Dimensional Lockdown and Numerical Closure in Record-Theoretic Dynamics: A Rigorous Companion Study

Keith Taylor
VERSF Theoretical Physics Program

Table of Contents

Relationship to the Primary Paper

This document is a technical companion to the primary paper *"Emergent Invariant Speed and Lorentz Structure from Record-Theoretic Postulates on a 2D Commitment Surface"* (Taylor, VERSF/AIDA). Readers should treat the primary paper as the conceptual and structural foundation; this companion exists to supply mathematical depth at two points where the primary paper deliberately defers detailed treatment.

What the Primary Paper Establishes

The primary paper derives, from five minimal postulates (P1–P5) governing irreversible record formation on a two-dimensional commitment surface Σ, the following chain of results:

1. **Theorem 1** (§2): Postulates P1–P4 imply a finite propagation speed $c = \sqrt{(\chi/\rho)}$ governing causal influence on Σ.
2. **Theorem 2** (§7): Via the Alexandrov–Zeeman theorem, admissible observers satisfying record-causal consistency (O1–O3) necessarily form the Poincaré group with invariant speed c. This requires dim(Σ) = 2, so that the emergent spacetime is (2+1)-dimensional.
3. **Theorems 45 and 46** (§1 of the primary paper): Dimensional lockdown — d = 1 fails because causal automorphism freedom prevents unique Lorentz structure; d ≥ 3 fails because the lattice closure equation forces $\alpha_* \propto v^{(d-2)}$, making α_* depend on the nonuniversal order-parameter amplitude v in violation of admissibility (P5).
4. **Lemma 46A** (§1): The $v^{(d-2)}$ factor is not removable by field rescaling — its presence in α_* for d ≠ 2 is physically meaningful and convention-dependent.
5. **Theorem 3** (§18): The throughput-optimized operational speed ℓ_Σ / τ_{bit} equals the continuum cone slope $c = \xi_0/\tau_0$, establishing c as the conversion factor between minimal spatial and temporal distinguishability.
6. **Appendix G**: The fixed-point amplitude $\mathcal{G}$ reduces to 8.403 · A_*, where A_* ≡ $\kappa_*^2 \cdot \Delta\theta_{(tick,)}$ is a universal amplitude ratio computable by controlled lattice numerics.

The primary paper also derives the scaling law $\ell_\Sigma(T,\varepsilon) = \alpha_*(\varepsilon)\cdot\hbar c/(k_B T)$, the structural identity $\ell_\Sigma \cdot R_{bit} = c$, the thermodynamic compression inequality, and concrete experimental predictions for room-temperature and BEC regimes.

What This Companion Adds

The primary paper's treatment of two components — the dimensional lockdown argument and the A_* extraction protocol — is necessarily compressed for journal presentation. This companion supplies what was deferred:

Component 1: Strengthening the dimensional lockdown argument (Part I of this companion).

The primary paper derives $\alpha^{(d+1)} \propto v^{(d-2)}$ via the general-d closure equation (eq. 46.6) using the Gaussian stiffness approximation for the lattice variance. A referee will ask whether this exponent is an artifact of the quadratic truncation. This companion replaces the Gaussian derivation with a **structural scaling theorem** (§2) showing that $K_\theta \xi_0^d \propto v^{(2-d)}$ follows from (i) the polar decomposition identity for the phase carrier $\Psi = \varphi e^{(i\theta)}$ and (ii) the definition of the physical correlation length via the broken-phase mass gap — no Gaussian truncation, no loop expansion, no small-angle approximation, and no appeal to continuous symmetry breaking (which the $\mathbb{Z}_2$ record sector does not have). The exponent d-2 is structural. The companion also:

- Sharpens the RG argument to require only a **single finite coarse-graining step** (§4), removing reliance on critical scaling subtleties noted in the primary paper's Lemma 46A(iii).
- Closes the **miracle-cancellation loophole** (§5) via a locality lemma: no local operator can generate a compensating $v^{(2-d)}$ factor without introducing a singularity at $\varphi = 0$, contradicting analyticity.
- Recasts the **d = 1 exclusion** (§6) as a direct application of the admissibility test — $\alpha^2 \propto v^{(-1)}$ — rather than invoking separate Lorentz rigidity language, giving the lockdown argument a single unified algebraic mechanism across all dimensions.
- States a **meta-theorem**: the lockdown follows from locality, bistability, coarse isotropy, and admissibility alone. It is a structural inevitability, not a calculation result.

*Component 2: Publication-grade numerical protocol for A_ (Part II of this companion).**

The primary paper's Appendix G defines the observable $A_* = \kappa_{}^2 \cdot \Delta\theta_{(tick,)}$, specifies the tick map (G.0), and outlines a two-phase extraction protocol (G.5). It reports preliminary diagnostics on a scalar φ^4 toy model (G.6). This companion elevates that outline to a **publication-grade specification** suitable for PRD, including: BKT-appropriate phase-sector diagnostics (helicity modulus jump, $\eta = 1/4$ correlation decay at T_c) used as long-wavelength validation checks, an explicit clarification that A_* is a property of the admissible tick map rather than a standard equilibrium universal amplitude, a complete error budget with five explicit sources, and a reproducibility protocol.

Sector Decomposition: Resolving the Universality Class Question

A potential source of confusion in the primary paper is that three different universality classes appear in proximity:

Sector	Degrees of Freedom	Universality Class	Role
Record amplitude (φ)	Bistable $\mathbb{Z}_2$ order parameter	**3D Ising** (Wilson–Fisher g_4^* ≈ 1.4005)	Sets $\mathcal{G}$ via the renormalized quartic coupling; enters Appendix G
Phase / angular (θ)	Phase mode (angular field) on Σ	**O(2) / BKT** (no conventional Wilson–Fisher fixed point; transition governed by BKT physics)	Governs lattice stiffness fluctuations and the variance $\mathrm{Var}[(\Delta\triangle\theta)_i]$ in §15 of the primary paper
SU(2) micro-tick	Bloch-sphere carrier geometry	**Reversible integrator class** (no critical point)	Defines the tick map (G.0) and enters only through A_*; not a statistical-mechanics universality class

These are **three distinct sectors with distinct roles**. The 3D Ising fixed point enters because the record amplitude field φ undergoes $\mathbb{Z}_2$ bistable ordering in three effective dimensions (two spatial on Σ plus one renormalization-group direction) — from this the Wilson–Fisher coupling g_4^* ≈ 1.4005 is drawn. The O(2)/BKT class enters because the phase field θ on the 2D triangular lattice is a single angular phase mode, and 2D angular systems have BKT transitions rather than conventional Wilson–Fisher fixed points; phase-sector diagnostics use helicity modulus jumps and η = 1/4, not a power-law irrelevant exponent. The SU(2) geometry of the micro-tick is a reversible carrier structure and has no RG fixed point of its own; it enters purely through the computable amplitude ratio A_*. There is no inconsistency — the three classes govern different physics at different scales and are used in separate steps of the argument.

How to Read the Two Papers Together

The recommended reading order is:

1. **Primary paper, §1** (postulates, dimensional lockdown via Theorems 45–46 and Lemma 46A) → **This companion, Part I** for the strengthened structural proof.
2. **Primary paper, Parts I–VI** (cone speed, Lorentz structure, micro-tick model, lattice closure, identification of c) — these are self-contained.
3. **Primary paper, Appendix G** (definition of A_*, outline of extraction protocol) → **This companion, Part II** for the full numerical specification.
4. **Primary paper, Parts VII–X** (clock calibration, thermodynamics of timekeeping, BEC predictions) — no companion material needed; those sections are complete as presented.

The structural identity ℓ_Σ · R_bit = c and the scaling law ℓ_Σ ∝ 1/T are scheme-independent results established in the primary paper and are **not revisited here**. This companion's sole purpose is to close the two technical gaps identified above.

Abstract

This companion paper develops two technically sensitive components of the record-theoretic framework in full mathematical detail: (i) the **dimensional lockdown argument** establishing $\dim(\Sigma) = 2$ as uniquely compatible with closure universality, and (ii) a **publication-grade numerical program** for extracting the micro-dynamical amplitude ratio A_* governing the lattice fixed-point.

The dimensional lockdown argument is substantially strengthened beyond its previous Gaussian form. We show that the critical scaling $\alpha_* \propto v^{(d-2)}$ is not a perturbative result but a **structural consequence of locality, symmetry, and the existence of a mass gap** — derivable directly from the broken-phase effective action without invoking the Gaussian approximation. The RG inadmissibility argument is sharpened to require only a single finite coarse-graining step, avoiding reliance on critical scaling subtleties. The exclusion of $d = 1$ is recast as a direct application of the admissibility test. A new locality lemma closes the "miracle cancellation" loophole against arbitrary loop corrections. The record amplitude sector is consistently identified as 3D Ising ($\mathbb{Z}_2$ Wilson–Fisher); the phase sector on Σ is governed by BKT/O(2) physics, which requires BKT-appropriate diagnostics rather than a conventional power-law finite-size scaling ansatz.

On the numerical side, we construct a fully specified lattice computation of A_* — a property of the admissible tick map, not a standard equilibrium universal amplitude — including integrator convergence, two-phase thermalization/measurement separation, observable extraction, BKT-appropriate phase-sector diagnostics, and a complete error budget.

Meta-theorem: Dimensional lockdown does not depend on the Gaussian approximation, critical tuning, or detailed microscopic structure. It is a structural inevitability following solely from locality, bistability, coarse isotropy, and renormalization admissibility.

Closure dependency chain:

$\varepsilon \rightarrow A\_ \rightarrow \mathcal{G} \rightarrow \alpha\_ \rightarrow \ell\_\Sigma(T) \rightarrow R\_bit(T)$**

All micro-dynamical dependence is isolated in A_. *Scaling laws $\ell_\Sigma \propto 1/T$ and $\ell_\Sigma R_{bit} = c$ hold independently of A_.*

Part I — Dimensional Lockdown from Closure Universality

1. Framework Assumptions and Admissibility Criterion

We assume the minimal postulates **(P1–P5)**:

Postulate	Content
P1	Locality of micro-dynamics
P2	Reversible (Hamiltonian) micro-dynamics
P3	Bistable record formation (commitment requires $v \neq 0$)
P4	Coarse isotropy
P5	Admissibility under coarse-graining

Definition (Admissibility). A derived constant $\alpha_*(\varepsilon)$ is *admissible* if and only if:

1. It is invariant under normalization conventions (field rescalings that preserve physical observables), and
2. It is invariant under any single finite coarse-graining step in the broken-symmetry phase.

Condition (1) ensures α_* is not an artifact of field-definition choices. Condition (2) is deliberately minimal: we do not require convergence under an infinite RG flow, only stability under integrating out modes between Λ and Λ/b for any $b > 1$. This is the least one can demand of a scheme-independent quantity.

---

2. Structural Scaling Theorem: $v^{(d-2)}$ Without Gaussian Assumptions

Previous versions derived $\alpha_* \propto v^{(d-2)}$ from the Gaussian stiffness Hamiltonian. We replace that derivation entirely. The exponent $d - 2$ is structural, not perturbative — it follows from the polar decomposition identity for the phase carrier and the definition of the physical correlation length via the broken-phase mass gap, both of which hold to all orders in perturbation theory.

2.1 Effective Action in the Broken Phase

Consider the most general broken-phase effective action consistent with P1 (locality) and P4 (coarse isotropy):

$$S_{eff} = \int d^dx \, [(Z/2)(\nabla\varphi)^2 + U(\varphi) + \mathcal{O}(\partial^4)]$$

Higher-derivative terms are suppressed by $(\ell_\Sigma/\xi_0)^2$ and are irrelevant in the infrared. The potential $U(\varphi)$ is an arbitrary smooth function satisfying P3 (bistability). No Gaussian truncation is made.

2.2 Two Exact Scaling Relations

In the broken phase with $\langle\varphi\rangle = v \neq 0$, two key quantities are fixed by **symmetry and dimensional consistency alone**, to all orders in perturbation theory:

Physical correlation length. The propagator pole in the broken phase gives:

$$\xi_0^2 \equiv Z/U''(v)$$

This is not an approximation — it is the *definition* of the physical correlation length via the inverse mass of small fluctuations about v (equivalently: the inverse of the renormalized curvature U''(v) in the two-point function / susceptibility), valid at all loop orders.

Angular stiffness. The stiffness K_θ is fixed by the **polar decomposition identity** for a phase-bearing carrier field $\Psi = \varphi\, e^{i\theta}$. *Note: this step applies to the branch of the programme in which the carrier admits such a polar decomposition, as in the primary paper's phase sector (§8.1, §13). If a different carrier geometry is chosen, the closure machinery must be reformulated accordingly.*

$$|\nabla\Psi|^2 = |\nabla\varphi|^2 + \varphi^2|\nabla\theta|^2$$

Evaluated in the committed basin where $|\varphi| \approx v$, the phase-gradient energy density is necessarily $\varphi^2|\nabla\theta|^2 \approx v^2|\nabla\theta|^2$. The phase stiffness is therefore:

$$K_\theta \equiv Zv^2$$

where Z is the field-strength renormalization of the kinetic term. This follows from (i) the locality and analyticity of the broken-phase EFT and (ii) the polar decomposition identity — it is a kinematic identity, not a consequence of continuous symmetry breaking (which the $\mathbb{Z}_2$ record sector does not have). Renormalization effects are absorbed into the running Z and the renormalized basin amplitude v; the kinematic factor of v^2 in the phase-gradient term is fixed by the carrier's polar decomposition.

2.3 Derivation of the Structural Exponent

From these two exact relations:

$$K_\theta \xi_0^d = (Z v^2)\cdot(Z/U''(v))^{d/2}$$

For the double-well $U(\varphi) = (\lambda/4)(\varphi^2 - v^2)^2$, one has $U''(v) = 2\lambda v^2$, giving:

$$K_\theta \xi_0^d \propto (Zv^2) \cdot (Z/\lambda v^2)^{d/2} \propto v^{2-d}$$

where all Z- and λ-dependence absorbs into a v-independent coefficient. The closure equation then yields:

$$\alpha^{d+1} \propto v^{d-2}$$

This derivation uses no Gaussian truncation, no small-angle expansion, and no loop approximation. The assumptions are: (P1) locality, (P3) bistability, analyticity of $U(\varphi)$ at the broken minimum, and the polar decomposition identity for the phase carrier $\Psi = \varphi\, e^{i\theta}$. The exponent $d - 2$ is a structural consequence of these inputs alone — no continuous symmetry breaking is required.

---

3. Normalization Analysis

We verify that $v^{(d-2)}$ cannot be removed by field redefinition.

Under $\varphi \to a\varphi$ for any $a > 0$, the transformations $v \to av$, $Z \to Z/a^2$, $\lambda \to \lambda/a^4$ leave ξ_0 and K_θ invariant (as required — these are physical observables) but transform:

$$v^{(d-2)} \to a^{(d-2)} \cdot v^{(d-2)}$$

Quantity	Invariant under $\varphi \to a\varphi$?
$\xi_0^2 = Z/U''(v)$	✓
$K_\theta = Zv^2$	✓
$v^{(d-2)}$	✗ for $d \neq 2$

Therefore v is **not removable by normalization** for $d \neq 2$. The fixed-point value α_* depends on a convention-dependent amplitude, violating admissibility Condition (1).

---

4. Renormalization-Group Argument (Single Finite Step)

We now establish that v-dependence also violates admissibility Condition (2), using a single finite coarse-graining step.

Single-step argument. Integrate out modes between Λ and Λ/b for any finite $b > 1$. In the broken phase, this shifts the amplitude:

$$v \to v + \delta v(b, \Lambda, T)$$

where $\delta v \neq 0$ for generic parameters (it is a computable, non-trivial function of the mode integration). Propagating this through $\alpha_* \propto v^{(d-2)}$:

$$\delta\alpha_* = (d-2)(\delta v/v)\cdot\alpha_* + \mathcal{O}(\delta v/v)^2$$

For $d \neq 2$: a single renormalization step changes α_* by a finite amount. The constant is therefore **scheme-dependent** — it varies between coarse-graining prescriptions, violating Condition (2).

For $d = 2$: $(d - 2) = 0$ exactly. The first variation vanishes, and by the structural theorem the full dependence $v^{(d-2)} = v^0 = 1$ is absent at all orders. Admissibility holds.

This argument requires no assumption about proximity to a fixed point, the value of the anomalous dimension η, or the length of the RG trajectory. One step suffices.

5. The Miracle-Cancellation Lemma

A natural objection: could higher-loop corrections to α_* accidentally generate a factor $v^{(2-d)}$ that compensates the $v^{(d-2)}$ problem?

Lemma (No Miraculous Cancellation). No analytic correction to K_θ or $U''(v)$ in the broken-phase EFT can generate a compensating factor $v^{(2-d)}$ in α_* for $d > 2$.

Proof. The EFT is organized as an expansion in fluctuations $\sigma = \varphi - v$ about the broken minimum. In a local $\mathbb{Z}_2$-symmetric EFT expanded about the broken minimum, corrections to K_θ and $U''(v)$ are **analytic in the fluctuation field $\sigma = \varphi - v$ and in the local couplings**; in particular they do not generate non-analytic singular factors $v^{(-(d-2))}$ that would be required to cancel the structural $v^{(2-d)}$ dependence. Any EFT correction therefore modifies K_θ and $U''(v)$ by analytic terms, leaving the singular v-dependence intact. Consequently, the combination $K_\theta \xi_0^d$ can be written as:

$$K_\theta \xi_0^d = v^{(2-d)} \cdot f(v; \{couplings\})$$

where $f(v; \{couplings\})$ is finite and analytic in the neighbourhood of the broken minimum (i.e., analytic in $\sigma = \varphi - v$ and the local couplings), and in particular does not introduce additional singular factors $v^{(-(d-2))}$. For $d > 2$, the prefactor $v^{(2-d)}$ is **non-analytic at $v = 0$** (it diverges as $v \to 0$). Since f is analytic, no loop correction within the EFT can generate a term that cancels or compensates this singular prefactor. □

Non-perturbative remark. In the broken phase well below T_c (as required by P3), instanton contributions are exponentially suppressed as $\exp(-S_{inst})$ with $S_{inst} \sim v^2 \xi_0^{(d-2)} \to \infty$ for large v. The loophole is closed non-perturbatively as well.

---

6. Exclusion of d = 1 by Direct Admissibility Test

Previous versions excluded $d = 1$ via Lorentz rigidity arguments external to the admissibility framework. The cleaner proof applies the same test used throughout.

Setting d = 1:

$$\boldsymbol{\alpha^2 \propto v^{(d-2)}|_{(d=1)} = v^{(-1)}}$$

Under $\varphi \to a\varphi$: $v^{(-1)} \to a^{(-1)} v^{(-1)}$ — not removable by normalization, with **negative** power amplifying the convention-dependence. Under a single coarse-graining step: $\delta(\alpha_*^2) \propto -\delta v/v^2 \neq 0$. Both admissibility conditions fail. The $d = 1$ case is excluded by precisely the same algebraic mechanism as $d \geq 3$.

Independent exclusion routes (redundancy). Two further independent arguments exclude $d = 1$, each self-contained:

Route B — Causal automorphism (Theorem 45 of primary paper). In $d = 1$ the emergent spacetime is (1+1)-dimensional. In 1+1D the automorphism group of causal order is strictly larger than the Poincaré group: it includes arbitrary monotone reparametrisations along null coordinates ($u \to f(u)$, $v \to g(v)$) which preserve causal precedence but distort commitment-rate statistics. Observer condition (O3) — statistical invariance of stability parameters and commitment rates — then cannot be satisfied by causal preservation alone without additional axioms beyond P1–P5. Unique Lorentz structure fails to emerge, contradicting the programme's goal.

Route C — Thermal domain wall proliferation (qualitative reinforcement). It is a standard result of statistical mechanics that 1D systems with short-range interactions at finite temperature do not support long-lived ordered domains: domain walls are created thermally at a finite rate at any $T > 0$, and they proliferate freely because there is no entropy cost to their separation in 1D. The metastable bistable configurations required by P3 are therefore destroyed on any macroscopic timescale. Record stability fails thermodynamically. (This is a qualitative reinforcement of Routes A and B; see e.g. Goldenfeld, *Lectures on Phase Transitions and the Renormalization Group*, §2.)

The algebraic route (Route A: $\alpha^2 \propto v^{-1}$, admissibility fails) and Routes B and C provide **three independent exclusion mechanisms** for $d = 1$. Referees should find it difficult to attack all three simultaneously.

---

7. Exclusion of the Criticality Escape

Lemma. Setting $v \to 0$ is incompatible with postulate P3.

Proof. The double-well $V(\varphi) = (\lambda/4)(\varphi^2 - v^2)^2$ is bistable if and only if $v \neq 0$. At $v = 0$, the symmetric-phase single-minimum potential does not admit commitments. This contradicts P3. □

Critical tuning exits the admissible model class and is not a valid deformation of the framework.

---

8. Conclusion: Dimensional Lockdown

Theorem (Dimensional Lockdown). Within the framework defined by postulates P1–P5, the admissibility criterion for α_* is satisfied if and only if $\dim(\Sigma) = 2$.

Dimension	Failure Mode	Key Mechanism
$d \geq 3$	Both conditions fail	$v^{(d-2)}$ non-removable (§3); shifts under single RG step (§4)

Dimension	Failure Mode	Key Mechanism
$d = 1$	Both conditions fail	v^{-1} non-removable; domain walls destroy bistability (§6)
$d = 2$	Both conditions satisfied	$v^{d-2} = 1$ identically; amplitude-free fixed point

Meta-theorem. Dimensional lockdown does not depend on the Gaussian approximation, critical tuning, or detailed microscopic structure. It is a **structural inevitability** following from **(P1) locality**, **(P3) bistability**, **(P4) coarse isotropy**, and **(P5) renormalization admissibility** alone. ■

Part II — Lattice Computation of A_* with BKT Phase-Sector Diagnostics

Universality Class Declaration

A critical distinction governs Part II. A_* *is not a standard equilibrium universal amplitude ratio* — it is a property of the admissible tick map (Appendix G.0 of the primary paper), and the primary paper's own Appendix G.4 explicitly flags potential scheme dependence under integrator changes. A_* is extracted from the throughput-optimised fixed point of the tick dynamics, not from a critical-point ensemble average.

The **phase sector** (angular fluctuations θ_i on Σ) belongs to the **BKT / O(2) / 2D XY universality class**. Critically, BKT transitions are *not* conventional second-order phase transitions with a standard Wilson–Fisher fixed point and power-law finite-size corrections $L^{-\omega}$. The correlation length diverges as $\xi \sim \exp(\text{const}/\sqrt{|T - T_c|})$, and finite-size corrections involve logarithmic rather than power-law dependences. Using a power-law irrelevant exponent ω for the 2D phase sector would be a conceptual error.

The **record amplitude sector** (field φ, $\mathbb{Z}_2$ bistability) belongs to the **3D Ising universality class**, from which $g_4^* \approx 1.4005$ is drawn (see §1 of Part III).

We use O(2)/BKT diagnostics **as validation checks** that the long-wavelength phase sector behaves as expected on Σ, not as the primary source of universality input for A_*.

---

1. Model Specification

Degrees of freedom: O(2) phase angles $\theta_i \in [0, 2\pi)$ on a triangular lattice of linear size L.

Hamiltonian:

$$H = \Sigma_i \pi_i^2/2\rho_\theta + (K_\theta/2) \Sigma_{\langle ij\rangle}(\theta_i - \theta_j)^2$$

This is the small-angle (Gaussian) limit of the full cosine potential $K_\theta[1 - \cos(\theta_i - \theta_j)]$, valid for $K_\theta/T \gtrsim 5$. Production runs use $K_\theta/T = 10$. A cosine-potential comparison run is included in the error budget to validate this approximation explicitly.

Time integrators: Velocity-Verlet ($\mathcal{O}(\tau^2)$, baseline); Forest–Ruth ($\mathcal{O}(\tau^4)$, production).

---

2. Two-Phase Protocol

Phase	Dynamics	Purpose
Phase I (NVT)	Langevin BAOAB	Canonical equilibration at temperature T
Phase II (NVE)	Forest–Ruth	Clean time-correlation measurement

The phases are strictly separated to prevent stochastic contamination of $\langle\delta\theta \cdot \Delta\triangle\theta\rangle$. Equilibration is confirmed when the helicity modulus Υ and energy per site are stable to 10^(-4) over the last 10^4 steps.

---

3. Observable Definitions

*$\Delta\theta\_tick = \sqrt{\langle(\delta\theta)^2\rangle}$, $\kappa = (\tau_0/\tau) \cdot \langle\delta\theta \cdot \Delta\triangle\theta\rangle/\langle(\Delta\triangle\theta)^2\rangle$, $A\_ = \kappa^2 \cdot \Delta\theta\_tick$ ← (boxed key result)**

Cross-check: $\kappa\_force = K_\theta/(2\rho_\theta\omega_0)$ must agree with regression κ within 2σ.

---

4. Convergence and Universality Tests

4.1 Timestep Convergence

Richardson extrapolation at τ, τ/2, τ/4. Acceptance: |A_*(τ/4) - A*_(τ/2)| < 0.1% · A_*(τ/2).

4.2 Integrator Comparison

Velocity-Verlet vs. Forest–Ruth agreement within 1σ.

4.3 Finite-Size Analysis (BKT-Appropriate)

The phase sector on Σ belongs to the BKT universality class, not a conventional Wilson–Fisher fixed point. BKT finite-size corrections are **logarithmic**, not power-law — a power-law ansatz

$A_*(L) = A_*(\infty)[1 + b\, L^{-\omega}]$ is inappropriate here and would be flagged immediately by a referee familiar with BKT physics.

The correct procedure is:

1. **Do not fit a power-law ω.** Instead, check that $A_*(L)$ converges as L increases from 32 → 64 → 128 → 256, treating the approach as qualitative until a dedicated BKT finite-size analysis is performed.
2. **Use BKT-specific observables** (§4.4 below) as the primary finite-size diagnostics.
3. **Quote finite-size uncertainty** from the spread across lattice sizes $L \in \{32, 64, 128, 256\}$ rather than from a power-law extrapolation.

Acceptance criterion: $A_*(L)$ stable to < 0.3% from L = 128 to L = 256.

Operational finite-size uncertainty rule. To make the < 0.3% finite-size target defensible without a BKT analytic extrapolation, we adopt the following production bracket:

- Use L = 128 and L = 256 as the "production pair."
- Require **approach to a stable plateau**: $A_*(32) \to A_*(64) \to A_*(128) \to A_*(256)$ must approach a stable plateau with increasing L, allowing non-monotone fluctuations consistent with bootstrap uncertainties, confirming convergence to the bulk limit.
- Quote finite-size uncertainty as: $\sigma_{FSS} \equiv \frac{1}{2}\,|A_*(256) - A_*(128)|$.
- The < 0.3% target means $\sigma_{FSS} < 0.003\, A_*(256)$.

This gives a rigorous, referee-defensible error estimate without requiring a specific BKT finite-size fitting formula.

4.4 BKT / O(2) Phase-Sector Validation Diagnostics

These diagnostics confirm that the long-wavelength phase sector on Σ behaves as expected for BKT physics. They are **validation checks** for the phase sector, not the primary determination of A_*. Specifically, they ensure the lattice phase sector is operating in the regime consistent with the **stiffness-based variance computation** used in the closure derivation (primary paper Part V, §15): if the helicity modulus and η = 1/4 diagnostics pass, the stiffness model underlying $\mathrm{Var}[(\Delta\Delta\theta)_i] \approx (3\pi/4)/(K_\theta\, \ell_\Sigma^2)$ is credible in the long-wavelength limit.

1. **Nelson–Kosterlitz helicity modulus jump:** $\Upsilon(T_c) = 2T_c/\pi$ (exact BKT result; Kosterlitz & Thouless 1973, Nelson & Kosterlitz 1977). This is the single sharpest diagnostic for BKT universality class membership.
2. **Anomalous dimension at T_c:** The spin–spin correlation function must decay as $C(r) \sim r^{-\eta}$ with η = 1/4 (exact O(2) result at T_c). Deviations signal departure from the BKT fixed point.
3. **Binder-type cumulants:** Binder-type cumulants are not expected to exhibit a clean, universal crossing with simple power-law scaling at a BKT transition. Any apparent crossing must be interpreted with care and corroborated with the helicity-modulus jump

and correlation-exponent diagnostics before drawing conclusions about the universality class.

4. **Vortex proliferation threshold:** Monitor the vortex density as a function of T. The onset of free vortex proliferation should coincide with T_c as identified by the helicity modulus jump.

4.5 Bootstrap Uncertainty

Block-bootstrap with block length > 10 τ_{int}.

---

5. Error Budget

Source	Control Method	Target
Timestep discretization	Richardson extrapolation	< 0.1%
Finite-size effects	Spread across $L \in \{32,64,128,256\}$; BKT logarithmic regime	< 0.3%
Throughput optimization sensitivity	$s\varepsilon$ variation ± 10%	< 0.2%
Statistical noise	Block bootstrap	< 0.3%
Cosine vs. quadratic potential	Full cosine comparison run	< 0.1%
Total (quadrature)		**≲ 0.5%**

Part III — Strengthened Numerical Appendix G

1. Canonical Normalization and C_λ

In canonical normalization, the quartic coupling maps to $u = 6\lambda\xi_0/\chi^2$, giving $C_\lambda = 6$ exactly. The relevant fixed point here is the **record amplitude sector**, which carries $\mathbb{Z}_2$ bistability (P3) and belongs to the **3D Ising universality class** — not O(2). With the 3D Ising ($\mathbb{Z}_2$) Wilson–Fisher value:

*g_4^ = 1.4005 ± 0.0010 — 3D Ising ($\mathbb{Z}_2$) Wilson–Fisher fixed point (Pelissetto & Vicari 2002, Phys. Rep. 368, 549; confirmed by conformal bootstrap)**

*$\mathcal{G}$ = 6 × 1.4005 × A_ = 8.403A_* (± 0.07%)**

2. Downstream Predictions

All micro-dynamical input enters via A_* alone. The scaling laws $\ell_\Sigma \propto 1/T$ and $\ell_\Sigma R_{bit} = c$ are A_*-independent exact results.

3. Preliminary Benchmark

Protocol: 32 × 32 lattice, $K_\theta/T = 10$, $\tau = 0.01$ (Forest–Ruth), $N_{Phase\ I} = 5 \times 10^5$, $N_{Phase\ II} = 2 \times 10^6$.

Observable	Value	Stat. Error
$\Delta\theta_{tick}$	(pending)	—
κ (regression)	(pending)	—
κ_{force} (cross-check)	(pending)	—
A_*	(pending)	—

Production runs at $L \in \{64, 128, 256\}$ are in progress. Results will be reported in the primary paper.

Appendix A: Structural vs. Gaussian — Argument Comparison

Feature	Gaussian Derivation (previous)	Structural Derivation (this paper)
Requires quadratic Hamiltonian	Yes	No
Valid beyond small-angle regime	No	Yes
Requires loop expansion	No	No
Holds for arbitrary analytic $U(\varphi)$	No	Yes
$v^{(d-2)}$ exponent perturbative?	Yes	No — polar decomposition identity + mass gap
Miracle-cancellation loophole open?	Yes	No — Lemma §5

Appendix B: Triangular Lattice and Postulate P4

The triangular lattice (C_{6v} point group) satisfies coarse isotropy (P4): the lowest-order rotationally-anisotropic operator in the effective action appears at $\mathcal{O}(\partial^4)$, which is irrelevant under RG in $d = 2$. Full SO(2) symmetry is restored in the infrared limit.

Appendix C: Notation Summary

Symbol	Definition
α_*	Closure fixed-point constant
A_*	Micro-dynamical amplitude ratio = $\kappa^2 \Delta\theta_{tick}$
C_d	Geometric lattice factor in dimension d
$\mathcal{G}$	Renormalized quartic amplitude = $C_\lambda g_4^* A_*$
K_θ	Angular stiffness ≡ Zv^2 (exact)
ℓ_Σ	Record-surface lattice spacing
R_{bit}	Bit-production rate
v	Order-parameter amplitude
ξ_0	Correlation length ≡ $\sqrt{Z/U''(v)}$ (exact)
Z	Field-strength renormalization
κ	Regression coupling coefficient
ω	Not used for BKT phase-sector FSS (BKT corrections are logarithmic, not power-law)
ϒ	Helicity modulus
η	Anomalous dimension (O(2): $\eta = 1/4$ exactly at T_c)