 + 6 \cdot 5 + 7 \cdot 4 + 8 \cdot 3 + 9 \cdot 2 + 10 \cdot 1 \\ &= 10 + 18 + 24 + 28 + 30 + 30 + 28 + 24 + 18 + 10 = 220. \end{aligned}$$

Hence the average segment length is $\bar{\ell} = \frac{220}{55} = 4.0$.

Expected Hamming distance. Starting from the all-zero reference, flipping a segment of length ℓ increases the Hamming distance d by exactly ℓ on the first flip of each site; treating increments as independent with mean $\bar{\ell} = 4$ (an approximation discussed in Section 6), the expected distance after t steps is $\mathbb{E}[d(t)] \approx \min(L, 4t)$. Setting $4t_{\text{sat}} = 10$ gives $t_{\text{sat}} = \frac{10}{4} = 2.5$, and since t is an integer step count,

$$t_{\text{sat}} = \lceil 2.5 \rceil = 3.$$

Result (labeled projection P0, not exact arithmetic). Under the independence approximation of Section 3.1 — increments treated as independent with mean $\bar{\ell} = 4$, ignoring re-flipping of already-inverted sites — the saturation time is $t_{\text{sat}} = \lceil 10/4 \rceil = 3$ steps. The only shown arithmetic here is the ceiling computation $\lceil 10/4 \rceil = 3$; everything else in this paragraph is projection P0. We state as an expected-direction hypothesis, not an exact inequality, that re-flipping reduces the per-step Hamming increment below $\bar{\ell}$ after the first few steps, so the exact expected saturation time is plausibly larger than 3; the approximation carries an uncertainty of order one step and is used only as an order-of-magnitude estimate. The exact expectation would require tracking the full overlap structure of successive segments and is deferred to future work.

4.2 Case study II: moment dynamics and steering time

First-moment rate. Differentiating M_1 and using $\dot{x}_\omega = \omega u(t)$:

$$\dot{M}_1 = \int_{-1}^1 \omega \cdot \omega u(t) \rho(\omega) d\omega = u(t) \int_{-1}^1 \omega^2 \rho(\omega) d\omega.$$

Compute the integral:

$$\int_{-1}^1 \omega^2 \rho(\omega) d\omega = \frac{1}{2} \int_{-1}^1 \omega^2 d\omega = \frac{1}{2} \left[\frac{\omega^3}{3} \right]_{-1}^1 = \frac{1}{2} \left(\frac{1}{3} + \frac{1}{3} \right) = \frac{1}{3}.$$

Hence $\dot{M}_1 = \frac{1}{3}u(t)$: the broadcast input drives the first moment at rate $\frac{1}{3}$ per unit input.

Immobility of the mean and exact variance. The exact solution is $x_\omega(t) = \omega U(t)$ where $U(t) = \int_0^t u(s) ds$. Then

$$\bar{x}(t) = U(t) \int_{-1}^1 \omega \rho(\omega) d\omega = U(t) \cdot \frac{1}{2} \left[\frac{\omega^2}{2} \right]_{-1}^1 = 0,$$

so the mean advance is identically zero for the sign-symmetric uniform distribution. The variance is

$$\sigma^2(t) = \int_{-1}^1 (\omega U(t))^2 \rho(\omega) d\omega = U(t)^2 \cdot \frac{1}{3}.$$

At $t = 0.3$ with $u \equiv +1$: $U = 0.3$, so $\sigma^2 = \frac{0.09}{3} = 0.03$ and $\sigma = \sqrt{0.03} \approx 0.1732$. The broadcast input can only widen the population, not translate its mean.

Minimum steering time. From $M_1(T) = \frac{U(T)}{3}$ and $|U(T)| \leq T$, reaching $M_1^{\text{target}} = 0.1$ requires $U(T) = 0.3$, hence with $u \equiv +1$:

$$T_{\min} = \frac{0.3}{1} = 0.3 \text{ time units.}$$

Check: $M_1(0.3) = \frac{1}{3} \times 0.3 = 0.1$. ✓

Skewed distribution. Take the normalized right-triangular density $\rho_s(\omega) = \frac{2}{3}(\omega + 1)$ on $[0, 1]$ (normalization: $\int_0^1 (\omega + 1) d\omega = \frac{3}{2}$, so the factor $\frac{2}{3}$ makes the integral 1). Then

$$c_2 = \int_0^1 \omega^2 \rho_s(\omega) d\omega = \frac{2}{3} \int_0^1 (\omega^3 + \omega^2) d\omega = \frac{2}{3} \left(\frac{1}{4} + \frac{1}{3} \right) = \frac{2}{3} \cdot \frac{7}{12} = \frac{7}{18} \approx 0.3889.$$

The same target requires $T_{\min}^{(s)} = \frac{0.1}{7/18} = \frac{1.8}{7} \approx 0.2571$ time units — faster, because the gain mass is concentrated away from the sign-symmetric center.

Cross-validation and closure. From $x_\omega = \omega U$, direct computation gives $M_1(t) = U(t) \int_{-1}^1 \omega^2 \rho d\omega = \frac{U(t)}{3}$, matching the integrated moment equation exactly. For the second moment, $\dot{M}_2 = u(t) \int_{-1}^1 \omega^3 \rho d\omega = 0$, and indeed $M_2(t) \equiv 0$ since $x_\omega \propto \omega$; the hierarchy closes exactly in the linear case. Adding a drift term $\dot{x}_\omega = \omega u + a x_\omega$ produces the term $\int \omega^3 x_\omega \rho d\omega$, which is not expressible in terms of M_1, M_2 alone once x_ω is not proportional to ω : the hierarchy fails to close, and truncation error is the operative uncertainty in any enriched model.

4.3 Case study III: ensemble leg drafting computations

Baseline (best single leg). Best leg is L_3 with $\alpha_3 = 0.7$; by the cost normalization the baseline cost per position is $c_* = 1$. Compute $\alpha^{\gamma+1} = 0.7^5$:

$$0.7^2 = 0.49, \quad 0.7^3 = 0.343, \quad 0.7^4 = 0.2401, \quad 0.7^5 = 0.16807.$$

$$\mathbb{E}[N] = \frac{1 - 0.16807}{1 - 0.7} = \frac{0.83193}{0.3} = 2.7731.$$

Cross-check via direct sum (rejection terms $k = 0, \dots, 3$, since the $k = 4$ rejection outcome is indistinguishable from full acceptance and its mass belongs to the full-acceptance branch): $k = 1$: $1 \times 0.7 \times 0.3 = 0.21$; $k = 2$: $2 \times 0.49 \times 0.3 = 0.294$; $k = 3$: $3 \times 0.343 \times 0.3 = 0.3087$. Sum: 0.8127 . Full acceptance: $4 \times 0.2401 = 0.9604$. Total: $0.8127 + 0.9604 + 1 = 2.7731$. Matches the closed form. Baseline throughput: $C_{\text{draft}} = 4 \times 1 = 4$, total cost 5:

$$T_{\text{base}} = \frac{2.7731}{5} = 0.55462.$$

Ensemble. Effective acceptance probability:

$$(1 - \alpha_1)(1 - \alpha_2)(1 - \alpha_3) = 0.5 \times 0.4 \times 0.3 = 0.06, \quad \alpha_{\text{eff}} = 1 - 0.06 = 0.94.$$

Compute 0.94^5 : $0.94^2 = 0.8836$; $0.94^3 = 0.830584$; $0.94^4 = 0.78074896$; $0.94^5 = 0.7339040224$. Then

$$\mathbb{E}[N_{\text{ens}}] = \frac{1 - 0.7339040224}{0.06} = \frac{0.2660959776}{0.06} = 4.43493296.$$

Cross-check via direct sum ($q = 0.06$): $k = 1$: $0.94 \times 0.06 = 0.0564$; $k = 2$: $2 \times 0.8836 \times 0.06 = 0.106032$; $k = 3$: $3 \times 0.830584 \times 0.06 = 0.14950512$; sum 0.31193712 ; full acceptance $4 \times 0.78074896 = 3.12299584$; total $0.31193712 + 3.12299584 + 1 = 4.43493296$. Matches.

Ensemble throughput: $C_{\text{draft}} = 4 \times (0.2 + 0.3 + 0.5) = 4.0$, total cost 5.0:

$$T_{\text{ens}} = \frac{4.43493296}{5.0} = 0.886986592.$$

Speedup: $0.55462 \times 1.6 = 0.887392$, which exceeds 0.886986592 by 0.000405408 ; so

$$S = 1.6 - \frac{0.000405408}{0.55462} = 1.6 - 0.000731 = 1.599269 \approx 1.5993.$$

Break-even. The ensemble wins iff $\frac{\mathbb{E}[N(\alpha_{\text{eff}})]}{\gamma C + 1} > T_{\text{base}}$. Solve $\gamma C^* + 1 = \frac{4.43493296}{0.55462}$: since $0.55462 \times 8 = 4.43696$ exceeds 4.43493296 by 0.00202704 ,

$$\gamma C^* + 1 = 8 - \frac{0.00202704}{0.55462} = 8 - 0.003655 = 7.996345, \quad C^* = \frac{7.996345 - 1}{4} = 1.749086.$$

The three-leg ensemble wins as long as $\sum_i c_i < 1.7491$; at the design point $\sum_i c_i = 1.0$, the leg-compute budget can rise by a factor 1.749086 (a 74.9% increase over the current $\sum_i c_i = 1.0$), equivalent to 42.8% headroom relative to the break-even budget, since $\frac{1.749086 - 1.0}{1.749086} \approx 0.4283$.

Projection P1 (labeled projection, correlated legs). Assumptions: exchangeable legs with common effective acceptance $\bar{\alpha} = \frac{0.5+0.6+0.7}{3} = 0.6$ and comonotone failure ($\rho = 1$, all legs fail together), so $\alpha_{\text{eff}}^{(\rho=1)} = 0.6$. With $0.6^5 = 0.07776$ (from $0.6^2 = 0.36$, $0.6^3 = 0.216$, $0.6^4 = 0.1296$, $0.6^5 = 0.07776$):

$$\mathbb{E}[N] = \frac{1 - 0.07776}{0.4} = \frac{0.92224}{0.4} = 2.3056, \quad T = \frac{2.3056}{5} = 0.46112.$$

Since $0.46112 <$

References

-
- [1] Chaos in DNA inversions (Draft paper). arXiv:1105.1512v1. <https://arxiv.org/abs/1105.1512v1> [2] DRAFT: A Formally Verified Constructive Proof of the Consistency of Peano Arithmetic Using Ordinal Assignments. arXiv:2603.00487v1. <https://arxiv.org/abs/2603.00487v1> [3] TABED: Test-Time Adaptive Ensemble Drafting for Robust Speculative Decoding in LVLMS. arXiv:2601.20357v1. <https://arxiv.org/abs/2601.20357v1> [4] Master Functions and Equations for Perturbations of Vacuum Spherically-Symmetric Spacetimes. arXiv:2108.08668v3. <https://arxiv.org/abs/2108.08668v3> [5] Moment-Based Ensemble Control. arXiv:2009.02646v1. <https://arxiv.org/abs/2009.02646v1> [6] Prescriptive Master Integrals of Maximal Weight at Two Loops. arXiv:2610.10672v1. <https://arxiv.org/abs/2610.10672v1> [7] Asymptotic of Bergman Kernel (Master Thesis). arXiv:2202.03383v1. <https://arxiv.org/abs/2202.03383v1> [8] Ensemble Control on Lie Groups. arXiv:2008.03243v1. <https://arxiv.org/abs/2008.03243v1> [9] QNFO: QNFO Consilient Synthesis v2.0: The Shared Adelic Kernel [10] QNFO: NUMERATA: A Multi-Axis Framework for Evaluating Numeral Systems [11] QNFO: Connecting Geometrogenesis and Biogenesis [12] QNFO: Kappa/SIIT Cross-Validation: Exact Reproduction of Standard Model 1-Loop β -Functions and Resolution of the Harmonic-Paradigm Tension