

The two-dimensional $O(3)$ sigma model: a volume-uniform mass-gap certificate and the one inequality that remains

Andrei - Sebastian Ursachi

contact@andreiursachi.eu

independent researcher <https://orcid.org/0009-0002-6114-5011>

Research Article

Keywords:

Posted Date: October 21st, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11440706/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

The two-dimensional $O(3)$ sigma model: a volume-uniform mass-gap certificate and the one inequality that remains

Andrei Ursachi

Independent Researcher ORCID 0009-0002-6114-5011 andreiursachi.eu

20 October 2026 — preprint, version 2.2 (certified range extended to $\beta = 0.85$; the remaining inequality restated; erratum to version 2.1)

Abstract

The two-dimensional lattice $O(3)$ sigma model is the standard warm-up for the Yang–Mills mass-gap problem: asymptotic freedom predicts a gap at every coupling, and no proof exists. We give a finite criterion for the gap — two partition-function ratios (“purities”) on small tori — and prove that when it is met at one seed size, the transfer matrix has a gap uniform in the volume at that coupling (Lemmas P, B, C). The criterion is checked by a certified computation (exact spherical-harmonic transfer matrices with rigorous truncation and rounding bounds: a reflection-positivity lower bound valid on tori of any side, a convexity upper bound, symmetry-reduced exact traces on rings of up to six sites, and Rayleigh bounds from exactly contracted matrix-product states on rings of up to twelve sites), giving a volume-uniform mass gap at eight couplings up to $\beta = 0.85$ (Theorem 1: for example $\gamma_L \geq 0.0195$ for every circle length $L \geq 6$ at $\beta = 0.85$, beyond the Aizenman–Simon bound $3/4$ and the fixed- β decay certificate 0.7701 of this paper). For arbitrary coupling the whole question reduces to a single explicit statement, C1'': at every β the doubling defect $\Delta(L, L) = 4 \log Z_{L \times L} - \log Z_{2L \times 2L}$ falls below $\Delta^* = 0.2239$ at some finite L . An exact dichotomy (Theorem 2) shows that this is sharp: at every coupling either $\Delta \geq \Delta^*$ for all L , or $\Delta \rightarrow 0$ doubly exponentially and the volume-uniform gap holds, so the interval $(0, \Delta^*)$ is empty for $\lim \Delta$. A gapless infrared theory forces $\Delta_\infty = 3 \log Z_{\text{IR}}(i) \geq \pi c/2 > \Delta^*$, which is why the two control models that must fail ($O(2)$ in its BKT phase and the easy-plane deformed $O(3)$) fail at every L (measured plateaus 4.3 and 5.7), while $O(3)$ drives Δ to zero at $\beta = 1.4$ and 1.6 (Monte Carlo, transfer matrices and DMRG; provisional). At weak coupling a rigorous Laplace expansion (Theorem 5) shows that the additive form of the inequality used in version 2.1 is false — the decrement vanishes as $\beta \rightarrow \infty$ at fixed L — and that the uniform drop lives in the running coupling, with the one-loop constant $3 \ln 2/(2\pi)$ for $O(3)$ and 0 for both controls; the rotor-ring deficit tends to $\ln 2/\pi$ from above. We prove a volume-uniform mass gap for $\beta \leq 0.85$ (at the eight certified couplings) and reduce the Polyakov conjecture at every coupling to a single explicit statement (C1''), which holds at every point we have tested — rigorously at the certified couplings, by Monte Carlo at $\beta = 1.4$ and through a proxy at 1.6 — and fails, as it must, for the two control models. The paper also contains the scale maps, the stiffness–Casimir identity, the rotor-chain certificates, an audit of the 2026 machine-generated proofs, and all adverse outcomes.

1 Introduction

The classical Heisenberg model on $\mathbb{Z}^2$, equivalently the lattice $O(3)$ non-linear sigma model, has Gibbs weight $\exp(\beta \sum_{\langle xy \rangle} s_x \cdot s_y)$ with $s_x \in S^2$. Polyakov’s 1975 one-loop calculation [1] makes the coupling asymptotically free and predicts a mass gap at every β , with correlation length $\xi \sim e^{2\pi\beta}/(2\pi\beta)$. Numerically this has been confirmed to $\xi \approx 10^5$ [3, 5], and the exact mass in units of the perturbative scale is known from the Bethe ansatz [2]. Rigorously, before October 2026, exponential decay of correlations was proved only at small β (Dobrushin uniqueness gives $\beta < \ln(5/3) \approx 0.511$ for this model; the Aizenman–Simon inequality [8] gives $\beta < 3/4$), while at all β one has only the absence of long-range order and power-law upper bounds [9]. The

abelian $O(2)$ model has a Berezinskii–Kosterlitz–Thouless (BKT) phase with power-law decay at large β [10], with a stiffness that stays bounded below by $2/\pi$ [11]; a proof of the gap for $O(3)$ must therefore use the non-abelian structure in an essential way. Garban and de Montgolfier [13] showed in October 2026 that arbitrarily small easy-plane deformations of $O(3)$ also have a BKT phase, so a proof must use the exact $O(3)$ invariance, not merely smoothness.

This problem is widely regarded as the warm-up for the Clay Millennium problem on four-dimensional Yang–Mills theory [15]: both theories are asymptotically free, both are expected to generate a mass dynamically, and in both the difficulty is the “middle band” of scales where the running coupling is of order one and no small parameter exists. On 6 October 2026, OpenAI released a large set of machine-generated manuscripts [16], among them a family of five papers on the $O(3)$ and $O(4)$ models, one of which claims exponential decay for all $O(n \geq 3)$ models at every β and comes with a Lean formalization of its finite-volume statement. These manuscripts are unrefereed.

What this paper does. It reports, from one day of work (19 October 2026), a family of finite-size measurements on the $O(3)$ model and its controls $O(2)$, $O(4)$ and λ -deformed $O(3)$, designed around a single question: is there a one-variable, β -independent map from scale L to scale $2L$ that strictly decreases the spin stiffness, and does it fail where it must ($O(2)$ BKT, $\lambda < 1$)? It rederives, in sector form, the exact relation between the square-torus helicity modulus and the twisted transfer trace (Section 5), states precisely the two missing pieces that would close the argument and what the standard tools cannot do (Section 6), reports a computer-assisted exponential-decay certificate at fixed β up to 0.785 (Section 4.7), proves by computer a positive spectral gap for finite $O(3)$ quantum rotor chains over a coupling grid and locates the cost wall of that certificate (Section 11), measures how the gap appears as a relaxation rate under reversible stochastic dynamics and derives an exact bound on that rate (Section 12), analyses the doubling step in the exact character basis, proving where positivity alone stops and isolating the first unproved inequality (Section 10), and gives an independent audit of the October 2026 machine-generated proofs (Section 13). Version 2.1 added the main result: a finite criterion whose two seed inequalities, certified on small tori, imply a mass gap uniform in the circle length (Section 7). Version 2.2 extends the certified couplings to $\beta = 0.85$, proves the dichotomy that makes the remaining statement C1'' sharp (Section 8), and settles the weak-coupling end rigorously at fixed size (Section 9). All simulations are reported with their pre-registrations, deviations and adverse outcomes.

What this paper does not do. It does not prove the mass gap, at any β beyond the certified region, nor asymptotic scaling, nor anything in the continuum. It claims no priority over the OpenAI lemmas; the stiffness-iteration lemma of their $O(4)$ paper is the pinned-square form of the inequality studied here, and the decay certificate of Section 4.7 is a human-checkable independent result on a small β range and a test of certificate machinery, not a first (the OpenAI manuscript 215-a, unrefereed and Lean-verified in finite volume, claims decay for all β). Pre-registered and post-hoc analyses are labelled throughout; a cubic fit that was pre-registered and failed is reported as failed.

Relation to Molke and Wolff (2001). The twist-defined finite-volume coupling studied here is not new. Molke and Wolff [31] defined the helicity-modulus coupling $g_Y^2(L)$ of the $O(3)$ model through boundary conditions periodic up to an $O(n)$ rotation, expressed the helicity modulus through the twisted transfer trace and the squared-generator expectation (their eqs. (2.7)–(2.11)) [details unverified], compared it with the finite-volume mass-gap coupling $\bar{g}_m^2(L) = 2m(L)L/(N - 1)$ of [19], and established the continuum relation between the two at all scales, with the weak-coupling matching $g_Y^2 = \bar{g}_m^2 + 0.1635 \bar{g}_m^4 + O(\bar{g}_m^6)$ at equal box size (coefficient as quoted in the referee report from [31]; not re-derived here). The sector identity of Section 5 is a regrouping of

that structure by $\text{SO}(N)$ representations; the contribution of this paper lies in the measurements, the controls ($\text{O}(2)$, λ -deformed $\text{O}(3)$, $\text{O}(4)$), the explicit bounds on sector weights, the decay certificate and the rotor-chain certificate, not in the twist observable itself. Version 1 of this paper identified $g_{\text{T}}^2 = 1/\rho_s$ with $\bar{g}_m^2$ exactly; that identification is withdrawn (Sections 4.6 and 3.4).

Changes in versions 1.1, 2.0, 2.1 and 2.2. Two sections are added relative to version 1: the finite-chain rotor certificate (Section 11) and the relaxation-dynamics study (Section 12), both from runs completed on 19 October 2026 after version 1 was frozen; their adverse outcomes are appended to Section 3.3 and their files to Section 15. Version 2.0 (same day) applies an external referee report in full — every change to a version-1 statement is listed in Section 3.4 — and adds the character-basis analysis of the doubling step (Section 10). Version 2.1 (20 October 2026) adds Section 7 (finite criterion and volume-uniform certificate), the doubling-defect seed and the statement of C1' (superseded in v2.2, E14), a transfer-matrix test of the first unproved step with a rotor-ring lemma and a transfer lemma (Section 10), the negative result on an N -uniform rotor-chain certificate (Section 11), the second-action run and direct sector weights (Section 4.6), an erratum to version 2.0 (Section 3.5), a response to the review of version 2.0 (Appendix B) and a data-and-code appendix for the new files (Appendix C). Version 2.2 (same day): Theorem 1 extended from two to eight certified couplings ($0.60 \leq \beta \leq 0.85$) with two new lemmas (U, T1') and a pass/fail map over seed shapes; new Section 8 (exact inequalities for the doubling defect, the dichotomy Theorem 2, the converse Theorem 3, the universal constant of a gapless infrared theory, the restated remaining statement C1'', Monte Carlo of the defect); new Section 9 (Theorem 5, the exact large- β expansion of the defect, the rotor-ring semiclassics, the running-coupling form, the obstruction); the DMRG rates folded into Section 10.1; erratum E14 to version 2.1 (Section 3.5): the additive inequality C1' is false at large β and is replaced by C1''; the seed threshold raised from 0.1 to 0.2239. Every number from the new tracks that rests on a truncation extrapolation, on exact diagonalization at small size or on a single Monte Carlo series is marked [provisional].

2 Model, observables and estimators

Model. $L \times L$ (or $L \times T$) periodic lattices; weight $\exp(\beta \sum_{\langle xy \rangle} s_x \cdot s_y)$ with $s \in S^{N-1}$, $N = 2, 3, 4$. The λ -deformed $\text{O}(3)$ model [13] uses $\langle \sigma, \sigma' \rangle_\lambda = \sigma^x \sigma'^x + \sigma^y \sigma'^y + \lambda \sigma^z \sigma'^z$ with $\sigma \in S^2$ and $\lambda \in (-1, 1)$; the definition was taken from the arXiv abstract.

Algorithm. Wolff single-cluster reflections [4] along random axes ($\text{O}(N)$) or symmetry axes only (controls), with a *fixed* number of clusters per measurement, set once after thermalization to $\text{round}(2V/\langle |C| \rangle)$ (Casimir map) or $\lfloor V/\langle |C| \rangle \rfloor$ (ξ study). A size-dependent stopping rule used earlier the same day biased observables and was removed (Section 3.3). For the λ -model one Metropolis sweep per measurement is added because reflections conserve $|s_z|$. Errors are leave-one-out jackknife over 50 blocks (20 blocks for the earlier Swendsen–Wang runs), with τ_{int} monitored; no block was shorter than $10\tau_{\text{int}}$.

Correlation length. $\xi_2(L) = \sqrt{\chi/F - 1}/(2 \sin(\pi/L))$ with $\chi = \langle |M|^2 \rangle/V$ and F the structure factor at $|p| = 2\pi/L$, both from Wolff improved estimators (direct estimators as a cross-check).

Helicity modulus and rotor stiffness. Let $e_x = \sum_{x\text{-bonds}} s \cdot s'$ and $J_{ab} = \sum_{x\text{-bonds}} (s_a s'_b - s_b s'_a)$. Twisting by $R_{ab}(\theta)$ across the x -seam and spreading the twist uniformly over the x -bonds gives,

by $O(N)$ invariance (Section 5),

$$\Upsilon_{\text{axis}} = \frac{1}{V} \left[\frac{2\beta}{N} \langle e_x \rangle - \beta^2 \langle J_{ab}^2 \rangle \right], \quad \rho_s = \frac{N}{2} \Upsilon_{\text{axis}} = \frac{1}{V} \left[\beta \langle e_x \rangle - \frac{\beta^2}{N-1} \sum_{a < b} \langle J_{ab}^2 \rangle \right]. \quad (1)$$

For $N = 3$ this is $\rho_s = [\beta \langle \sum s \cdot s' \rangle - \frac{1}{2} \beta^2 \langle |\sum s \times s'|^2 \rangle] / V$. At tree level ρ_s is the coefficient of $|\dot{s}|^2$ in the collective-rotor action, so that the moment of inertia of an $L \times L$ box is $\Theta = \rho_s L$ and the rotor levels are $E_\ell = \ell(\ell + N - 2)/(2\Theta)$ [17]; beyond tree level the square-torus helicity stiffness and the rotor inertia differ (for an ideal thermal rotor on a square, $\rho_{\text{hel}} = \Theta/L - \frac{1}{6} + \dots$, Section 4.6), and the helicity coupling $1/\rho_s$ and the mass-gap coupling $2mL/(N-1)$ are distinct couplings that agree at leading order only [31]. On square tori both lattice directions and all axes are averaged. On the rectangular $L \times T$ lattices of Section 4.2 the spatial and temporal twist responses are inequivalent; the record of that exploratory run does not state whether the quoted $\rho_s(L)$ is the spatial-twist response alone or the two-direction average, and this is left as an open item (Section 3.4). For $O(2)$ and for the λ -model we quote the $U(1)$ stiffness $\Upsilon_z = [\beta \langle A_z \rangle - \beta^2 \text{Var } B_z] / V$, normalized so that $\Upsilon = \langle q^2 \rangle$ over winding sectors.

Topological charge. The geometric charge Q of Berg and Lüscher [6]: each lattice plaquette is split into two triangles, the signed solid angle of the spherical triangle spanned by the three spins is summed, and $Q = \sum / 4\pi \in \mathbb{Z}$. $\chi_t = \langle Q^2 \rangle / V$. “Lattice-cell wraps” are plaquettes whose image area exceeds π (dislocations in the sense of [7]).

Block realignment. Each $b \times b$ block B is rotated into its own frame, defined by the mean direction $\hat{m}_B = m_B / |m_B|$, $m_B = \sum_{x \in B} s_x$. The inner wobble is $1 - |m_B|/b^2$; the frame misalignment between neighbouring blocks is $a(b) = \langle 1 - \hat{m}_B \cdot \hat{m}_{B'} \rangle$. The realignment map is $a(2b) = G(a(b))$. A pair $(b, 2b)$ is “clean” if $b \geq 4$ and the lattice contains at least 8 blocks of size $2b$ per side. The growth per doubling is $\ln[a(2b)/a(b)]$. For $O(3)$ the sub-frame geometric phase Ω is the solid angle enclosed by the four sub-block frames of a block.

Exact twisted partition function. $Z(\pi)/Z(0)$ for the π -rotation about an axis is measured exactly by the FK-winding estimator: the π -rotation is the product of two reflections; given the moduli, the two embedded Ising models are conditionally independent, and $Z_{\text{AP}}/Z_{\text{P}} = P_{\text{FK}}(\text{no cluster with odd } x\text{-winding})$ for each. The product of the two indicators is therefore unbiased for $Z(\pi)/Z(0)$. Reweighting of the uniformly spread twist from the untwisted ensemble was tried first and found unreliable for $\theta \gtrsim 2\pi/3$ (Section 3.3). The sector weights p_0, p_1, p_2 are solved from $\sum p_\ell = 1$, $\sum p_\ell \ell(\ell + 1)/2 = \rho_s$ and $\sum p_\ell \chi_\ell(\pi)/(2\ell + 1) = Z(\pi)/Z(0)$ under the truncation $\ell \leq 2$. This truncation is a model assumption: the same three numbers bound the singlet weight from both sides without it (Eq. (2)), and the truncated p_0 is a lower bound on the true singlet weight.

3 Validation, corrections and adverse outcomes

3.1 Cross-checks

[pre-registered] for the ξ study: (a) high-temperature series $E = t + 2t^3 + O(t^5)$, $\chi = 1 + 4t + 12t^2 + 36t^3 + O(t^4)$, $t = \coth \beta - 1/\beta$, at $\beta = 0.1, 0.2, 0.3$ ($L = 32, 10^5$ measurements): all $|z| \leq 1.4$ including the truncation budget; (b) $O(3)$ invariance: 228 pairwise comparisons of $\langle M_a^2 \rangle$ over 76 runs, rms $z = 1.06$, max $|z| = 2.79$; (c) hot versus cold starts at $(\beta, L) = (1.5, 96)$ and $(1.8, 256)$: $|z| \leq 1.41$; (d) improved versus direct ξ_2 : rms $z = 0.79$, max 2.62. An independent heat-bath code run through the same measurement function at $L = 8$, $\beta = 1.0$ (4 seeds $\times 4 \cdot 10^5$ sweeps) gave $\chi_{\text{HB}} = 8.730(8)$ against $\chi_{\text{Wolff,imp}} = 8.737(3)$ and $E_{\text{HB}} = 0.38558(5)$ against $E_{\text{Wolff}} = 0.38575(10)$.

Table 1: Correlation lengths: this work versus literature. Direct values at $L/\xi \geq 7.4$; $\beta \geq 2.0$ by the step-scaling extrapolation of [3] (reduced statistics).

β	this work	reference	value	z
1.6	19.026 ± 0.026	Wolff 1990, ξ_{exp}	19.07 ± 0.06	-0.68
1.7	34.547 ± 0.088	Wolff 1990, ξ_{exp}	34.57 ± 0.07	-0.20
1.8	64.792 ± 0.16	Wolff 1990, ξ_{exp}	64.78 ± 0.15	0.06
1.9	122.33 ± 0.15	Wolff 1990, ξ_{exp}	121.2 ± 0.6	1.83
1.9	122.33 ± 0.15	CEPS 1995, $\xi_{2,\infty}$	122.34 ± 0.29	-0.02
2.0	230.03 ± 1.8	CEPS 1995	228.54 ± 0.67	0.87
2.2	795.95 ± 23	CEPS 1995	780.04 ± 4.8	0.98
2.4	2635.6 ± 160	CEPS 1995	2568.8 ± 21	0.68
2.6	8689.9 ± 790	CEPS 1995	8569 ± 92	0.21
2.8	26489 ± 3800	CEPS 1995	28737 ± 390	-0.90

For the Casimir map [pre-registered] gates V1–V4 were run through the production code path: V1, Wolff versus an independent Metropolis chain at $L = 8$: O(3) $\beta = 1.6$, $E = 0.6501(8)$ vs $0.6514(6)$, $\rho_s = 0.7368(38)$ vs $0.7393(48)$; $\lambda = 0.5$, $\beta = 1.0$: $E = 0.3330(10)$ vs $0.3333(13)$, $\Upsilon_z = 0.0606(22)$ vs $0.0588(42)$; O(4) $\beta = 2.5$: $E = 0.6850(4)$ both, $\rho_s = 1.232(4)$ vs $1.234(5)$, $Z(\pi)/Z(0) = 0.315(2)$ vs $0.320(2)$. V2, exact $Z(\pi)/Z(0)$ versus the independent Swendsen–Wang implementation of Section 5: 14 overlapping points, none beyond 2σ , max pull 1.7. V4, per-axis $\Upsilon_x = \Upsilon_y = \Upsilon_z$ for O(3): max spread pull 2.9 over 22×3 values. V3 (rotor identity) is reported in Section 4.6.

3.2 Literature correlation lengths

Table 1 compares the direct and step-scaling ξ_2 of this work with Wolff [5] and Caracciolo et al. [3]. All $|z| < 2$ (statistical errors only; Wolff quotes ξ_{exp} , expected ratio $\xi_2/\xi_{\text{exp}} \approx 0.999$).

3.3 Corrections and adverse outcomes

1. **Wolff measurement-schedule bias (found and fixed).** An early version of the cluster code chose the number of clusters per measurement by a size-dependent stopping rule. This biased observables toward order: at $L = 8$, $\beta = 1.0$ it gave $E = 0.399$ against the heat-bath value $0.3856(4)$. With a fixed number of clusters per measurement (set once after thermalization) the same code gives $E = 0.3864(9)$ and $\chi = 8.747(58)$ against heat-bath $8.756(47)$. All results in this paper were produced with the fixed schedule; the ξ -study code fixed the count from the start and was not affected.
2. **Helicity normalization (corrected).** A first working formula, $\beta\langle e \rangle - \beta^2\langle J_a^2 \rangle/3$, mixed the normalizations of Υ_{axis} and ρ_s . Equation (1) is the corrected form, $\rho_s = (N/2)\Upsilon_{\text{axis}}$, which follows from $\langle s_1^2 + s_2^2 \rangle = 2/N$ in the symmetric finite-volume state and from the frozen-rotor check (Section 5). With the wrong normalization the rotor prediction of $Z(\pi)/Z(0)$ misses the exact value by up to 0.144; with the corrected one by at most 0.020. The stiffness numbers of Section 4.2 were recomputed after the correction.
3. **Withdrawn flag on eq. (2.14) of the OpenAI decay paper.** A first reading of the PDF flagged the identity $S(\bar{f}_j, \bar{f}_i) = -S_{ij}$ as false for complex test functions. The TeX source reads $S(\bar{f}_j, \bar{f}_i) = -\bar{S}_{ij}$, which is correct; the flag was a PDF-reading artefact and is withdrawn (Section 13).
4. **Twist reweighting fails for sector weights (replaced).** The pre-registered estimator of $Z(\pi)/Z(0)$ by reweighting the uniformly spread twist reproduces only the Gaussian part

of $Z(\theta)$: for $O(2)$ at $\beta = 1.2$ it gives $0.017 = e^{-\Upsilon\pi^2/2}$ where the exact FK value is 0.032. The exponent has spread $\sigma \approx 2.6$ at $\theta = \pi$, so the integer-sector interference sits in tails the untwisted sample never visits. The first Casimir-map launch was aborted after two points for this reason, the exact FK-winding estimator was added, gate V3 was replaced by the one-parameter rotor test, and all runs were restarted with new seeds, before any target data were analysed.

5. **Pre-registered hypothesis H1 of the Casimir map failed** ($\chi^2 = 123/10$, $p \approx 10^{-21}$); the post-hoc fit restricted to the running regime does not pass either ($\chi^2 = 24.1/9$, $p = 0.004$; with the full covariance of the drops $17.6/9$, $p = 0.04$; Section 4.6). Version 1 described the restricted fit as passing; corrected (Section 3.4).
6. **Pre-registered AF-versus-KT test not met as stated.** On the pre-registered window $\beta \in [1.4, 2.8]$ every fitted form (asymptotic freedom and Kosterlitz–Thouless) is rejected, because the non-monotone dip of $\xi/\xi_{3\text{-loop}}$ near $\beta = 1.5\text{--}1.6$ is described by none of them; the discrimination reported in Section 4.3 uses $\beta_{\min} = 1.8$ and is post hoc.
7. **Prediction 2 of the frame-drift test not met as stated** ($O(4)/O(3)$ growth ratio 1.36 ± 0.21 against the predicted 2; Section 4.5).
8. **Step-scaling tension.** The extrapolation from $L = 512$ at $\beta = 1.9$ gives $123.5(5)$ against the direct $L = 1024$ value $122.33(15)$ (2.1σ , about 1%), indicating a systematic in the reduced-statistics step scaling beyond its bootstrap error.
9. **Earlier Monte Carlo value of the certificate quantity was low.** An earlier report quoted $\varphi_{\text{pin}}(0.77) = 0.896 \pm 0.008$; the certified value of the truncated model is 0.9134, confirmed by two new heat-bath runs (0.9150 ± 0.0006 and 0.91356 ± 0.00044). The old value was about 2σ low, so the slack is smaller than first reported (Section 4.7).
10. **Pre-registration erratum of the ξ study.** Five throughput benchmarks (up to 4000 measurements) had printed E and ξ_2 before the pre-registration was frozen; none enters a result.
11. **Rotor certificate: cost wall reached (v1.1).** The certificate of Section 11 fails at $N = 4$, $g = 0.1$ (lower bound -0.176 at $\ell_{\max} = 4$; $\ell_{\max} = 5$ over the pre-registered dense cap) and at $N = 5$, $g \leq 0.2$; $N = 5$, $g = 0.3$ is certified only barely ($[0.006, 0.125]$), so prediction P7 (gap $\approx g/N$) is inconclusive at $N = 5$. 276 of 1920 rows are not certified, all by a non-positive lower bound, none by a failed positivity check.
12. **Rotor certificate: run-time patches (v1.1).** The η retry schedule left E_0^{low} undefined in 10 rows at $g = 50$; the five cases were re-run with a patched schedule and replace the old rows. The eigenvalue locator was switched from shift-invert LU to Lanczos after a stall (search value only; every bound is verified independently of it). The independent ball-arithmetic cross-check covers three small cases, not every row, and the entry-error budget ε_A was not logged per row.
13. **Langevin integrator replaced before target runs (v1.1).** The pre-registered Jacobi geodesic Euler–Maruyama scheme has an equilibrium bias of -7.5% in E at $dt = 0.04$; it was replaced by Metropolis-adjusted geodesic Langevin in a dated addendum written after validation and before any target run (Section 12). The pre-registered M -based free-rotor test was underpowered and a site-averaged estimator was added in the same addendum.
14. **Langevin hypothesis H1 inconclusive, H3 failed (v1.1).** The Langevin dynamic exponent $z = 1.78(4)$ is neither a pass nor a fail under the registered rule (heat bath $z = 2.04(3)$ passes); the energy channel relaxes 3–9 times faster than the predicted two-mode threshold $2\gamma_M$.

15. **Langevin precision below target at $\beta = 1.4, 1.5$ (v1.1).** Relative errors 6% and 8% against a 5% stopping target (wall caps hit); these rows, the z fits and the saturation values 0.759, 0.738 are provisional. The heat-bath $m_{\text{dyn}}\xi_2$ range quoted in the methods summary (0.87–1.03) is 0.86–1.06 when recomputed from the results file.

3.4 Errata relative to version 1.0 (external referee report, 19 October 2026)

An external referee report on version 1.0 was received the same day; its points were verified independently (recomputation of the rotor series, the fit statistics, the sector bounds and the annulus geometry) and are all accepted. Each item below names the version-1 statement, the correction and where it is applied.

- E1 **Helicity coupling $\neq$ mass-gap coupling.** Version 1 (Section 4.6, H5) identified $\bar{g}^2 = 2mL/(N-1)$ with $1/\rho_s$ exactly. False: these are two different finite-volume couplings that agree at leading order, $g_\Upsilon^2 = \bar{g}_m^2 + 0.1635\bar{g}_m^4 + \dots$ [31]; an ideal thermal rotor on a square has $\rho_{\text{hel}} = 0.826333$ at $mL = 1$ ($g_\Upsilon^2 = 1.210$ against $\bar{g}_m^2 = 1$). Corrected in Sections 2, 4.2, 4.6; the two-loop step formula is retained, justified by matching of the first two β -function coefficients. Molke–Wolff [31] added and positioned in the Introduction.
- E2 **O(4) sector formula demoted.** $\rho_s = \langle \ell(\ell+2) \rangle / 3$ is not exact for the full SO(4) transfer spectrum ((j_L, j_R) representations, $C_2 = 2[j_L(j_L+1) + j_R(j_R+1)]$; (1, 0) and (0, 1) have $C_2 = 4$). The general identity (5) stands; the O(4) form is now a collective-rotor approximation (Sections 5, 4.6, Table 3).
- E3 **Annulus paragraph rewritten; one sentence retracted.** Version 1 said that annulus twist bounds $Z(\pi)/Z(0) \geq 1 - C_\beta/\log N$, “including the one used in the OpenAI decay paper”, are N -blind. The cited proposition uses fixed-ratio annuli $N \leq \|x\|_\infty \leq 2N$, where the spin-wave cost $\pi\rho\theta^2/\log 2$ does not vanish for finite stiffness; a vanishing cost there is strictly stronger than the capacity bound and already contains the non-abelian step. “The annulus lemma cannot be the carrier” is retracted (Section 6; the same statement in the project records `attack_plan_03_gap_2026-10-19.md` §7 and `stiffness_theory.md` is superseded by this erratum).
- E4 **Restricted fit does not pass; covariance-aware refit added.** $\chi^2 = 24.1/9$, $p = 0.004$ (version 1 wrote “passes”; pre-registered cubic 123/10, $p \approx 10^{-21}$). Adjacent drops are correlated by construction; the recurrence fitted directly to ρ_s with the propagated covariance gives 17.6/9, $p = 0.04$ (Section 4.6). The same caveat is attached to the step-scaling and block-map χ^2 values (not refitted).
- E5 **Sector weights replaced by brackets.** The $\ell \leq 2$ reconstruction is a lower bound on p_0 ; Eq. (2) and Table 4 give model-independent two-sided bounds with propagated errors. The rise at $\beta = 2.0$, $L = 8 \rightarrow 256$ is $\geq 0.207(4)$ without truncation; stepwise monotonicity is not established where consecutive brackets overlap.
- E6 **Three claims separated.** Decreasing $\langle C_2 \rangle$, eventual singlet concentration and stepwise monotonicity of p_0 are distinct (counterexample (0.5, 0, 0.5) vs (0.3, 0.7, 0)); Sections 4.6, 6.
- E7 **Topology conclusion narrowed.** Section 4.1 and the abstract now claim only that bare $\chi_t \xi^2$ does not follow constant-amplitude scaling on the measured interval; exponent 0.763(17) with window and finite-volume sensitivity ($\xi_2(1.8)$: 62.1 in that run vs 64.79 in the larger-volume study, 4.2%; exponent 0.779(16) with the latter); [32] cited; the causal question is left open.
- E8 **Two missing pieces.** Abstract and Section 6 now state both tasks: the uniform stiffness inequality and the small-twist-response $\Rightarrow$ decay criterion.

E9 **Rigorous bounds rounded down.** $0.0296 \rightarrow 0.02956$ and the other mass bounds in Table 6.

E10 **Minor items.** Fig. 3(a) label relabelled (the version-1 label “no proof for $\beta > 0.511$ ” contradicted the text; the panel was relabelled on the rendered figure, the underlying plot is unchanged); Section 4.3 “small u ” $\rightarrow$ large u ; Section 4.4 clean domain $a \in [0.21, 0.44]$ with the iteration to 0.65 labelled as extrapolation; rectangular-lattice helicity direction flagged as an open item (Section 2); uniform integrability for the converse and proper $\text{SO}(3)$ rotations in Section 5; “detect masslessness” $\rightarrow$ “detect loss of stiffness”; the rotor Hamiltonian of Section 14 qualified as the continuum-time limit of $-\log \mathcal{T}$ and the centre-flux comparison as analogy; minimum observed drops stated as central estimates, not bounds; the denominators of the “1%” collapse statements defined.

E11 **Certificate made reviewable.** Appendix A gives the theorem and the lemma chain with hypotheses, attributes the amplitude-FKG and wired-comparison ingredients of Lemmas 2–3 to Campbell–Chayes [30], lists the 18 covering intervals (Table 14), states the arithmetic error model and describes exactly what the independent replay recomputes. The referee’s independent recomputation of ε , κ , the truncation errors and $\varphi = 0.91511728$ agrees with the printed values. A versioned archive of the code and data is to accompany the deposit (Section 15).

Not changed by the report: the reported adverse outcomes, the general Casimir identity, the controls, and the static Lean audit (the referee’s independent static check matched the 117 files, 12,907 lines, token search and comparator correspondence; the local build and checker runs remain reported by this paper only).

3.5 Erratum to version 2.0 (version 2.1)

E12 **Factor 2 in the first unproved step.** Version 2.0 wrote the first unproved step of Section 10 in two forms presented as equivalent: the ratio form $\Delta_1(2L) \geq \frac{1}{2}\Delta_1(L)[1 + c_1 L \Delta_1(L)]$ and the inertia form $\Theta_{2L} \leq 2\Theta_L - c_1 L$ (equivalently $\rho_H(2L) \leq \rho_H(L) - c_1/2$). They are not equivalent: with $\Theta = 1/\Delta_1$, the inertia form is $\Delta_1(2L) \geq \frac{1}{2}\Delta_1(L)/(1 - \frac{1}{2}c_1 L \Delta_1(L)) \approx \frac{1}{2}\Delta_1(L)[1 + \frac{1}{2}c_1 L \Delta_1(L)]$, so the ratio form as written double-counts c_1 . One loop fixes $c_1 = \ln 2/\pi$ in the inertia (Θ) form (drop $c_1/2 = \ln 2/(2\pi) = 0.110$ in ρ_H per doubling). Equation (12) is restated in the Θ -form; the pre-registered ratio predictions 0.58–0.66 of the transfer-matrix test (file `gap_step_test.md`, §0) used the written ratio form and are superseded by the Θ -form values 0.54–0.58 quoted in Section 10.1. Found by the transfer-matrix track on 19 October 2026 (`gap_step_test.md`, “Bookkeeping finding”).

E13 **Withdrawn link between Gibbs purity and singlet weight (project record, not the version-2.0 text).** The version-1 proof-chain record (`proof_chain_closure_2026-10-19.md`, v1) stated $T(L, L) = \sum_k w_k^2 \geq w_0^2 \geq p_0^2$, i.e. that a small square-torus stiffness bounds the bootstrap seed through $1 - T \leq 2\rho_s$. This is false in direction: w_0 is the weight of the *top eigenvalue* while p_0 is the weight of the whole singlet *sector*, so $w_0 \leq p_0$; the excited-singlet weight $p_0 - w_0$ is not controlled by ρ_s . Counterexample with the same $\text{SO}(3)$ action: $\text{O}(3)$ at $\beta = 0$ ($p_0 = 1$, $\rho_s = 0$) stacked with a low-temperature Ising layer whose two near-degenerate top eigenvalues are both singlets gives $T \approx \frac{1}{2}$. Version 2.0 of this paper did not print the false link (its Section 6 used only $1 - p_0 \leq \rho_s$, which stands); the correction is recorded here because the record is deposited with the paper. The repair is the doubling-defect seed of Section 7 (Lemma 6).

E14 **Erratum to version 2.1: the additive inequality C1’ is false at large β and is replaced by C1’’.** Version 2.1 stated the remaining inequality as $\Delta(2L, 2L) \leq \Delta(L, L) - c(\Delta(L, L))$ for $L \geq L_0(\beta)$ with $c > 0$ independent of β . This cannot hold. By the rigorous fixed- L Laplace bounds of Theorem 5, the decrement $D(L, \beta) = \Delta(L, L) - \Delta(2L, 2L)$ tends,

as $\beta \rightarrow \infty$ at fixed L , to the free-field constant $\gamma(L) = -4 \log \tau_L + 5 \log \tau_{2L} - \log \tau_{4L}$ (τ the number of spanning trees), which is positive but behaves as $\gamma(L) \approx 1.41/L^2 \rightarrow 0$; the β -dependent part of the decrement is $d(L)/\beta$ with $d(L) \rightarrow 3 \ln 2/(2\pi)$. Hence any c with $D \geq c$ for all L and β must satisfy $c \leq \inf_L \gamma(L) = 0$. The additive form also asks for more than is needed: Δ itself is $3 \log(2\beta) + O(1)$ at weak coupling, so its per-doubling drop is $\approx 0.331/\beta$, while the uniform drop lives in the running-coupling variable $\hat{\beta}(L) = \frac{1}{2} \exp[(\Delta - \delta_0)/3]$, where it is $\ln 2/(2\pi)$ per doubling. The statement that remains is the pointwise “reach” form C1’’ of Section 8.4: for every β , $\inf_L \Delta(L, L; \beta) < \Delta^* = 0.2239$, which by the dichotomy Theorem 2 is equivalent to the full conclusion and needs neither a rate nor uniformity in β . All statements of version 2.1 that used the additive form (its abstract, Sections 6, 7.5, 8.2 and Appendix B, item 1) are superseded by Sections 8.4 and 9; the seed threshold 0.1 of version 2.1 is raised to $\Delta^* = 0.2239$ (Section 8.4). The Θ -form step of Section 10, $\Theta_{2L} \leq 2\Theta_L - c_1 L$, is unaffected: it is the running-coupling form, and the weak-coupling end confirms its constant ($d(\infty) = \frac{3}{2}c_1$, Section 9).

4 Results

4.1 The bare topological susceptibility does not follow constant-amplitude scaling

[**exploratory**] Wolff runs at $\beta = 0.8$ – 1.8 , $L = 64$ – 256 (file `o3_sphere_map_results.csv`; Fig. 1). The scaled topological susceptibility $\chi_t \xi_2^2$ rises from $0.097(14)$ at $\beta = 1.2$ ($\xi_2 = 3.20$) to $1.031(36)$ at $\beta = 1.8$ ($\xi_2 = 62.1$ in this run). A weighted power-law fit over $\beta \geq 1.2$ (largest L at each β , 7 points) gives $\chi_t \xi_2^2 \propto \xi^\alpha$ with $\alpha = 0.763 \pm 0.017$ ($\chi^2 = 2.1/5$; 0.762 ± 0.023 when the ξ_2 errors are propagated); the fit window matters little ($\alpha = 0.760 \pm 0.018$ for $\beta \geq 1.4$, 0.736 ± 0.028 for $\beta \geq 1.5$), a logarithmic form $a + b \ln \xi$ is excluded ($\chi^2 = 130/5$), and the value 0.80 printed in version 1 came from a fit whose weighting is not recoverable from the record and is superseded. Finite-volume sensitivity: this run’s $\xi_2(1.8) = 62.13(94)$ at $L = 256$ is 4.2% below the $L = 1024$ value $64.792(160)$ of Table 1; using the latter moves the last point to $\chi_t \xi_2^2 = 1.121$ and the exponent to $\alpha = 0.779 \pm 0.016$. All of these say the same thing: on $3 \lesssim \xi \lesssim 65$ the bare $\chi_t \xi^2$ is not constant, growing by a factor of 10. Lattice-cell wraps (plaquettes with image area $> \pi$) fall as $e^{-5.6\beta}$ while $1/\xi^2$ falls as $e^{-9.9\beta}$; per correlation area the number of such wraps grows from 0.41 ($\beta = 1.2$) to 5.5 ($\beta = 1.8$). The bare geometric charge is therefore dominated by lattice-scale dislocations [7], in line with the known ultraviolet divergence of χ_t in this model, which persists under substantial gradient-flow smoothing [32]. What this observable establishes is that the bare $\chi_t \xi^2$ does not follow constant-amplitude scaling on the measured interval; it does *not* establish that topological excitations play no role in mass generation (an observable can carry large short-distance contributions while other aspects of the same configurations act at long distance), and that causal question is left open. In the same runs, Wolff reflection clusters never wrap the torus for $L/\xi \geq 7$, and the site-weighted mean reflection-cluster size is 0.75–0.78 χ at every β .

4.2 Finite-box rotor relation and stiffness running

[**exploratory**] On $L \times T$ lattices ($T = 128$ – 512 , $L = 4$ – 64 , $\beta = 1.4$ – 2.0 ; file `o3_rotor_box.csv`; Fig. 2) the time-slice gap $m(L)$ and the stiffness $\rho_s(L)$ of Eq. (1) were measured together. The tree-level rigid-rotor relation $m L \rho_s = 1$ ($E_1 = 1/\Theta$, $\Theta = \rho_s L$) holds to about 8% in the smallest boxes, $m L \rho_s = 0.918(6)$ at $\beta = 1.8$ and $0.919(8)$ at $\beta = 2.0$ for $L = 4$, and fades continuously as L approaches ξ : $0.735(13)$ at $L/\xi = 0.86$, $0.634(34)$ at $L/\xi = 1.7$, $0.364(58)$ at $L/\xi = 2.6$ ($\beta = 1.6$), and $0.12(5)$ at $L/\xi = 3.5$ ($\beta = 1.4$). The stiffness runs down at every doubling: $\beta = 1.6$: $\rho_s = 0.923, 0.725, 0.508, 0.311, 0.112$ for $L = 4, 8, 16, 32, 64$ (drops 0.198, 0.217, 0.197, 0.199); $\beta = 1.8$: $1.149, 0.964, 0.812, 0.644$ (drops 0.185, 0.152, 0.168); $\beta = 2.0$: $1.357, 1.190, 1.048, 0.893$




Figure 1: (a) A spin configuration at $\beta = 1.5$, $L = 128$. (b) Its signed covering of the sphere ($Q = 3$). (c) $\chi_t \xi_2^2$ versus ξ_2 : the bare susceptibility does not follow constant-amplitude scaling (weighted refit $\alpha = 0.76(2)$; the drawn line is the version-1 fit $\propto \xi^{0.80}$, superseded in the text). (d) Lattice-cell wraps per site fall as $e^{-5.6\beta}$, slower than $1/\xi^2 \sim e^{-9.9\beta}$.

(drops 0.167, 0.142, 0.155); $\beta = 1.4$: 0.688, 0.446, 0.185, 0.026 (drops 0.242, 0.261, 0.159). The per-doubling drops on this rectangular geometry lie in 0.14–0.27 (0.14–0.23 for $\beta \geq 1.6$); they are not directly comparable with the square-torus values of Section 4.6. The quantity $mL\rho_s$ compares two different finite-volume couplings, $\bar{g}_m^2 = 2mL/(N - 1)$ and $g_\Upsilon^2 = 1/\rho_s$, which coincide only at tree level [31]; the 8% deficit in small boxes is of the size of their one-loop mismatch and is not by itself a measure of non-rotor physics. Version 1 read this relation as exact; see Section 3.4.

4.3 Correlation length and step scaling

[pre-registered] (infinite-volume ξ study; file `o3_sim_results.csv`; Fig. 3). ξ_2 is finite at every β simulated, measured directly at $L/\xi \geq 7.4$ up to $\beta = 1.9$ ($\xi_2 = 122.33(15)$ at $L = 1024$) and by the step-scaling extrapolation of [3] for $2.0 \leq \beta \leq 2.8$ (Table 1). The ratio to the three-loop asymptotic-scaling prediction with the Bethe-ansatz normalization [2, 3] (no free parameter) is 0.70 at $\beta = 1.6$, 0.85 at 2.0 and 0.95(6) at 2.6, approaching 1 from below as in [3]. Exponential clustering is *proved* only for $\beta < \ln(5/3) = 0.5108$ by the Dobrushin condition computed here ($C(\beta) = \tanh(\beta/2)$, attained at zero background field and antipodal neighbour; $4C < 1$), for $\beta < 3/4$ by [8], and for $\beta \leq 0.7701$ by Section 4.7; everything above is numerics plus perturbation




Figure 2: (a) $mL\rho_s$ versus L/ξ : the tree-level rigid-ball relation ($= 1$) holds to 8% in small boxes and fades gradually; the two couplings compared here differ already at one loop. (b) $\rho_s(L)$ on $L \times T$ lattices falls at every doubling for every β .

theory.

[post hoc] With $\beta_{\min} = 1.8$ the asymptotic-freedom form with fixed normalization and two free $1/\beta^2, 1/\beta^3$ terms gives $\chi^2/\text{dof} = 2.5/8$, while the three-parameter Kosterlitz–Thouless form $\xi = A \exp(b/\sqrt{\beta_c - \beta})$ gives $101.2/7$ with $\beta_c = 12.9(2.0)$ ($\Delta\text{AIC} = +101$). Fitting on $[1.8, 2.2]$ only, the AF form predicts the CEPS values at $\beta = 2.5$ and 3.0 to 2.2% and 1.8%; the KT form over-predicts them by factors 1.41 and 2.64. This disfavors the simple KT form under the standard finite-size-scaling assumptions [14]; it cannot exclude a transition at β_c beyond the simulated range, and the $\beta \geq 2.0$ points rest on the same extrapolation method whose justification the CEPS reply called “considerably more subtle”.

[exploratory] (finite-size step scaling; file `o3_step_scaling.csv`; Fig. 4). With $u = \xi(L)/L$ and $\sigma(u) = \xi(2L)/\xi(L)$, 9 values of $\beta \in [1.2, 2.0]$ and $L = 8\text{--}128$ (4000 measurements each), all pairs with $L \geq 16$ fall near one curve (degree-4 fit, $\chi^2/\text{dof} = 2.64$; rms pull 1.6 at $L = 16, 32$ and 1.0 at $L = 64$; residuals in σ of order 0.02, i.e. about 1% of σ). The $L = 8$ pairs deviate systematically (mean pull -1.65 , rms 3.8). The zoom-out map $u \mapsto u\sigma(u)/2$ crosses the measured band $u \in [0.2, 0.72]$ in 6 doublings independently of β , with per-step contraction $\sigma/2 \leq 0.932$ (margin ≥ 0.068 from the fixed point $\sigma = 2$ that a stiff phase would require). Limits: the fit is valid only on the data range; the hand-off to perturbation theory, which with $u = \xi(L)/L$ lies at large u ($L \ll \xi$), is not reached. The quoted χ^2/dof uses diagonal weights; $\sigma(u)$ and u share the measurement $\xi(L)$, so this value is indicative only (the covariance-aware treatment of Section 4.6 was not repeated here).

4.4 Block-realignment map

[exploratory] Single large lattices ($\beta = 1.6, L = 128$, 1500 configurations; $\beta = 1.8, 2.0, L = 256$, 600 each; files `o3_realign_blocks.csv`, `o3_realign_map.csv`; Fig. 5). The inner wobble grows roughly linearly in $\log b$: $+0.07$ per doubling at $\beta = 2.0$, $+0.09\text{--}0.10$ at 1.8, $+0.11\text{--}0.13$ at 1.6. The frame misalignment stays small and grows slowly ($\beta = 2.0$: $a = 0.207 \rightarrow 0.273$ for $b = 4 \rightarrow 32$). All clean pairs from the three β (input misalignment $a(b) \in [0.207, 0.445]$, outputs up to 0.564) fall on one quadratic map G with residuals $|G(a) - a(2b)| \leq 0.008$ ($\chi^2/\text{dof} = 4.7$ against statistical errors of 0.001–0.005: a collapse to about 1–3% of a , not within statistics; $a(b)$ and $a(2b)$ come from the same configurations, so this diagonal χ^2 is indicative only). The excluded pairs deviate




Figure 3: (a) $\xi_2(\beta)$ with the three-loop prediction, the AF and KT fits, and the proved region $\beta < \ln(5/3)$ (green). (b) Ratio to the three-loop prediction. (c) Probability that a reflection cluster wraps the torus versus $L/\xi_{2,\infty}$: one curve, zero for $L/\xi \gtrsim 10$.




Figure 4: (a) Step-scaling function $\sigma(u)$: all β with $L \geq 16$ on one curve. (b) The zoom-out map crosses the band $u \in [0.2, 0.72]$ in six steps.

systematically: $b = 2$ by -0.025 to -0.038 (lattice), $L/(2b) = 4$ by -0.027 to -0.053 (finite size). On the measured input band $a \in [0.21, 0.44]$, $G(a) - a \geq 0.011$: the map has no fixed point in the data; iterating from $a = 0.207$ reaches 0.65 in 7 steps, of which the steps beyond $a \approx 0.44$ are extrapolations of the fitted quadratic outside the clean domain. The realignment recursion is thus a one-variable, approximately β -independent map measured on a single lattice without coarse-graining the action; the residual β -dependence (residuals up to 0.008 in a , i.e. 2–3% of a) is the footprint of the other couplings.

4.5 Frame drift in O(2), O(3) and O(4)

Predictions written before the run (no separate pre-registration file): (1) for O(2), a single rotation axis, the misalignment does not grow; (2) O(4) grows about twice as fast as O(3), from the factor $N - 2$; (3) O(3) growth tracks the sub-frame geometric phase $\langle \Omega^2 \rangle$. Runs: O(2) $\beta = 1.5, 2.0$; O(3) $\beta = 2.0, 2.4$; O(4) $\beta = 3.0, 3.6$; $L = 256$, 500 configurations each; growth per doubling $\ln[a(2b)/a(b)]$ for $b = 4, 8, 16$ (files `o3_axis_spin_ON_blocks.csv`, `o3_axis_spin_ON_growth.csv`; Fig. 6). Prediction 1 is supported, more strongly than stated:




Figure 5: (a) Frame misalignment (solid) and inner wobble (dashed) versus block size. (b) The realignment map $a(2b)$ versus $a(b)$: one quadratic map for all β (clean pairs filled; excluded $b = 2$ and $L/(2b) = 4$ pairs open).

O(2) frames settle (growth -0.02 to -0.05 per doubling; growth/ $a = -0.41 \pm 0.08$), while O(3) and O(4) frames drift apart at every scale. Prediction 2 is not met as stated: the growth/ a ratio O(4)/O(3) is 1.36 ± 0.21 . [post hoc] Counting transverse directions ($a \propto (N-1)g$, growth $\propto (N-2)g$) gives $(N-2)/(N-1)$, i.e. a ratio $4/3$, consistent with the data but derived after seeing them. Prediction 3: growth/ $\langle \Omega^2 \rangle = 0.91, 1.01, 0.92$ ($\beta = 2.0$; $b = 4, 8, 16$) and $0.97, 1.68, 1.27$ ($\beta = 2.4$), but $\langle \Omega^2 \rangle$ scales with a^2 , so this does not discriminate the geometric phase from plain misalignment squared. Supported: the drift of realigned frames with scale is present for O(3) and O(4) and absent (reversed) for O(2). Unresolved: whether the geometric phase itself is the carrier.

4.6 Casimir map: stiffness per doubling with controls and sector weights

[pre-registered] (file `casimir_map_methods.md`, Section 1, written before the target runs; data `casimir_map_results.csv`; Fig. 7). Square tori; O(3) at $\beta \in \{1.6, 1.8, 2.0, 2.2\}$, $L = 8\text{--}128$ plus $L = 256$ at $\beta = 2.0, 2.2$; O(2) at $\beta \in \{0.9, 1.2, 1.4\}$; λ -model at $(\lambda, \beta) \in \{(0.5, 2.0), (0.5, 1.6), (0.9, 2.0)\}$; statistics fixed in advance ($n_{\text{meas}} = 32000$ at $L = 8, 16$; 16000 at 32, 64; 8000 at 128; 4000 at 256); O(4) at $\beta = 2.5, 3.0$ added during the runs at the coordinator's request. Primary observable: the drop per doubling $c = \rho_s(L) - \rho_s(2L)$ as a function of $\rho_s(L)$.

H1, β -independence of the map: pre-registered test FAILED; the post-hoc fit does not pass either. The cubic in ρ fitted to all O(3) pairs with $L \geq 16$ (14 pairs, all β pooled) gives $\chi^2 = 123.4$ for 10 degrees of freedom ($p \approx 10^{-21}$), max residual 0.037. The pre-registered form cannot accommodate the crossover pair ($\beta = 1.6$, $L = 64 \rightarrow 128$, ρ 0.10 $\rightarrow$ 0.00, where necessarily $c \leq \rho$) together with the running regime. [post hoc] Restricted to the running regime $\rho_s(L) \geq 0.45$, $L \geq 16$ (12 pairs, four β), the form $c = a + b/\rho + c'/\rho^2$ gives $\chi^2 = 24.1$ for 9 degrees of freedom ($p = 0.004$) with the diagonal weights $\sigma_c^2 = \sigma_{\rho(L)}^2 + \sigma_{\rho(2L)}^2$ used in version 1, max residual 0.025 (the pair $\beta = 2.0$, $L = 128 \rightarrow 256$, pull -3.6 ; all other pulls ≤ 1.7). **Covariance-aware refit** (added in v2.0, from `casimir_map_results.csv`). The drops are correlated by construction: $d_i = x_i - x_{i+1}$ with $x_i = \rho_s(L_i)$ gives $\text{Cov}(d_i, d_{i+1}) = -\text{Var}(x_{i+1})$ and $\text{Cov}(x_i, d_i) = \text{Var}(x_i)$. Fitting the recurrence $x_{i+1} = x_i - c(x_i)$ directly to the stiffness measurements, with the linearized residual covariance $\text{Var}(r_i) = (1 - c'(x_i))^2 \sigma_{x_i}^2 + \sigma_{x_{i+1}}^2$, $\text{Cov}(r_i, r_{i+1}) = -(1 - c'(x_{i+1})) \sigma_{x_{i+1}}^2$ along




Figure 6: (a) Frame misalignment between neighbouring blocks: O(2) settles, O(3) and O(4) drift. (b) Growth per doubling divided by misalignment; the dashed $(N-2)/(N-1)$ line is post hoc.

each β chain (generalized least squares, iterated to convergence), gives $\chi^2 = 17.6/9$ ($p = 0.04$; parameters $a = 0.168$, $b = -0.099$, $c' = 0.067$ against $0.172, -0.105, 0.070$ diagonal; max pull 3.2, same pair) for the restricted set and $\chi^2 = 139.6/10$ for the pre-registered cubic. Neither fit passes; the data are an approximate one-variable curve with residuals beyond the statistical errors, not a statistically established β -independent map. The same correlation structure affects the step-scaling and block-map fits of Sections 4.3 and 4.4, whose quoted χ^2 values are diagonal and indicative only (not refitted; declared). Per- β mean residuals $+0.008 \pm 0.005$, 0.000 ± 0.003 , -0.001 ± 0.003 , -0.002 ± 0.004 for $\beta = 1.6, 1.8, 2.0, 2.2$: no β ordering. Reading: one curve to 0.01–0.02 in the drop (about 10% of c), not within statistics, the same quality of collapse as the ξ step scaling. The drop is positive at every O(3) point: the smallest central estimate is 0.098(8) (the crossover pair) and the smallest in the running regime is 0.123(6); these are minimum observed central values with their statistical errors, not rigorous lower bounds. The outlier pair was re-run with new seeds ([**post hoc**], declared): $L = 256$ gave 0.4739(74) and 0.4410(99) (2.7σ apart; merged 0.4622(59)), $L = 128$ gave 0.6238(45) and 0.6143(68); the merged pair still sits 0.025 below the curve. Whether this residual is a second coupling or an under-estimated error at $L \geq 128$ is unresolved.

H2, lattice artefacts: separated. The $L = 8 \rightarrow 16$ pairs lie above the $L \geq 16$ curve by $+0.041, +0.023, +0.010, +0.007$ at $\beta = 1.6, 1.8, 2.0, 2.2$ (pulls 11.2, 5.7, 2.8, 1.9); the artefact shrinks with β as $\xi(8)/a$ grows.

H5, perturbative limit. The two-loop step quoted here, $c(\rho) = (N-2)\frac{\ln 2}{2\pi}[1 + \frac{1}{2\pi\rho}]$, i.e. $0.1103 + 0.01756/\rho$ for $N = 3$, was derived in `casimir_map_methods.md`, Section 3, by inserting the two-loop running of the mass-gap coupling $\bar{g}_m^2(L) = 2m(L)L/(N-1)$ [19] into the identification $\bar{g}_m^2 = 1/\rho_s$. That identification is false as an exact statement (**correction, v2.0**): $g_\Gamma^2 = 1/\rho_s$ and $\bar{g}_m^2$ are different finite-volume couplings, and Molke and Wolff [31] give $g_\Gamma^2 = \bar{g}_m^2 + 0.1635\bar{g}_m^4 + \dots$ at equal L . Because a change of coupling definition of this form preserves the first two β -function coefficients, the two-loop step in g_Γ^2 has the same leading and next-to-leading terms, so the formula stands at two loops, justified by matching rather than by equality. An ideal rotor shows the gap between the two notions directly: on a square torus the thermal rotor with inertia Θ has $p_\ell \propto (2\ell+1)e^{-\ell(\ell+1)L/(2\Theta)}$ and $\rho_{\text{hel}} = \frac{1}{2}\sum p_\ell \ell(\ell+1)$, which at $r_\Theta = \Theta/L = 1$ equals 0.826333426 (recomputed) while $mL = 1/r_\Theta = 1$; numerically

Table 2: O(3) Casimir map on square tori (merged data; three points carry an independent replicate). Two-loop $c(\rho) = (\ln 2/2\pi)[1 + 1/(2\pi\rho)]$. $Z(\pi)/Z(0)$ from the exact FK-winding estimator; p_ℓ from the three exact equations under the truncation $\ell \leq 2$ (model-dependent; p_0 is a lower bound on the true singlet weight, see Table 4).

β	L	$\rho_s(L)$	drop to $2L$	two-loop	$Z(\pi)/Z(0)$	p_0	p_1	p_2
1.6	8	0.7439(24)	0.1977(36)	0.134	0.3262(19)	0.455	0.446	0.099
1.6	16	0.5462(28)	0.2215(51)	0.143	0.4249(20)	0.550	0.403	0.048
1.6	32	0.3247(42)	0.2273(63)	0.164	0.6091(34)	0.702	0.285	0.013
1.6	64	0.0975(47)	0.0982(75)	—	0.8788(19)	0.908	0.089	0.003
1.6	128	−0.0007(59)			0.9957(4)	0.997	0.004	−0.002
1.8	8	0.9802(24)	0.1606(40)	0.128	0.2441(23)	0.364	0.464	0.172
1.8	16	0.8196(32)	0.1504(44)	0.132	0.2923(22)	0.421	0.459	0.120
1.8	32	0.6692(30)	0.1686(48)	0.137	0.3532(24)	0.484	0.439	0.077
1.8	64	0.5006(37)	0.2393(83)	0.145	0.4671(31)	0.584	0.374	0.042
1.8	128	0.2613(74)			0.6679(40)	0.749	0.246	0.005
2.0	8	1.2004(25)	0.1432(37)	0.125	0.1993(16)	0.300	0.451	0.250
2.0	16	1.0571(27)	0.1413(48)	0.127	0.2273(20)	0.341	0.460	0.199
2.0	32	0.9158(40)	0.1362(57)	0.130	0.2576(30)	0.383	0.467	0.150
2.0	64	0.7796(41)	0.1588(55)	0.133	0.3025(28)	0.434	0.459	0.107
2.0	128	0.6209(38)	0.1586(70)	0.139	0.3708(37)	0.503	0.435	0.062
2.0	256	0.4622(59)			0.4942(55)	0.607	0.359	0.035
2.2	8	1.4186(23)	0.1397(35)	0.123	0.1693(16)	0.244	0.424	0.332
2.2	16	1.2789(27)	0.1233(57)	0.124	0.1849(16)	0.277	0.445	0.278
2.2	32	1.1556(50)	0.1402(67)	0.126	0.2126(28)	0.315	0.449	0.235
2.2	64	1.0155(45)	0.1324(79)	0.128	0.2310(24)	0.350	0.467	0.183
2.2	128	0.8831(64)	0.1509(123)	0.130	0.2627(42)	0.392	0.471	0.138
2.2	256	0.7322(104)			0.3121(66)	0.448	0.462	0.090

$\rho_{\text{hel}}(r_\Theta) = r_\Theta - \frac{1}{6} - 0.0056/r_\Theta + \dots$ (0.3126, 0.8263, 1.8303, 3.8319, 7.8326 at $r_\Theta = \frac{1}{2}, 1, 2, 4, 8$). Stiffness is a thermal average over charge sectors; inertia sets the energies in that average. For $\rho \geq 1$ the measured drops are 0.123(6)–0.141(5) against two-loop 0.123–0.128: data = two-loop + (0.00...0.015). Below $\rho \approx 0.9$ the drop rises above the perturbative curve (0.159(6) at $\rho = 0.78$, 0.169(5) at 0.67, 0.239(8) at 0.50, 0.227(6) at 0.32): the measured total drop exceeds the two-loop prediction, by 0.094 at $\rho = 0.50$ (two-loop 0.145). Version 1 called the excess “the non-perturbative part”; the supported statement is only that the total exceeds two loops, since higher perturbative orders, the $g_{\text{T}}\text{--}\bar{g}_m$ matching and lattice effects can all contribute to the residual.

Sector weights: model-independent brackets (v2.0, replaces the $\ell \leq 2$ reconstruction as the primary result). Normalization, one Casimir moment and one twist ratio do not determine an unrestricted sector distribution; the assumption $\ell \leq 2$ did essential work in version 1. The same three numbers bound the singlet weight from both sides. Write $R_\pi = Z(\pi)/Z(0)$ and $\rho = \rho_s$. Solving the three equations under $\ell \leq 2$ gives $p_0^{(2)} = 1 - \rho/6 - \frac{5}{8}(1 - R_\pi)$; applied to the exact moments of an arbitrary distribution this equals $\sum_\ell p_\ell a_\ell$ with $a_0 = 1$, $a_1 = a_2 = 0$ and $a_\ell = 1 - \ell(\ell+1)/12 - \frac{5}{8}(1 - (-1)^\ell/(2\ell+1)) \leq 0$ for $\ell \geq 3$ ($a_3 = -5/7$), so $p_0^{(2)} \leq p_0$. From $1 - (-1)^\ell/(2\ell+1) \in [\frac{4}{5}, \frac{4}{3}]$ for $\ell \geq 1$ one also has $1 - \frac{5}{4}(1 - R_\pi) \leq p_0 \leq 1 - \frac{3}{4}(1 - R_\pi)$. Hence

$$\max\left(0, 1 - \frac{\rho}{6} - \frac{5}{8}(1 - R_\pi), 1 - \frac{5}{4}(1 - R_\pi)\right) \leq p_0 \leq 1 - \frac{3}{4}(1 - R_\pi), \quad (2)$$

with no assumption on the sector content. Table 4 gives the brackets at every O(3) point with errors propagated from ρ_s and $Z(\pi)/Z(0)$ (the first lower bound is the active one everywhere). At $\beta = 2.0$: $0.2995(11) \leq p_0 \leq 0.3995(12)$ at $L = 8$ and $0.6068(33) \leq p_0 \leq 0.6207(38)$ at $L = 256$, so the singlet weight rises by at least 0.207(4) over five doublings without any truncation; at $\beta = 1.6$ it rises from ≤ 0.495 to ≥ 0.908 between $L = 8$ and 64. Stepwise monotonicity, however, is *not*

established everywhere: the lower bound at $2L$ exceeds the upper bound at L at every doubling for $\beta = 1.6$, for $L \geq 16$ at $\beta = 1.8$ and for $L \geq 64$ at $\beta = 2.0$, while at $\beta = 2.2$ and at the smaller boxes of $\beta = 1.8, 2.0$ consecutive brackets overlap. The $\ell \leq 2$ values p_1, p_2 of Table 2 are model-dependent: a true $p_3 = 0.05$ would lower $p_0^{(2)}$ by 0.036 relative to the truth and inflate the reconstructed p_2 by 0.107. Three claims must be kept apart (referee item E6): a decreasing Casimir mean $\langle C_2 \rangle$, which is measured; eventual singlet concentration $p_0 \rightarrow 1$, which follows from $\rho_s \rightarrow 0$ by Eq. (2); and stepwise monotonicity of p_0 , which is a separate statement and does not follow from the first (the distributions $(p_0, p_1, p_2) = (0.5, 0, 0.5)$ and $(0.3, 0.7, 0)$ have $\rho_s = 1.5$ and 0.7 with a *decreasing* singlet weight).

V3 (replacement), rotor identity. Fixing the one rotor parameter from $\rho_s = \langle C_2 \rangle_\theta / (N - 1)$ and predicting $Z(\pi)/Z(0)$: O(3), 22 points, max deviation 0.020 ($\beta = 1.8, L = 64$), mean +0.003, 19 of 22 within 0.01; O(4), 10 points, max 0.006; O(2) BKT ≤ 0.008 (disordered $\beta = 0.9$: up to 0.07, not a rotor regime). With Υ_{axis} in place of ρ_s the O(3) deviations reach 0.144: the normalization $\rho_s = (3/2)\Upsilon_{\text{axis}}$ is confirmed. For O(4) the one-parameter prediction uses the collective-rotor spectrum $(j_L, j_R) = (\ell/2, \ell/2)$ only, i.e. $\rho_s \approx \langle \ell(\ell + 2) \rangle / 3$; its agreement to 0.006 shows that this approximation works for these observables and parameters, not that the restricted spectrum is exact (Section 5).

Control A, O(2): PASS (H3). $\beta = 1.2$ (BKT phase): $\rho_s = 0.855, 0.840, 0.834, 0.832, 0.828$ for $L = 8-128$, drops 0.0148(13), 0.0066(10), 0.0020(10), 0.0036(11), i.e. ≤ 0.004 from $L = 32$ on, with $\rho = 0.83 > 2/\pi = 0.637$ [11]; $\beta = 1.4$: $\rho_s = 1.090 \rightarrow 1.077$, drops 0.009, 0.003, $-0.000, 0.001$. $\beta = 0.9$ (disordered): $\rho_s = 0.383, 0.209, 0.039, 0.010, -0.007$, drops 0.174, 0.170, 0.029, 0.017, $Z(\pi)/Z(0) \rightarrow 1.000$: an O(3)-like collapse. The square-torus map does not see N by itself; the inequality has to supply it.

Control B, λ -deformed O(3): PASS (H4). $(\lambda, \beta) = (0.5, 2.0)$: $\Upsilon_z = 1.400, 1.389, 1.389, 1.388, 1.387$ for $L = 8-128$, drop $64 \rightarrow 128$ equal to 0.0003(7), $Z(\pi)/Z(0) = 0.002$: map on the diagonal. $(0.5, 1.6)$: $\Upsilon_z = 0.945, 0.927, 0.919, 0.916, 0.917$, $\Upsilon_z(128) = 0.917 > 2/\pi$. $(0.9, 2.0)$: $\Upsilon_z = 1.274, 1.263, 1.261, 1.261, 1.259$, drops ≤ 0.011 ; the map leaves the O(3) curve already at $L = 8$, where a 10% anisotropy forces $\langle s_z^2 \rangle = 0.176$ (O(3): $1/3$) and $\Upsilon_z(8) = 1.27$ against the O(3) value 0.80. The broken symmetry shows as $\Upsilon_x = \Upsilon_y \neq \Upsilon_z$ (0.57 vs 1.39 at $\lambda = 0.5$). The U(1) rotor identity $\Upsilon_z = \langle q^2 \rangle$ still holds (rotor prediction of $Z(\pi)/Z(0)$ within 0.007 at all 15 λ -points). Any proof of $\rho_s(2L) \leq \rho_s(L) - c$ must use the full O(3) invariance: the λ -model satisfies every N -blind bound and its U(1) stiffness does not run at all.

O(4) and the $(N - 2)$ scaling. Table 3. The per-doubling drops are 0.25–0.30 against the one-loop $\ln 2/\pi = 0.221$ and the two-loop 0.24–0.26: as for O(3), the measured drop exceeds the perturbative one at every size. At matched $\rho \approx 1.0$ the O(4) drop 0.2686(86) ($\beta = 3.0, L = 64, \rho = 0.994$) against the O(3) drops 0.1324(79) ($\beta = 2.2, L = 64, \rho = 1.016$) and 0.1413(48) ($\beta = 2.0, L = 16, \rho = 1.057$) gives a ratio 2.0 ± 0.15 ; O(2) in the BKT phase gives 0.000–0.004. The coefficient therefore scales as $0 : 1 : 2$ for $N = 2, 3, 4$, i.e. as $N - 2$, the Ricci curvature of S^{N-1} . This replaces a pinned-square Cauchy–Schwarz saturation test requested during the runs, which could not be defined from the available information and was not run (declared deviation).

Second and third lattice action: no second coupling (v2.1). [pre-registered] (file `second_action_methods.md`, §1, written 19 October 2026 19:50 before any target run; data `second_action_results.csv`; Fig. 8a). The question left open above — whether the 0.01–0.02 residual of the one-variable map is a second coupling of the lattice action — was tested by repeating the Casimir map with two further actions on square tori: NN plus diagonal couplings

Table 3: O(4) on square tori. Two-loop $c = (2 \ln 2/2\pi)[1 + 1/(2\pi\rho)]$. Rotor prediction from the collective-rotor approximation $\rho_s \approx \langle \ell(\ell + 2) \rangle / 3$ (not exact for the full SO(4) spectrum).

β	L	ρ_s	drop to $2L$	two-loop	$Z(\pi)/Z(0)$ exact	rotor
2.5	8	1.2330(31)	0.2717(46)	0.249	0.3214(24)	0.319
2.5	16	0.9613(34)	0.2856(67)	0.257	0.3863(27)	0.385
2.5	32	0.6758(57)	0.3041(77)	0.272	0.4901(33)	0.489
2.5	64	0.3716(52)	0.2716(84)	—	0.6664(30)	0.666
2.5	128	0.1000(67)	—	—	0.9067(22)	0.901
3.0	8	1.7609(44)	0.2517(59)	0.241	0.2394(19)	0.239
3.0	16	1.5093(39)	0.2506(72)	0.244	0.2680(24)	0.271
3.0	32	1.2587(60)	0.2645(85)	0.248	0.3144(36)	0.314
3.0	64	0.9942(61)	0.2686(86)	0.256	0.3772(36)	0.376
3.0	128	0.7256(62)	—	—	0.4689(49)	0.467




Figure 7: (a) $\rho_s(2L)$ versus $\rho_s(L)$: O(3) and O(4) lose stiffness at every doubling; O(2)-BKT and λ -O(3) stay on the diagonal. (b) Drop per doubling versus $\rho_s(L)$ with the two-loop curves and the post-hoc restricted fit (diagonal weights; $\chi^2 = 24.1/9$, see text); $L = 8$ pairs (open rings) lie above. (c) O(3) sector weights $p_\ell(L)$ under the $\ell \leq 2$ truncation (model-dependent; $p_0^{(2)}$ is a lower bound on p_0 , Table 4).

with $\beta_2/\beta = 0.3$ (“NNN”) and the tree-level Symanzik action $\beta[\frac{4}{3} \sum_{nn} - \frac{1}{12} \sum_{2\text{-step}}]$ (“SYM”). For a general action the uniform-twist estimator is $\Upsilon_a = [\beta(\sum_b \kappa_b d_{x,b}^2 A_b) - \beta^2 \text{Var} \sum_b \kappa_b d_{x,b} B_b]/V$ with $A_b = s \cdot P_\perp(a)s'$, $B_b = s \cdot G_a s'$ and $d_{x,b}$ the x -extent of bond b ; it reduces to Eq. (1) for the NN action and was validated against the standard-action data at $(\beta, L) = (2.0, 8)$ ($\rho_s = 1.2052(28)$ vs $1.2004(25)$) and against an independent Metropolis chain for NNN and SYM (all pulls ≤ 1.4). The couplings were chosen from an $L = 16$ scan so that $\rho_s(16)$ matches the four standard-action values (0.534/0.827/1.060/1.279 NNN; 0.554/0.816/1.057/1.272 SYM; standard 0.546/0.820/1.057/1.279), and $L = 8$ –256 were run with the statistics of the standard map. Tested against the fixed standard-action curve $c_{\text{std}}(\rho) = 0.17225 - 0.10491/\rho + 0.06958/\rho^2$ (its own 12 pairs: $\chi^2 = 17.7/12$, $p = 0.13$) with the full covariance of adjacent pairs: NNN $\chi^2 = 39.9/12$ ($p = 7 \times 10^{-5}$), mean offset -0.0007 ± 0.0018 , max residual 0.027; SYM $\chi^2 = 21.3/12$ ($p = 0.047$), offset $+0.0051 \pm 0.0019$, max residual 0.028. Restricted to $\rho_s(L) \geq 0.7$ (8 pairs): NNN 22.4 ($p = 0.004$), offset $-0.001(2)$; SYM 8.1 ($p = 0.42$), offset $+0.003(2)$. By the pre-registered rule neither branch fires: “same curve” fails on the letter (one residual > 0.02 per action; NNN $\chi^2/\text{dof} \geq 3$) and “second coupling” (mean offset > 0.03 at $> 3\sigma$) fails by a factor of six or more. The largest pulls of all three actions sit on the steep part of the curve, $\rho_s(L) \approx 0.49$ – 0.62 at $L = 64$ – 128 . Replicas with new seeds ([**post hoc**], declared): the standard-action outlier $\beta = 2.0$, $L = 128 \rightarrow 256$ gave drop 0.184(8) against the curve value 0.180 (original 0.159(7); the two differ by 0.025 ± 0.010), and with replicas averaged in the standard-action χ^2 becomes 11.4/12 ($p = 0.50$, max residual 0.011); the NNN $\beta = 1.21$ chain reproduces (pulls < 1) and its deviation from the curve (± 0.012 per step, summing to zero over three steps) is real at that

Table 4: Model-independent brackets on the singlet weight $p_0(L)$ for O(3) on square tori, Eq. (2), with errors propagated from ρ_s and $Z(\pi)/Z(0)$ (last digits in parentheses); the column $p_0^{(2)}$ is the $\ell \leq 2$ reconstruction of version 1, which coincides with the lower bound. At $\beta = 1.6$, $L = 128$ the measured ρ_s is $-0.0007(59)$ and the two bounds cross within errors.

β	L	$p_0 \geq$	$p_0 \leq$	$p_0^{(2)}$	β	L	$p_0 \geq$	$p_0 \leq$	$p_0^{(2)}$
1.6	8	0.4549(12)	0.4947(14)	0.455	2.0	8	0.2995(11)	0.3995(12)	0.300
1.6	16	0.5495(13)	0.5687(15)	0.550	2.0	16	0.3409(14)	0.4204(16)	0.341
1.6	32	0.7015(22)	0.7068(26)	0.702	2.0	32	0.3834(21)	0.4432(24)	0.383
1.6	64	0.9080(14)	0.9091(15)	0.908	2.0	64	0.4341(19)	0.4769(21)	0.434
1.6	128	0.9974(10)	0.9968(3)	0.997	2.0	128	0.5033(21)	0.5281(24)	0.503
1.8	8	0.3642(15)	0.4331(17)	0.364	2.0	256	0.6068(33)	0.6207(38)	0.607
1.8	16	0.4211(15)	0.4692(17)	0.421	2.2	8	0.2444(11)	0.3770(12)	0.244
1.8	32	0.4842(15)	0.5149(17)	0.484	2.2	16	0.2774(14)	0.3887(16)	0.277
1.8	64	0.5835(20)	0.6003(23)	0.584	2.2	32	0.3153(21)	0.4094(23)	0.315
1.8	128	0.7489(28)	0.7510(30)	0.749	2.2	64	0.3501(17)	0.4233(18)	0.350
					2.2	128	0.3920(26)	0.4470(28)	0.392
					2.2	256	0.4480(47)	0.4841(53)	0.448

level. Reading: there is no action-dependent shift of the map at matched ρ_s to ± 0.005 ; the residual 0.02–0.03 scatter is action-blind and behaves like statistics (or mildly under-estimated errors on the steep part). An action-independent second coupling that moves points *along* the curve is not excluded by this test. Lattice artefacts: the $L = 8 \rightarrow 16$ pairs lie above the curve by $+0.040/+0.027/+0.014/+0.001$ (NNN) and $+0.014/+0.005/+0.006/-0.006$ (SYM) against $+0.041/+0.023/+0.010/+0.007$ for NN; the Symanzik action removes most of the $L = 8$ artefact, as it should. Declared deviation: the FK odd-winding estimator of $Z(\pi)/Z(0)$ is not valid for the negative Symanzik couplings (frustrated FK configurations are missed), so SYM contributes to the stiffness map only.

Direct sector weights by seam twist (v2.1; review angle 5). [pre-registered] (same methods file, part 2). For the NN action the twisted partition function was measured at 13 angles $\theta_k = k\pi/12$ by a seam twist $s \cdot R_z(\theta)s'$ on the bonds crossing the x -boundary, with $g(\theta) = \partial_\theta \log Z = \beta \langle \sum_{\text{seam}} (-\sin \theta A_b + \cos \theta B_b) \rangle_\theta$ integrated by Simpson’s rule (heat bath + over-relaxation + z -reflecting Wolff clusters; Bennett acceptance ratios between neighbouring angles as a cross-check, agreeing to ≤ 0.004 in $\log Z$). The weights $p_0, \dots, p_3$ are the weighted least-squares solution of $Z(\theta)/Z(0) = \sum_\ell p_\ell \chi_\ell(\theta)/(2\ell+1)$ at $\theta \in \{\pi/3, \pi/2, 2\pi/3, \pi\}$ with $\sum p_\ell = 1$. Validation: $Z(\pi)/Z(0)$ from the integration agrees with the exact FK value of the standard runs at all six points (max pull 1.0); $\sum_\ell p_\ell \ell(\ell+1)/2$ agrees with the helicity ρ_s within 1.4σ at all six points; $-\frac{3}{2}g'(0)$ reproduces ρ_s . Table 5 gives the results. Every direct p_0 lies inside the rigorous bracket of Table 4 and $0.014/0.011/0.006$ ($\beta = 2.0$) and $0.002/0.002/0.002$ ($\beta = 1.8$) above the $\ell \leq 2$ lower bound, so the $\ell \geq 3$ weight is $0.017(2)$ at $\rho_s \approx 1.06$ and ≤ 0.01 below $\rho_s \approx 0.9$; a secondary $\ell \leq 4$ fit to all 11 interior angles agrees to ≤ 0.002 in p_0 with $|p_4| \leq 0.004$. The singlet log-odds $\Lambda = \log[p_0/(1-p_0)]$ from the direct weights rise at every doubling: increments $+0.169(16)$, $+0.186(16)$ ($\beta = 2.0$, $16 \rightarrow 32 \rightarrow 64$) and $+0.256(17)$, $+0.397(20)$ ($\beta = 1.8$), minimum $> 10\sigma$ above zero, and the four increments lie on one monotone curve in p_0 (quadratic residuals ≤ 0.014) that agrees with the $\ell \leq 2$ indirect curve to -0.015 , -0.046 , -0.003 , -0.006 (Fig. 8b). Caveat: the two β do not overlap in p_0 , so this is consistency with a single curve, not an independent two- β collapse test. The $\beta = 2.0$, $L = 32$ chain shows angle-to-angle scatter $1.8\times$ its quoted errors; its errors were inflated by 1.79 in the Λ test (declared). A sign error in the seam observable was caught by the $L = 8$ validation before any target run (declared).

Table 5: Direct SO(3) sector weights on square tori (NN action) from the seam-twist thermodynamic integration over 13 angles; file `second_action_results.csv`, block `sector_weights`. “bracket” is Eq. (2); “ $\ell \leq 2$ ” is the three-equation lower bound of Table 2.

β	L	$Z(\pi)/Z(0)$ (TI)	exact FK	p_0	p_1	p_2	p_3	bracket for p_0	$\ell \leq 2$
2.0	16	0.2256(14)	0.2273(22)	0.3549(16)	0.4733(16)	0.1548(15)	0.0171(19)	[0.340, 0.419]	0.3409
2.0	32	0.2607(18)	0.2576(33)	0.3944(19)	0.4711(20)	0.1264(19)	0.0081(22)	[0.385, 0.446]	0.3834
2.0	64	0.3044(21)	0.3025(28)	0.4397(21)	0.4594(21)	0.0930(22)	0.0078(23)	[0.435, 0.478]	0.4341
1.8	16	0.2894(24)	0.2923(22)	0.4232(25)	0.4602(24)	0.1064(24)	0.0102(28)	[0.419, 0.467]	0.4211
1.8	32	0.3531(32)	0.3532(23)	0.4866(33)	0.4418(28)	0.0698(33)	0.0017(31)	[0.484, 0.515]	0.4842
1.8	64	0.4670(41)	0.4671(31)	0.5852(37)	0.3776(37)	0.0380(39)	-0.0008(36)	[0.584, 0.600]	0.5835




Figure 8: (a) Drop per doubling versus $\rho_s(L)$ for three lattice actions (nearest-neighbour, NN+diagonal with $\beta_2/\beta = 0.3$, Symanzik two-step); the line is the standard-action restricted fit, the star the new-seed replica of the standard-action outlier pair. (b) Singlet log-odds increment per doubling versus $p_0(L)$: direct seam-twist weights (filled) against the $\ell \leq 2$ reconstruction (crosses).

4.7 A computer-assisted decay certificate for $\beta \leq 0.7701$ and at isolated points up to 0.785

[pre-registered] (file `decay077_preregistration.json`, written 2026-10-19T15:31Z before any kernel-truncation $L \geq 5$ run at $\beta \geq 0.76$; methods and lemma chain in `decay077_methods.md`; results `decay077_results.json`, `coverage_summary.json`; code `decay077_certificate.py`; replay `decay077_replay.py`). Here L denotes the spherical-harmonic truncation degree, not a lattice size.

Criterion. Write $G_1(x, y) = \langle s_x^1 s_y^1 \rangle = \langle s_x \cdot s_y \rangle / 3$. The Aizenman–Simon inequality [8] (Lemma 1), for finite $A \ni i$ and $j \notin A$, $G_1(i, j) \leq \beta \sum_{l \in A, k \notin A, l \sim k} G_1(i, l) G_1(k, j)$, with $A = [-1, 1]^2 + i$ gives exponential decay whenever $\varphi = \beta[4G_1(i, i + e) + 8G_1(i, i + d)] < 1$ (e an edge neighbour, d a diagonal one). The Wolff embedding $s = (r\tau, q \cos \theta, q \sin \theta)$ with the Campbell–Chayes amplitude association [30] gives $G_1(x, y) = E[r_x r_y \mathbf{1}\{x \leftrightarrow y\}]$ with (r, n) positively associated (Lemma 2). Pinning the ring $|x - y|_\infty = 3$ to e_1 can only increase this expectation (Lemma 3), so $G_1(y, l) \leq \langle s_y^1 s_l^1 \rangle^{\text{pin}}$ in the 7×7 box with 25 free spins. The certified quantity is $\varphi_{\text{pin}}(\beta) = \beta[4c_e + 8c_d]$ with c_e, c_d the pinned-box correlations. Lemmas 2 and 3 are comparison statements whose ingredients — the amplitude-FKG property of the

Wolff-embedded (r, n) variables, the ordering under axial fields, the wired-boundary comparison and the amplitude-times-connectivity representation — are the results of Campbell and Chayes [30]; Appendix A states both lemmas with their hypotheses and marks which steps are taken from [30] and which are assembled here. Lemma 1 was checked against the reproduction of its proof in [13], Appendix A.2 (Theorem A.3 at $\lambda = 1$, with the coefficient β and no extra factor 3), because the original could not be retrieved.

Computation. Each bond kernel $e^{\beta x}$ (40 internal, 20 to the ring) is replaced by its character expansion truncated at degree $L = 6$, $k_L(x) = \sum_{l \leq L} (2l+1) i_l(\beta) P_l(x)$; the truncated model is then computed *exactly* by a site-by-site transfer in the y_{lm} basis (Funk–Hecke diagonal horizontal integration, Gaunt-coefficient vertical bonds, exact degree projections, $M = 0$ sector with 6.2×10^7 slots). Rigour: Gaunt coefficients exact as sign- $\sqrt{\text{rational}}$; $i_l(\beta)$ and the tail $T_L = e^\beta - \sum_{l \leq L} (2l+1) i_l(\beta)$ as 256-bit arb enclosures; every float64/BLAS product carries a componentwise a-priori error bound (a γ_n majorant, valid for any summation order and with FMA); the truncation lemma gives $|\langle F \rangle - \langle F \rangle_L| \leq \Delta = \kappa(1 + c_L)/(1 - \kappa)$ with $\kappa = ((1 + \varepsilon)^{60} - (1 - \varepsilon)^{60})/(2(1 - \varepsilon)^{60})$, $\varepsilon = T_L e^\beta$. At $\beta = 0.77$: $\varepsilon = 2.750 \times 10^{-6}$, $\kappa = 1.650 \times 10^{-4}$, $\Delta_e = 1.85 \times 10^{-4}$, $\Delta_d = 1.79 \times 10^{-4}$; the truncation bound dominates the budget (the float error is $\sim 10^{-11}$ on correlators). This is float64 with a-priori bounds, not full ball arithmetic. Validation: 2×2 , 3×3 and 5×5 pinned boxes agree with independent heat-bath Monte Carlo to within errors (e.g. 5×5 , $\beta = 0.77$: φ 0.91344 vs 0.91500 ± 0.00062 and 0.91356 ± 0.00044); D_4 symmetry to 12 digits; L -convergence $c_e = 0.12305225400, 0.12305226187, 0.12305226191$ at $L = 4, 5, 6$; small- β limits $c_e/\beta \rightarrow 1/9$, $c_d/\beta^2 \rightarrow 2/27$. An independent mpmath interval replay (`decay077_replay.py`) reproduces every φ_{up} to 10^{-11} and every certified flag; what it does and does not recompute is stated in Appendix A.

Results (Table 6). At $\beta = 77/100$: $c_e \in [0.122867, 0.123238]$, $c_d \in [0.086580, 0.086939]$, hence $\varphi_{\text{pin}} \in [0.911765, 0.915117] < 1$; the pre-registered acceptance rule $\varphi_{\text{up}} < 1$ is met with margin 0.085 (cross-check at $L = 5$: $\varphi_{\text{up}} = 0.94218$). Iterating Lemma 1 with Lemma 3 (Simon’s argument, Lemma 5) gives, on every torus of side ≥ 8 and in free boxes far from the boundary, $\langle s_0 \cdot s_x \rangle \leq \varphi_{\text{up}}^{\lfloor d(0,x)/3 \rfloor}$, i.e. a mass $m \geq -\ln(0.915117)/3 \geq 0.02956$ per lattice step (rounded down). Isolated certificates hold at $\beta = 0.78$ ($\varphi_{\text{up}} = 0.96296$) and 0.785 (0.98765). At $\beta = 0.79$ the certified *lower* bound is 1.0086: this criterion (3×3 block, 7×7 pinned box) itself fails there, not its precision. **Interval coverage** (declared in the pre-registration as a secondary aim): a change-of-measure lemma, $|\langle F \rangle_{\beta+\delta} - \langle F \rangle_\beta| \leq \frac{1}{2}(e^{120|\delta|} - 1)\sqrt{\langle F^2 \rangle_\beta}$ for $|F| \leq 1$ and $|H| \leq 60$, with 18 certified grid points (listed in Table 14), covers every $\beta \in [0.7491374, 0.7701917]$ with no gaps (replay agrees). Together with Aizenman–Simon for $\beta < 3/4$, and conditional on Lemmas 2–3 as stated in Appendix A, exponential decay of the two-point function holds for every $\beta \leq 0.7701$ (Theorem 6). Covering $[0.7702, 0.785]$ would need about 40 more grid points.

Scope. This is a fixed- β statement about the isovector two-point function at couplings just beyond the 1980 bound. It is not a gap uniform in β , says nothing about scaling or the continuum, does not claim uniqueness of the Gibbs state, and is a human-checkable independent result on a small range, obtained as a test of certificate machinery that could in principle be pushed further (an $R = 4$ pinned box needs about 6×10^9 states).

5 Exact identities

The derivations below are taken from the project record `stiffness_theory.md` (19 October 2026); they hold for all $N \geq 2$ and are elementary. They are a sector-by-sector regrouping of the relation between the helicity modulus and the twisted transfer trace written by Molke and

Table 6: Certified pinned-box criterion, kernel truncation L (6 unless stated). Mass bounds are rounded down.

β	c_e	c_d	φ_{pin} interval	$\varphi < 1$	mass $\geq -\ln \varphi_{\text{up}}/3$
0.75	≤ 0.116439	≤ 0.079350	≤ 0.82542	yes	0.06395
0.77	[0.122867, 0.123238]	[0.086580, 0.086939]	[0.911765, 0.915117]	yes	0.02956
0.77 ($L = 5$)	[0.119873, 0.126232]	[0.083683, 0.089836]	[0.88470, 0.94218]	yes	0.01985
0.78	[0.126380, 0.126792]	[0.090527, 0.090925]	[0.95919, 0.96296]	yes	0.01258
0.785	[0.128174, 0.128608]	[0.092546, 0.092966]	[0.98365, 0.98765]	yes	0.00414
0.79	[0.129993, 0.130450]	[0.094595, 0.095038]	[1.00862, 1.01286]	no	—

Wolff [31] (their eqs. (2.7)–(2.11) [details unverified]); the general identity (5) is an independent rederivation, not a new construction.

5.1 Helicity identity with the rotor normalization

On an $L \times T$ torus, twist by $g = R_{12}(\theta)$ across the x -seam and gauge the twist uniformly onto all x -bonds ($R(\theta/L)$ per bond); Z is unchanged. Then $\partial_\theta^2(\beta \sum s \cdot R(\theta/L)s')|_0 = -(\beta/L^2) \sum (s_1 s'_1 + s_2 s'_2)$, the first-derivative term has zero mean by reflection symmetry, and $O(N)$ invariance gives $\langle \sum (s_1 s'_1 + s_2 s'_2) \rangle = (2/N) \langle e_x \rangle$, so $\partial_\theta^2 \ln Z|_0 = L^{-2} [-(2\beta/N) \langle e_x \rangle + \beta^2 \langle J_{12}^2 \rangle]$. Defining Υ by $F(\theta) - F(0) = \Upsilon \theta^2 T / (2L) + O(\theta^4)$ gives Eq. (1). Classical check: a frozen rotor (s constant in space, rotating in time) has action $(\beta/2)L \int |\dot{s}|^2$, so $\Theta = \beta L$, while a twist about a fixed axis costs $(\beta/2)(\theta/L)^2(1 - s_3^2)$ with $\langle 1 - s_3^2 \rangle = 2/3$; hence $\rho_s = (3/2)\Upsilon_{\text{axis}}$ for $N = 3$.

5.2 Character expansion and positivity

By Funk–Hecke, $e^{\beta s \cdot s'} = \sum_\ell a_\ell(\beta) P_\ell(s \cdot s')$ with $a_\ell > 0$, so the bond kernel is positive definite on S^{N-1} ; the column transfer matrix $\mathcal{T} = D^{1/2} K^{\otimes T} D^{1/2}$ is positive, trace class and commutes with the global rotation $U(g)$. Hence

$$Z_{L,T}(g) = \text{Tr } U(g) \mathcal{T}^L = \sum_\rho \chi_\rho(g) n_\rho(L), \quad n_\rho(L) = \text{Tr}_{M_\rho} T_\rho^L = \sum_k \lambda_{\rho,k}^L \geq 0. \quad (3)$$

Consequences: Z is a positive-definite class function, so $Z(g) \leq Z(e)$ and $\Upsilon \geq 0$; each $n_\rho(L)$ is log-convex in L ; at fixed T , $Z(g)/Z(e) \rightarrow 1$ exponentially in L (Perron–Frobenius, singlet ground state), which is a one-dimensional statement and not running; $\ln Z$ is not concave in θ in general ($Z(\theta)$ has its minimum at π); and there is no monotonicity at fixed aspect ratio, because $L \times L \rightarrow 2L \times 2L$ compares traces on different Hilbert spaces. Reflection positivity, Cauchy–Schwarz and chessboard estimates act within one torus and give only $|Z(gh)| \leq Z(e)$ -type bounds, all of which hold verbatim for $O(2)$.

5.3 Square-torus identity: stiffness equals the mean Casimir

On the $L \times L$ torus the 90° lattice rotation maps a spatial twist to a temporal twist. Transferring in t with spatial circle L , $Z_{\text{sq}}(g) = \sum_\rho \chi_\rho(g) n_\rho$ with $n_\rho = \sum_k e^{-LE_{\rho,k}(L)}$ over the isospin- ρ eigenstates of the circle- L Hamiltonian. With $p_\rho = d_\rho n_\rho / Z_{\text{sq}}(e)$ and $-\chi''_\rho(0)/d_\rho = \langle J_{12}^2 \rangle_\rho = C_2(\rho) / \dim \mathfrak{so}(N)$ (Schur’s lemma over equivalent generators),

$$\Upsilon_{\text{axis}}(L \times L) = \sum_\rho p_\rho \frac{C_2(\rho)}{\dim \mathfrak{so}(N)}, \quad \rho_s = \frac{\langle C_2 \rangle}{N-1}, \quad (4)$$

$$\text{SO}(3) : \rho_s = \frac{1}{2} \langle \ell(\ell+1) \rangle, \quad \text{O}(2) : \Upsilon = \langle q^2 \rangle, \quad \text{SO}(4) : \rho_s = \frac{2}{3} \langle j_L(j_L+1) + j_R(j_R+1) \rangle \quad (5)$$




Figure 9: Swendsen–Wang runs on $L \times L$ tori (8000 measurements per point). (a) $\rho_s(L)$: only $O(2)$ in its BKT phase stays stiff; $O(3)$ and disordered $O(2)$ run to zero. (b) Exact $Z(\pi)/Z(0)$ with the one-parameter rotor prediction (dotted).

The expectation runs over *all* transfer-matrix eigenstates with their multiplicities; no rigid-rotor spectrum is assumed. For $SO(3)$ every integer-spin representation has $C_2 = \ell(\ell + 1)$, so the first form is exact. For $SO(4)$ the representations are labelled (j_L, j_R) with $C_2 = 2[j_L(j_L + 1) + j_R(j_R + 1)]$, and the single-rotor harmonics on S^3 are only the subset $(j_L, j_R) = (\ell/2, \ell/2)$ with $C_2 = \ell(\ell + 2)$; already two vector sites contain $(\frac{1}{2}, \frac{1}{2}) \otimes (\frac{1}{2}, \frac{1}{2}) = (0, 0) \oplus (1, 0) \oplus (0, 1) \oplus (1, 1)$, and $(1, 0), (0, 1)$ have $C_2 = 4$, which is not of the form $\ell(\ell + 2)$. The $O(4)$ formula $\rho_s = \frac{1}{3}\langle \ell(\ell + 2) \rangle$ used in version 1 is therefore a *collective-rotor approximation*, not an identity (**correction, v2.0**); its success in predicting $Z(\pi)/Z(0)$ to 0.006 (Section 4.6) is evidence that the approximation is good for those observables, nothing more. For $SO(3)$, $\chi_\ell(\pi)/d_\ell = (-1)^\ell/(2\ell + 1)$ and $1 - (-1)^\ell/(2\ell + 1) \in [4/5, 4/3]$ for $\ell \geq 1$, so $1 - Z(\pi)/Z(0) \in [\frac{4}{5}(1 - p_0), \frac{4}{3}(1 - p_0)]$ and $\Upsilon_{\text{axis}} \geq \frac{2}{3}(1 - p_0)$; hence

$$1 - Z(\pi)/Z(0) \leq 2 \Upsilon_{\text{axis}}(L \times L), \quad 1 - \frac{Z(g)}{Z(e)} \leq 2(1 - p_0) \leq \frac{5}{2} \left(1 - Z(\pi)/Z(0) \right) \quad \text{for every } g \in SO(3), \quad (6)$$

and the same first bound holds for $O(2)$ (rotational singlets can transform non-trivially under improper elements of $O(3)$, so the statement is for $SO(3)$). Thus $\Upsilon_{\text{sq}} \rightarrow 0 \Rightarrow p_0 \rightarrow 1 \Leftrightarrow Z(\pi)/Z(0) \rightarrow 1 \Leftrightarrow$ all proper twists free; the converse, $p_0 \rightarrow 1 \Rightarrow \langle C_2 \rangle \rightarrow 0$, needs uniform integrability of C_2 under the sector weights (tightness alone does not suffice). The identity was tested independently of the Casimir-map runs with a Swendsen–Wang code (two embedded Ising components, exact FK-winding $Z(\pi)/Z(0)$; Fig. 9): one fitted rotor parameter reproduces $Z(\pi)/Z(0)$ to 0.01–0.03 at every point, and in the Casimir-map runs to 0.020 ($O(3)$) and 0.006 ($O(4)$) (Section 4.6). The character-expansion consequence used in Section 4.6 is that the three exact numbers $\sum p_\ell$, $\sum p_\ell \ell(\ell + 1)/2$ and $\sum p_\ell \chi_\ell(\pi)/(2\ell + 1)$ determine p_0, p_1, p_2 under the truncation $\ell \leq 2$.

6 The missing inequality and its controls

Statement. For the nearest-neighbour $O(3)$ model on $L \times L$ tori there exist $L_0(\beta)$ and a function $c : (0, \infty) \rightarrow (0, \infty)$, bounded below on compact sets and *independent of β* , such that

$$\rho_s(2L) \leq \rho_s(L) - c(\rho_s(L)) \quad \text{for all } L \geq L_0(\beta), \quad (7)$$

with $\rho_s = \frac{1}{2}\langle \ell(\ell + 1) \rangle$ the square-torus stiffness of Eq. (5); one loop gives $c(\rho) \rightarrow \ln 2/(2\pi) \approx 0.110$ as $\rho \rightarrow \infty$. In sector language the inequality says that the Casimir mean $\langle C_2 \rangle$ decreases by a

controlled amount at each doubling; it does *not* say that p_0 increases at every step (Section 4.6, three claims kept apart). Iterating gives $\rho_s \rightarrow 0$ and, by Eq. (6), $p_0 \rightarrow 1$ and $Z(\pi)/Z(0) \rightarrow 1$. **Two pieces are missing, not one (v2.0):** (i) inequality (7) with its stated uniformity in β ; (ii) a finite-size criterion of Dobrushin–Shlosman or Simon–Lieb type stating that sufficiently small *torus twist response* at scale L implies exponential decay. The certificate of Section 4.7 is a finite-box *correlation* criterion at small β and does not supply the implication from small square-torus stiffness to that criterion. Section 4.6 shows the inequality holding numerically with every measured drop positive (smallest central estimate 0.098(8)) and a curve that is β -independent to 0.01–0.02 but not within statistics.

Status in versions 2.1–2.2. Piece (ii) is supplied: the finite criterion of Section 7 converts two partition-function ratios on tori of a single size into a gap uniform in the circle length, and exponential decay in the plane follows (Lemma 3). The criterion detects every gapped β (Tasaki’s converse, Proposition D.1 of [36], and its classical analogue Theorem 3), so the obstruction is not the criterion. Piece (i) remains, in its sharp form: the seed observable is the doubling defect $\Delta(L, L) = 4 \log Z_{L \times L} - \log Z_{2L \times 2L}$ (Lemma 6), and the single statement that remains, C1'', is that at every β this defect falls below $\Delta^* = 0.2239$ at some finite L (Section 8.4); by the dichotomy Theorem 2 nothing weaker suffices and nothing stronger is needed — in particular no β -uniform decrement (version 2.1’s additive C1', withdrawn in erratum E14) and no monotonicity. The stiffness inequality (7) and the Θ -form step of Section 10 are sufficient routes to C1'' in the running-coupling variable, where the weak-coupling end (Section 9) shows the uniform drop actually lives; the link $1 - p_0 \leq \rho_s$ stands but is no longer needed for the seed.

What reflection positivity and Mermin–Wagner cannot do. Every inequality obtained from reflection positivity, chessboard estimates or Cauchy–Schwarz within one torus holds verbatim for $O(2)$, whose BKT phase has $\Upsilon(L) \rightarrow \Upsilon_\infty \geq 2/\pi$ [11]; such inequalities cannot prove $\Upsilon(L) \rightarrow 0$. The spin-wave (Jensen) bound $F(\theta) - F(0) \leq (\beta/N) \langle e_{\text{bond}} \rangle \sum_{\text{bonds}} (\Delta\phi)^2$ gives on a torus $\Upsilon_{\text{axis}} \leq 2\beta/N$ (no decay in L) and on an annulus $1 < r < R$ a cost $\leq C\beta\theta^2/\ln R$ for every $N \geq 2$ [9]; in the $O(2)$ BKT phase the annulus cost is $\approx \pi\rho_\infty\theta^2/\ln(R/r)$, which goes to zero exactly as the bound allows while Υ stays finite. **Two annulus limits must be distinguished (correction, v2.0; version 1 conflated them).** With inner radius fixed and $R \rightarrow \infty$ the spin-wave cost vanishes as $1/\ln R$ for every $N \geq 2$, so a bound of the form $Z(\pi)/Z(0) \geq 1 - C_\beta/\log R$ on such annuli is N -blind and cannot by itself imply running. The annulus proposition of the machine-generated decay proof [16] (215-a, Prop. 3.1, file `annulus.tex` at the cited revision) is *not* of this kind: it concerns the fixed-ratio annuli $D_N = \{N \leq \|x\|_\infty \leq 2N\}$, where the spin-wave denominator is $\log 2$ and a finite stiffness gives a cost $\pi\rho\theta^2/\log 2$ that does not vanish as N grows; in the $O(2)$ BKT phase it stays finite. A vanishing twist cost on $N \rightarrow 2N$ annuli is therefore strictly stronger than the capacity bound, and its proof in 215-a explicitly uses the preceding non-abelian fourth-variation estimate (Prop. 2.2), consistent with Section 13. The version-1 sentence “the annulus lemma cannot be the carrier” is retracted; the correct statement is that expanding-ratio annulus bounds are N -blind, while the fixed-ratio statement already contains the non-abelian step. Section 4.6 confirms numerically that the square-torus diagnostics ($\rho_s \rightarrow 0$, $Z(\pi)/Z(0) \rightarrow 1$) detect loss of stiffness, not N : disordered $O(2)$ at $\beta = 0.9$ shows the same collapse as $O(3)$, and loss of stiffness is not by itself a mass-gap criterion.

Where N enters. In perturbation theory $\rho_s(2L) = \rho_s(L) - (N-2) \ln 2/(2\pi) + O(1/\rho_s)$; $N-2$ is the Ricci curvature of S^{N-1} . In the helicity formula it arises because a twist about axis 3 couples to fluctuations in the (1, 3) and (2, 3) planes through $[J_{13}, J_{23}] \propto J_{12}$; for $N = 2$ the bond-energy deficit and the current fluctuation cancel at order $\ln L$. None of the exact identities of Section 5 carries an N -dependent sign. There is no rigorous non-perturbative counterpart: perturbation theory is proved asymptotic only for $N = 2$ in two dimensions; Aru, Garban and Sepúlveda [12]

proved that local $O(3)$ fluctuations at high β are a two-component Gaussian free field and built an example showing that “mostly stiff” does not imply “stays stiff”.

The Garban–de Montgolfier constraint. Since the λ -deformed model has power-law decay at large β for every $\lambda \in (-1, 1)$ [13], inequality (7) must fail for it, as it must for $O(2)$ at $\beta > \beta_{KT}$. Any proof must therefore use the exact $O(3)$ invariance (the curvature, or the non-abelian commutator) and not merely smoothness or near-symmetry. Both required failures were observed (Section 4.6): $O(2)$ BKT drops ≤ 0.004 , λ -model drop 0.0003(7).

7 A finite criterion and a volume-uniform certificate

Records: `purity_bootstrap_03.md` (lemmas, the $n_0 = 4$ computation, controls; v2.1), `seed_push.md` and `seed_push_results.json` (the $n_0 = 5, 6$ computation, Lemmas U and T1', pass/fail map, controls; v2.2; code `seed_push.py`, modules `seed_push_code.tar.gz`, pre-registration `preregistration.md`), `purity_results.json`, `purity_bootstrap_03.py`, `purity_ising.py` (exact Ising validation), `aspect_pocket_ising_2026-10-20.md` (seed aspect on exact Ising), `proof_chain_closure_2026-10-19.md` (v3, the chain), `c1_attack.md` (§3, the doubling-defect seed), `tasaki_verification_2026-10-19.md`. This section supplies piece (ii) of Section 6: a finite-size criterion that turns two ratios of partition functions on tori of one size into a spectral gap uniform in the circle length, and exponential decay in the plane. The criterion is imported from the machine-generated manuscript on the periodic spin-one Haldane gap (OpenAI family 268, file `build/bootstrap.tex` at commit `adc7f1241b`: Lemma “Squaring probabilities”, Lemma “Coupled update”, Proposition “Uniform gap from two initial purities”) [35], read in full. Only its *analytic* part is used; the computer-assisted seed estimates of that manuscript are not. The analytic part has been independently re-derived in expository notes by Tasaki [36] (Section 7.6). Everything specific to the classical $O(3)$ model — Lemma P, the iterated Lemma B, Lemma C, the truncation lemmas U and T1', the seed computation and its error bounds — is new here and is the part a referee should check. (The lemma letters follow the records `purity_bootstrap_03.md` and `seed_push.md` and are unrelated to Lemmas A–D of Section 10; the numbered labels are unambiguous.)

7.1 Setting and transfer to the classical model

Let $Z_{L \times T} = \int \prod_x d\mu(s_x) \prod_{\langle xy \rangle} e^{\beta s_x \cdot s_y}$ on the $L \times T$ torus (L the circle or “space” length, T the “time” length, μ the normalized uniform measure on S^2). Define the spatial and Gibbs purities

$$S(n, b) = \frac{Z_{2n \times b}}{Z_{n \times b}^2}, \quad T(n, b) = \frac{Z_{n \times 2b}}{Z_{n \times b}^2}. \quad (8)$$

Lemma 1 (P: positivity, trace class, all L). *Let $E_L(\sigma) = \exp(\beta \sum_x \sigma_x \cdot \sigma_{x+1})$ on $(S^2)^L$ and $V_L = E_L^{1/2}(\otimes_x K)E_L^{1/2}$ on $L^2((S^2)^L)$, where K has kernel $e^{\beta s \cdot s'}$. Then V_L is positive, trace class and of infinite rank, and $Z_{L \times T} = \text{Tr } V_L^T$ for all $T \geq 1$. The same holds for the column operator W_T , with $Z_{L \times T} = \text{Tr } W_T^L$ for all $L \geq 1$.*

Proof. By Funk–Hecke $K = \sum_\ell a_\ell \Pi_\ell$ with $a_\ell = i_\ell(\beta) > 0$ (modified spherical Bessel functions), so K and $\otimes K$ are positive and injective; congruence by the bounded invertible $E^{1/2}$ preserves both; the kernel is continuous on a compact space, so Mercer gives trace class; the trace formula is the transfer-matrix identity, each horizontal energy appearing once. $\square$ Hence $Z_{L \times T} = \sum_i \lambda_i(T)^L = \sum_k \mu_k(L)^T$ with all $\lambda_i, \mu_k > 0$: exactly the hypothesis of the family-268 Proposition, with every L admitted (the family-268 statement is for even L ; the positivity of Lemma P removes that restriction here). With $\pi_i = \lambda_i(b)^n / Z_{n \times b}$ and $w_k(L, b) = \mu_k(L)^b / Z_{L \times b}$, both S and T are purities $\sum(\text{prob.})^2$ of probability vectors, in $(0, 1]$. The gap is $\gamma_L = \log(\mu_0(L)/\mu_1(L))$.

7.2 The bootstrap (Lemma B)

Let $f(u) = u^2/(2(1-u)^2)$, $G(u) = \frac{1}{2}((1-u)^{-1} + (1-u)^{-2} + (1-u)^{-3})^2$, and $u^* = 0.12731$ the root of $uG(u) = 1$. The ingredients, all verbatim from [35] except B4 and the iteration:

B1 (squaring). A probability vector of purity P has normalized squares of purity P' with $1 - P' \leq f(1 - P)$; applied to π ($n \rightarrow 2n$) and w ($b \rightarrow 2b$): $1 - S(2n, b) \leq f(1 - S(n, b))$, $1 - T(n, 2b) \leq f(1 - T(n, b))$.

B2 (exact cancellations). Definition-level identities, checked line by line:

$$S(n, 2b) = \frac{S(n, b)^2 T(2n, b)}{T(n, b)^2}, \quad T(4n, b) = \frac{T(2n, b)^2 S(2n, 2b)}{S(2n, b)^2}.$$

B3 (coupled update). If $1 - S(n, b) \leq p$ and $1 - T(2n, b) \leq q$ then $1 - S(2n, 2b) \leq p_+ := f(1 - (1 - p)^2(1 - q))$ and $1 - T(4n, 2b) \leq q_+ := f(1 - (1 - q)^2(1 - p_+))$.

B4 (interpolation, all integers). For $n \leq L \leq 2n$: $T(L, b) \geq [T(2n, b)S(n, b)^2]^{L/2n}$, from $\sum \lambda^L \leq (\sum \lambda^n)^{L/n}$ at time b and $\sum \lambda^L \geq (\sum \lambda^{2n})^{L/2n}$ at time $2b$.

B5 (spectrum). If $T(L, b) > \frac{1}{2}$ then $\max_k w_k \geq \sum w_k^2 > \frac{1}{2}$, so $\mu_0(L)$ is simple and $e^{-b\gamma_L} = w_1/w_0 \leq (1 - w_0)/w_0$.

Lemma 2 (B: iterated two-purity bootstrap for classical $O(3)$). *Let $n_0, b_0 \geq 1$ and $0 \leq p, q < 1$ with $1 - S(n_0, b_0) \leq p$, $1 - T(2n_0, b_0) \leq q$. Define (p_{j+1}, q_{j+1}) from (p_j, q_j) by B3, $n_j = 2^j n_0$, $b_j = 2^j b_0$, $P_j = (1 - q_j)(1 - p_j)^2$. Suppose that for some $K \geq 0$: $P_j > \frac{1}{2}$ for all $j < K$, and $u := \max(p_K, q_K)$ satisfies $u < \frac{1}{6}$, $C := G(u)$, $r := Cu < 1$ and $C(1 - 3u) > 3$. Then for every $L \geq n_0$ the top eigenvalue $\mu_0(L)$ of V_L is simple and*

$$\gamma_L \geq \gamma^* := \min \left(\min_{j < K} \frac{1}{b_j} \log \frac{P_j}{1 - P_j}, \frac{\log(1/r)}{b_K} \right). \quad (9)$$

Proof. B3 propagates the bounds to every (n_j, b_j) . For $L \in [n_j, 2n_j]$, $j < K$, B4 gives $T(L, b_j) \geq P_j > \frac{1}{2}$ and B5 the interval bound. For $L \geq n_K$ apply the Proposition of [35] with seed (n_K, b_K) ; its proof uses only B1–B5, the recursion $u_{i+1} = Cu_i^2$ and Hölder, all valid here by Lemma P. The dyadic intervals cover $[n_0, \infty)$. $\square$ The case $K = 0$ is the family-268 Proposition itself ($\gamma_L > \log(1/r)/b_0$). The iterated form matters: it needs the seed defects neither below u^* nor equal to each other, and $\beta = 0.65$ below is certified only through it.

7.3 From the torus to the plane (Lemma C)

Lemma 3 (C). *Let $\gamma = \gamma^*$ from Lemma B and $L \geq n_0$. (a) On the cylinder $L \times \infty$ (the $T \rightarrow \infty$ limit of $L \times T$ tori, which exists because μ_0 is simple), $|\langle s_{(x,0)} \cdot s_{(x',t)} \rangle| \leq e^{-\gamma|t|}$. (b) On the $L \times T$ torus, $|\langle s_0 \cdot s_{(a,t)} \rangle| \leq e^{-\gamma t} + e^{-\gamma(T-t)} + 3(1 - w_0(L, T))$, where w_0 is increasing in T and $1 - w_0 \leq 1 - P_j$ for $L \in [n_j, 2n_j]$, $T \geq b_j$. (c) On square tori with $n_0 = b_0$ the 90° symmetry gives the same bound with $|a|$ in place of t , so every weak limit point of the square-torus Gibbs measures satisfies $|\langle s_0 \cdot s_x \rangle| \leq e^{-\gamma\|x\|_\infty} \leq e^{-\gamma|x|/\sqrt{2}}$.*

Proof. $SO(3)$ commutes with V_L and has no non-trivial one-dimensional representations, so the top eigenvector Ω is invariant and $\langle \Omega, s_x^c \Omega \rangle = 0$. (a) $\langle s_{(x,0)} \cdot s_{(x',t)} \rangle = \sum_c \langle s_x^c \Omega, (V/\mu_0)^t P_\perp s_{x'}^c \Omega \rangle$ and $\|s_x^c \Omega\| \leq 1$. (b) Write $V^t = \mu_0^t P + R_t$; the PP term vanishes, the cross terms are bounded by $\mu_0^T e^{-\gamma(T-t)}$ and $\mu_0^T e^{-\gamma t}$, and Hölder in Schatten norms with exponents T/t , $T/(T-t)$ bounds the RR term by $Z - \mu_0^T$. (c) is (b) plus the lattice rotation. $\square$

Table 7: Certified seeds and volume-uniform gaps. Rows 1–2 are the version-2.1 certificates at $n_0 = 4$ (`purity_results.json`, truncation bound T2’); rows 3–9 are the version-2.2 certificates at $n_0 = 5, 6$ (`seed_push_results.json`, Lemmas U and T1’). p, q are the rigorous seed bounds rounded *up* to five decimals; γ^* is rounded *down*; K is the number of coupled-update steps of Lemma B before the family-268 Proposition is applied; “trunc.” gives the non-rigorous central values of the truncated network for orientation.

β	seed (n_0, b_0)	$p: 1 - S(n_0, b_0) \leq$	$q: 1 - T(2n_0, b_0) \leq$	trunc. p/q	K	$\gamma_L \geq$ (all $L \geq n_0$)	version
0.60	(4,4)	0.05785	0.10727	0.0532 / 0.0985	7	0.2358	2.1
0.65	(4,4)	0.08637	0.15691	0.0760 / 0.1377	10	0.0342	2.1
0.60	(5,5)	0.01818	0.03478	0.0166 / 0.0317	6	0.4611	2.2
0.70	(5,5)	0.04554	0.08413	0.0413 / 0.0759	7	0.2552	2.2
0.75	(5,5)	0.06855	0.12389	0.0620 / 0.1115	8	0.1411	2.2
0.78	(5,5)	0.08643	0.15398	0.0779 / 0.1381	9	0.0392	2.2
0.80	(6,6)	0.05681	0.10465	0.0416 / 0.0755	7	0.1628	2.2
0.82	(6,6)	0.06776	0.12366	0.0498 / 0.0896	8	0.1188	2.2
0.85	(6,6)	0.08736	0.15703	0.0646 / 0.1144	10	0.0195	2.2
0.70	(4,4)	0.12836	0.22781	0.1054 / 0.1865	—	chain diverges at $j = 1$	2.1 (FAIL)
0.80	(4,4)	—	—	$1 - S_{\text{trunc}} = 0.185 > u^*$	—	FAIL	2.1

7.4 The certified statement

Theorem 1 (computer-assisted volume-uniform gap, $0.60 \leq \beta \leq 0.85$). *For each row of Table 7 — a coupling β , a square seed $n_0 = b_0$ and two numbers (p, q) — the seed inequalities*

$$1 - S(n_0, n_0) \leq p, \quad 1 - T(2n_0, n_0) \leq q$$

hold for the exact (untruncated) nearest-neighbour $O(3)$ model. Consequently, by Lemmas P, B and C, for every circle length $L \geq n_0$ the transfer operator V_L has a simple top eigenvalue and $\gamma_L \geq \gamma^*$ with the value in the last column; in particular

$$\gamma_L \geq 0.0195 \quad \text{for every } L \geq 6 \quad \text{at } \beta = 0.85,$$

and every weak limit point of the square-torus Gibbs measures at these couplings satisfies $|\langle s_0 \cdot s_x \rangle| \leq e^{-\gamma^* \|x\|_\infty}$. The certified couplings are $\beta \in \{0.60, 0.65, 0.70, 0.75, 0.78, 0.80, 0.82, 0.85\}$; the largest, 0.85, lies beyond the Aizenman–Simon bound $3/4$ [8] and beyond the fixed- β decay certificate $\beta \leq 0.7701$ of Section 4.7.

The statement is at grid points: no covering argument in β (like Lemma S of Appendix A) has been carried out for the purities, so “ $\beta \leq 0.85$ ” in the abstract means “at the eight certified couplings up to 0.85”. The certified γ^* are far below the physical ring gaps (≈ 1.24 and 1.14 at $\beta = 0.6, 0.65$, non-rigorous); they are existence bounds, not estimates.

What is proved by hand and what by computer. Lemmas P, B (given the family-268 lemmas B1–B3, B5 and the Proposition’s proof, which were read and re-derived in [36]), C, U and T1’ are analytic. The seed inequalities are computer-assisted: each needs (a) upper bounds on $Z_{n_0 \times b_0}$ and $Z_{2n_0 \times b_0}$, (b) a lower bound on $Z_{2n_0 \times b_0}$ and (c) a lower bound on $Z_{2n_0 \times 2b_0}$. In version 2.1 ($n_0 = 4$) these came from an exact $\ell \leq 2$ network with the two-sided truncation bound T2’ and a ring-8 power iteration; in version 2.2 ($n_0 = 5, 6$) they come from the pipeline below. The evaluation of Lemma B from the seed bounds is done in mpmath interval arithmetic (40 digits, $K \leq 24$), and γ^* is attained on the first dyadic intervals (B5), not at the final Proposition. The direct Proposition ($K = 0$) is admissible only at $\beta = 0.6$, $n_0 = 4$ ($u_0 = 0.1073$, $C = 7.146$, $r = 0.766$, $\gamma \geq 0.0665$); everywhere else the iterated form is essential.

Two new lemmas (v2.2).

Lemma 4 (T1’: reflection-positivity monotonicity on tori of any side ≥ 3). *Let every bond kernel K_b depend only on the bond direction, be pointwise positive and have non-negative Funk–Hecke*

coefficients, and let $K'_b - K_b$ have non-negative coefficients. Then $Z(K) \leq Z(K')$. In particular $Z_{\text{trunc}} \leq Z$ for the $\ell \leq 2$ (state bonds) and $\ell \leq 3$ (in-ring bonds) truncations, on tori with even or odd sides.

Proof. Interpolate $K_t = K + t(K' - K)$; then $dZ/dt = \sum_b \int (K' - K)(s_i \cdot s_j) F_b(s_i, s_j)$. Reflect through the plane bisecting b . For an even side the plane bisects a second bond (bond reflection); for an odd side it passes through the fixed column $x_f = (L + 1)/2$ on the far side (site reflection there). The weight factorizes as $A(\sigma_+, \sigma_0) \theta A(\sigma_-, \sigma_0) C(\sigma_0) \prod_{\text{crossing}} K_t$ with $C \geq 0$ collecting the bonds inside the fixed column, and every crossing kernel expands as $\sum_\alpha c_\alpha \phi_\alpha \phi_\alpha$ with $c_\alpha \geq 0$; hence $F_b(s, s') = \int d\sigma_0 C(\sigma_0) \sum_\alpha c_\alpha G_\alpha(s, \sigma_0) G_\alpha(s', \sigma_0)$ is a positive semidefinite kernel, and pairing it with the positive definite $K' - K$ gives $dZ/dt \geq 0$. $\square$ Check: $Z_{\ell \leq 2} \leq Z_{(2,3)} \leq Z_{\ell \leq 3}$ on the 3×3 , 3×4 and 3×5 tori at $\beta = 0.8$ and 1.5 . This removes the even-side restriction of T1 (v2.1).

Lemma 5 (U: convexity upper bound). *Let $\tilde{K}_b > 0$ pointwise and $u_b = \log(K/\tilde{K}_b)$. Then*

$$\log Z(K) \leq \log Z(\tilde{K}) + \sum_b \rho_b, \quad \rho_b = \sup_{|N|, |N'| \leq 3} R_b(N, N'),$$

$$R_b(N, N') = \frac{\iint u_b(s \cdot s') K(s \cdot s') e^{\beta s \cdot N} e^{\beta s' \cdot N'}}{\iint K(s \cdot s') e^{\beta s \cdot N} e^{\beta s' \cdot N'}}.$$

Proof. $Z(t) = \int \prod_b \tilde{K}_b e^{t u_b}$ is an exponential family, so $\log Z(t)$ is convex and $\log Z(1) \leq \log Z(0) + \sum_b \langle u_b \rangle_{\text{exact}}$. Condition on every spin except the two endpoints i, j of b (DLR): in the exact model the three other bonds at i give exactly $e^{\beta s_i \cdot N_i}$ with $|N_i| \leq 3$ (sides ≥ 3 exclude double bonds), so $\langle u_b \rangle_{\text{exact}}$ is an average of R_b over environments and $R_b \leq \rho_b$. $\square$ In Funk–Hecke form $R = \sum_\ell f_\ell(2\ell + 1) i_\ell(y) i_\ell(y') P_\ell(c) / \sum_\ell a_\ell(2\ell + 1) i_\ell(y) i_\ell(y') P_\ell(c)$ with $f_\ell = \frac{1}{2} \int u e^{\beta x} P_\ell$, $y = \beta |N| \in [0, 3\beta]$, $c = \hat{N} \cdot \hat{N}'$. *Rigorous evaluation (sp_rho.py)*: f_ℓ by 160-node Gauss–Legendre with error $\leq 4 \times 10^{-19}$ from the Bernstein-ellipse bound ($\rho = 1.2$; $\text{Re } \tilde{K} > 0$ on the ellipse checked in interval arithmetic, so $\log \tilde{K}$ is analytic there); the $\ell > 16$ tail bounded by $\sup |F| \sum (2\ell + 1) i_\ell(3\beta)^2$; the supremum over the box by branch-and-bound in mpmath interval arithmetic to 1% above the best point value; it sits at $y = y' = 3\beta$, $c \approx -0.5$ ($\ell \leq 2$) and $c = -1$ ($\ell \leq 3$). Values: $\rho_2 = 3.17 \times 10^{-5}, 8.86 \times 10^{-5}, 1.389 \times 10^{-4}, 1.829 \times 10^{-4}, 2.187 \times 10^{-4}, 2.603 \times 10^{-4}, 3.354 \times 10^{-4}$ and $\rho_3 = 2.11 \times 10^{-7}, 9.69 \times 10^{-7}, 1.94 \times 10^{-6}, 2.88 \times 10^{-6}, 3.72 \times 10^{-6}, 4.78 \times 10^{-6}, 6.86 \times 10^{-6}$ at $\beta = 0.60, 0.70, 0.75, 0.78, 0.80, 0.82, 0.85$, against the v2.1 bound T2' of $7.6 \times 10^{-5}, 4.07 \times 10^{-4}, 1.09 \times 10^{-3}, 2.11 \times 10^{-3}, 3.89 \times 10^{-3}$ at $\beta \leq 0.80$: Lemma U is $18\times$ smaller per bond at $\beta = 0.8$, and T2' breaks down at $\beta \approx 0.85$ (its denominator turns negative) while Lemma U does not. A first-order “best $\ell \leq 2$ kernel” variant gave no gain (the supremum moves to the aligned-field corner). T2' is superseded; it remains the bound behind rows 1–2 of Table 7.

Certified seed computation (v2.2 pipeline).

Network. Complex spherical-harmonic basis; state bonds (along the transfer direction) carry $\ell \leq 2$ (9 labels), in-ring bonds carry $\ell \leq 3$ (16 labels, entering only the MPO bond dimension); the vertex tensor $G[a, b, h, h'] = \int \bar{Y}_a Y_b \bar{Y}_h Y_{h'} d\mu$ is exact (sympy Gaunt coefficients summed at 40 digits, rounded once); M is exactly real symmetric ($G[a, b, h, h'] = G[b, a, h', h]$ and $A^{ab} = S(A^{ab})^T S$ with $S_{h\bar{h}} = (-1)^{m_h}$).

(a) **Upper bounds.** $Z_{n \times b_0} \leq Z_{\text{trunc}} e^{\sum_b \rho_b}$ by Lemma U, with $Z_{\text{trunc}} = \text{Tr } M_{b_0}^n$ an exact trace of the ring- b_0 column transfer matrix for $n \in \{n_0, 2n_0\}$; $b_0 \leq 6$ is the hard limit (ring 6: $9^6 = 531441$ states). *Symmetry reduction.* M_L commutes with $U(1)$ (total weight W), translations (k) and $SO(3)$; every integer-spin irrep has exactly one $W = 0$ state, so $\text{Tr } M^p = \sum_k \sum_J (2J + 1) \text{Tr}(B_{0,k}|_J)^p$ with $J^2 = J_+ J_-$ on $W = 0$ assembled from the sparse ladder matrix. Enclosures per J : Davis–Kahan ($\|\sin \Theta_J\|_F \leq \|J^2 W_J - J(J+1)W_J\|_F / \text{gap}_J$, $\text{gap}_J = 2J$, exact integer Casimir spectrum), Ostrowski for the orthonormalization factor, Weyl on $Q_J^* \hat{B} Q_J$ with an a-posteriori residual and the build error $\|\hat{B} - B\|_F \leq$

$\gamma_{4H^2+200}\|B_{\text{abs}}\|_F$ ($|G|$ network), and $|\text{Tr}X^p - \text{Tr}Y^p| \leq p\|\cdot\|^{p-1}\|X - Y\|_*$. For rings 4 and 5 an all- W route (Weyl+Ostrowski per (W, k) block) agrees to 6.6×10^{-15} . About 25 minutes per β at ring 6.

- (b) **Lower bound on $Z_{2n_0 \times b_0}$.** $Z \geq Z_{\text{trunc}}$ by Lemma T1', with the same exact trace.
- (c) **Lower bound on $Z_{2n_0 \times 2b_0}$.** $Z \geq Z_{\text{trunc}} \geq \lambda_{\max}(M_L)^T \geq (\langle v|M_L|v \rangle / \langle v|v \rangle)^T$ (Rayleigh, M_L positive semidefinite) on the larger of ring $2n_0$ and ring $2b_0$ ($L = 10$ or 12), with v a periodic uniform matrix-product state of bond dimension $\chi = 24$ from canonical-form power iteration (any v is admissible); $\langle v|M|v \rangle = \text{Tr} E_M^L$ and $\langle v|v \rangle = \text{Tr} E_0^L$ are contracted exactly, with a sequential normwise rounding bound and the construction error propagated through $\|\cdot\|_*$ and $\|E\|_2 \leq (\|E\|_1\|E\|_\infty)^{1/2}$. Relative slack of the Rayleigh bound: 1.4×10^{-8} (ring 10, $\beta = 0.78$), 3.1×10^{-8} and 4.3×10^{-8} (ring 12, $\beta = 0.80, 0.85$).

Error budget. Truncated $\rightarrow$ rigorous defects: $1-S$: $0.0779 \rightarrow 0.0864$, $1-T$: $0.1381 \rightarrow 0.1540$ at $\beta = 0.78$ (Lemma U log-corrections 4.64×10^{-3} on $Z_{5 \times 5}$, 9.29×10^{-3} on $Z_{10 \times 5}$; trace relative widths 1.5×10^{-8} , 3.0×10^{-8}); $0.0416 \rightarrow 0.0568$, $0.0755 \rightarrow 0.1046$ at 0.80 (8.01×10^{-3} , 1.60×10^{-2} ; widths 4.4×10^{-7} , 8.4×10^{-7}); $0.0498 \rightarrow 0.0678$, $0.0896 \rightarrow 0.1237$ at 0.82 ; $0.0646 \rightarrow 0.0874$, $0.1144 \rightarrow 0.1570$ at 0.85 (1.23×10^{-2} , 2.46×10^{-2}). The Lemma U term on the state bonds dominates; rounding is $\leq 10^{-6}$ relative. Without the Lemma U cost the truncated values would give $\gamma \approx 0.098$ (0.78) and 0.138 (0.85). Arithmetic: IEEE-754 double precision with a-priori and a-posteriori bounds for the network, interval arithmetic for ρ_b and for Lemma B; not ball arithmetic throughout.

Validation. Complex-basis blocks against the v2.1 real-basis dense rings 3 and 4: full spectra agree to 4.2×10^{-15} ($Z_{4 \times 4} = 7.5856142515459855$ against 7.58561425154599); the v2.1 ring-8 power iterates reproduced to $\leq 9.4 \times 10^{-14}$; transposition checks $Z_{4 \times 5}$ (ring 5 vs ring 4) to 10^{-16} and $Z_{6 \times 3}$, $Z_{6 \times 4}$ (ring 6 vs rings 3, 4 with swapped truncation) to 5×10^{-15} ; the MPS top eigenvalue on ring 6 at $\beta = 0.7$ equals the exact $W=0, k=0$ block value 2.76697038 ; on ring 8 the MPS gives 3.878606003 , above the v2.1 certified 3.878605889 , as it must; the $\lambda = 1$ limit of the deformed-model code reproduces $O(3)$ to 10^{-13} .

Pass/fail map, seed shape and pre-registration (v2.2). Table 8 gives the certified γ^* over the scanned seeds (n_0, b_0) . Readings: at fixed b_0 , $1 - S$ falls with n_0 while $1 - T(2n_0, b_0)$ rises with n_0 (both terms are extensive); the best seeds inside the feasible box $b_0 \leq 6$ are square, $(5, 5)$ and $(6, 6)$; at fixed $n_0 = 4$ the aspect-1.25 seed $(4, 5)$ certifies at $\beta = 0.70$ where $(4, 4)$ fails, while aspect 1.5, $(4, 6)$, fails at $\beta = 0.80-0.85$ on $1 - S(4, 6) \geq 0.25$. The same ‘‘aspect pocket’’ exists in the exactly solvable Ising model (`aspect_pocket_ising_2026-10-20.md`, Kaufman partition functions at $K = 0.38$, $\xi = 3.93$): at $n_0 = 12$ the criterion certifies only for $T_0 = 14-22$ (aspect 1.17–1.83) and fails at aspect 1 and ≥ 2 ; the square seed needs $n_0 = 13$ while $T_0 = n_0 + 2$ certifies already at $n_0 = 11$; the defect $-2 \log S - \log T(2n_0, T_0)$ is $0.322, 0.273, 0.254, 0.253, 0.262, 0.296$ at $T_0 = 12, 14, 16, 18, 20, 24$, minimal near aspect 1.4, because $1 - S$ increases and $1 - T$ decreases with T_0 . *Pre-registration scored* (`preregistration.md`, written 05:34 CEST before any new seed number): $(4, 5)$ at $\beta = 0.70$ was predicted a borderline FAIL and passes ($\gamma \geq 0.138$); $\beta = 0.85$ was predicted to FAIL for all $b_0 \leq 6$ and $(6, 6)$ passes; $(6, 6)$ at $\beta = 0.80$ was predicted a marginal PASS and is a clear pass; both wrong predictions erred in the pessimistic direction. The predicted $(4, 6)$ PASS at $\beta = 0.70$ was not assessed (ring-6 traces at 0.70 were not run before the stop). *Not assessed*: $b_0 = 6$ at $\beta \notin \{0.80, 0.82, 0.85\}$ (stopped on instruction, about 45 minutes per β); $b_0 \geq 7$ at every β (exact ring- ≥ 7 traces infeasible here); $T_0 \in \{8, 10, 12\}$ (needs exact ring- ≥ 8 traces). With $b_0 \leq 5$ every seed fails at $\beta \geq 0.80$. The trend at $(6, 6)$, $q = 0.105 \rightarrow 0.124 \rightarrow 0.157$ for $\beta = 0.80 \rightarrow 0.82 \rightarrow 0.85$, makes 0.88 doubtful at $b_0 = 6$; that trend is not a result.

Reach and cost. Version 2.1 reached $\beta = 0.65$ with $n_0 = 4$ (ring 8, 4.3×10^7 states, power iteration) and failed at 0.70 and 0.80 ; version 2.2 reaches 0.85 with $n_0 = 6$ because the exact traces

Table 8: Pass/fail map of the certificate over seed shapes (n_0, b_0) (`seed_push_results.json`, `pass_fail_map`). Entries are certified γ^* as recorded in the results file (four decimals, nearest rounding; Table 7 gives the rounded-down values for the best seeds); $\times$ = fails; n.a. = not assessed. Rows $\beta \geq 0.88$ fail at every assessed seed with $b_0 \leq 5$.

β	(3,4)	(4,4)	(5,4)	(6,4)	(7,4)	(8,4)	(3,5)	(4,5)	(5,5)	(6,5)	(7,5)	(8,5)	(3,6)	(4,6)	(5,6)	(6,6)	(7,6)	(8,6)
0.60	$\times$	0.2599	0.2471	0.1923	0.1329	0.0628	$\times$	0.3321	0.4611	0.4567	0.4287	0.4008	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.
0.70	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	0.1384	0.2553	0.2499	0.2190	0.1855	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.
0.75	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	0.1412	0.1353	0.0951	0.0370	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.
0.78	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	0.0392	0.0305	$\times$	$\times$	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.
0.80	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	0.1105	0.1629	0.1555	0.1320
0.82	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	0.0575	0.1189	0.1107	0.0826
0.85	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	0.0196	$\times$	$\times$
0.88–1.0	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	$\times$	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.

are reduced to the $W = 0$ sector of ring 6 and the large-ring lower bounds use exactly contracted matrix-product states instead of power iteration. The remaining wall is the exact trace: $b_0 = 7$ needs ring-7 traces ($9^7 = 4.8 \times 10^6$ states before reduction) and the pre-registered aspect points $T_0 \in \{8, 10, 12\}$ need exact traces on rings of ≥ 8 sites. The certified range now exceeds the fixed- β decay results (3/4 [8]; 0.7701 in Section 4.7), and the statement is of a stronger kind (a gap of the circle transfer operator for every $L \geq n_0$, hence decay on every cylinder and in the plane).

Controls. *Exact Ising (own Kaufman code, `purity_ising.py`, `mpmath` 60 digits, verified against brute-force enumeration to 10^{-14} for $L \leq 4$; v2.1).* The smallest square seed satisfying the Proposition is $n_0/\xi = 3.67, 3.80, 3.51, 3.19, 3.19, 3.14$ at $K = 0.2, 0.3, 0.35, 0.4, 0.42, 0.43$ ($\xi = 0.82\text{--}23.2$); the iterated chain lowers this to $2.9\text{--}3.7\xi$. The best certified gap (seed 9.4–11.4 ξ) is $0.903/0.480/0.289/0.121/0.061/0.031$ against the exact $1.223/0.633/0.390/0.168/0.084/0.043$ (factor 1.35–1.4). Required failure: at $K = K_c = 0.4407$ and at $K = 0.5$ (ordered) the minimal defect over $n \leq 256$ is 0.54 and 0.50, far above u^* . Our $G(u)$ reproduces Tasaki’s eq. (4.15), $G(1/8) = 913952/117649$. *O(2) (v2.1, FK π -twist estimator, Monte Carlo).* At $\beta = 1.3$ (BKT phase) $Z(\pi)/Z(0) = 0.0208, 0.0178, 0.0179, 0.0174$ for $L = 8\text{--}64$ is scale-invariant, so $p_0 \leq (1 + Z(\pi)/Z(0))/2 \approx 0.51$ and the seed fails at every L ; at $\beta = 0.7$ (massive) the seed proxy passes from $L = 16$ on. *O(2) and $\lambda = 0.5$ in the v2.2 pipeline (truncated network, non-rigorous; by Lemma T1’ the rigorous defects are at least as large; seeds with $b_0 \leq 5$ only).* O(2) (U(1) charges $|n| \leq 2$ on state bonds, ≤ 3 in-ring): PASS at $\beta = 0.4$ at every seed and at $\beta = 0.5$ at (4, 5), (5, 5) (massive phase: the pipeline does not fail trivially on O(2)); FAIL at every seed for $\beta = 0.6, 0.7, 0.8, 1.0$ and in the BKT phase at $\beta = 1.2, 1.5$ — the BKT control fails as required, while in the massive window 0.6–1.0 O(2) also fails at these sizes, so the control does not discriminate there. $\lambda = 0.5$ (bond kernel compressed to $\ell \leq 2/3$; Lemmas T1’/U do not apply to this compression): PASS at $\beta = 0.75$ (high-temperature massive regime, $\gamma_{\text{trunc}} = 0.22$ at (5, 5)); FAIL at $\beta = 1.5$, the large- β regime where [13] prove power-law decay “for β large enough” — whether 1.5 is already large enough is not established. *Estimator check (v2.1).* On the exact 4×4 torus at $\beta = 0.6$, $Z(\pi)/Z(0) = 0.96404$ (transfer matrix) against 0.96396 ± 0.0005 (Monte Carlo); the singlet-dominance proxy $p_0^2 = 0.9467$ against the exact $T(4, 4) = 0.9468$.

Physical coupling, $\beta = 1.6$ (Monte Carlo, non-rigorous, [provisional]; v2.1). The purities were estimated with the exact FK odd-winding estimator of $Z(\pi)/Z(0)$ on the relevant seam (Wolff updates; $(5z - 1)/4 \leq p_0 \leq (3z + 1)/4$ from the SO(3) characters) and the proxy $T \approx p_0^2$, which assumes singlet-sector dominance by the vacuum (exact at 4×4 to 10^{-4} , see above, but unproved in general and precisely the gap in E13). Results: $S(64, 64)$ from $Z(\pi)/Z(0) = 0.8785(20)$ gives $1 - S \approx 0.17\text{--}0.28$ (FAIL); $T(128, 64)$ from 0.7981(28) gives $0.28\text{--}0.44$ (FAIL); $S(128, 128)$ from 0.9947(5) gives $0.008\text{--}0.013$ (PASS); $T(256, 128)$ from 0.9897(7) gives $0.015\text{--}0.026$ (PASS). The worst case $u_0 = 0.0257$ at seed (128, 128) would give $\gamma \geq 0.016$ (Proposition)

or 0.021 (chain) if these purities held rigorously; the true mass is $1/\xi \approx 0.05$; $n_0 = 128 \approx 6.7\xi$ against $\approx 3\xi$ for Ising (three species, every one-particle multiplet weighs $3\times$). Section 8 measures the defect Δ itself at $\beta = 1.4$ and 1.6 with a dedicated estimator.

7.5 The seed observable and its threshold

Lemma 6 (doubling-defect seed). *Define $\Delta(n, w) := 4 \log Z_{n \times w} - \log Z_{2n \times 2w}$. Then*

$$\Delta(n, w) = -2 \log S(n, w) - \log T(2n, w), \quad \text{both terms} \geq 0,$$

so $1 - S(n, w) \leq -\log S \leq \Delta/2$ and $1 - T(2n, w) \leq \Delta$. Over the worst split of Δ into its two terms, the iterated Lemma B is admissible whenever $\Delta(n_0, w_0) < \Delta^* = 0.2239$; hence $\Delta(n_0, w_0) < \Delta^*$ is a sufficient seed.

Proof. $4 \log Z(n, w) - \log Z(2n, 2w) = [4 \log Z(n, w) - 2 \log Z(2n, w)] + [2 \log Z(2n, w) - \log Z(2n, 2w)]$; by Lemma P all transfer eigenvalues are positive, so $S, T \in (0, 1]$ and $-\log S, -\log T \geq 0$. For the threshold put $a = -2 \log S$, $b = -\log T$, $p = 1 - e^{-a/2}$, $q = 1 - e^{-b}$ and minimize over $a + b = \Delta$ the admissibility of Lemma B (`liminf_formulation.py`): the binding constraint is $P_1 > \frac{1}{2}$, which gives $\Delta^* = 0.22391$; the basin of the coupled update B3 has the same threshold $\Delta^{**} = 0.22391$. The basin part is certified at 0.221 in interval arithmetic (monotone box cover); the value 0.2239 itself is a floating-point computation. At $\Delta = 0.1$ the worst-case certified gap is $\gamma \geq 0.441/b_0$, at 0.2 it is $0.0019/b_0$; on $u \leq 0.05$ the map contracts by ≤ 0.277 per step and then quadratically. $\square$ The seed observable

is thus the free-energy defect of one simultaneous doubling, the classical analogue of Tasaki's $D_L = E_0(2L) - 2E_0(L)$ (and $Z_{2L} \leq Z_L^2$, i.e. $E_0(2L) \geq 2E_0(L)$, follows from the spatial transfer representation alone [36]). It replaces the withdrawn link through ρ_s (E13); version 2.1 used the threshold 0.1, which is superseded by Δ^* . The rigorous $\beta = 0.6/0.65$, $n_0 = 4$ certificates lie in the admissible region $\mathcal{A}$ of Lemma B but not in $\{\Delta < \Delta^*\}$ (their bounds give only $\Delta(4, 4) \leq 0.233$ and ≤ 0.352): the admissible region, not the scalar threshold, is the sharp seed condition. The machine-generated O(3) continuum manuscript (215-c, Proposition “Preliminary mixing and rectangular doubling”, part (ii)) asserts $\Delta(n, w) \leq C(1 + b + n + w)^p e^{-c \min(n, w)/X(b)}$ for even, aspect-bounded $n, w \geq CX(b)$; if that estimate held it would supply the seed at every β , but it is unrefereed and proved from the same pinned iteration as the stiffness lemma of Section 13. The statement that remains for arbitrary β is formulated, and shown to be sharp, in Section 8.

7.6 External status of the imported criterion

Tasaki's two expository notes [36] (continuous- and discrete-time versions, 34 pages each, published 18 October 2026 and revised 19 October; fetched and read on 19 October) give a self-contained account of the analytic argument of the family-268 manuscript — the gap criterion in terms of partition functions, the spatial transfer representation and the extension to all large rings — and state explicitly that the computer-assisted finite inputs of that manuscript were not reproduced or checked there. They are cited here as an independent re-derivation of the analytic part that this section imports, not as a referee report, and not as a check of anything computed in this paper. Two further items from those notes are used: (i) Appendix D, Proposition D.1 (converse): for a family with a spatial transfer representation, unique ground states, gap $\geq \Delta$ and $E_0(2L) - 2E_0(L) = o(1/L)$, the two initial purity inequalities are satisfied at $(L, \beta_0 = cL)$ for all large even L and $c > 2 \log \nu / \Delta$. Consequence for O(3): the finite criterion detects *every* gapped coupling (a seed exists at $n_0 \sim L$, $b_0 \sim cL$), so the only obstruction to closing the chain is the reach statement C1'', not the criterion; Theorem 3 is the classical-torus form of this converse. (ii) The identity $Z_{2L} \leq Z_L^2$ from the transfer representation alone (their footnote 17). The $\beta = 1.6$ Monte Carlo seeds above are the empirical form of (i): the criterion is met at $n_0 = 128$.

Scope. Theorem 1 is a statement about eight fixed couplings. It does not give a gap uniform in β and says nothing about scaling or the continuum. Its value is that it is the first torus-native, volume-uniform rigorous statement in this programme, that it now reaches beyond the fixed- β decay results, that it is obtained from finite computations whose every step is listed above, and that it fixes the seed observable (Δ) and the single statement (C1'', Section 8) on which the general case rests.

8 The dichotomy and the remaining statement C1''

Records: `liminf_formulation.md` (exact inequalities, Theorems D and Conv, thresholds, universal constant, Monte Carlo; code `liminf_formulation.py`, data `liminf_results.csv`, pre-registration `PREREG_liminf.md` frozen 2026-10-20T06:02:54Z before any $O(2)/O(3)/\lambda$ run), `proof_chain_closure_2026-10-19.md` (v3). Everything in Sections 8.1–8.4 is exact and uses only Lemma P, the 90° lattice symmetry (the column operator W_T is unitarily equivalent to V_T) and the family-268 bootstrap; the Monte Carlo of Section 8.5 is **[provisional]**.

8.1 Exact inequalities for the doubling defect

Notation: $e_0(L) = -\log \mu_0(L)$; $F = \lim L^{-2} \log Z(L, L)$; $\varepsilon(L) = e_0(L) + FL$ (Casimir energy), $C(L) = L\varepsilon(L)$; $D_L = e_0(2L) - 2e_0(L)$; $R(L, T) = \sum_{k \geq 1} (\mu_k(L)/\mu_0(L))^T$ (excitation sum of the ring of circumference L run for time T); $\ell(x) = \log(1+x)$; $c(L) = \log Z(L, L) - FL^2$ (torus constant); $\Delta(L) = \Delta(L, L)$.

Proposition 1 (exact sign and monotonicity laws). *For lattice $O(N)$ on tori, any N , any $\beta \geq 0$:*

- (a) $\Delta(n, w) \geq 0$ (Lemma 6).
- (b) For fixed w , $n \mapsto Z(n, w)^{1/n} = \|\lambda(w)\|_{\ell^n}$ is non-increasing, so $Z(m, w) \leq Z(n, w)^{m/n}$ for all $m \geq n$; hence $\log Z(L, L)/L^2$ is non-increasing, $c(L) \geq 0$, the generalized defect $(L'/L)^2 \log Z(L, L) - \log Z(L', L') \geq 0$, and $\log Z(n, w)$ is convex in n and in w separately. Lower bounds: $Z(2n, w) \geq Z(n, w)^2(1 + R(w, n))^{-2}$ and $Z(n, 2w) \geq Z(n, w)^2(1 + R(n, w))^{-2}$.
- (c) $2e_0(L) \leq e_0(2L) \leq 2e_0(L) + \min\{4\beta, \ell(R(2L, L))/L\}$; hence $e_0(L)/L$ is non-decreasing, $\varepsilon(L) \leq 0$ (the Casimir energy is attractive) and $C(L) \leq 0$.
- (d) Spectral identity: $\Delta(L) = 2LD_L + 4\ell(R(L, L)) - \ell(R(2L, 2L))$. Sandwich: $S(L, L) \in [\pi_0^2, \pi_0]$ and $T(2L, L) \in [w_0^2, w_0]$ with $\pi_0 = 1/(1 + R(L, L))$, $w_0 = 1/(1 + R(2L, L))$; hence

$$2\ell(R(L, L)) + \ell(R(2L, L)) \leq \Delta(L, L) \leq 4\ell(R(L, L)) + 2\ell(R(2L, L)).$$

- (e) 4-adic identities: $c(L) = \sum_{j \geq 0} 4^{-(j+1)} \Delta(2^j L)$ and $C(L) = -\sum_{j \geq 0} 4^{-(j+1)} \delta(2^j L)$ with $\delta(L) = 2LD_L$; the torus constant at L is a weighted average (total weight $\frac{1}{3}$) of all future defects, so $\Delta \rightarrow \Delta_\infty \Leftrightarrow c \rightarrow \Delta_\infty/3$.

Proof sketch. (b): $\|\cdot\|_{\ell^n}$ norms of the positive vector $\lambda(w)$ (Lemma P) decrease in n ; the lower bounds are $S = \sum \pi_i^2 \geq \pi_0^2$ by the 90° symmetry and its time-transfer analogue. (c): the left inequality is $Z(2L, T) \leq Z(L, T)^2$ as $T \rightarrow \infty$; the 4β bound is variational, $V_{2L} = X^{1/2}(V_L \otimes V_L)X^{1/2}$ with X the multiplication operator of the four seam bonds, trial vector $X^{-1/2}(\Omega_L \otimes \Omega_L)$, $X^{-1} \leq e^{4\beta}$; the $\ell(R)/L$ bound follows from $-\log S(L, L) = LD_L + 2\ell(R(L, L)) - \ell(R(2L, L))$ with $S \geq \pi_0^2$, and $S \leq \pi_0$ gives the matching lower bound $\ell(R(2L, L)) - \ell(R(L, L)) \leq LD_L$. (d), (e): algebra from the definitions and $c(L)/L^2 \rightarrow 0$. $\square$

Physical meaning of (d): $\Delta(L)$ is small exactly when the rings of circumference L and $2L$, run for imaginary time L , have relaxed into their vacua; the bulk free energy and the Casimir term drop

out. A necessary condition for a seed is $L m_L \geq \log(g/0.0776)$, i.e. ≥ 3.66 for an $\text{SO}(3)$ triplet ($g = 3$) and ≥ 2.56 for Ising ($g = 1$). All of (a)–(e) were checked on exact Ising (Kaufman, 60 digits) at $K \in \{0.3, 0.4, K_c, 0.47, 0.5, 0.6\}$ and $L = 1, \dots, 128$ (96 rows, all pass; the sign checks $c \geq 0$, $\varepsilon \leq 0$ fail only at the 10^{-27} level, the precision of the bulk integral times L^2).

Monotonicity of Δ cannot be proved from positivity alone (closed). Exact Ising, which satisfies every structural hypothesis used above, has $\Delta(4) = 2.041 < \Delta(16) = 2.081$ at $K = 0.47$ and $\Delta(1) = 1.859 < \Delta(2) = 1.892$ at $K = 0.40$, which is massive. Any proof of decrease of Δ must use the non-abelian structure. (In the $\text{O}(3)$ Monte Carlo below, Δ decreases at every tested L .)

8.2 The dichotomy

Theorem 2 (D: dichotomy). *Assume Lemma P and the 90° symmetry (lattice $\text{O}(N)$), any N , any β ; also Ising). If $\Delta(L_0) < \Delta^* = 0.2239$ for some L_0 (more generally, if the seed pair $(1 - S(L_0, L_0), 1 - T(2L_0, L_0))$ lies in the admissible region $\mathcal{A}$ of Lemma B), then (i) $\gamma_L \geq \gamma^* > 0$ for all $L \geq L_0$; (ii) $\Delta(2^j L_0) \leq -3 \log(1 - u_j)$ with $u_{j+1} \lesssim C u_j^2$, i.e. doubly exponential decay along the dyadic sequence; (iii) $\Delta(L) \rightarrow 0$ exponentially over all integers L . Consequently, at every β exactly one of the following holds: either $\Delta(L) \geq \Delta^*$ for every $L \geq 1$, or $\Delta(L) \rightarrow 0$ and the volume-uniform gap holds; equivalently $\inf_L \Delta(L) < \Delta^* \Leftrightarrow \lim_L \Delta(L) = 0$, and $\Delta_\infty = \lim \Delta \notin (0, \Delta^*)$.*

Proof. (i) is Lemma B. For (ii), B3 propagates $1 - S(n_j, b_j) \leq p_j$ and $1 - T(2n_j, b_j) \leq q_j$, and $\Delta(n_j, b_j) = -2 \log S - \log T$. For (iii): by the sandwich (d), (ii) gives $LD_L \rightarrow 0$ along $2^j L_0$; the 4-adic identity (e) then gives $\varepsilon(2^j L_0) \rightarrow 0$, so $|\varepsilon(T_0)| < \gamma^*/2$ for some T_0 , and Theorem 3 applies. $\square$

Theorem 3 (Conv: converse). *Suppose $\gamma_L \geq \gamma > 0$ for $L \geq L_0$ and $|\varepsilon(T_0)| < \gamma/2$ for some $T_0 \geq L_0$. Then $R(L, L) \leq (1 + R(T_0, T_0))e^{-\gamma(L-T_0)+L|\varepsilon(T_0)|}$ and $R(2L, L) \leq (1 + R(T_0, T_0))e^{-\gamma(L-T_0)+2L|\varepsilon(T_0)|}$, so $\Delta(L) \rightarrow 0$ at rate $\geq \gamma - 2|\varepsilon(T_0)|$.*

Proof. $R(L, L) \leq e^{-\gamma(L-T_0)}R(L, T_0)$ and $R(L, T_0) \leq Z(L, T_0)/\mu_0(L)^{T_0} = \exp(T_0\varepsilon(L) - L\varepsilon(T_0) + \ell(R(T_0, L)))$ by the 90° symmetry; then use $\varepsilon(L) \leq 0$. $\square$ This is the classical analogue of Tasaki's Proposition D.1 [36], with $o(1/L)$ weakened to " $|\varepsilon(T_0)| < \gamma/2$ at a single T_0 ".

Equivalence chain (dyadic sense where marked). $[\forall \beta : \inf_L \Delta < \Delta^*] \Leftrightarrow [\forall \beta : \Delta \rightarrow 0] \Leftrightarrow [\forall \beta : c(L) \rightarrow 0] \Leftrightarrow [\forall \beta : R(L, L), R(2L, L) \rightarrow 0] \Leftrightarrow [\forall \beta : \text{uniform gap and } |\varepsilon(T_0)| < \gamma/2 \text{ at one } T_0] \Rightarrow [\forall \beta : L\varepsilon(L) \rightarrow 0]$. The last arrow is not reversible: low-temperature Ising has $2LD_L \rightarrow 0$ and $\Delta \rightarrow 3 \log 2$. The version-2.1 additive C1' sits strictly above the first bracket and is false (E14). Let $\mathcal{A}$ be the admissible region of Lemma B (open, since all its conditions are strict) and $B_{\text{good}} = \{\beta : \exists(n, w), (1 - S, 1 - T) \in \mathcal{A}\} \supseteq \{\beta : \exists L, \Delta(L; \beta) < \Delta^*\}$; B_{good} is open (finite-volume Z is analytic in β) and contains 0 and the eight certified couplings. If β_c is the first point of its complement, the seed scale $L_\eta^*(\beta) := \min\{L : \Delta(L; \beta) \leq \Delta^* - \eta\}$ tends to infinity as $\beta \uparrow \beta_c$, for every η .

8.3 The universal constant of a gapless infrared theory

Along the second branch of Theorem 2, $\Delta_\infty = 3 \log Z_{\text{IR}}(\tau = i)$, where Z_{IR} is the modular-invariant torus partition function of the infrared theory including the Casimir factor $e^{\pi c/6}$. For a unitary CFT $Z(i) \geq e^{\pi c/6}$, so $\Delta_\infty \geq \pi c/2 \geq \pi/4 > \Delta^*$, and its Casimir part is $2LD_L \rightarrow \pi c/2$ (from $e_0 = fL - \pi c/(6L)$). Checks and controls: *Ising at K_c* (exact): $\Delta(128) = 1.9197 = 3 \log[(2\theta_2 + \theta_3)/(2\eta)](i) = 3 \times 0.63991$ and $2LD_L(128) = 0.785 = \pi/4$. *Two ordered vacua* give

$\Delta_\infty = 3 \log 2$. *Compact boson at stiffness K* : $Z(i) = \sqrt{2\pi K} \theta_K^2 / \eta(i)^2$ with $\theta_K = \sum_m e^{-2\pi^2 K m^2}$ and $\eta(i) = \Gamma(\frac{1}{4}) / (2\pi^{3/4})$, so $3 \log Z = 3.66, 4.10, 4.28, 4.34$ at $K = 2/\pi, 0.85, 0.96, 1.0$, always $> \Delta^*$ in the BKT phase ($K \geq 2/\pi$): **the O(2) control fails at every L** , and the measured O(2) plateau at $\beta = 1.3$ (4.30–4.42, Table 9) matches $3 \log Z_{\text{CB}}(i; \rho_{\text{measured}}) = 4.28\text{--}4.37$ to ≤ 0.14 . *Easy-plane $\lambda = 0.5$ model* at $\beta = 3.0$: plateau 5.71–5.84 (the record labels this control a substitute because the original definition was not accessible to that track; it coincides with the Garban–de Montgolfier form $\langle \sigma, \sigma' \rangle_\lambda = \sigma^x \sigma'^x + \sigma^y \sigma'^y + \lambda \sigma^z \sigma'^z$ used throughout this paper, Section 2). *Spin-wave form for $O(N)$ ($L \ll \xi$)*: $c(L) \approx \log |S^{N-1}| + \frac{N-1}{2} \log(\rho_L/2\pi) - 2(N-1) \log \eta(i)$ (lattice check: $\log \det' \Delta_{L \times L} - L^2 \cdot 4G/\pi - \log L^2 = -1.05470$ at $L = 256$ against $4 \log \eta(i) = -1.05469$); for O(3), $\Delta_{\text{sw}} = 4 \log(2\rho_L) - \log(2\rho_{2L}) + 3.164$, so in $Y = e^{\Delta/3} \approx 5.74\rho$ the one-loop drop per doubling is $5.74 \ln 2 / (2\pi) = 0.63$, uniform, while in Δ itself it is $\approx 3\kappa e^{-\Delta/3}$, non-uniform. This is the precise shape the additive C1' would have needed, and the reason it fails.

8.4 The remaining statement C1''

C1'' (reach): for every $\beta \geq 0$ there is $L^*(\beta) < \infty$ with $\Delta(L^*, L^*; \beta) < \Delta^* = 0.2239$ (10)

(sufficient: a seed pair in the admissible region $\mathcal{A}$ of Lemma B). By Theorem 2 and the equivalence chain, C1'' is equivalent to each of: $\forall \beta, \lim_L \Delta(L) = 0$; $\forall \beta$, no universal torus constant ($c(L) \rightarrow 0$); $\forall \beta$, the scaled ring excitation weights $R(L, L), R(2L, L) \rightarrow 0$; the seed scale $L_\eta^*(\beta)$ is locally bounded on $[0, \infty)$, i.e. there is no finite β_c at which the seed scale diverges. By the sandwich, the complementary branch $\inf_L \Delta \geq \Delta^*$ means $R(L, L)$ or $R(2L, L) \geq e^{\Delta^*/6} - 1 = 0.038$ at every scale — a scale-invariant excitation weight, i.e. a critical or degenerate infrared theory; Mermin–Wagner excludes degenerate SO(3)-breaking vacua, so what C1'' asks is the exclusion of a massless (Patrascioiu–Seiler) phase [14] at every β . In the language of finite-volume couplings: the LWW coupling $L m_L$ has no infrared fixed point. This is exactly the remaining physical content of the Polyakov conjecture, and nothing else is needed: *not* a decrement uniform in β (false, E14), *not* a rate in Δ , *not* monotonicity of Δ (false in general, Section 8.1).

Theorem 4 (conditional). *Let $\beta \geq 0$. If $\liminf_L \Delta(L, L; \beta) < 0.2239$ — or, weaker, $\Delta(n_0, w_0; \beta) < 0.2239$ for a single torus, or a seed pair in $\mathcal{A}$ — then (i) for every $L \geq n_0$ the circle- L transfer operator has a simple top eigenvalue and gap $\gamma_L \geq \gamma^*(\text{seed}) > 0$; (ii) $|\langle s_0 \cdot s_x \rangle| \leq e^{-\gamma^* \|x\|_\infty}$ for every weak limit of square tori, and the cylinder analogue holds (Lemma C); (iii) $\Delta(L) \rightarrow 0$, $c(L) \rightarrow 0$ and $L(e_0(L) - fL) \rightarrow 0$.*

Proof. Lemma P, Lemma B with the thresholds of Lemma 6, Lemma C, Theorems 2 and 3. $\square$ **With C1'' proved, Theorem 4 gives the lattice O(3) mass gap at every β ; without it the rigorous content of this paper is Theorem 1 and the fixed- β decay for $\beta \leq 0.7701$.**

What an unconditional proof must supply. For every β , one of the equivalent inputs: (1) a single scale with $4\ell(R(L, L)) + 2\ell(R(2L, L)) < 0.2239$ (vacuum dominance of the rings L and $2L$ after time L); or (2) exclusion of the branch $\inf_L \Delta \geq 0.2239$; or (3) local boundedness of $L_\eta^*(\beta)$. None of reflection positivity, chessboard estimates, Ginibre or Griffiths inequalities can do this on their own: they are satisfied by Ising and O(2), where the conclusion is false (Sections 8.1, 8.3). The input must use the non-abelian bracket quantitatively, as already found for the Θ -form (Section 10.2); the weak-coupling end (Section 9) shows where the uniform drop lives and what obstructs the fixed-size methods. Routes: (a) a rigorous renormalization step in β_{eff} , $\beta_{\text{eff}}(2L) \leq \beta_{\text{eff}}(L) - \kappa$ with $\kappa > 0$ uniform, from $\log L \sim \beta$ down to $\beta_{\text{eff}} \sim 1$; (b) a direct exclusion of a massless phase at every β ; (c) an independent referee of the 215-c estimate **prelim-weighted-current**, which with Lemma T gives an averaged β -uniform step conditional on a verified rather than unrefereed result.

Table 9: Doubling defect $\Delta(L, L)$ (`liminf_results.csv`). O(3), O(2) and λ rows are Monte Carlo, $\pm 1\sigma$, [provisional]; Ising rows are exact (Kaufman). $\Delta^* = 0.2239$. At $\beta = 1.6$, $L = 128$ the error ± 0.40 makes the point uninformative. [†]Combined value of the pre-registered run 0.197(99) and three declared extra seeds.

model	coupling	$L = 4$	8	16	32	64	128
O(3)	$\beta = 1.4$ ($\xi \approx 6.9$)	4.641(21)	3.419(38)	1.659(51)	0.229(36) [†]	-0.05(16)	—
O(3)	$\beta = 1.6$ ($\xi \approx 19$)	5.446(26)	4.717(38)	3.810(65)	2.574(137)	0.72(23)	-0.05(40)
O(2)	$\beta = 1.3$	4.419(22)	4.304(32)	4.324(47)	4.419(101)	4.33(20)	—
$\lambda = 0.5$	$\beta = 3.0$	5.770(26)	5.729(39)	5.707(55)	5.844(124)	—	—
Ising exact	$K = 0.3$ ($\xi = 1.58$)	0.533	0.054	4.4×10^{-4}	2.4×10^{-8}	5.3×10^{-17}	1.9×10^{-34}
Ising exact	$K = 0.4$ ($\xi = 5.96$)	1.645	1.169	0.437	0.039	2.4×10^{-4}	7.2×10^{-9}
Ising exact	K_c	1.946	1.925	1.921	1.920	1.920	1.920
Ising exact	$K = 0.5$	2.077	2.082	2.080	2.079	2.079	2.079

8.5 Monte Carlo of the defect at $\beta = 1.4$ and 1.6 ([provisional])

[pre-registered] (`PREREG_liminf.md`). *Estimator (new)*. The $2L \times 2L$ torus and the four $L \times L$ tori have the same sites and differ only in $8L$ seam bonds; with weights $t\beta$ on the new seam bonds and $(1-t)\beta$ on the old wraps, $\Delta = -\beta \int_0^1 \langle \sum_{\text{new}} s \cdot s' - \sum_{\text{old}} s \cdot s' \rangle_t dt$, evaluated by 24-node Gauss–Legendre quadrature in u with $t = u^3$, Wolff reflection clusters with a fixed cluster count per measurement and 20 bins per node. The variance grows like L , not like L^4 as in β -integration. A state-dependent stopping rule (“clusters until flipped size $\geq N$ ”) biased the sampled distribution by up to 20σ and was removed before production (the same pitfall as item 1 of Section 3.3). Validation on exact Ising through the same code path: $K = 0.4407, 0.5, 0.4$ at $L = 8$ give $z = -0.42, +0.51, -0.24$; $K = 0.4$ at $L = 16$ gives $z = -1.64$.

Readings (Table 9). O(3) drives Δ to zero at both couplings, decreasing at every tested L : at $\beta = 1.4$ the seed sits at $L \approx 32$ –64 (4.6 – 9ξ ; $\Delta(32) = 0.229(36)$ is on the threshold); at $\beta = 1.6$, $L = 64$ fails as a seed ($0.72(23)$) and $L = 128$ is inconclusive. Measured seed scales: Ising $K = 0.3 \rightarrow L = 6$ (3.8ξ), $K = 0.4 \rightarrow L = 24$ (4.0ξ), O(3) $\beta = 1.4 \rightarrow L \approx 32$ –64; the necessity bounds of Proposition 1(d) are 2.56ξ (Ising) and 3.66ξ (O(3)). The controls plateau far above Δ^* , at the universal constants of Section 8.3. *Pre-registration scorecard*: P-O2-1,2,3, P-L-1, P-O3-1,2,3 pass; **P-O3-4 fails on magnitude**: the drop of $Y = e^{\Delta/3}$ per doubling is 1.32, 1.26, 1.20 ($\beta = 1.6$) and 1.57 ($\beta = 1.4$, $4 \rightarrow 8$), outside the registered $[0.3, 1.2]$ and about $2\times$ one loop — the same $2\times$ excess as in ρ itself ($0.947 \rightarrow 0.742 \rightarrow 0.548 \rightarrow 0.340$) and in the rotor-ring drops of Section 10.2; its qualitative part (nearly constant Y -drops while the Δ -drops grow, 0.73, 0.91, 1.24, 1.85) passes. *Deviations*: correlation lengths $\xi(1.4) \approx 6.9$, $\xi(1.6) \approx 19$ are literature values; three extra seeds at $\beta = 1.4$, $L = 32$ were run beyond the stopping rule for seed-verdict precision (no verdict changed); Monte Carlo errors grow like $L^{1.3}$. Smallest next discriminating test (not run): $\beta = 1.4$ at $L = 32, 48$ with $4\times$ statistics and $\beta = 1.5$ ($\xi \approx 11$) at $L = 32$ –64, to test whether L^*/ξ stays in $[3.7, 6]$ as β grows.

9 The weak-coupling end: Laplace bounds, rotor semiclassics and the running-coupling form

Records: `laplace_betainf.md` (Theorem A, exact expansion, controls, Monte Carlo; code `laplace_betainf.py` with the embedded sampler `o3mc.c`; data `laplace_betainf_results.csv`; pre-registered predictions `preregistered_predictions.csv`), `rotor_semiclassic.md` (ring semiclassics; code `rotor_semiclassic.py`, `run_scan.py`, `analyze_semiclassic.py`, `fb_check_N2.py`; data `rotor_semiclassic_results.csv`, `rotor_semiclassic_fits.csv`, `fb_check_N2.csv`; pre-registration `prereg_semiclassic.md`, 2026-10-20T05:46:15Z).

Both tracks attack the $\beta \rightarrow +\infty$ end of the doubling defect, where everything is computable; they agree on the constant and on the obstruction.

9.1 Exact large- β expansion of the defect and Theorem A

Since $|s| = 1$, $\sum_{\langle xy \rangle} s_x \cdot s_y = E - \frac{1}{2} \sum_e |s_x - s_y|^2$ exactly ($E = 2N$ bonds, $N = L^2$ sites), and rotation invariance fixes $s_0 = e_3$: $Z(L, T) = (4\pi)^{-(N-1)} e^{2N\beta} J(\beta)$ with $J = \int \exp(-\frac{\beta}{2} \sum_e |s_x - s_y|^2) \prod_{x \neq 0} d\Omega_x$. After this gauge fixing the maximizer is unique and non-degenerate, the collective S^2 has been integrated exactly (no Faddeev–Popov Jacobian remains). Laplace expansion in the chart $s_x = (\pi_x, \sigma_x)$ with Hessian $\beta \Delta_0 \otimes I_2$ (Δ_0 the torus Laplacian with the row and column of site 0 removed, $\det \Delta_0 = \tau(L, T)$ spanning trees) gives

$$\log Z = 2N\beta - (N-1) \log(2\beta) - \log \tau + \frac{B}{\beta} + O(\beta^{-2}),$$

$$B = \sum_x G_{xx} - \frac{1}{8} \sum_e [6W_e^2 + 2U_e V_e],$$

with $G = \Delta_0^{-1}$, $W = G_{xx} - G_{yy}$, $U = G_{xx} + G_{yy} - 2G_{xy}$, $V = G_{xx} + G_{yy} + 2G_{xy}$ (general $O(m+1)$: $B = \frac{m}{2} \sum G_{xx} - \frac{1}{8} \sum_e [m(m+1)W^2 + mUV]$). Exact checks: on the cycles C_3, C_4, C_6 against Bessel sums for $O(2)$ and $O(3)$ ($\beta(\text{exact} - \text{order } 0) \rightarrow B$, e.g. $O(3)$ C_6 : $0.6252 \rightarrow 0.5590$ against $B = 0.5556$); $O(2)$ on tori reproduces the angle-variable result $B_2 = (N-1)^2/(16N)$ to machine precision. Structure: $B_3(L) = N/8 - (1/2\pi) \log L + \text{const} + o(1)$ (slope per $\ln 2$ at $L = 512$: -0.159163 against $-1/(2\pi) = -0.159155$); the extensive $N \log L$ parts of the measure term and the quartic vertex cancel, and the $-(1/2\pi) \log L$ left over is the one-loop running $\beta_{\text{eff}}(L) = \beta - (1/2\pi) \log L$ entering through the zero-mode factor.

The defect and its decrement at large β (exact algebra). $\Delta(L, L) = 3 \log(2\beta) + \delta_0(L) + b(L)/\beta + O(\beta^{-2})$: the energy $8L^2\beta - 8L^2\beta = 0$ cancels exactly, the β -power *does not* ($4(L^2 - 1) - (4L^2 - 1) = -3$: one zero mode per torus, $4 - 1 = 3$; for $O(n)$ the coefficient is $3(n-1)/2$, i.e. 3 for $O(3)$ and $\frac{3}{2}$ for $O(2)$ and the λ -model), $\delta_0(L) = \log \tau_{2L} - 4 \log \tau_L = 3.701, 3.457, 3.320 \rightarrow 3.164$ at $L = 2, 3, 4, \infty$, and $b(L) = 4B(L) - B(2L) = -0.990, -1.254, -1.414 \approx -(3/2\pi) \log L + \text{const}$. Decrement $D(L, \beta) = \Delta(L) - \Delta(2L) = \gamma(L) + d(L)/\beta + O(\beta^{-2})$ with

$$\gamma(L) = -4 \log \tau_L + 5 \log \tau_{2L} - \log \tau_{4L}, \quad d(L) = 4B(L) - 5B(2L) + B(4L);$$

$\gamma(L) = 0.3817, 0.2338, 0.1241, 0.0239, 5.6 \times 10^{-3}, 1.38 \times 10^{-3}, 8.6 \times 10^{-5}$ at $L = 2, 3, 4, 8, 16, 32, 128$ ($L^2\gamma \rightarrow 1.411$; positive at every computed L), and $d(L) = 0.4247, 0.3723, 0.3556, 0.3376, 0.3326, 0.33138, 0.33098$, tending to $3 \ln 2/(2\pi) = 0.330953$ (positive at every L). In D every term proportional to N cancels at every order ($4 - 20 + 16 = 0$) and every L -independent constant cancels ($4 - 5 + 1 = 0$); what survives is sub-extensive and L -dependent. γ is the free Gaussian field's contribution, shared with $O(2)$ up to the factor $(n-1)/2$, and carries no information about asymptotic freedom; the discriminating information is $d(\infty)$: the factor 3 is the number of zero modes lost in doubling, and $d(\infty) = \frac{3}{2}c_1$ with $c_1 = \ln 2/\pi$ the Θ -form constant of Section 10 (the rotor zero-mode factor $\approx 2\Theta_L/L$ with $\Theta_L/L = \beta_{\text{eff}}(L)$ gives $D - \gamma \approx 3[\ln \beta_{\text{eff}}(L) - \ln \beta_{\text{eff}}(2L)] \approx 3 \ln 2/(2\pi\beta)$ and $2\Theta_L - \Theta_{2L} = (\ln 2/\pi)L$; a consistency dictionary, not a derivation on the ring).

Theorem 5 (A: Laplace bounds with explicit remainder, fixed L). *Let $G = \Delta_0^{-1}$ with eigenvalues μ_i , $R_x = G_{xx}$ (effective resistance to 0), $S = \sum_e [6W_e^2 + 2U_e V_e] \geq 0$, $J_G = (2\pi/\beta)^{N-1}/\tau$. For $c \in (0, 1)$ put $r^2 = 1 - c^2$, $a = -\ln c/(1 - c^2)$, $P_A = 1 - \sum_x \exp(-\beta r^2/(2R_x))$, $\varepsilon_- = -\ln P_A + S/(8c^2\beta P_A)$ and $\varepsilon_+ = \ln [\prod_i (1 - 2a\mu_i/\beta)^{-1} + \tau 2^{N-1} \sum_x \exp(-\beta(1-c)/(2R_x))]$. Then for every $c \in (0, 1)$ and every $\beta > 2a\mu_{\max}$ with $P_A > 0$,*

$$-\varepsilon_-(\beta, c) \leq \ln J(\beta) - \ln J_G \leq \varepsilon_+(\beta, c), \quad \varepsilon_{\pm} = O(1/\beta).$$

Proof sketch. Split $J = J_U + J_{U^c}$ with $U = \{s_x \cdot e_3 \geq c \forall x\}$, covered by the chart $|\pi_x| \leq r$. Upper bound on J_U : drop $-\frac{\beta}{2} \sum (\sigma_x - \sigma_y)^2 \leq 0$, use the chord bound $-\ln \sigma \leq a|\pi|^2$ on $[0, r^2]$ (convexity of $-\frac{1}{2} \ln(1-p)$) and extend to $\mathbb{R}^{2(N-1)}$: $J_U \leq J_G \det(I - (2a/\beta)G)^{-1}$. Upper bound on

J_U : if $s_x \cdot e_3 < c$ then $|s_x - s_0|^2 > 2(1 - c)$; the Dirichlet principle gives $\sum_e |\nabla s|^2 \geq |s_x - s_0|^2 / R_x$; split the exponent in half and bound $J(\beta/2)$ by a spanning-tree integration ($\int d\Omega e^{-(\beta'/2)|s-s'|^2} \leq 2\pi/\beta'$ per leaf), giving $J(\beta/2) \leq \tau 2^{N-1} J_G$; union bound over x . Lower bound: $J \geq J_U$, $1/\sigma \geq 1$, $\sigma_x + \sigma_y \geq 2c$ gives $-\frac{\beta}{2} \sum (\sigma_x - \sigma_y)^2 \geq -\frac{\beta}{8c^2} Q$ with $Q = \sum_e (|\pi_x|^2 - |\pi_y|^2)^2$; conditional Jensen on the event $A = \{|\pi_x| \leq r \ \forall x\}$, Wick ($E_G Q = S/\beta^2$) and $|\pi_x|^2 \sim (R_x/\beta)\chi_2^2$ with a union bound give $P(A) \geq P_A$. $\square$

Corollary A (fixed- L decrement at large β ; rigorous). $\gamma(L) - [4\varepsilon_-(L) + 5\varepsilon_+(2L) + \varepsilon_-(4L)] \leq D(L, \beta) \leq \gamma(L) + [4\varepsilon_+(L) + 5\varepsilon_-(2L) + \varepsilon_-(4L)]$, so $D(L, \beta) \rightarrow \gamma(L)$ as $\beta \rightarrow \infty$, and $D(L, \beta) \geq \gamma(L)/2 > 0$ for $\beta \geq \beta_0(L)$ with $\beta_0 = 5.0 \times 10^2, 2.2 \times 10^3, 8.1 \times 10^3, 2.4 \times 10^4, 5.7 \times 10^4, 2.1 \times 10^5$ at $L = 2, 3, 4, 5, 6, 8$ ($\beta_0 \sim L^4 \log L$: remainder $\propto \text{Tr } G \sim (N/2\pi) \log L$, margin $\gamma \sim 1.41/L^2$). Rigour status: the inequalities are proved; the constants are evaluated in double precision (dense eigendecompositions, $N \leq 1024$) without interval arithmetic, with a factor-2 margin in γ ; $\gamma(L)$ is an exact finite sum of logarithms. Not done: a rigorous $O(1/\beta^2)$ remainder, i.e. a theorem containing $d(L)/\beta$; $B(L)$ is verified numerically (exact cycles, Monte Carlo), not with rigorous error bounds. **Consequence (E14):** since $D(L, \infty) = \gamma(L) \rightarrow 0$, no β -uniform additive decrement of Δ exists.

Running-coupling form and the obstruction. Define $\hat{\beta}(L) := \frac{1}{2} \exp[(\Delta(L, L) - \delta_0(L))/3] \approx \beta_{\text{eff}}(L)$ up to $O(1/\beta)$. The $\beta \rightarrow \infty$ end gives $\hat{\beta}(L) - \hat{\beta}(2L) \rightarrow \ln 2/(2\pi)$ per doubling, uniform in β : this is the Θ -form statement $\Theta_{2L} \leq 2\Theta_L - c_1 L$, and the large- β end supports it while refuting the additive form. Obstructions to uniformity in L : (O1, genuine) the large- β series at fixed L is an expansion in $(\log L)/(2\pi\beta)$ — the sub-extensive part of the k -th coefficient grows like $(\log L)^k$, so the expansion is one in $1/\beta_{\text{eff}}(L)$ and cannot be uniform beyond $L \sim e^{2\pi\beta} \sim \xi(\beta)$, exactly the scale where $C1''$ must deliver $\Delta < \Delta^*$; (O2, technical) the rigorous remainder of Theorem 5 is $\propto \text{Tr } G \sim (N/2\pi) \log L$, not $\propto N$ with an L -independent coefficient, so it does not cancel by extensivity — the $N \log L$ pieces cancel in the true B only through the $O(3)$ rotation Ward identity (measure term against quartic vertex), which term-by-term Gaussian bounds do not see; a local-frame or multiscale (Balaban-type) treatment would be needed; (O3) the free-field margin $\gamma(L) \sim 1/L^2$ is itself non-uniform, so uniform positivity must come from d/β . The large- β end controls $C1''$ only in the window $\log L \ll \beta$; the remaining window is the renormalization-group regime.

Controls (exact Laplace algebra). $O(2)$: $\log Z = 2N\beta - \frac{N-1}{2} \log(2\pi\beta) - \frac{1}{2} \log \tau + B_2/\beta$ with $B_2 = (N-1)^2/(16N)$, so $D_2 = \gamma/2 + d_2/\beta$ with $d_2 = 45/(256L^2) \rightarrow 0$: no $\ln 2/\pi$ term, as required; the $O(1)$ term $\gamma/2$ is the same free field. λ -model ($\lambda \in (-1, 1)$): the maximizers form the equatorial S^1 ; θ is massless with one zero mode, z is massive with quadratic form $\frac{1}{2} z^T [\lambda \Delta + 4(1-\lambda)] z$; the β -power in Δ is $\frac{3}{2} \log \beta$ (losing the full S^2 zero mode halves the coefficient); $D_\lambda = \gamma/2 + \gamma_z(L) + d_\lambda(L)/\beta$ with γ_z exponentially small for $L \gg \xi_z = (|\lambda|/4(1-\lambda))^{1/2}$ (1.3×10^{-5} at $L = 8$, $\lambda = 0.5$) and $d_\lambda(L) \rightarrow 0$ (equal to the $O(2)$ value for $L \geq 32$); near $\lambda \rightarrow 1$ the z -mode mimics the missing zero modes for $L \lesssim \xi_z$ ($\gamma_z = 2.17, 1.78, 1.07, 0.28$ at $L = 2, 4, 8, 16$ for $\lambda = 0.99$), a finite total drop and no running; the order of limits $\beta \rightarrow \infty$ versus $\lambda \rightarrow 1$ matters (the Laplace formula needs $\beta(1-\lambda) \gg 1$). Both controls give $d(\infty) = 0$; only $O(3)$ has $d(\infty) = 3 \ln 2/(2\pi) > 0$.

Monte Carlo test of the expansion (pre-registered; pass). Heat-bath plus over-relaxation sampler validated on the exact 2×2 torus; $\Delta(L, L) = 8L^2 \int_0^\beta (e_L - e_{2L}) d\beta'$ by thermodynamic integration (6 segments $\times$ 5 Gauss-Legendre nodes), $d\Delta/d\beta$ measured directly, $L = 2, 3, 4$, $\beta \leq 8$, 420 runs. Predictions frozen before any run: (i) the slope of $\Delta_{\text{MC}} - \delta_0 - b/\beta$ against $\log \beta$ on $\beta \in \{4, 6, 8\}$ is 3.023, 3.042, 3.052 for $L = 2, 3, 4$, inside 3 ± 0.15 ; the $O(2)$ -like value $\frac{3}{2}$ is excluded;

(ii) $\beta^2 \times \text{residual}$ is bounded ($|\cdot| \leq 0.93$ at $\beta = 6, 8$) and $\beta^3 \times (\text{derivative residual})$ settles to a constant (0.64/1.03/1.37 at $\beta = 8$), i.e. $O(\beta^{-3})$; the β^{-2} coefficient grows with L , as O1 predicts; (iii) $D(2, \beta) = 0.5050(4), 0.4544(4), 0.4321(4)$ at $\beta = 4, 6, 8$ against $\gamma(2) + d(2)/\beta = 0.4879, 0.4525, 0.4348$, approaching $\gamma(2) = 0.3817$ from above as Theorem 5 requires. Declared deviation: the 6×6 torus (used for $L = 3$) ran at half the planned statistics because of machine load (errors $\approx \sqrt{2}$ larger).

9.2 Rotor ring at weak coupling

Conventions as in Section 10.1: $H_N = \frac{g}{2} \sum_i L_i^2 - \sum_i n_i \cdot n_{i+1}$ periodic, $g = 1/\beta^2$, $\Theta = \beta^2$, $a_\tau = \beta$, $\Theta_N = 1/(E_1 - E_0)$, softening $\delta_N = N\Theta - \Theta_N > 0$ (Lemma R), deficit $D_N = (2\Theta_N - \Theta_{2N})/(N\beta) = (\delta_{2N} - 2\delta_N)/(N\beta) = 2[\rho_H(N) - \rho_H(2N)]$, one-loop value $\ln 2/\pi$.

Semiclassical expansion (derived; confirmed by exact diagonalization). The minima of V form the diagonal S^2 (Morse–Bott); with $n_i = \sqrt{1 - |\pi_i|^2} \hat{n} + \pi_i$, $\hat{n} = m/|m|$, $m = \sum n_i$, one has $\sum_i \pi_i = 0$ exactly, $N - 1$ relative modes $\omega_k = 2\sqrt{g} \sin(\pi k/N)$ with two polarizations, and the zero mode $\hat{n} \in S^2$ as a rigid rotor of moment $N\Theta$; the Gaussian ground state has $\sum_i \langle (e \cdot \pi_i)^2 \rangle = \sqrt{g} S_N/4$ for any unit $e \perp \hat{n}$, with $S_N = \sum_{j=1}^{N-1} \csc(\pi j/N)$. The inertia of a bent configuration about e is $I_{ee} = N\Theta - \Theta \sum_i (e \cdot \pi_i)^2$ (even in π : the harmonic $(\partial I)^2$, cubic-anharmonic and Coriolis contributions of the vibration–rotation expansion vanish), so

$$\Theta_N = \frac{N}{g} - \frac{\beta}{4} S_N + O(1), \quad D_N \rightarrow D_N^{(0)} := \frac{S_{2N} - 2S_N}{4N} \quad (\beta \rightarrow \infty),$$

with $D_N^{(0)} = 0.2500, 0.2286, 0.2242, 0.2227, 0.2219$ for $N = 1, \dots, 5$, tending to $\ln 2/\pi = 0.2206$. *Lemma (alternating sum; proved).* With $a_j = \csc(\pi j/2N)$ decreasing on $1 \leq j \leq N$, $S_{2N} - 2S_N = \sum_{j=1}^{2N-1} (-1)^{j+1} a_j \geq 2(a_1 - a_2)$, hence $D_N^{(0)} \geq (2 \cos(\pi/2N) - 1)/(2\pi) \geq (\sqrt{2} - 1)/(2\pi) = 0.066$ for all $N \geq 2$, and $D_N^{(0)} \rightarrow \ln 2/\pi$ (the $1/x$ part of $\csc x$ gives the alternating harmonic series). Numerically $N^2(D_N^{(0)} - \ln 2/\pi) = 0.0326781, 0.0327245, 0.0327249$ at $N = 10, 100, 1000$, consistent with $\pi/96 = 0.0327249$ (identification, not proved): $D_N^{(0)}$ never drops below $\ln 2/\pi$. *Where the commutator enters.* The whole softening comes from the *unbroken* generator: rotation about $\hat{n}$ acquires a lever arm $|\pi_i|$ on every tilted spin because the stabilizer $n_i \cdot L_i = c_i \hat{n} \cdot L_i + \pi_i \cdot L_i$ of a tilted spin is the stabilizer of $\hat{n}$ conjugated by the broken generators, and $e^{-i\theta \cdot L^\perp} (\hat{n} \cdot L) e^{i\theta \cdot L^\perp} = \hat{n} \cdot L - i[\theta \cdot L^\perp, \hat{n} \cdot L] + \dots$ is non-trivial exactly because $[L^a, L^b] = i\varepsilon_{abc} L^c$; in $O(2)$ there is no unbroken generator and the same algebra gives zero. *Exact diagonalization* (kinetic cutoff $\sum \ell_i(\ell_i + 1) \leq \Lambda$, dihedral-trivial sector, Rayleigh–Ritz monotone in Λ ; Λ raised until the relative change of Θ was $< 2 \times 10^{-8}$ or the size cap of $1.2\text{--}2.6 \times 10^6$ states; validated against unsymmetrized ED to 4×10^{-15} and a Legendre-basis solver to 5×10^{-13}): $A_N := \sqrt{g}(N/g - \Theta_N)$ extrapolated to $g \rightarrow 0$ from the three smallest converged g gives $0.25001/0.57747/0.96128$ for $N = 2/3/4$ against $S_N/4 = 0.25000/0.57735/0.95711$; the rival $S_N/2$ (plain Feynman–Bijl with $A = \sum n^z$) and “no softening” are excluded by wide margins; $D_2(g \rightarrow 0) = 0.2298$ against 0.2286 , approached from above (slope $+0.065$ in $\sqrt{g}$). **[provisional]**: $N \leq 4$ fully converged ($N = 5$ only at $g \geq 0.2$, $N = 6$ at $g = 0.3$); the pre-registered fit window $g \in \{0.05, \dots, 0.3\}$ was too coarse (P1 fails its 3% rule there: $+3.1/+3.9/+4.4\%$; P4 fails its 5% rule) and the small- g points $g = 0.002\text{--}0.03$ and three-point local fits were added post hoc; with them P1 passes ($+0.01/+0.02/+0.44\%$).

Rigorous half, conditional half. *Rigorous (N side).* The Feynman–Bijl inequality with the order-parameter trial $A = \hat{n}^z$ gives, for every $N \geq 3$ and g (exact operator algebra; at $N = 2$ with a regularized $h(|m|)$),

$$\frac{1}{\Theta_N} = E_1 - E_0 \leq \frac{g}{2} \left\langle \sum_i \frac{1 + (\hat{n} \cdot n_i)^2}{|\sum_j n_j|^2} \right\rangle_0,$$

where the “1” is the unbroken-generator lever arm and the “ c_i^2 ” the broken generators; expanding, $E_1 - E_0 \leq (g/N)[1 + \sqrt{g}S_N/(4N)] + O(g^2)$, the same law as above, and the bound is asymptotically sharp (at $N = 2$ it exceeds the exact gap by $4.8 \times 10^{-5}g^2$ at $g = 0.002$, an excess vanishing faster than g^2). This is the direction needed on the N side. The plain choice $A = \sum n^z$ loses a factor 2 in the softening ($\Theta_N \geq N/g - (\beta/2)S_N$). *Conditional (2N side)*. A lower bound on the $2N$ gap is reduced to Temple’s inequality with $\varphi = A\psi_0$: $E_1 \geq \langle H \rangle_\varphi - \text{Var}_\varphi(H)/(\mu - \langle H \rangle_\varphi)$; rotation invariance of ψ_0 cancels the leading part of $[H, A]\psi_0$, leaving $\text{Var}_\varphi = O(g^{5/2})$, so with $\mu - E_0 \geq c\sqrt{g}$ the Temple loss is $O(g^2)$, below the $O(g^{3/2})$ signal; the $\ell \geq 2$ states do not interfere ($E_\ell - E_0 \geq \ell(\ell + 1)g/(2N)$ by Lemma R(i), rigorous). This removes the circularity noted in Section 10.2: μ only needs to be crude. Two lemmas are assumed, not proved: **G1** (harmonic ground-state moments for a Morse–Bott well: $\langle \sum |\pi_i|^2 \rangle = \sqrt{g}S_N/2 + O(g)$, quartic moments $O(g)$, residual $O(g^{5/2})$ — of the type of Simon [39] and Helffer–Sjöstrand [40], located but not read in full; their applicability is an assessment) and **G2** (inside the $\ell = 1$ sector every eigenvalue except the lowest is $\geq E_0 + c_N\sqrt{g}$; expected $c_N\sqrt{g} \approx 2-4\sqrt{g}\sin(\pi/N)$). *Theorem W (fixed N ; conditional on G1, G2 for rings of N and $2N$ sites)*. $E_1 - E_0 = (g/N)[1 + \sqrt{g}S_N/(4N)] + O(g^2)$, hence $\Theta_N = N/g - (\beta/4)S_N + O(1)$ and $D_N(g) = D_N^{(0)} + O(\sqrt{g})$; for every $c < D_N^{(0)}$ (in particular $c = \ln 2/\pi - \epsilon$ or $c = (2\cos(\pi/2N) - 1)/(2\pi)$) there is $g_0(N, c) > 0$ with $\Theta_{2N} \leq 2\Theta_N - cNa_\tau$ for all $g \leq g_0$. The Feynman–Bijl inequality, the alternating-sum lemma, the $\ell \geq 2$ exclusion and the Temple step are proved; G1, G2 are open; $g_0(N)$ is not explicit.

Controls (same code path). O(2): the only generator is broken and there is no stabilizer, $I = N\Theta$ identically, Lemma R(iii) exact; ED over $N = 2-6$, $g \leq 0.3$: $|A| \leq 2.8 \times 10^{-4}$, $|D| \leq 1.9 \times 10^{-4}$ ($\mathbb{Z}_N$ tunnelling). $\lambda = 0.5$ (easy-plane, U(1) about z): the softening mode is the massive out-of-plane field, $\omega_k^2 = 2g(1 - \lambda \cos k)$ including $k = 0$, δ_N is extensive and the doubling compares two trapezoid rules of an analytic periodic function, so the leading deficit is *negative and exponentially small*: $D_N^{\lambda, (0)} = -0.211, -0.041, -0.009, -0.002$ for $N = 1-4$ (no logarithm); ED: $A_2^\lambda(0) = 0.78883$ against 0.78868 , $D_2^\lambda < 0$ at every g (-0.062 to -0.155). Lemma R(iv) noted that the λ -model also has $\Theta_z < N\Theta$; the expansion shows why that does not help.

Obstruction to N -uniformity (precise form). The leading constant is N -uniform ($D_N^{(0)} \in [\ln 2/\pi, \frac{1}{4}]$); the expansion parameter is not: the anharmonic corrections are controlled by $\langle \pi_i^2 \rangle = \sqrt{g}S_N/(2N) \approx (\sqrt{g}/\pi)[\ln(2N/\pi) + \gamma_E] = (\ln N)/(\pi\beta) + O(1/\beta)$, the Mermin–Wagner growth of spin-wave fluctuations, so the k -th coefficient carries up to $(\ln N)^{k-1}$ — the fixed- N re-expansion of the running coupling $1/\rho(N) = 1/(\beta - \ln N/2\pi)$. Measured: the $\sqrt{g}$ -slope a_N of A_N is $0.093, 0.209, 0.296$ ($a_N/N = 0.047, 0.070, 0.074$, increasing; P5 pass, weak). The fixed- N result therefore covers $N \leq e^{2\pi\theta\beta}$ only with $\theta \ll 1$; in the rigorous Temple route the protecting gap $4\sqrt{g}\sin(\pi/N)$ closes like $1/N$, so any explicit $g_0(N)$ is at least polynomially small in N . This is the same obstruction as O1 above, seen from the ring: the window $\ln N \approx 2\pi\theta\beta$, $\theta \in (\theta_0, 1)$, is the renormalization-group regime, and no fixed-size argument reaches it.

What the two ends give together. Both pictures place the one-loop constant exactly: $d(\infty) = 3\ln 2/(2\pi)$ in the torus defect at order $1/\beta$, $D_N^{(0)} \rightarrow \ln 2/\pi$ in the ring inertia, related by the zero-mode count ($d(\infty) = \frac{3}{2}c_1$). Both are zero for O(2) and for the λ -model. Both are rigorous at fixed size (Theorem 5; the Feynman–Bijl half) and both stop at the same place: the expansion parameter $(\log L)/(2\pi\beta)$ must be carried up to $L \sim \xi \sim e^{2\pi\beta}$, where C1’’ needs $\Delta < \Delta^*$. The uniform statement lives in β_{eff} , not in Δ , which is the content of erratum E14.

10 Character-basis analysis of the doubling step

[exploratory] (record `character_contraction.md`, code `character_contraction.py`, table `character_contraction_check.csv`; all computed from `casimir_map_results.csv` and from exact diagonalization of a two-rotor model). This section asks where, in the exact sector structure of Section 5, the doubling $L \rightarrow 2L$ can produce the stiffness decrease of Eq. (7), and where positivity alone stops. Every statement is labelled proved, numerically supported or conjectured, and every lemma is run through the $O(2)$ and λ -deformed controls. It supersedes an earlier, unpublished “Lyapunov / sector-weight” draft: the singlet log-odds $\Lambda(L) = \ln[p_0/(1-p_0)]$ appear below only through the rigorous bracket (2) and through provisional increments.

Set-up. $\mathcal{T}_L = D_L^{1/2} K^{\otimes L} D_L^{1/2}$ on $L^2(S^2)^{\otimes L}$ (positive, trace class, strictly positive kernel, commuting with $U(g)$); $n_\ell(L, T) = \text{Tr } \mathcal{T}_{L,\ell}^T = \sum_k \lambda_{\ell,k}^T$, $p_\ell = (2\ell+1)n_\ell/Z$, $E_{\ell,k} = -\ln \lambda_{\ell,k}$, $\Delta_\ell(L) = E_{\ell,0} - E_{0,0} > 0$ (Perron–Frobenius). Two stiffnesses are kept distinct: the helicity stiffness $\rho_{\text{hel}}(L, T) = (T/L)\langle C_2 \rangle / (N-1)$ (on the square torus $\rho_{\text{hel}} = \langle \ell(\ell+1) \rangle / 2$, the quantity measured in Section 4.6) and the rotor inertia $\Theta_L = 1/\Delta_1(L)$ with $\rho_H(L) = \Theta_L/L$.

Lemma 0 (ideal-rotor helicity map; proved, elementary). If $n_\ell(L, T) = e^{-TE_0} e^{-T\ell(\ell+1)/(2\Theta)}$ exactly, then $\rho_{\text{hel}}(L, T) = (T/L)F_3(\Theta/T)$ with $F_3(x) = \sum_\ell (2\ell+1)\ell(\ell+1)e^{-\ell(\ell+1)/2x} / (2\sum_\ell (2\ell+1)e^{-\ell(\ell+1)/2x})$, and Mulholland’s expansion gives

$$F_3(x) = x - \frac{1}{6} - \frac{1}{180x} + O(x^{-2}), \quad F_3(1) = 0.8263, \quad F_3(8) - 8 = -0.1674; \quad (11)$$

for $O(2)$ the same construction gives $F_2(x) = x$ up to $O(e^{-2\pi^2 x})$ ($F_2(1) = 1.0000$). On the square torus an ideal $O(3)$ rotor therefore has $\rho_{\text{hel}} = \rho_H - \frac{1}{6} + O(1/\Theta)$: helicity stiffness and inertia differ by a constant curvature $(N-2)$ offset, so their per-doubling drops agree to $O(1/\Theta)$ while their values do not. This is the correct replacement for the identification $\bar{g}^2 = 1/\rho_s$ withdrawn in erratum E1 (the referee’s counterexample $\rho_{\text{hel}} = 0.826333$ at $mL = 1$ is $F_3(1)$; the direct series evaluation in Section 4.6 gives 0.826333426 and the numerical coefficient -0.0056 there agrees with $-1/180 = -0.00556$).

Lemma A (time doubling is exact and N -blind; proved). With $r_\ell = p_\ell/p_0$ and $\varepsilon_\ell(L, T) = \sum_{k \geq 1} (\lambda_{\ell,k}/\lambda_{\ell,0})^T \geq 0$ (decreasing in T),

$$r_\ell(L, 2T) = e^{-T\Delta_\ell(L)} r_\ell(L, T) \frac{(1 + \varepsilon_\ell(2T))(1 + \varepsilon_0(T))}{(1 + \varepsilon_\ell(T))(1 + \varepsilon_0(2T))},$$

$$\frac{e^{-T\Delta_\ell} r_\ell(L, T)}{1 + \varepsilon_\ell(L, T)} \leq r_\ell(L, 2T) \leq e^{-T\Delta_\ell} (1 + \varepsilon_0(L, T)) r_\ell(L, T).$$

Proof. $n_\ell(T) = e^{-TE_\ell}(1 + \varepsilon_\ell(T))$; substitute into $r_\ell = (2\ell+1)n_\ell/n_0$ and use $\varepsilon_\ell(2T) \leq \varepsilon_\ell(T)$, $\varepsilon_0 \geq 0$. $\square$ Time doubling multiplies the $\ell:0$ odds by the sector Boltzmann factor; on the square torus this is $e^{-\ell(\ell+1)/(2\rho_H)}$, an $O(1)$ contraction exactly when $\rho_H = O(1)$. It is monotone, exact and identical for $O(2)$ (sectors $m \in \mathbb{Z}$) and for the λ -model ($U(1)$ sectors); it does not touch Θ_L or the gaps, so it cannot produce running. All running sits in the space doubling.

Lemma B (block/seam decomposition; proved). Label the $2L$ circle $A_1 \dots A_L B_1 \dots B_L$; let $\mathcal{T}_L^o$ be the open-chain transfer operator, $D_{\text{in}} = e^{\beta\sigma_L \cdot \sigma_1}$ the closing bond of one block and $D_s = \exp[\beta(\sigma_L^A \cdot \sigma_1^B + \sigma_L^B \cdot \sigma_1^A)]$ the two seam bonds. (i) $\mathcal{H}_{2L} = \mathcal{H}_A \otimes \mathcal{H}_B$, $\mathcal{T}_{2L} = D_s^{1/2} (\mathcal{T}_L^o \otimes \mathcal{T}_L^o) D_s^{1/2}$ and $\mathcal{T}_L \otimes \mathcal{T}_L = (D_{\text{in}}^A D_{\text{in}}^B)^{1/2} (\mathcal{T}_L^o \otimes \mathcal{T}_L^o) (D_{\text{in}}^A D_{\text{in}}^B)^{1/2}$: the two differ only in the wiring of two bonds per time slice. (ii) $V_{\ell_1} \otimes V_{\ell_2} = \bigoplus_{\ell=|\ell_1-\ell_2|}^{\ell_1+\ell_2} V_\ell$, the composite singlet space is $\bigoplus_j (V_j \otimes V_j)^{\text{SO}(3)} \otimes M_j^A \otimes M_j^B$, one singlet per equal-isospin pair. (iii) $C_2^{\text{tot}} = C_A + C_B + 2L_A \cdot L_B = 2(C_A + C_B) - (L_A - L_B)^2$; in

the $A \leftrightarrow B$ -symmetric thermal state $\langle C_2^{\text{tot}} \rangle = 4\langle C_A \rangle - \langle (L_A - L_B)^2 \rangle$. (iv) Decoupled comparison: $\text{Tr}[U(g)(\mathcal{T}_L \otimes \mathcal{T}_L)^T] = Z_{L \times T}(g)^2$ has composite singlet weight $p_0^{\text{dec}} = \sum_j p_j^2 / (2j+1)^2$ (O(2): $\sum_m p_m^2$). Numbers (O(3), $\beta = 1.8$, $L = 16$, four-sector weights from the data file $p_0, \dots, p_3 = 0.428, 0.464, 0.095, 0.018$): $p_0^{\text{dec}} = 0.183 + 0.024 + 0.0004 + \dots \approx 0.21 < p_0(L) = 0.43 < p_0(2L) = 0.47$. Two free blocks are *less* singlet than one block; the two seam bonds must raise the singlet weight by ≈ 0.26 to break even and by a further ≈ 0.04 to produce the observed contraction.

Lemma C (positive seam character expansion; proved). By Funk–Hecke $e^{\beta s \cdot s'} = \sum_j a_j(\beta) P_j(s \cdot s')$ with $a_j > 0$. With $J = (j_t, j'_t)_{t \leq T}$ the representation labels on the $2T$ seam bonds and g a twist across the time seam,

$$Z_{2L \times T}(g) = \sum_J w_J(\beta) Q_J(g), \quad w_J = \prod_t \frac{(4\pi)^2 a_{j_t} a_{j'_t}}{(2j_t+1)(2j'_t+1)} > 0,$$

$$Q_J(g) = \sum_{m, m'} |\text{Tr}_A[U_A(g) X_{J, m, m'}^A]|^2 \geq 0,$$

where X^A is the open-chain- A time-ordered product with the seam insertions $Y_{j_t m_t}(\sigma_{L, t}^A) \bar{Y}_{j'_t m'_t}(\sigma_{1, t}^A)$; the $J = 0$ term is $a_0^{2T} Z_{L \times T}^o(g)^2$. *Proof sketch.* Expand each seam bond; the B factor is $\text{Tr}_B[U_B(g) X^B]$ with conjugate insertions; the spatial reflection $B_x \mapsto B_{L+1-x}$ composed with complex conjugation is an antiunitary commuting with $U(g)$ that maps X^B to X^A , so the B factor is the complex conjugate of the A factor. $\square$ Each Q_J is a class function. O(2) has the identical expansion with $a_j \rightarrow I_j(\beta)$.

No-go for positivity alone (proved). No inequality $p_0(2L \times T) \geq \Phi(\{p_j(L, T)\})$ with $\Phi > p_0$ on the relevant range can be derived from Lemmas A–C and the Casimir identity by manipulations valid for every compact group. *Proof.* Such manipulations hold verbatim for O(2) (sectors m , Casimir m^2 , $m_1 \otimes m_2 = m_1 + m_2$, singlet pairs $(m, -m)$). For O(2) at $\beta = 1.2 > \beta_{\text{KT}}$ on square tori the Casimir-map data give $p_0(L) = 0.448, 0.435, 0.436, 0.436, 0.437$ for $L = 8, \dots, 128$ — no gain — and $\rho_{\text{hel}} \rightarrow \rho_\infty \geq 2/\pi$ [11, 10] forbids $p_0 \rightarrow 1$; Φ would be violated there. $\square$ The gain must come from a step that is false for O(2): the recombination of non-zero seam fluxes $J \neq 0$ into composite sectors through $6j$ symbols, i.e. from the non-commutativity of the generators.

Lemma D (two-rotor reduction of the seam; proved). Model: two rigid rotors $n_A, n_B \in S^2$ with block inertia Θ coupled by the seam, $H_2 = (L_A^2 + L_B^2)/(2\Theta) - K n_A \cdot n_B$, $K = 2\beta$. Then $H_2 = L_{\text{tot}}^2/(4\Theta) + H_{\text{rel}}$, $H_{\text{rel}} = (L_A - L_B)^2/(4\Theta) - K n_A \cdot n_B$, with $[L_{\text{tot}}^a, H_{\text{rel}}] = 0$, so $E_\ell = \ell(\ell+1)/(4\Theta) + \epsilon_\ell$ with $\epsilon_\ell = \min \text{spec}(H_{\text{rel}}|_\ell)$. (O3) On $|\ell_1 \ell_2; \ell m\rangle$: $(L_A - L_B)^2 = 2\ell_1(\ell_1+1) + 2\ell_2(\ell_2+1) - \ell(\ell+1)$ and $\langle \ell'_1 \ell'_2; \ell | n_A \cdot n_B | \ell_1 \ell_2; \ell \rangle = (-1)^{\ell_1 + \ell'_2 + \ell} \begin{Bmatrix} \ell'_1 & \ell'_2 & \ell \\ \ell_2 & \ell_1 & 1 \end{Bmatrix} \langle \ell'_1 || n || \ell_1 \rangle \langle \ell'_2 || n || \ell_2 \rangle$: both depend on ℓ beyond the additive term, and nothing forces $\epsilon_\ell = \epsilon_0$. (O2) For planar rotors $H_2 = (m_A^2 + m_B^2)/(2\Theta) - K \cos(\phi_A - \phi_B)$ with $M = m_A + m_B$, $k = m_A - m_B \equiv M \pmod{2}$: $H_2 = M^2/(4\Theta) + [k^2/(4\Theta) - \frac{K}{2}(\text{shift } k \text{ by } \pm 2)]$, and the bracket acts on the parity class of M only, so $E(M) = M^2/(4\Theta) + \epsilon_{M \bmod 2}$: the composite inertia is exactly 2Θ up to the tunnelling splitting $\epsilon_{\text{odd}} - \epsilon_{\text{even}}$. *Proof.* Algebra; for (O2) the potential shifts k by ± 2 and commutes with M .

$\square$ Where the commutator enters: the in-plane inertia of a planar pair is 2Θ for every bend angle, while a bent pair on S^2 has inertia tensor $(2\Theta, 2\Theta \cos^2 \frac{\theta}{2}, 2\Theta \sin^2 \frac{\theta}{2})$, bend-dependent because $[L^a, n^b] = i\epsilon_{abc} n^c$; the zero-point bend couples to the total rotation (vibration–rotation coupling), which is the ℓ -dependence of ϵ_ℓ .

Proposition D' (numerically supported; semiclassical derivation). For O(3), $E_\ell - E_0 = \ell(\ell+1)/(2\Theta_{\text{eff}})$ with $\Theta_{\text{eff}} = 2\Theta - \sqrt{\Theta/(8K)}(1 + O((K\Theta)^{-1/2}))$. *Derivation.* Harmonic bend

with reduced inertia $\Theta/2$ and frequency $\omega = \sqrt{2K/\Theta}$ has zero-point $\langle \theta^2 \rangle = \sqrt{2/(K\Theta)}$; the two large principal inertias give a rotational constant $\approx (1/4\Theta)(1 + \theta^2/8)$, hence $\delta = 2\Theta - \Theta_{\text{eff}} = (\Theta/4)\langle \theta^2 \rangle = \sqrt{\Theta/(8K)}$. *Exact diagonalization* (product basis to $\ell_{\text{max}} = 10\text{--}17$, $M = 0$ subspace, projection on L_{tot}^2 ; Δ_1 converged to 2×10^{-4} relative): $(\Theta, K) = (4, 1.6)$: $\Theta_{\text{eff}} = 7.2595$, deficit 0.7405 vs $\sqrt{\Theta/8K} = 0.5590$; $(4, 12.8)$: 0.2141 vs 0.1976; $(8, 3.2)$: 0.6287 vs 0.5590; $(8, 12.8)$: 0.2954 vs 0.2795; $(16, 3.2)$: 0.8565 vs 0.7906; $(32, 3.2)$: 1.1815 vs 1.1180; $(32, 12.8)$: 0.5863 vs 0.5590; the composite is an exact rotor to 10^{-4} ($(E_2 - E_0)/3(E_1 - E_0) = 0.9945\text{--}1.0000$) and the deficit/semiclassical ratio falls from 1.32 at $K\Theta = 6.4$ to 1.05 at $K\Theta = 410$. **O(2) control**: deficit 7.9×10^{-4} (4, 1.6), 2.8×10^{-9} (8, 3.2), 6.0×10^{-2} (8, 0.4) — exponentially small in $\sqrt{K\Theta}$, as Lemma D(O2) requires. λ **control** (easy-plane term $D n_z^2$ per rotor, U(1) inertia Θ_z from $E(M=1) - E(M=0)$; $\Theta = 8$, $K = 3.2$, $\lambda = 0.5$): $\delta_z = 0.314$ at $D = 0$ (half the isotropic value), -0.227 at $D = 4$, -0.030 at $D = 16$, $+0.024$ at $D = 64$ (the last point converged to ≈ 0.01 only; provisional): the easy-plane term freezes the out-of-plane bend and quenches the gain by an order of magnitude towards the planar value 0.

Honest statement of the toy’s scaling. Translating to the lattice step with $\Theta_L = \rho_H L$, $K = 2\beta$ gives $\rho_H(2L) = \rho_H(L) - \delta/2L = \rho_H(L) - \frac{1}{8}\sqrt{\rho_H/(\beta L)}$: at $\beta = 1.6$, $L = 8$, $\rho = 0.73$ this is 0.030 against the measured 0.198, and it decays like $L^{-1/2}$ whereas the data (and one loop) are L -independent. The single-seam-mode model reproduces the sign and the mechanism, not the rate. The missing factor is the whole band of relative-orientation modes of the $2L$ system with wavelengths between L and $2L$: in the spin-wave approximation $\sum_{L < 2\pi/|k| < 2L} 1/(\rho k^2)$ per unit area gives the $(N-2) \ln 2/(2\pi)$ of one loop, which is exactly the band sum $T = \sum_{M \leq |k| \leq K_1} |k|^{-2} = \log(K_1/M)/2\pi$ with $K_1/M = 2$ of the machine-generated stiffness-iteration lemma (Section 13) and the $\log 2$ of a fixed-ratio annulus (Section 6). The seam/Clebsch–Gordan step is the torus-native, one-mode truncation of that lemma’s rotation-score Cauchy–Schwarz gain: the score of a relative rotation of block B against A acts on the composite through $[L_B^a, \cdot]$, and its coupling to the total-rotation sector is the ℓ -dependence of ϵ_ℓ — the same algebraic object, $[L^a, L^b] = i\varepsilon_{abc}L^c$, with structure constant $c = 1$ on S^2 , 2 on S^3 in quaternion variables, 0 for O(2).

The first unproved step (exact statement; Θ -form, corrected in v2.1). Let $\Delta_1(\cdot)$ be the $\ell = 1$ sector gap of the circle transfer Hamiltonian $-\ln \mathcal{T}$ and $\Theta_L = 1/\Delta_1(L)$. Needed: there exist $c_1 > 0$ independent of β and $L_0(\beta)$ such that for $L \geq L_0(\beta)$

$$\begin{aligned} \Theta_{2L} &\leq 2\Theta_L - c_1 L \quad (\Leftrightarrow \rho_H(2L) \leq \rho_H(L) - c_1/2), \\ \text{i.e. } \Delta_1(2L) &\geq \frac{\frac{1}{2}\Delta_1(L)}{1 - \frac{1}{2}c_1 L \Delta_1(L)} \approx \frac{1}{2}\Delta_1(L)[1 + \frac{1}{2}c_1 L \Delta_1(L)], \end{aligned} \tag{12}$$

together with the uniform rotor-regime statement $\varepsilon_0(L, L) \leq C$ of Lemma A so that gaps control weights; one loop predicts $c_1 = \ln 2/\pi \approx 0.221$, i.e. a drop $\ln 2/(2\pi) = 0.110$ in ρ_H per doubling. Version 2.0 wrote the bracket as $[1 + c_1 L \Delta_1(L)]$, which is not equivalent and double-counts c_1 (erratum E12, Section 3.5); the quantity tested in Section 10.1 is the deficit $D(L) := 2\Theta_L - \Theta_{2L}$ against $c_1 L$. Why nothing so far reaches it: (i) direction — Lemma B gives variational *upper* bounds on $E_0(2L)$, while (12) needs a *lower* bound on $E_1(2L)$ relative to $E_0(2L)$ (Temple/Kato bounds need a priori knowledge of the $\ell = 1$ sector’s own gap); (ii) model dependence — Proposition D’ proves nothing about $\mathcal{T}_{2L}$, whose blocks are open chains with internal spin-wave modes and whose seam is two bonds; (iii) uniformity — c_1 must not depend on β , and every finite- ℓ_{max} or perturbative argument carries β . Controls the step must pass: O(2) in the BKT phase has $\Delta_1(2L) = \frac{1}{2}\Delta_1(L)(1 + O(e^{-c\sqrt{\beta L}}))$ (Lemma D(O2) and [11]), so the inequality must fail there; the $\lambda < 1$ power-law phase [13] forces failure at large L , so a proof must use the full SO(3) recombination, not U(1) plus smoothness.

Conjecture (sector contraction under doubling) and data check. For the nearest-neighbour O(3) model on square tori, with $\Lambda(L) = \ln[p_0(L)/(1 - p_0(L))]$: $\forall \beta > 0 \exists L_0(\beta) \forall L \geq L_0(\beta): \Lambda(2L) - \Lambda(L) \geq f(\rho_{\text{hel}}(L))$ with $f: (0, \infty) \rightarrow (0, \infty)$ continuous, independent of β , bounded below on compacts, $f(\rho) = \ln 2/(2\pi\rho)(1 + o(1))$ as $\rho \rightarrow \infty$. Required failures: O(2) for $\beta > \beta_{\text{KT}}$ and $\lambda < 1$ ($\Lambda \rightarrow \Lambda_\infty$ finite). **Rigorous part** (from the bracket (2)): Λ lies in $[\text{logit } p_0^{\text{lo}}, \text{logit } p_0^{\text{up}}]$, and the increment over the full measured range is bounded below by $\Lambda^{\text{lo}}(L_{\text{max}}) - \Lambda^{\text{up}}(L_{\text{min}}) = 5.98$ ($\beta = 1.6, 8 \rightarrow 128$), 1.36 ($1.8, 8 \rightarrow 128$), 0.84 ($2.0, 8 \rightarrow 256$), 0.29 ($2.2, 8 \rightarrow 256$); per doubling the bracket increment is positive at every step for $\beta = 1.6$ (0.22, 0.58, 1.41, 3.66), for $L \geq 16$ at $\beta = 1.8$ (0.06, 0.28, 0.69) and for $L \geq 64$ at $\beta = 2.0$ (0.11, 0.32), and not resolved elsewhere. **Provisional part** (single Monte Carlo series per point, p_0 from the four-sector solution of the data file with scatter ≈ 0.01 , i.e. ≈ 0.05 in Λ at $p_0 \approx 0.4$; no lattice-artefact separation): $\Delta\Lambda > 0$ on all 17 O(3) doublings, minimum 0.061 ($\beta = 2.0, 32 \rightarrow 64$, next to 0.209 and 0.252), median 0.21; one-loop $f(\rho) = 0.110/\rho$ gives 0.078, 0.092, 0.110 at $\rho = 1.42, 1.20, 0.98$ against measured 0.102, 0.158, 0.246; O(2) BKT ($\beta = 1.2, 1.4, 8$ doublings): $|\Delta\Lambda| \leq 0.052$, compatible with zero (the required failure); λ -model (12 doublings): $|\Delta\Lambda| \leq 0.011$ at $\rho_{\text{hel}} = 0.9\text{--}1.4$ (the required failure, already at $L = 8\text{--}64$); O(2) disordered ($\beta = 0.9$): $\Delta\Lambda$ large, allowed. The ideal-rotor map (11), inverted to predict Λ from ρ_{hel} , reproduces the O(3) values within 0.1 and the O(2) BKT values within 0.07, and is off by 1.1 for the λ -model (-0.47 vs -1.57): the O(3) Casimir map does not describe U(1) sectors, and the “sectors split” control is visible in the data themselves. The conjecture is not contradicted; its lower envelope is not resolved by these data.

10.1 Transfer-matrix test of the first unproved step (v2.1; review angle 1)

[pre-registered] (§0 of `gap_step_test.md`, written 19 October 2026 before any production run and not edited afterwards; code `gap_step_test.py`; tables `gap_step_test.csv`, `gap_step_test_best.csv`, `gap_step_test_convergence.csv`, `gap_step_test_all.csv`; validation `v1_mc_check.json`). The circle- L transfer matrix $\mathcal{T}_L = D^{1/2} K^{\otimes L} D^{1/2}$ ($D = \exp \beta \sum_x n_x \cdot n_{x+1}$, K diagonal in $Y_{\ell m}$ with eigenvalue $4\pi i_\ell(\beta)$; $\text{Tr } \mathcal{T}_L^T = Z_{L \times T}$ exactly) was built in the per-site truncation $\ell \leq \ell_{\text{max}}$ and diagonalized in the sectors (total M , parity $\Pi = (-1)^{\sum \ell}$, momentum $k = 0$, flip f); λ_0 is the top of ($M=0, \Pi=+, f=+$), the $\ell = 1$ leading eigenvalue the top of $M = 1$ (checked absent from $M = 2$), $\Delta_\ell = \ln(\lambda_0/\lambda_\ell)$. Grid: $\beta \in \{1.4, 1.6, 1.8, 2.0\}$, $L \in \{2, 3, 4, 6\}$ with $\ell_{\text{max}} \leq 6, 6, 5, 4$ and $L = 8$ at $\ell_{\text{max}} = 2$ only; O(2) control ($m_{\text{max}} \leq 5\text{--}6$, $L \leq 8$); $\lambda = 0.5$ control ($L \leq 6$); the rotor chain $H = \frac{g}{2} \sum L_i^2 - \sum n_i \cdot n_{i+1}$ at $g = 1/\beta^2$ (the Hamiltonian-limit map $g = 1/(\beta_t \beta_s)$, qualitative at these β). *Validation (V1–V3)*. $\text{Tr } \mathcal{T}_L^T$ reproduces the heat-bath Monte Carlo energy per bond of the 2×2 , 2×4 and 3×3 tori within 1σ ($\leq 2.4 \times 10^{-4}$) wherever ℓ_{max} is converged (the 3×3 , $\beta = 2.0$ energy at $\ell_{\text{max}} = 4$ is 0.0027 low, a pure truncation effect; gaps converge about 100× faster than the energy); Δ_1 is converged to $\leq 4 \times 10^{-4}$ relative at $L \leq 4$, to 0.9–6% at $L = 6$ (geometric extrapolation applied, residual $\lesssim 1\%$), and $L = 8$ is corrected by the L -extrapolated $\ell_{\text{max}} = 2$ truncation factor ($f_8 = 1.13\text{--}1.48$; $\pm 1\text{--}3\%$ on $\Delta_1(8)$, so every $4 \rightarrow 8$ number is [provisional]); the rotor code reproduces the `rotor_cert` gaps at $g = 2$ to all digits (1.5142, 1.4117, 1.3853, 1.3780 for $N = 2\text{--}5$). The rotor-regime check $\Delta_2/\Delta_1 \in [2.7, 3.3]$ holds on all lattice doublings except ($\beta = 1.4, 2 \rightarrow 4$) and the truncated $L = 8$ at $\beta = 1.4$.

Readings (Table 10). (a) $\Delta_1(2L)/\Delta_1(L) > \frac{1}{2}$ on all 12 O(3) lattice doublings and all 8 rotor-chain doublings in the rotor regime; the excess over $\frac{1}{2}$ grows with L at fixed β (0.062, 0.078, 0.095 at $\beta = 1.6$) and shrinks with β . (b) The deficit per unit length $D(L)/L$ is *not* $\propto L^{-1/2}$: it rises monotonically with L on every β series (0.327 $\rightarrow$ 0.369 $\rightarrow$ 0.404 at $\beta = 1.4$, where the single-seam toy of Proposition D' predicts 0.327 $\rightarrow$ 0.267 $\rightarrow$ 0.231), so the toy scaling is excluded in direction and magnitude. It is not L -independent within the pre-registered 20% either ($4 \rightarrow 8$ against

Table 10: Doubling test on circle transfer matrices (best estimates: $\ell_{\max}$ -extrapolated, $L = 8$ truncation-corrected; file `gap_step_test_best.csv`). “one-loop ratio” is the Θ -form prediction $\frac{1}{2}/(1 - 0.110/\rho_H(L))$; $D/L \geq c_1$ tests Eq. (12) with $c_1 = 0.221$. All $4 \rightarrow 8$ rows are **[provisional]**. The λ rows use the U(1) inertia $\Theta_z = 1/(2\Delta_1)$; the $3 \rightarrow 6$ λ doublings are not converged and omitted.

model	β	$L \rightarrow 2L$	$\Delta_1(L)$	$\Delta_1(2L)$	ratio	one-loop ratio	$D/L \geq 0.221$	$\rho_H(L)$	D/L
O(3)	1.4	$2 \rightarrow 4$	0.51593	0.31035	0.6015	0.5642	yes	0.969	0.327
O(3)	1.4	$3 \rightarrow 6$	0.37966	0.24037	0.6331	0.5719	yes	0.878	0.369
O(3)	1.4	$4 \rightarrow 8$	0.31035	0.20705	0.6671	0.5793	yes	0.806	0.404
O(3)	1.6	$2 \rightarrow 4$	0.43247	0.24305	0.5620	0.5527	yes	1.156	0.255
O(3)	1.6	$3 \rightarrow 6$	0.30683	0.17747	0.5784	0.5565	yes	1.086	0.295
O(3)	1.6	$4 \rightarrow 8$	0.24305	0.14469	0.5953	0.5601	yes	1.029	0.329
O(3)	1.8	$2 \rightarrow 4$	0.37025	0.20009	0.5404	0.5445	no	1.350	0.202
O(3)	1.8	$3 \rightarrow 6$	0.25694	0.14171	0.5515	0.5465	yes	1.297	0.242
O(3)	1.8	$4 \rightarrow 8$	0.20009	0.11239	0.5617	0.5484	yes	1.249	0.275
O(3)	2.0	$2 \rightarrow 4$	0.32279	0.17063	0.5286	0.5383	no	1.549	0.168
O(3)	2.0	$3 \rightarrow 6$	0.22113	0.11872	0.5369	0.5395	no	1.507	0.207
O(3)	2.0	$4 \rightarrow 8$	0.17063	0.09282	0.5440	0.5407	yes	1.465	0.237
O(2)	1.4	$2 \rightarrow 4$	0.25659	0.11901	0.4638	0.5	—	0.974	−0.152
O(2)	1.4	$4 \rightarrow 8$	0.11901	0.05836	0.4903	0.5	—	1.050	−0.041
O(2)	2.0	$2 \rightarrow 4$	0.15561	0.07387	0.4747	0.5	—	1.607	−0.171
O(2)	2.0	$4 \rightarrow 8$	0.07387	0.03656	0.4949	0.5	—	1.692	−0.035
$\lambda=0.5$	1.4	$2 \rightarrow 4$	0.39169	0.19661	0.5020	0.5	—	0.638	+0.005
$\lambda=0.5$	1.6	$2 \rightarrow 4$	0.30460	0.14089	0.4625	0.5	—	0.821	−0.133
$\lambda=0.5$	1.8	$2 \rightarrow 4$	0.24441	0.11066	0.4527	0.5	—	1.023	−0.214
$\lambda=0.5$	2.0	$2 \rightarrow 4$	0.20221	0.09204	0.4552	0.5	—	1.236	−0.244

$2 \rightarrow 4$: 1.23–1.41; $3 \rightarrow 6$ against $4 \rightarrow 8$: 9–15%), hence the pre-registered verdict “undecided (lattice regime)” for strict L -independence; the $L = 2$ doubling (doubled bonds) is the main outlier. In terms of ρ_H drops per doubling ($D/2L$): 0.16–0.20 ($\beta = 1.4$), 0.13–0.16 (1.6), 0.10–0.14 (1.8), 0.08–0.12 (2.0) against one loop 0.110 and against the Monte Carlo square-torus drops at $L = 8 \rightarrow 16$ of 0.198, 0.161, 0.143 ($\beta = 1.6, 1.8, 2.0$; Table 2); the $4 \rightarrow 8$ transfer-matrix drops (0.16, 0.14, 0.12) sit just below the $8 \rightarrow 16$ Monte Carlo values and are rising, consistent with one continuous curve above one loop from $L \approx 8$. Against the one-loop $c_1 = 0.221$ the Θ -form inequality holds in 9 of 12 lattice doublings; the three failures are at $\beta \geq 1.8$ with $L \leq 3$ ($\rho_H \geq 1.3$), where $D/L = 0.17$ –0.21 approaches 0.221 from below as L grows. D/L at fixed L falls with β (0.40 $\rightarrow$ 0.24 at $4 \rightarrow 8$): c_1 is not β -independent on this grid, but it is ≥ 0.221 for $L \geq 4$ at every β in the grid. (c) *O(2) control: closed*. The ratio converges to $\frac{1}{2}$ from below (0.464–0.475, 0.482–0.490, 0.490–0.495 for the three doublings) with $D/L < 0$, β -independent to $\pm 10\%$ and decaying roughly like L^{-2} : the doubled circle is slightly stiffer than two blocks, a lattice correction of the wrong sign for a gain; $\Delta_2/\Delta_1 \rightarrow 4$ (U(1) rotor). The pre-registered threshold $|D/L| < 0.03$ at $L \geq 4$ is missed by 0.005–0.011, with the wrong sign; no gain, the required failure. (d) $\lambda = 0.5$ *control: closed*. In the converged rows ($L \leq 4$, $\ell_{\max} = 5$) the ratio is 0.502, 0.463, 0.453, 0.455 and $D_z/L = +0.005, -0.133, -0.214, -0.244$ at $\beta = 1.4, 1.6, 1.8, 2.0$: the composite is *stiffer* than two blocks for $\beta \geq 1.6$, the same reversal as $\delta_z < 0$ in the two-rotor λ -control of Proposition D'; $\Delta(M=2)/\Delta(M=1)$ interpolates from 3.0–3.6 ($L = 2$) to 3.4–3.9 ($L = 4$), SO(3) toward U(1). The step inequality ($D > 0$ with a β -independent margin) fails in the λ model, as required. (e) *Rotor chain*. Ratios above the Θ -form one-loop reference on all doublings; the velocity-normalized deficit $D\sqrt{g}/L = 0.38/0.39/0.41$ ($g = 0.51$), 0.33/0.33/0.34 (0.39), 0.31/0.30/0.31 (0.31), 0.29/0.28/0.28 (0.25) for $L = 2/3/4 \rightarrow 2L$ is ≥ 0.221 everywhere and flat in L within 5% (the heuristic normalization $c_{\text{bare}} = \sqrt{g}$ is stated as such); the $g = 2$ doubling (ratio 0.91, $\Delta_2/\Delta_1 = 2.2$ –2.5) is a magnon system, outside the test’s domain. *What this test gives and does not give*. The margin required by Eq. (12) is present at every coupling and size, grows with L rather than shrinking, is at or above one loop for $L \geq 4$, and the two controls

fail as the step demands. It is not a proof, and it does not yet show a size-independent constant: D/L still rises between $L = 4$ and 8, and the $L = 8$ values rest on a truncation correction. Deviations: $\ell_{\max} = 3$ at $L = 8$ exceeded the 6 GB budget; the planned λ , $L = 6$, $\ell_{\max} = 4$ run was dropped for time; three drivers briefly used about 4 cores instead of ≤ 3 .

DMRG rates on open rotor chains (v2.2; [provisional]). Record `rotor_rate.md` (pre-registration `prereg.md`, 2026-10-20T03:19:03Z; code `rotor_rate.py`, `run2.py`, `analyze.py`, `ed_obc_check.py`, `partB.py`; data `rotor_rate_results.csv`, `rotor_rate_fits.csv`, `dmrg_energies.csv`, `rotor_rate_partB.csv`). Two-site DMRG (TeNPy, $Y_{\ell m}$ basis with $\ell \leq 3$, L_z conserved, $\chi \leq 128$, energies extrapolated in the truncation error) of $H = \frac{g}{2} \sum L_i^2 - \sum n_i \cdot n_{i+1}$ at $g = 1/\beta^2$ on *open* chains (periodic DMRG was too slow; declared deviation). On an open chain $\Theta_N^{\text{OBC}} = \Theta_N^{\text{bulk}} + c_b$, so the literal deficit carries $c_b/(Na_\tau)$; the two-point estimator $\kappa_N = (D_{2N} - D_N)/(Na_\tau) = (3\Theta_{2N} - \Theta_{4N} - 2\Theta_N)/(Na_\tau)$ cancels c_b exactly and equals the Θ -form rate for logarithmic running ($\Theta_N = N\beta(\beta - a \ln N) + c_b$ gives $\kappa = 2a \ln 2$, i.e. $\ln 2/\pi$ at one loop). Validation: open-chain exact diagonalization at $N = 4$ agrees with DMRG to $\leq 9 \times 10^{-11}$ in all three sectors. Results: $\kappa_4 = 0.342/0.335/0.303/0.282$ at $\beta = 1.4/1.6/1.8/2.0$ ($N = 4, 8, 16$; DMRG errors ≤ 0.004), all above $\ln 2/\pi = 0.2206$, and within 0.01 of the periodic transfer-matrix rotor-chain values 0.33/0.31/0.28 of Table 10 for $\beta \geq 1.6$; the $\ell_{\max} = 3 \rightarrow 4$ check at $\beta = 2.0$, $N = 4$ lowers $\hat{D}$ by 9%, so the value 0.282 is likely biased high by 0.02–0.03 and still above one loop. At $\beta = 1.4$ the rate falls ($\kappa_4 = 0.342 \rightarrow \kappa_8 = 0.290$) as the chain leaves the rotor band ($R = (E_2 - E_0)/(E_1 - E_0) = 2.73 \rightarrow 2.64 \rightarrow 2.50$ at $N = 4, 8, 16$), the massive crossover where any Θ -form step must stop holding (pre-registered P2, confirmed). Controls: $O(2)$ $|\kappa| \leq 0.0058$ at all β , N (≤ 0.0024 for $\beta \geq 1.6$); $\lambda = 0.5$: $\kappa_z \in [-0.093, +0.013]$, no positive gain, and the uncorrected $\hat{D}_z$ negative at $\beta \geq 1.6$. *Not reached*: the pre-registered window of doublings $8 \rightarrow 16 \rightarrow 32 \rightarrow 64$ — $O(3)$ at $N = 32$ converged only at $\beta = 1.4$, $N = 64$ not run — so the rules “ $\kappa \geq 0.2206$ β -uniformly on rotor-regime doublings” and “plateau over $N \geq 8$ ” are evaluated on $N = 4$ only (OPEN; pass where testable) and not testable as registered, respectively. Fits with DMRG-statistical errors only: at $\beta = 2.0$ a plateau is rejected ($\chi^2 = 225/2$) and both $c + a/N$ ($c = 0.282(4)$, $p = 0.84$) and $c + a \ln N$ ($p = 0.95$) fit; three points cannot separate them.

10.2 Rotor-ring lemma, transfer lemma and the routes to the Θ -form step (v2.1; review angles 1–2)

Record `c1_attack.md`; code `c1_attack.py`, `lam_job.py`, `c1_tables.py`, `h0_check.py`; data `ed_small.csv`, `ed_big.csv`, `ed_lam.csv`, `c1_doubling_table.csv`, `c1_rotor_ring_03.csv`. The inputs read were 215-b `preliminary.tex` (lines 1–880) and 215-c `preliminary.tex` (lines 1–699) at commit `adc7f1241b`; both are unrefereed machine-generated preprints.

Lemma R (rotor ring of N rotors; proved). For $H = \sum_i L_i^2/(2\Theta) - K \sum_{\langle ij \rangle} n_i \cdot n_j$ on a ring of N sites: (i) $N \sum_i L_i^2 - (\sum_i L_i)^2 = \sum_{i < j} (L_i - L_j)^2$ as operators, so $H = L_{\text{tot}}^2/(2N\Theta) + H'$ with $H' = (2N\Theta)^{-1} \sum_{i < j} (L_i - L_j)^2 + V$ commuting with every L_{tot}^a ; hence $E_\ell = \ell(\ell + 1)/(2N\Theta) + \epsilon_\ell$, $\epsilon_\ell = \inf \text{spec } H'|_\ell$, and $\Delta_1 = 1/(N\Theta) + (\epsilon_1 - \epsilon_0)$. (ii) $SO(3)$: the relative fields $Y_{ija} = X_a^{(i)} - X_a^{(j)}$ are Haar-divergence-free and bracket-generating ($[Y_{ija}, Y_{ijb}] = \varepsilon_{abc}(X_c^{(i)} + X_c^{(j)})$), so the Lie hull contains every $X_a^{(i)}$, so H' is hypoelliptic with compact resolvent and strictly positive heat kernel (Hörmander; Chow–Rashevskii connectivity and Bony’s strong maximum principle) [37]; $e^{-tH'}$ is positivity-improving, its ground state is unique, positive and rotation-invariant, hence a singlet, so $\epsilon_1 > \epsilon_0$ and $\Theta_N := 1/\Delta_1 < N\Theta$ *strictly*; moreover $\lim_N \Delta_1(N) = \lim_N (\epsilon_1 - \epsilon_0)$, the whole mass gap sits in H' . (iii) $O(2)$: the Y ’s commute and the shift $m_i \rightarrow m_i + 1$ maps the M -sector unitarily onto $M + N$ preserving H' , so $E(M) = M^2/(2N\Theta) + \epsilon_{M \bmod N}$ exactly: the ring is rigid up to $\mathbb{Z}_N$ tunnelling. (iv) Limitation: the λ -deformed ring also satisfies the bracket condition

through its single-site transverse generators, so Lemma R is qualitative and does not separate the λ control; the discriminating quantity is the doubling ratio $\Theta_{2N}/(2\Theta_N)$.

Mode sum (derived, semiclassical; review angle 2). Rigidly rotating a bent configuration in the plane $(\bar{n}, a)$ has inertia $\Theta \sum_i (1 - \sum_{c \geq 3} \pi_{i,c}^2)$, so $\delta_N = N\Theta - \Theta_N = \Theta \frac{N_c-2}{N_c-1} \sum_i \langle \pi_i^2 \rangle$ ($N = 2$ reproduces Lemma D'). Zero-point motion of the ring modes $\omega_j = 2\sqrt{K/\Theta} \sin(\pi j/N)$ gives $\delta_N = \frac{N_c-2}{4} \sqrt{\Theta/K} S_N$ with $S_N = \sum_{j=1}^{N-1} 1/\sin(\pi j/N)$, and $S_N/(4N) = (2\pi)^{-1} [\ln(2N/\pi) + \gamma_E] + O(N^{-2})$, so $\rho_H(2N) - \rho_H(N) \rightarrow -(N_c - 2) \ln 2/(2\pi)$: summing the $N - 1$ relative-orientation modes turns the single-mode $N^{-1/2}$ deficit of Proposition D' into an N -independent drop per doubling (0.1143, 0.1121, 0.1113, 0.1106, 0.1103 at $N = 2, 3, 4, 8, \geq 64$). The logarithm is the band of relative modes between scales N and $2N$, the torus counterpart of T in 215-b/c; with $c = 1$ the rotation-score lemma captures half of it. *Exact diagonalization* ($M = 0, 1$ sectors, $\ell_{\max} \leq 5$; [provisional]: $N \leq 6$, one K , $\ell_{\max} = 3$ at $N = 5, 6$): $O(3)$ drops in ρ_H of 0.253/0.249 ($\beta_{\text{eff}} = \sqrt{K\Theta} = 1$; $2 \rightarrow 4 / 3 \rightarrow 6$), 0.188/0.196 (1.41), 0.147 (2; $2 \rightarrow 4$) — N -independent at fixed β_{eff} and approaching one loop from above as β_{eff} grows, where the single-mode model gives only 0.057–0.085; $\Theta_N < N\Theta$ in all 13 $O(3)$ points. Controls: $O(2)$ |drop| $\leq 4 \times 10^{-4}$ (ratio $\Theta_{2N}/2\Theta_N = 1.0003$); $\lambda = 0.5$ ratio 0.9965–1.0000 ($D = 4$) and 0.980–0.998 ($D = 0$) against 0.825–0.847 for isotropic $O(3)$ at the same β_{eff} : the gain is quenched by a factor ≥ 9 . A Temple/Kato lower bound on $\Delta_1(2N)$ was attempted and not obtained: Temple needs an a-priori bound on the second eigenvalue of $H'|_{\ell=1}$, Kato a residual with an a-priori gap — both circular — and Feynman–Bijl bounds point the wrong way.

Lemma T (transfer from pinned cells to tori; proved). Let $\mathbb{T} = \mathbb{Z}_{S_1} \times \mathbb{Z}_{S_2}$ with $S_\mu = \ell_\mu R$, homogeneous strength $\beta \leq b$; rings are the sites with $x_1 \equiv 0$ or $x_2 \equiv 0 \pmod{R}$ and cells the translates of the pinned square $Q(R)$. Let $H_{\mathbb{T}}(F, G)$ be the torus form of the 215-c stiffness quadratic form (contact term minus current covariance) on constant 3×2 matrices, and $h(R)$ the supremum over all pins and strengths in $[0, b]$ of $\lambda_{\max}(H_R)$ (the induction quantity of the 215-c “Stiffness iteration” lemma). Then for every constant A ,

$$H_{\mathbb{T}}(A, A) = (\ell_1 \ell_2)^{-1} \sum_c E[\mathcal{H}_c^\omega(A, A)] - \text{Var } j(A) \leq h(R)|A|^2,$$

with $j(A) = E[X_{\mathbb{T}}(A) \mid \omega]$. *Proof.* Every bond has an endpoint in a cell interior or lies on a ring line; ring-line bonds are tangent to exactly two cells and the half weights add to the full weight, so the action is a sum of cell actions. Given the ring configuration ω the cell interiors are independent pinned squares; contact and current are sums of per-cell terms; the law of total variance gives $\text{Var } X_{\mathbb{T}} = E \text{Var}(X_{\mathbb{T}} \mid \omega) + \text{Var } j$, the conditional variance is additive over cells, and $\mathcal{H}_c^\omega(A, A) \leq h(R)|A|^2$ by definition. $\square$ (No reflection positivity and no positivity of H_c is used.) *Link to sector weights (proved).* A twist by θ about a across the seam normal to μ is gauge-equivalent to rotating every μ -bond by θ/S_μ , so $-\partial_\theta^2 \log Z|_0 = (S_\nu/S_\mu) H_{\mathbb{T}}(a \otimes e_\mu, a \otimes e_\mu)$, and by Schur $-\partial_\theta^2 \log Z|_0 = \langle C_2 \rangle/3$; hence $\langle C_2 \rangle \leq 3(S_\nu/S_\mu)h(R)$, $\rho_s \leq \frac{3}{2}h(R)$ on square tori, and $1 - p_0 \leq \langle C_2 \rangle/2$ (normalization check: $H_{\mathbb{T}}(a \otimes e_\mu) = \Upsilon_{\text{axis}}$ with tree value $2\beta/3$, $\rho_s = \frac{3}{2}\Upsilon_{\text{axis}}$ with tree value β , as in Section 5).

Corollary T1 (conditional on the unrefereed 215-c Lemmas “Stiffness iteration” and “Initial reduction”). For $b \geq b_0$ every torus with sides in $R^*\mathbb{Z}$ has $\rho_s \leq \frac{3}{2}b^{-D}$ with $\log R^* \leq (4\pi - \pi u)b + o(b)$; at the k -th scale R_k of the iteration, $\Upsilon_{\text{axis}}(\mathbb{T}) \leq h_k$ for sides in $R_k\mathbb{Z}$ with $\log(R_k/R_0) \leq 4\pi(b - h_k) + C\sqrt{B \log B}$. This is a β -uniform, *averaged* form of the stiffness inequality at rate $\log 2/(4\pi)$ per doubling in the h normalization — half of one loop, consistent with Section 13 — and it is conditional on the imported estimate chain, of which the weighted current estimate (215-b eqs. (6-6)/(6-7), 215-c **prelim-weighted-current**) is the first step not verified by us.

Routes closed (structural, negative). (i) *The 215-b/c rotation-score method cannot produce the Θ -form step in its torus-to-torus, per-doubling form.* The band $T = \sum_{M \leq |k| \leq K_1} |k|^{-2}$ needs $M = B^J \leq K_1 = \ell/B^{4J}$, i.e. an enlargement $\ell \geq B^{5J}$ per step (the papers take $\ell \geq B^{10J}$); at $\ell = 2$ the band is empty and the gain is zero; each step also loses an additive $C \log B$, so the “ $\log 2/(4\pi)$ per doubling” exists only as an average over $\gtrsim 10J \log_2 B$ doublings; and conditioning a torus on cell rings yields pinned cells, never tori, so the self-reproducing hypothesis is a supremum over pins and no induction closes inside the torus class. In the torus construction ($\psi \equiv 1$, exact periodic Ward identity, no boundary term) the leading algebra is the same as 215-c ($\sum EY_i^2 = vT/2$, $\text{Cov} = cvT/2$, $\text{Hessian} \leq v^2T/2$) and gives $\text{Var}J(A) \geq (c^2T/2)[1 - O(T/v + T_w/v)]$ with $T = \log(K_1/M)/(2\pi) + O(1/M)$, $c = 1$ on S^2 . Of the nine estimate groups E0–E8 of the imported proof, the exact identities (E0) and the score moment bound (E1, Haar invariance) were proved on the torus; E2–E6 were verified at the level of the written argument; the weighted current estimate (E7) and the crossed term (E8) were not verified (E8’s sub-ingredient (7-7) was proved and (7-4) checked numerically, [provisional]). Reflection positivity, the Casimir identity and Funk–Hecke positivity prove only E0 and the links above; none of E2–E8 follows from them (pins break reflection symmetry; the Casimir identity is a torus statement; Funk–Hecke positivity is N -blind). Controls: for $O(2)$, $c = 0$ and $\text{Cov}(J(A), S) = 0$, so Lemma T transfers a pinned bound that does not decrease in the BKT phase; for $\lambda < 1$ the Ward identity needs rotations about $a \perp$ the twist colour to be a symmetry, which fails for the easy-plane action. (ii) *The $\ell_{\max} = 1$ Schur certificate is never N -uniform* (Section 11). *Conditional route:* Corollary T1. *First unproved step for the torus-native route:* Eq. (12) in Θ -form, uniform in β , which by Lemma R(i) is a statement about H' alone (v2.2: its weak-coupling end is settled at fixed N , with constant $D_N^{(0)} \geq \ln 2/\pi$ and one rigorous half, in Section 9.2; the Θ -form is a sufficient route to C1'' through the running coupling, Section 9); it must fail for $O(2)$ (Lemma R(iii)) and for $\lambda < 1$ (ratio ≥ 0.98 in exact diagonalization), so any proof must use $[L^a, L^b] = i\varepsilon_{abc}L^c$ quantitatively, not only Hörmander’s qualitative positivity. Next discriminating test (not run): DMRG of the $O(3)$, $O(2)$ and $\lambda = 0.5$ rotor rings at $\beta_{\text{eff}} = 1.4\text{--}2.0$ for $N \leq 64$, checking whether $\Theta_{2N}/(2\Theta_N) \leq 1 - \kappa/\rho_H$ holds uniformly (one loop $\kappa = \ln 2/\pi$).

11 A rigorous finite-chain gap certificate for the $O(3)$ quantum rotor chain

[pre-registered] (pre-registration in §0 of `rotor_cert_methods.md`, written 19 October 2026 before any target run and not edited afterwards; code `rotor_cert.py`; results `rotor_cert_results.csv` (one row per case $\times g \times \text{floor}$) and `rotor_cert_sectors.csv`; Fig. 10). This section transposes to the rotor chain the Schur-complement certificate used in the author’s local Weil-positivity work (Proposition 7.3 of [21]): keep a finite block of the Hilbert space exactly, price the infinite complement by one rigorous floor, and verify positive definiteness of one Schur complement with a rounding-error budget.

Model and sectors. $H_N(g, \sigma) = \frac{g}{2} \sum_{i=1}^N L_i^2 - \sigma \sum_{\langle ij \rangle} n_i \cdot n_j$ on $L^2((S^2)^N)$, with $\sigma = +1$ the ferromagnetic target and $\sigma = -1$ a wrong-sign control; bonds periodic ($N = 2$ periodic counts the bond (1, 2) twice) or open. Up to an overall factor and the relabelling $g \rightarrow g^2$ this is the Hamiltonian of Section 14, i.e. the transfer-matrix (Hamiltonian) form of the lattice $O(3)$ model with N sites on the spatial circle. The product basis is $|\ell_1 m_1 \cdots \ell_N m_N\rangle$ with the bond operator’s matrix elements given by Wigner–Eckart (stretched Clebsch–Gordan coefficients; machine checks $\sum_q (-1)^q n_q n_{-q} = 1$ to 2×10^{-16} and agreement with sympy Gaunt integrals). Every bond term changes ℓ_i and ℓ_j by ± 1 , so the parity $\Pi = (-1)^{\sum \ell_i}$ is conserved. We work in total $M = 0$ (every integer- J multiplet has an $M = 0$ member) and split by Π , translation momentum k (periodic) or reflection r (open), and the $m \rightarrow -m$ permutation F . Since e^{-tH} is positivity improving (heat

semigroup on a connected compact manifold with bounded potential [22]), the ground state is unique, has $J = 0$ and lies in the vacuum sector ($\Pi = +, F = +, k = 0$); hence

$$\text{gap} = \lambda_2(H) - \lambda_1(H) = \min\{\lambda_2(\text{vacuum sector}), \min_{\text{other sectors}} \lambda_1\} - E_0. \quad (13)$$

(Machine check: the union of sector spectra equals the unreduced $M = 0$ spectrum to 10^{-14} for $N = 2, 3, 4$.)

Certificate. Let P be the span of all basis states with every $\ell_i \leq \ell_{\max}$, $Q = 1 - P$, B the number of bonds, $L = PHP$ and $W = QHP$ ($= QVP$, finite rank, exactly computable: it maps P into the first shell $\max \ell = \ell_{\max} + 1$).

- *Floor F0.* $QHQ \geq d_0 Q$ with $d_0 = \frac{g}{2}(\ell_{\max} + 1)(\ell_{\max} + 2) - B$ (the kinetic term is diagonal and $|n_i \cdot n_j| \leq 1$).
- *Floor F1 (two-rotor bootstrap).* For $0 \leq b \leq g/(2 \max_i \deg i)$, $H = \sum_i c_i L_i^2 + \sum_{\langle ij \rangle} [b(L_i^2 + L_j^2) - \sigma n_i \cdot n_j]$ with $c_i = \frac{g}{2} - b \deg i \geq 0$, so $QHQ \geq d_1(b)Q$, $d_1(b) = \min_i c_i (\ell_{\max} + 1)(\ell_{\max} + 2) + B e_2^{\text{low}}(b)$, where $e_2(b) = \inf \text{spec}[b(L_1^2 + L_2^2) - n_1 \cdot n_2]$ (independent of σ). By Perron–Frobenius the two-rotor ground state is a function of $x = n_1 \cdot n_2$, so e_2 is the lowest eigenvalue of the Jacobi matrix $2b\ell(\ell + 1) - x$, $\langle \ell + 1 | x | \ell \rangle = (\ell + 1)/\sqrt{(2\ell + 1)(2\ell + 3)}$, bounded below by the same Schur certificate (truncation 60) in exact-rational ball arithmetic (python-flint arb, 120 bits; cross-check $e_2^{\text{low}}(0.5) = -0.157663483155$ against -0.157663483138 by diagonalization). The best b on a 40-point grid is used; $d_M = \max(d_0, \max_b d_1(b))$.
- *Schur lemma.* Let $H_{\text{aug}} = PHP + PHQ + QHP + d_M Q$. Then $H - H_{\text{aug}} = Q(H - d_M)Q \geq 0$, so by min-max $\lambda_k(H) \geq \lambda_k(H_{\text{aug}})$ for every k . For $\lambda < d_M$ and $x = p + q$,

$$\begin{aligned} \langle x, (H_{\text{aug}} - \lambda)x \rangle &\geq \langle p, (L - \lambda)p \rangle - 2\|q\| \|Wp\| + (d_M - \lambda)\|q\|^2 \geq \langle p, S(\lambda)p \rangle, \\ S(\lambda) &= L - \lambda - \frac{W^*W}{d_M - \lambda}. \end{aligned} \quad (14)$$

(a) If $S(\lambda) > 0$ then $\lambda_1(H) > \lambda$. (b) If $S(\lambda) + ww^* > 0$ for some $w \in P$, then on $\{x : \langle w, p \rangle = 0\}$ the form $H - \lambda$ is positive, so $\lambda_2(H) > \lambda$. Both statements hold sector by sector.

- *Upper bounds.* For any $p \in P$ in a sector, $\lambda_1(\text{sector}) \leq \langle p, Lp \rangle / \|p\|^2$, evaluated with an a-priori dot-product error bound and rounded outward; E_0^{up} from the vacuum sector.
- *Verification.* Positive definiteness of $\tilde{A} = \text{fl}(S(\lambda_c) - sI)$ is established by a completed LAPACK Cholesky factorization with the backward-error bound $\|\Delta A\|_2 \leq \delta = \gamma_{n+1} \text{tr} \tilde{A} / (1 - \gamma_{n+1})$ ([23], Thm. 10.3; complex case γ_{4n+8}), plus an explicit entry-error budget ε_A for the assembly of S ; the certificate requires $s > \delta + \varepsilon_A + 2u \text{tr} \tilde{A}$. The search value λ^* comes from a non-rigorous eigensolver; verification is at $\lambda_c = \lambda^* - \eta$, $\eta = 10^{-10}(1 + |\lambda^*|)$, multiplied by 10 on failure (at most five attempts). A fully independent exact-entry ball-arithmetic LDL on the unreduced sector (separate builder, no symmetry code) confirmed positive definiteness at the float-certified λ_c and failure at $\lambda_c + 10^{-6}$ in **three** cases only ($N = 3$ periodic $\ell_{\max} = 2$ $g = 1$; $N = 2$ open $\ell_{\max} = 4$ $g = 0.5$; $N = 3$ open $\ell_{\max} = 2$ $g = 2$; both Π sectors each).

A case is *certified* iff $\text{gap}_{\text{low}} = \min_s \lambda_{\text{low}}(s) - E_0^{\text{up}} > 0$ with every Cholesky check passed; no tolerance.

Grid. $g \in \{50, 8, 4, 2, 1, 0.7, 0.5, 0.3, 0.2, 0.1\}$; $\ell_{\max}$ up to 4 ($N = 1$), 14 ($N = 2$), 7 ($N = 3$), 5 ($N = 4$), 3 ($N = 5$); periodic and open; $\sigma = \pm 1$; dense sector dimension capped at 8000 (real) or 6000 (complex), larger cases reported as “cap” and not run.

Table 11: Smallest certified g as a function of the truncation $\ell_{\max}$ (periodic chain, $\sigma = +1$, floor F1). “—”: beyond the pre-registered $\ell_{\max}$ grid or the dense cap. With F0 alone the table is the same except $N = 3$, $\ell_{\max} = 1 \rightarrow 2.0$; $N = 5$, $\ell_{\max} = 2 \rightarrow 1.0$ and $\ell_{\max} = 3 \rightarrow 0.5$. Right block: smallest $\ell_{\max}$ certifying a positive gap, $\ell_{\text{needed}}(N, g)$.

$\ell_{\max}$	$N = 2$	$N = 3$	$N = 4$	$N = 5$	g	$N = 2$	$N = 3$	$N = 4$	$N = 5$
1	1.0	1.0	2.0	2.0	50, 8, 4	1	1	1	1
2	0.5	0.5	0.7	0.7	2	1	1	1	1
3	0.2	0.3	0.3	0.3	1	1	1	2	2
4	0.1	0.2	0.2	—	0.7	2	2	2	2
5	0.1	0.1	—	—	0.5	2	2	3	3
6–7	0.1	0.1	—	—	0.3	3	3	3	3
8–14	0.1	—	—	—	0.2	3	4	4	—
					0.1	4	5	—	—

Results. 1920 rows: **1604 certified**, 276 not certified, 40 over the cap. All 276 failures are of one kind: a finite but non-positive lower bound (Schur bound too loose). No positive-definiteness check failed at the final η . Controls: (P2) $N = 1$, $H = \frac{g}{2}L^2$: certified gap = g to 5.1×10^{-9} . (P1) $g = 50$: certified gap/ $g \in [0.9867, 1.0000]$ over all N , bc, σ . (P3) every certified interval contains the plain diagonalization gap at the largest $\ell_{\max}$ run (0 violations in 1604 rows). (P4) on bipartite chains $\sigma = -1$ reproduces the $\sigma = +1$ ground energy to 4×10^{-16} and the gaps to 4×10^{-12} (excitation moves from $k = 0$ to $k = \pi$, as the sublattice unitary predicts); frustrated rings differ ($N = 3$: E_0 by up to 1.06, excitation at $k = 2\pi/3$; $N = 5$: by 0.61, at $k = 4\pi/5$). (P6) F1 certifies cases F0 does not ($N = 3$ periodic $g = 1$ at $\ell_{\max} = 1$; $N = 4$ open $g = 1$ at $\ell_{\max} = 1$; $N = 5$ periodic $g = 0.7$ at $\ell_{\max} = 2$ and $g = 0.3$ at $\ell_{\max} = 3$); the gain is $\max_b d_1 - d_0 \approx B^2/(2(\ell_{\max} + 1)(\ell_{\max} + 2))$. Table 11 gives the smallest certified g per $\ell_{\max}$ and N (periodic, ferromagnetic, F1): a positive gap is *proven* for $g \geq 0.1$ at $N = 2, 3$, $g \geq 0.2$ at $N = 4$ and $g \geq 0.3$ at $N = 5$.

Both ends of the coupling range. At strong coupling the gap is N -independent and equals $g - \frac{2}{3}$ up to small corrections: one rotor is lifted to $\ell = 1$ and hops along the ring with first-order amplitude $\langle 10|n \cdot n'|01 \rangle = \frac{1}{3}$, giving a band bottom at $g - 2 \cdot \frac{1}{3}$ (deficit 0.67 at $g = 50$; at $g = 8$ the diagonalization gaps exceed $g - \frac{2}{3}$ by 0.047, 0.012, 0.010, 0.010 for $N = 2, \dots, 5$, at $g = 2$ by 0.181, 0.078, 0.052, 0.045). At weak coupling the finite chain rotates as one rigid body with moment of inertia N/g : the pre-registered prediction P7, $\text{gap} \approx g/N$, is met with $\text{gap} \cdot N/g = 1.047$ ($N = 2$, $g = 0.1$), 1.075 ($N = 3$, $g = 0.1$), [1.077, 1.212] ($N = 4$, $g = 0.2$); for $N = 5$ at $g = 0.3$ the certified interval [0.101, 2.088] is too wide to test (inconclusive). The lowest excitation always lies in the sector $(\Pi = -, F = +, k = 0)$, i.e. odd J . Truncation convergence is fast once the floor clears the spectrum: for $N = 3$ periodic $g = 1$ the certified gap interval is [0.0172, 0.9219] at $\ell_{\max} = 1$, [0.5253, 0.5405] at 2, [0.533612, 0.533697] at 3 and [0.533653, 0.533653] from $\ell_{\max} = 4$ on; for $g = 0.1$ the lower bound is $-1.46, -0.89, -0.43, -0.032$ at $\ell_{\max} = 1-4$ and becomes positive only at $\ell_{\max} = 5$ ([0.033854, 0.037369]), converging to [0.035829, 0.035831] at $\ell_{\max} = 7$.

First failure and cost wall. The first failure is $N = 4$ at $g = 0.1$: not certified at $\ell_{\max} = 4$ ($\text{gap}_{\text{low}} = -0.176$ with F1, while plain diagonalization gives 0.029), and $\ell_{\max} = 5$ has complex sectors of dimension $8.4-8.8 \times 10^3$, above the pre-registered cap. $N = 5$ fails at $g \leq 0.2$ ($\ell_{\max} = 3$, dimension 5.9×10^3 ; $\ell_{\max} = 4$ is estimated at 6×10^4 per sector) and is certified only barely at $g = 0.3$: $\text{gap} \in [0.006, 0.125]$ (diagonalization 0.077). The failure is never a broken inequality; it is a lower bound too loose to be positive because the floor clears the spectrum by too little ($d_M - E_0^{\text{up}} \leq 1.04$ in every failed row, ≥ 0.65 in every certified row) relative to $\|Wp\|^2/\text{gap}$. Pre-registered cost prediction P5 holds: ℓ_{needed} is non-decreasing as g decreases in every family and, from log-log fits over $g \leq 0.5$ (3–4 integer points each, *provisional*), grows as $g^{-0.40}$ ($N = 2$),




Figure 10: Rotor-chain certificate (periodic, $\sigma = +1$, F1). (a) Certified lower bound on the gap versus g for $N = 2$ –5; dashed: $\text{gap} = g$; dotted: $g/3$. Open circles: not certified (plain diagonalization value at the largest $\ell_{\max}$). (b) Smallest $\ell_{\max}$ with a certified positive gap; the dashed guide is $\propto g^{-1/4}$. Open triangles: not certified at any feasible $\ell_{\max}$.

$g^{-0.57}$ ($N = 3$), $g^{-0.30}$ ($N = 4$), pooled $g^{-0.39}$ over 13 points, at least as steep as the predicted $g^{-1/4}$; the largest sector dimension grows as $\ell_{\max}^{2N-2}/(2N)$ (381, 812, 1558 at $\ell_{\max} = 5, 6, 7$ for $N = 3$; 567, 2526, 8778 at $\ell_{\max} = 3, 4, 5$ for $N = 4$), so the dense cost $\propto \ell_{\max}^{6N-6}$ is exponential in N at fixed g and a power of $1/g$ at fixed N .

Not frustration-free. The finite-size criteria that convert a local gap into a gap uniform in the chain length (Knabe [24], Gosset–Mozgunov [25], the martingale method [26], Lemm–Mozgunov [27]) require a frustration-free Hamiltonian, $H = \sum_X h_X$ with $h_X \geq 0$ and a common ground state annihilated by every h_X . $H_N(g)$ is not frustration-free: with $h_b = \frac{g}{4}(L_i^2 + L_j^2) - n_i \cdot n_j$ one has $E_0 > \sum_b \inf h_b$ strictly, e.g. $N = 3$, $g = 1$: $E_0 \in [-0.59083, -0.59083]$ against $3e_2(\frac{1}{4}) = 3 \times (-0.27864) = -0.83593$. These criteria therefore cannot be fed with the finite- N certificates.

Deviations and arithmetic caveats. (i) The λ^* locator was switched from shift-invert LU to Lanczos after a stall on $N = 5$, $\ell_{\max} = 3$; λ^* is a search value only and every bound is still verified. (ii) The η retry schedule could not exceed the floating shift s in large-trace cases, so E_0^{low} came out undefined in 10 rows ($g = 50$ at the largest $\ell_{\max}$); the five affected cases were re-run with the patched schedule and the re-run rows replace the old ones. (iii) The $N = 5$ batch was interrupted once and restarted; duplicate rows were removed (last run kept). (iv) The tightness column ℓ_{90} of the methods file is post hoc. (v) **Arithmetic.** The verification is IEEE-754 double precision with an a-priori backward-error bound, not ball arithmetic; it relies on LAPACK `potrf` obeying the standard componentwise bound (valid for any summation order). The entry-error constants are deliberately generous but ε_A was not logged per row (spot value for an $N = 5$ complex sector of dimension 5866: $\varepsilon_A = 1.4 \times 10^{-11}$, $\delta = 2.9 \times 10^{-7}$, $s = 5.7 \times 10^{-7}$). The fully independent ball-arithmetic LDL was run on the three small cases listed above and on no other row.

Scope and the N -uniform target. Each certificate is a statement about $H_N(g)$ at one N and gives no bound on $\lim_{N \rightarrow \infty} \text{gap}$; the data show the finite- N gap closing like g/N at small g , which is the rigid-rotor tower of Section 4.2, not a mass gap. The analogue of “a certificate on a fixed support is not a certificate on all supports” applies verbatim. The one place where the certificate is already N -independent is strong coupling: for $g \geq 2$ every N is certified at $\ell_{\max} = 1$.

and the certified gap is $g - \frac{2}{3}$ plus a correction that *decreases* with N (0.181, 0.078, 0.052, 0.045 at $g = 2$). The natural next target was an $\ell_{\max} = 1$ Schur certificate whose floor and cross-term bound are uniform in N at $g \geq 2$, which would give a gap uniform in the chain length in the strong-coupling region where a gap is already known by cluster expansion. Version 2.1 shows that no such certificate exists in this class (Section 11). A second direction is a two-shell complement (exact treatment of $\max \ell = \ell_{\max} + 1$) to push the cost wall.

No N -uniform $\ell_{\max} = 1$ certificate (v2.1; review angle 4; closed, negative). Record `nuniform_cert.md`; code `nuniform_cert.py`; data `part2_rigorous.csv` (rigorous F0/F1 rows, $N = 2-10$, $g \in \{2, 3, 4, 8, 50\}$), `part2_oracle.csv`, `part2_qweights.csv`, `part2_Wnorm.json`, `part2_nstar.json`. Re-running the pipeline at $\ell_{\max} = 1$ reproduces the $N \leq 5$ rows of `rotor_cert_results.csv` to all printed digits (16/16) and extends them to $N = 10$; the result is that for every g the certified $\ell = 1$ gap goes negative at a finite $N^*(g)$, for the rigorous floors F0 and F1 and even for the best conceivable scalar floor (“oracle” $d = \inf \text{spec } QHQ$, a diagnostic that is not itself a valid bound). Two independent losses are responsible, both proved in structure and measured. (i) *Vacuum Q -dressing*. The certificate subtracts a lower bound for the excited sector, which contains the vacuum’s coupling to the discarded shells $-\|W\psi_0\|^2/(d - E_0)$, from an upper bound for the vacuum, which does not. By the exact variational lemma $\lambda_0(H_{\text{aug}}) \leq$ lowest eigenvalue of $\begin{pmatrix} E_0^P & \sqrt{w} \\ \sqrt{w} & d \end{pmatrix}$ with $w = \|W\psi_0^P\|^2$, the loss is $\delta_{\text{vac}} \geq \Phi(w, d - E_0^P)$, $\Phi(w, \Delta) = \frac{\Delta}{2}[\sqrt{1 + 4w/\Delta^2} - 1] \approx w/\Delta$, and $w = N q_{\text{vac}}(g)$ is *extensive*: $q_{\text{vac}} = 0.02168, 0.01020, 0.005857, 0.001494, 3.9 \times 10^{-5}$ at $g = 2, 3, 4, 8, 50$, i.e. $q_{\text{vac}}g^2 \rightarrow 0.0975$ (second-order perturbation theory; measured at $N \leq 9$, converged to four digits by $N = 6$). Since $d - E_0^P \leq 3g + O(1)$ for any admissible floor, δ_{vac} grows linearly in N while the gap of H_{aug} itself converges to an N -independent limit (3.3357 at $g = 4$, oracle floor), so $\text{gap}_{\text{low}} \rightarrow -\infty$. Measured δ_{vac} agrees with $Nq_{\text{vac}}/(d - E_0)$ to 3% ($N = 6$, $g = 4$: 0.00295 vs 0.00292). (ii) *Floor collapse (rigorous floors only)*. d_M is an absolute energy, F0 = $3g - N$ and F1 $\approx 3g - 2\sqrt{N}$ for $36 \lesssim N \leq 9g^2$ (from $e_2(b) = -1/(12b) + O(b^{-3})$, optimum $b^* = \sqrt{N}/12$), while $E_0 \approx -N/(6g)$; the clearance $d_M - E_0$ vanishes at $N \approx 3.1-3.3g^2$ (measured 16, 31, 53, 200 at $g = 2, 3, 4, 8$). *Cross term*. $W = QHP = \sum_b Qv_bP$ with exactly vanishing products over disjoint bonds, so $W^*W = \sum_b A_b + \sum_{b \sim b'} (C_{bb'} + C_{bb'}^*)$ with $\|A_b\| = 9/25$, $\|C\| = 2/9$, $\langle A_b \rangle = 74/225$ on $\otimes|1, 0\rangle$; hence $(74/225)N \leq \|W\|^2 \leq 0.804N$, measured $\|W\|^2 \approx 0.40N$ for $N = 3-8$: $\|QHP\| \propto \sqrt{N}$, not bounded. *Momentum resolution changes nothing*: a build with the translation group removed gives identical bounds ($N = 3$, $g = 2$: 1.26667 both ways), because $S(\lambda)$ is block-diagonal in k and both losses are k -blind. *Scale of the failure* (Table 12): with F1 at $g = 2$ the certified bound falls 1.367, 1.267, $\dots$, 0.388 for $N = 2, \dots, 10$ with an accelerating slope ($-0.19/\text{site}$ at $9 \rightarrow 10$), so the certificate ends at $N^* \approx 12$ while the P -space diagonalization gap is N -uniform ($1.643 \rightarrow 1.448$); for $g \geq 3$, $N^* \lesssim 3.3g^2$; with the oracle floor $N_{\text{or}}^* \approx 27g^4$ (381, 6.9×10^3 , 1.2×10^5 at $g = 2, 4, 8$). *Proposition (no N -uniform scalar-floor certificate at $\ell_{\max} = 1$)*. For any admissible floors $d_N \leq \inf \text{spec } QH_NQ$, $\text{gap}_{\text{low}}(N) \leq \text{gap}_{\text{aug}}(d_N) - \Phi(\|W\psi_0^P(N)\|^2, d_N - E_0^P(N))$; if $\|W\psi_0^P\|^2 \geq cN$ and gap_{aug} stays bounded, then $\text{gap}_{\text{low}}(N) < 0$ for all $N > N^*(g) < \infty$. Proof: the variational lemma, the Rayleigh principle in P , and $d_N - E_0^P \leq 3g + O(1)$ (one $\ell = 2$ site on the P ground state). Status: proved conditional on the two measured inputs (extensivity of the vacuum Q -weight, boundedness of gap_{aug}), both second-order perturbation-theory facts at large g and confirmed at every g of the grid. *What would restore uniformity* is not a better scalar or diagonal floor from local operator inequalities (their energy density is strictly below $e_0(g)$, so the clearance loss is extensive by construction) but a *relative* certificate that bounds $E_1 - E_0$ directly, with the Q -dressing cancelling between the two sectors inside one inequality (Feshbach/Schur relative to the dressed vacuum, or a Yarotsky-type perturbative gap-stability argument around $H_0 = \frac{g}{2} \sum L_i^2$ [38], whose g_0 would be explicit in principle but large). Deviations: the rigorous series stops at $N = 10$ (dense-Cholesky cap), the oracle diagnostic at $N = 7$ (from the $\ell \leq 3$

Table 12: Certified $\ell = 1$ gap lower bound at $\ell_{\max} = 1$ versus chain length (periodic ferromagnetic ring; floor F1 rigorous; “oracle” uses $d = \inf \text{spec } QHQ$ and is a diagnostic, not a valid bound; last column: diagonalization gap of the P -space at $g = 2$). From `part2_rigorous.csv`, `part2_oracle.csv`. Asterisk: oracle floor from the $\ell \leq 2$ space.

N	F1, $g=2$	F1, $g=3$	F1, $g=4$	F1, $g=8$	F1, $g=50$	oracle, $g=2$	oracle, $g=4$	oracle, $g=8$	P -ED, $g=2$
2	1.3673	2.3744	3.3711	7.3570	49.3377	1.4558	3.3947	7.3630	1.6427
3	1.2665	2.3107	3.3227	7.3314	49.3333	1.3710	3.3442	7.3362	1.4942
4	1.1886	2.2814	3.3073	7.3279	49.3332	1.3413	3.3355	7.3339	1.4604
5	1.1080	2.2583	3.2974	7.3263	49.3332	1.3296	3.3336	7.3336	1.4517
6	1.0090	2.2339	3.2878	7.3249	49.3332	1.3230	3.3328	7.3335	1.4493
7	0.8879	2.2059	3.2775	7.3235	49.3331	1.3191*	3.3323*	7.3335*	1.4486
8	0.7437	2.1718	3.2659	7.3221	49.3331	—	—	—	1.4484
9	0.5798	2.1334	3.2539	7.3207	49.3331	—	—	—	1.4484
10	0.3885	2.0869	3.2403	7.3194	49.3331	—	—	—	1.4484

space for $N \leq 6$, $\ell \leq 2$ at $N = 7$); the g^{-2} law of q_{vac} is read from the grid, not proved. *First test of the relative certificate (v2.2, [provisional], $N \leq 5$; rotor_rate.md Part B).* With the oracle floor and the $\ell_{\max} = 3$ spectrum as truth proxy, the dressing differences $\delta_1 - \delta_0$ ($\delta_M = E_M - E_M^{\text{aug}}$) are 0.01682, 0.01687, 0.01708 at $g = 4$ and 0.0267, 0.0280, 0.0294 at $g = 2$ for $N = 3, 4, 5$: nearly N -flat at $g = 4$, growing by 0.0014 per site at $g = 2$; δ_0 and δ_1 each grow with N , so the cancellation is partial and not yet shown to be uniform (gap_{aug} reproduces Table 12). Since both the true gap and gap_{aug} converge in N , $\delta_1 - \delta_0$ converges; the open question is a bound uniform in N , not the limit.

12 Relaxation dynamics: the gap as a forgetting rate

[pre-registered] (pre-registration §1 and addendum §1b of `langevin_methods.md`, both written before any target- β run; code `langevin_o3.py`; results `langevin_results.csv`, `langevin_local_forgetting.csv`; Fig. 11). The question tested is whether the mass gap of the classical model can be read as the rate at which spin orientations forget their initial values under a reversible stochastic dynamics with the Gibbs measure as its stationary state.

Dynamics. Overdamped Langevin diffusion on $(S^2)^V$ with unit diffusion constant, generator $\mathcal{L} = \sum_x (\Delta_x + \beta h_x \cdot \nabla_x)$, $h_x = \sum_{y \sim x} s_y$ (D1). The pre-registered Jacobi geodesic Euler–Maruyama integrator was found in validation to have a large equilibrium bias ($E = 0.3566/0.3715/0.3787$ and $\chi = 7.73/8.24/8.50$ at $dt = 0.04/0.02/0.01$ against 0.3857, 8.756 at $L = 8$, $\beta = 1$); its quadratic $dt \rightarrow 0$ extrapolation passes ($E = 0.38578(47)$, $\chi = 8.784(58)$) but a three- dt extrapolation at every target was not affordable, so the addendum replaced it, before any target run, by Metropolis-adjusted geodesic Langevin (MALA, checkerboard, one sweep = dt ; exact stationary measure at every dt , same generator as $dt \rightarrow 0$), $dt = 0.04$ (acceptance ≈ 0.99). Rate checks: $\gamma_M = 0.2091(35)$ vs 0.2099(34) at $dt = 0.04/0.02$ ($L = 8$, $\beta = 1$) and 0.0832(24) vs 0.0791(46) ($L = 32$, $\beta = 1.2$); $dt = 0.08$ is biased by +9% and excluded. A second, discrete-time dynamics is the exact single-site heat bath with checkerboard sweeps (D2; rates per sweep). Statics (E , χ , ξ_2) come from an independent Wolff code path. Validation: at $\beta = 0$ the site-averaged $\ell = 1$ and $\ell = 2$ rates are 2.0155(54) and 6.040(11) at $dt = 0.01$ (exact 2 and 6; within 1%); MALA at $L = 8$, $\beta = 1$ gives $E = 0.38585(13)$, $\chi = 8.757(17)$ against heat bath 0.38568(13), 8.748(14); at every target β the dynamic E and χ agree with Wolff (36 comparisons, $|\text{pull}| \leq 2.5$).

Observables and estimator. $M = \sum_x s_x$, $C_M(t) = \langle M(0) \cdot M(t) \rangle / \langle |M|^2 \rangle$; lowest mode $\phi_p = \sum_x s_x e^{ip \cdot x}$, $|p| = 2\pi/L$, $C_p(t)$ averaged over both axes; energy $C_E(t)$ (connected); geometric charge $C_Q(t)$; site-averaged $C_1(t) = \langle s_x(0) \cdot s_x(t) \rangle$ (the literal single-spin forgetting). Rates γ_χ

Table 13: Relaxation rates (MALA Langevin, $dt = 0.04$, time units; heat bath per sweep). ξ_2 and χ from Wolff. $R = (\gamma_p - \gamma_M)/(\gamma_M \hat{p}^2 \xi_2^2)$ tests the dispersion $\gamma_p = \kappa(\hat{p}^2 + m^2)$; $m_{\text{dyn}} \xi_2 = R^{-1/2}$. Last column: saturation of the exact bound (15). Rows $\beta = 1.4, 1.5$ are *provisional* (precision 6%, 8%, above the 5% target; wall caps hit).

β	L	ξ_2	χ	γ_M (Lang.)	γ_p (Lang.)	γ_M (HB)	R (Lang.)	R (HB)	γ_E/γ_M	$\gamma_M \chi/2$
0.6	16	0.727(23)	2.786	0.7128(77)	0.7691(47)	0.7224(59)	0.98(18)	1.34(15)	3.1	0.993
0.8	16	1.160(12)	4.711	0.4223(48)	0.5053(40)	0.3505(27)	0.96(8)	1.08(6)	3.4	0.995
1.0	16	1.798(11)	9.085	0.2158(44)	0.3151(25)	0.1439(20)	0.94(7)	1.03(5)	3.7	0.980
1.2	32	3.215(24)	22.34	0.0832(24)	0.1150(18)	0.0465(13)	0.96(12)	0.89(10)	4.7	0.929
1.4	64	6.770(65)	78.28	0.0194(12)	0.0308(8)	0.00876(53)	1.33(24)	1.14(22)	7.1	0.759
1.5	96	11.16(8)	177.9	0.00829(68)	0.01308(38)	0.00350(23)	1.08(26)	0.93(20)	8.6	0.738

are the slopes of weighted linear fits of $-\log C_X$ over the window $C_X \in [0.05, 0.6]$ ($[0.02, 0.2]$ for E), fixed in advance; jackknife over 24 blocks, 3 replicas. Targets $\beta \in \{0.6, 0.8, 1.0, 1.2, 1.4, 1.5\}$ at $L = 16, 16, 16, 32, 64, 96$ ($L/\xi_2 \geq 8.6$ everywhere); small- β scan $\beta \in \{0, 0.05, 0.125, 0.25, 0.4\}$ at $L = 16$.

An exact bound. [post hoc] (found during analysis). For the reversible generator $\mathcal{L}$, $-C'_M(0) = \langle M \cdot (-\mathcal{L}M) \rangle / \langle |M|^2 \rangle$, and the Dirichlet form gives $\langle M_a(-\mathcal{L}M_a) \rangle = \sum_x \langle |\nabla_x s_{x,a}|^2 \rangle$; summing over the three components, $\sum_a |\nabla s_a|^2 = 2$ on S^2 pointwise, so $\sum_a \langle M_a(-\mathcal{L}M_a) \rangle = 2V$ and

$$-C'_M(0) = \frac{2V}{\langle |M|^2 \rangle} = \frac{2}{\chi} \quad \text{exactly, at every } \beta \text{ and } L. \quad (15)$$

Since $C_M(t) = \int e^{-\lambda t} d\mu(\lambda)$ for a probability measure μ on the spectrum of $-\mathcal{L}$, Jensen gives $C_M(t) \geq e^{-2t/\chi}$, so the asymptotic decay rate γ_M of the magnetization (which is at least the spectral gap of $\mathcal{L}$) satisfies $\gamma_M \leq 2/\chi$. This is the dynamic analogue of a Rayleigh upper bound, and it ties the forgetting rate to a static quantity: a slow magnetization requires a large susceptibility. The only rigorous lower bound is Bakry–Émery [28]: $\text{Ric}(S^2) = 1$ and the Hessian of the interaction is $\geq -8\beta$, so the spectral gap is $\geq 1 - 8\beta$, non-trivial only for $\beta < 1/8$.

Results (Table 13).

- **The rate goes as the gap squared.** From $\beta = 1.0$ to 1.5 , ξ_2 grows from 1.80 to 11.2 and γ_M falls from 0.216 to 0.0083. H1 (dynamic scaling $\gamma_M = Am^z$, $m = 1/\xi_2$, over $\beta \in \{1.0, 1.2, 1.4, 1.5\}$, prediction $z \approx 2$): heat bath $z = 2.044 \pm 0.030$ **pass**; Langevin $z = 1.776 \pm 0.035$ ($\chi^2/\text{dof} = 5.5/2$; local slopes 1.64, 1.96, 1.70), within the registered 0.3 of 2 but 6σ from it, hence **inconclusive** by the registered rule. [post hoc]: γ_M tracks the exact value $2/\chi$ ($\gamma_M \propto \chi^{-1.10}$) and χ/ξ_2^2 drifts from 2.81 to 1.43 over the window (a logarithmic correction from asymptotic freedom), so a pure power in m is the wrong ansatz at these ξ and its fitted exponent depends on the dynamics.
- **The static mass is recoverable from the dynamics.** H2 (Ornstein–Zernike dispersion $\gamma_p = \kappa(\hat{p}^2 + m_{\text{dyn}}^2)$, $\hat{p}^2 = 4\sin^2(\pi/L)$, prediction $R = 1$, pass window $[0.7, 1.3]$ at $\beta \geq 1.2$): Langevin $R = 0.96, 1.33, 1.08$ and heat bath $R = 0.89, 1.14, 0.93$ at $\beta = 1.2, 1.4, 1.5$, all within the window at 1σ (the 1.33 central value just above) — **pass within errors**. Over $\beta = 0.6$ – 1.5 , $m_{\text{dyn}} \xi_2 = R^{-1/2} = 1.01, 1.02, 1.03, 1.02, 0.87, 0.96$ (Langevin) and $0.86, 0.96, 0.99, 1.06, 0.93, 1.04$ (heat bath), i.e. the mass read off the dynamics as $m = \sqrt{\gamma_0/\kappa}$ agrees with the static $1/\xi_2$ to within 14%. (The methods file summarizes both dynamics as 0.87–1.03; the heat-bath values recomputed from the results file span 0.86–1.06.)
- **Saturation of the exact bound.** $\gamma_M \chi/2 = 0.993, 0.995, 0.980$ at $\beta = 0.6, 0.8, 1.0$ and $1.025, 1.031, 1.002, 0.996, 1.025$ in the small- β scan (a +1.3% bias of order dt is expected there), so up to $\beta = 1$ the magnetization decays as a single exponential at the Rayleigh




Figure 11: (a) Relaxation rate γ_M versus the static mass $1/\xi_2$ for Langevin (time units) and heat bath (sweeps), with the z fits over $\beta \in \{1.0, 1.2, 1.4, 1.5\}$. (b) γ_M versus β with the exact upper bound $2/\chi$ and the Bakry-Émery lower bound $1 - 8\beta$; the shaded band is the only region where a lower bound is rigorous.

rate; at $\beta = 1.2, 1.4, 1.5$ the saturation drops to 0.929, 0.759, 0.738. The last two values rest on the provisional rows; the departure from saturation at $\beta = 1.2$ is 2.7σ and is not explained here.

- **H3 failed.** The singlet (energy) channel was predicted to relax at the two-particle threshold $2\gamma_M$ (consistency window $[1.6, 2.4]$); measured $\gamma_E/\gamma_M = 3.1, 3.4, 3.7, 4.7, 7.1, 8.6$ (Langevin; heat bath 3.8, 5.2, 7.2, 9.2 at $\beta \geq 1.0$), i.e. energy forgets 3–9 times faster than the two-mode tail would, through short-wavelength pairs; the asymptotic $2\gamma_M$ tail has too small an amplitude to resolve in the window.
- **Bakry-Émery region is a limit of the technique, not of the physics (H4).** Every measured $\gamma_M \geq \max(0, 1 - 8\beta)$ (sanity pass). The rate falls smoothly from the free-rotor value 2: $\gamma_M = 2.049, 1.914, 1.695, 1.375, 1.097$ at $\beta = 0, 0.05, 0.125, 0.25, 0.4$, with no feature at $\beta = 1/8$, and stays $O(1)$ to $\beta \approx 0.6$; it collapses ($\propto 1/\chi$) only where $\xi \gg 1$.
- **Topology carries no slow relaxation (H5, [exploratory]).** The geometric charge is integer to 10^{-11} and changes in 83–94% of measurement intervals at every β ; $\tau_{\text{int}}(Q) \leq 0.14$ ($\beta \leq 1.0$), 0.30, 0.68, 1.27 ($\beta = 1.2, 1.4, 1.5$; at or below the measurement spacing, so upper bounds) against $\tau_{\text{int}}(M) = 4.6, 11.9, 48.8, 118$. The direction $\hat{M}$ relaxes at $\approx \gamma_M$ and $|M|^2$ at $\approx 2\gamma_M$ (Gaussian M): the slow mode is the long-wavelength rotation of the magnetization. This is the dynamic counterpart of Section 4.1.
- **Single-spin forgetting.** The parameter-free Ornstein-Zernike prediction $C_1(t) = V^{-1} \sum_p S(p) e^{-\gamma_p t}$, built from the measured χ , ξ_2 , γ_M , matches the measured C_1 at $t = 1/\gamma_M$ and $3/\gamma_M$ within 3–20%. The fraction of a single spin's orientation memory held in modes with $\hat{p}\xi < 1$ is 0.33, 0.25, 0.20, 0.13, 0.11, 0.07 ($\beta = 0.6$ –1.5); the rest is forgotten in $O(1)$ time; the late-time rate of tC_1 is 0.75–0.93 γ_M .

Deviations. MALA replaced the pre-registered integrator (declared in the addendum before target runs). The full three- dt extrapolation was done only at $L = 8$, $\beta = 1$; elsewhere $dt = 0.04$ was checked against 0.02 at two points. The pre-registered M -based $\beta = 0$ estimator was underpowered ($\pm 2.4\%$ per dt against a 2% criterion; 1.99(5)/1.98(4)/1.96(5)) and the site-averaged estimator was added in the addendum. γ_M precision at $\beta = 1.4/1.5$ is 6%/8% (wall caps of 40 and 45 minutes hit), above the 5% stopping target, so those rows and everything derived

from them (the z fits, the saturation values 0.759, 0.738, R at $\beta \geq 1.4$) are provisional; the z fit uses inverse-variance weights so it is not dominated by them.

What this adds and does not add. The reading “the gap is the rate at which spin angles forget their start” is correct up to the square: for a reversible dynamics with non-conserved order parameter (model A in the classification of [29]) the slow rates are $\kappa(p^2 + m^2)$ and the $p = 0$ rate is κm^2 , so the static mass is $\sqrt{\gamma_0/\kappa}$ and the single rigorous dynamic inequality available, $\gamma_M \leq 2/\chi$, bounds the rate by the susceptibility rather than the other way round. No statement about the infinite-volume gap follows: a dynamic lower bound uniform in L at $\beta > 1/8$ would be a proof of the gap in disguise and none is claimed.

13 Audit of the October 2026 machine-generated proofs

Source: github.com/openai/math, branch `main`; all family-215 files unchanged since the initial public commit `adc7f1241b` (2026-10-06T21:58:50Z) and untouched by the 7 October fixes and withdrawals (which concerned the Hodge/Weil families). Status of every item: unrefereed, AI-generated preprint. Full reports: `o3_t2_openai_audit.md`, `v8_halftrate_audit.md`; numerical checks `o3_t2_checks.py`, `v8_checks.py`.

13.1 Inventory

The release contains no lattice gauge theory or Yang–Mills manuscript (keyword search of `CONTENTS.md`); every sigma-model item is in family 215, “Canonical $O(3)$ continuum limit and exact $O(4)$ mass asymptotics”:

- **215-a.** *Exponential decay in two-dimensional classical $O(n)$ models* (23 Sep 2026). Lean: `lean/docs/215.md` $\rightarrow$ `ComparatorChallenges/ClassicalON.lean`, declaration `OAI.ClassicalON.main`.
- **215-b.** *Sharp mass bounds for the two-dimensional $O(4)$ model* (23 Sep 2026). No Lean. Full transfer gap $\asymp \sqrt{\beta}e^{-\pi\beta}$ for $\beta \geq \beta_0$; Corollary: gap $\geq c \exp[-\pi\beta - C\sqrt{(1+\beta)\log(2+\beta)}]$ at every β .
- **215-c.** *The canonical massive continuum limit of the two-dimensional $O(3)$ model* (4 Oct 2026). No Lean. OS axioms, unique vacuum, gap, non-Gaussian, along $\beta_N = H + (\log L/2\pi)N + O(\log N)$.
- **215-d.** *An Isolated Particle Pole for the Two-Dimensional $O(3)$ Spin Field* (4 Oct 2026). No Lean. Isolated atom in the Källén–Lehmann measure.
- **215-e.** *Exact mass asymptotics for the two-dimensional $O(4)$ lattice model* (5 Oct 2026). No Lean. $\lim e^{\pi\beta}\beta^{-1/2}m_{\text{lat}}(\beta) = 32e^{\pi/4-1/2}$.

Dependency graph from the bibliographies: $215\text{-a} \rightarrow 215\text{-b} \rightarrow 215\text{-c} \rightarrow 215\text{-d} \rightarrow 215\text{-e}$; 215-c imports estimates of 215-b through 358 cross-reference locators and proves several key lemmas as “the changes to” $O(4)$ proofs. Related: family 216 (five planar-XY papers, the BKT control family), 268 (spin-one Haldane gap, two papers, no Lean), 270 ($SU(N)$ BFSS matrix models). Of these only the analytic bootstrap of family 268 (`build/bootstrap.tex`) was read in full, for Section 7; its computer-assisted seed estimates were not audited and are not used. A non-OpenAI Zenodo paper (10.5281/zenodo.23349319) builds an $O(4)$ continuum limit explicitly conditional on 215-b and 215-c.

13.2 215-a: no error found; Lean compiled

The full proof (Sections 2–4, about 25 pages) was read step by step. Theorem 1.1: for all $n \geq 3$, $\beta > 0$ there are $A(n, \beta) < \infty$, $m(n, \beta) > 0$ with $0 \leq \langle \sigma_x \cdot \sigma_y \rangle \leq Ae^{-m\|x-y\|_2}$ on every finite nearest-neighbour subgraph of $\mathbb{Z}^2$ with free boundary and edge strengths in $[0, \beta]$. Architecture: an abelian gauge identity (Lemma 2.1); a fourth-order non-abelian rotation inequality $R(Q, \bar{Q}) \leq K_{\text{rot}}\|T\|_{\text{HS}}^2 + \beta \sum \min\{\dots\}$, $K_{\text{rot}} = 6\beta^2 + 4\beta^{3/2} + \beta$ (Prop. 2.2, from $W(dF) = W(B)$ with $B_e = \frac{1}{2}[F_x, F_y] + O(F^3)$ and Faà di Bruno); a vanishing annular twist cost $\min Z / \max Z \geq 1 - C_\beta / \log N$ on $\{N \leq \|x\|_\infty \leq 2N\}$ (Prop. 3.1, multiscale Fourier family with $\lambda_N \geq \frac{1}{2}H_N$); Wolff/Edwards–Sokal sign clusters with the exact identity $P_{\text{pinned}}(\text{crossing}) = 1 - Z(\pi)/Z(0)$, pinning monotonicity and a coarse-path union bound with 25 residue classes ($p_N < 4^{-25}$ needed); reduction $n \rightarrow 3$ by spherical disintegration. Numerical checks: the fourth-order pairing identity holds to $\leq 1.4 \times 10^{-14}$ at $\beta = 0.5, 2.0$ with random *complex* test functions (exact Gauss quadrature on $S^2 \times S^2$); $R(h, df) = 0$ to 3×10^{-17} ; the multiscale identity $\sum S(f_i, \bar{f}_i) = -i\lambda_N \tilde{h}$ to 2×10^{-14} with $\lambda_N \geq \frac{1}{2}H_N$ ($N = 24$: 4.435 vs 1.888) and $(\|T\|^2 + \text{mincost})/\lambda_N^2$ decreasing from 2.49×10^5 to 1.82×10^5 for $N = 4 \dots 24$. The proof genuinely uses $n \geq 3$ (Prop. 2.2 needs two generators with $[a, b] \neq 0$; for $n = 2$, $B \equiv 0$ and no response is produced, consistent with Fröhlich–Spencer) and $d = 2$ ($\|T\|_{\text{HS}}^2 \sim N^2$ against $R \sim N$ in $d = 3$). It does not prove too much: the implied rate is $m_0 \sim 1/N$ with $\log N \gtrsim 10^{20}\beta^2$, far below the physical $e^{-2\pi\beta}$. Our own crossing-probability Monte Carlo (pinned annuli, $\beta = 1.0, 1.25, 1.5$) gives half-points at $N \approx 5\text{--}6\xi$ and requires $C_\beta \geq 1.77, 2.46, 3.45$; no contradiction with $C_\beta \propto \beta^2$. The flag on eq. (2.14) was withdrawn (Section 3.3).

Lean. 117 files, 12,907 lines; `grep` finds no `sorry`, `axiom`, `admit`, `native_decide`, `implemented_by` or `extern`. The comparator statement `ExponentialDecay` encodes Theorem 1.1 faithfully (finite `LatticeGraph` in $\mathbb{Z}^2$, spins on the unit sphere with normalized measure, weight $\exp \sum_e b_e \langle \sigma_x, \sigma_y \rangle$, Euclidean `siteDistance`); the vacuity check passes ($Z \neq 0$ since the measure is finite positive and the integrand positive). The formalization was **compiled locally** (Lean v4.34.1, Mathlib d13f23b, the 117 ClassicalON files plus the comparator challenge file; 9042 jobs, “Build completed successfully”, 20 min); `#print axioms OAI.ClassicalON.main` gives `[propext, Classical.choice, Quot.sound]`; all 11 statement-level definitions in `Model.lean` are textually identical to the comparator file; the only `sorry` warning is the challenge stub itself. A `leanchecker` replay of the `Main` module completed with no errors (1.6 min); a full replay of all 117 modules was stopped after 30 min without finishing and is *not* reported as done. OpenAI’s Comparator binary and nanoda were not run. *Not formalized:* Corollary 1.2 (subsequential limits), infinite-volume uniqueness, any transfer gap, the rate, and all of 215-b to 215-e.

13.3 215-b/215-c: the half-rate resolution

Both proofs lower-bound the current variance on a pinned square by one Cauchy–Schwarz step, $\text{Var } J(A) \geq \text{Cov}(J(A), S)^2 / ES^2$, with S the score of slowly varying rotations about a single colour axis a . With $T = \sum_{M \leq |k| \leq K_1} |k|^{-2} = \log(K_1/M)/(2\pi)$ the three inputs are $\sum_i EY_i^2 = vI_\chi T/2$, $ES^2 \approx v^2 I_\chi T/2$ and $\text{Cov} = cvI_\chi T/2$, where c is the coefficient of the *midpoint* rotation angle in the bracket term of the bond-current derivative; the gain is $c^2 T/2$. For $O(4)$ with spins as unit quaternions rotated by left multiplication $q \mapsto e^{a\eta} q$, $D_\eta s_e = t_e(\eta_y - \eta_x) + (\eta_y + \eta_x) \times s_e$ (215-b eq. pre:1-4), so $c = 2$ and $\text{Var } J(A) \geq 2T = \log L/\pi$: the full one-loop $O(4)$ rate. For $O(3)$ with the vector action $R_a(u)$ on S^2 , $D_u s_e = E_e(u_y - u_x) + \frac{1}{2}(u_y + u_x) \times s_e$ (215-c eq. prelim-bond-derivative), so $c = 1$ and $\text{Var } J(A) \geq T/2 = \log \ell/(4\pi)$, half the one-loop $O(3)$ rate. Algebraically c is the structure constant of the colour algebra in the normalization where the generator rotates the spin at unit speed: $[L_a, L_b] = 2L_{a \times b}$ for left multiplication on S^3 , $[L_a, L_b] = L_{a \times b}$ for the vector action on S^2 , and $c = 0$ for $O(2)$ (one axis, no gain). Both bond-derivative formulas were verified by finite differences (residuals $\leq 2 \times 10^{-10}$ over 200 random trials; swapping the coefficients gives residuals of order 1); T and $\sum |k|^{-2} \hat{k} \hat{k}^T = (T/2)I$ were verified numerically.

Neither bound exceeds the one-loop value $(N-2)T$, as a Cauchy–Schwarz lower bound must not. The asymmetry is therefore **explained and not an error**: 215-c states the factor $\frac{1}{2}$ explicitly and compensates with an independent “Initial reduction” lemma that makes the total exponent strictly less than 4π ; its sharp rates come from the RG, so the half rate costs nothing there provided $\eta > 0$. In 215-b the full rate is *load-bearing* (the box comparison in `assembly.tex` needs $\log \Xi(\beta_K) = K \log L + o(K)$), a design dependence rather than an error. Our square-torus drops (0.123–0.24 per doubling for $\rho_s \geq 0.3$) exceed the lemma’s leading-order requirement $\log 2/(4\pi) = 0.0552$ by a factor 2.2–4.4 and the one-loop 0.110 by 1.1–2.2; this is a plausibility check, not a test of the lemma (pinned squares with signed couplings and $h \geq A \log B$ versus periodic boxes at $\rho_s \leq 1.4$).

13.4 Unverified links

(a) The analytic error body shared by 215-b and 215-c (band estimates, weighted current estimate, crossed term) and the Ward-identity error terms; (b) the Gaussian-augmentation hypothesis $C_c = vI - H_c - c_0 P_a \geq B^{-D'} I$; (c) the one-site Initial-reduction lemma of 215-c (read, not re-derived); (d) the $O(3)$ -specific current identities and Gaussian normalization in 215-c, proved as modifications of $O(4)$ proofs; (e) the Brownian-regulator $\leftrightarrow$ Wilson partition-function comparison of 215-e (the step flagged publicly by P. Orland), which must control absolute differences with a complex Brownian density to an error smaller than an exponentially small odd trace; its headline constant $32e^{\pi/4-1/2} = 42.56933$ equals $(m/\Lambda_{\overline{\text{MS}}})(\Lambda_{\overline{\text{MS}}}/\Lambda_L)\sqrt{\pi}$ from [2] exactly, and the finite coupling shift $\beta - b = \frac{1}{4} - (1 + \log 2)/(2\pi) = -0.0194727$ was verified to 4×10^{-11} ; (f) 215-d, conditional on 215-c, not refereed; (g) the full Lean replay of 215-a (see above). First failure point overall: none located. The one result with both a readable proof and a complete kernel-checked formalization is 215-a, and what it proves is exponential decay of the spin two-point function at fixed β with non-explicit, doubly-exponentially weak constants: not uniqueness of the Gibbs state, not a transfer gap, not the rate, not the continuum.

14 Discussion: relation to the Yang–Mills gap

In the continuum-time (anisotropic, Hamiltonian) limit the lattice $O(3)$ model becomes a chain of coupled quantum rotors, $H = \frac{g}{2} \sum_x L_x^2 - \frac{1}{g} \sum_x n_x \cdot n_{x+1}$ (the chain certified at finite length in Section 11); the exact isotropic lattice Hamiltonian is $-\log \mathcal{T}$, of which this is the leading form; the Kogut–Susskind Hamiltonian of lattice gauge theory [20] has the same two-term structure, an electric Casimir term $\frac{g^2}{2} \sum E^2$ on links and a magnetic plaquette term $\frac{1}{g^2} \sum \text{Re Tr } U_\square$. The objects of this paper have the following counterparts. The square-torus identity (5) decomposes the twisted partition function over isospin sectors of the circle Hamiltonian; in Yang–Mills the twisted partition functions are ’t Hooft’s twisted-boundary (’t Hooft flux) sectors, and a loose analogue of $\rho_s = \langle C_2 \rangle / (N-1)$ is an electric-flux energy averaged over centre sectors of a box; this is an analogy only, since centre sectors do not determine quadratic Casimir energies. The finite-volume helicity coupling $g_Y^2(L) = 1/\rho_s(L)$ [31] is the analogue of the Schrödinger-functional or twisted-boundary running coupling. The per-doubling map $\rho_s \mapsto \rho_s - c(\rho_s)$ is the analogue of a step-scaling function $\sigma(u)$ for the gauge coupling, which lattice gauge theory measures but does not prove. The $O(2)$ control, where the stiffness stays finite, corresponds to compact $U(1)$ in four dimensions, which has a Coulomb phase; the $(N-2)$ scaling of c corresponds to the group-theoretic factor in the one-loop gauge β -function.

What would transfer: the positivity and character structure of twisted partition functions (Osterwalder–Seiler reflection positivity and transfer matrix hold for lattice gauge theory); the formulation of the missing inequality as a β -independent, strictly decreasing one-variable map on a compact band of couplings, and the requirement that it fail for the abelian control; the certificate template of Section 4.7 (finite block, rigorously priced complement, interval-controlled computa-

tion) as a way to extend a strong-coupling region. What would not transfer: the Wolff embedding and the Campbell–Chayes association, which make the $O(3)$ problem a random-coupling Ising problem and underlie both the certificate and the sign-cluster part of 215-a; the Aizenman–Simon/Simon–Lieb inequality (no GKS-type inequalities for gauge theory); the λ -deformation control, which has no analogue once the gauge group is fixed; and the two-dimensional conformal modulus $\ln R$ that makes annulus twist bounds vanish. The spin-1 antiferromagnetic Heisenberg chain, the $\theta = 2\pi$ Hamiltonian regularization of the $O(3)$ model, has a claimed machine-generated gap proof in family 268 that builds on Tasaki’s conditional theorem [18]; a gap at fixed spin says nothing about the $g \rightarrow 0$ scaling of the rotor chain and does not transfer to the classical model on $\mathbb{Z}^2$. Its finite *criterion*, however, does transfer: the two-purity bootstrap needs only a positive transfer representation, which lattice gauge theory also has (Osterwalder–Seiler), so the volume-uniform certificate of Section 7 is a template whose gauge analogue would be a seed computed on small tori; what would not transfer is the character network used to certify the seed. The gauge problem also carries gauge invariance and a four-dimensional band in which the running is slower per doubling. The warm-up therefore delivers a template and a list of controls, not a method that carries over by itself.

15 Data and code availability

All data, code and records are deposited with this preprint under CC-BY-4.0, as a versioned archive (one DOI per version) so that every file named below is retrievable at the revision used here; version 1 named the files without a retrievable location (referee item E11). Version 2.0 adds `p0_brackets_v2.csv` (Table 4), the covariance-aware refit (Section 4.6, reproducible from `casimir_map_results.csv`) and the errata file `errata_v1_0_from_external_referee_2026-10-19.md`. Versions 2.1 and 2.2 add the files of Appendix C.

- Section 4.1: `o3_sphere_map.py`, `o3_sphere_map_results.csv`, `o3_sphere_map.png`.
- Section 4.2: `o3_rotor_box.csv`, `o3_rotor_box.png`.
- Section 4.3: `o3_sim_methods.md`, `o3_sim_results.csv`, `o3_xi_vs_beta.png`, `o3_step_scaling.csv`, `o3_step_scaling.png`.
- Section 4.4: `o3_realign.py`, `o3_realign_blocks.csv`, `o3_realign_map.csv`, `o3_realign.png`.
- Section 4.5: `o3_axis_spin_ON.py`, `o3_axis_spin_ON_blocks.csv`, `o3_axis_spin_ON_growth.csv`, `o3_axis_spin_ON.png`.
- Section 4.6: `casimir_map.py`, `casimir_map_methods.md`, `casimir_map_results.csv`, `casimir_map.png`.
- Section 4.7: `decay077_certificate.py`, `decay077_replay.py`, `decay077_methods.md`, `decay077_preregistration.json`, `decay077_results.json`, `coverage_summary.json`.
- Section 5: `stiffness_theory.md`, `twist_scan.png`.
- Section 10 (v2.0): `character_contraction.md`, `character_contraction.py`, `character_contraction_check.csv`. The second-action run announced in v2.0 is reported in Section 4.6 (v2.1).
- Versions 2.1 and 2.2 (Sections 4.6, 7, 8, 9, 10.1, 10.2, 11): see Appendix C, which lists every new file with its content.
- Section 11 (v1.1): `rotor_cert.py`, `rotor_cert_methods.md` (pre-registration in §0), `rotor_cert_results.csv`, `rotor_cert_sectors.csv`, `rotor_cert_gap_vs_g.png`.
- Section 12 (v1.1): `langevin_o3.py`, `langevin_methods.md` (pre-registration §1, addendum §1b), `langevin_results.csv`, `langevin_local_forgetting.csv`, `langevin_rate_vs_mass.png`.

- Literature and route surveys: `o3_t1_literature_status.md`, `o3_t4_routes_to_failure.md`.
- Section 13: `o3_t2_openai_audit.md`, `v8_halfrate_audit.md`.
- Research record: `decode_spinning_ball_cue_2026-10-19.md`, `mechanism_concept_03_v0.1.md`, `attack_plan_03_gap_2026-10-19.md`.

AI-assistance statement

AI assistance was used throughout the research and preparation of this manuscript, including conceptual development, mathematical derivations, simulations, and exposition. Claude (Claude Science, model Claude Fable 5.1, developed by Anthropic) performed the simulations, the stiffness–Casimir derivation, the rotor-chain certificate, the relaxation-dynamics study and the audit of the machine-generated proofs; the version-2.1 and 2.2 tracks (purity bootstrap and seed certification to $\beta = 0.85$, doubling-defect seed and dichotomy, Laplace bounds, rotor semiclassics, transfer-matrix and DMRG tests, rotor-ring and transfer lemmas, N -uniformity analysis, second-action run) were likewise AI-executed from the author’s direction, with the lemma proofs written out in the records named in Appendix C; the author supplied the cues that directed the seed-aspect scan and the $\beta \rightarrow +\infty$ attack. The author directed the research, contributed the motivating intuitions, connected ideas, and coordinated AI-assisted development. The version-2.0 corrections follow an external referee report whose points were verified by recomputation before being applied; the version-2.1 response (Appendix B) addresses a second, informal review of version 2.0. The author takes full responsibility for the manuscript, including its mathematical correctness, computational results, and references.

A The decay certificate: theorem, lemma chain and computation

Setting. Λ is a finite graph, either the torus $\mathbb{T}_N = (\mathbb{Z}/N)^2$ with $N \geq 8$ or a box in $\mathbb{Z}^2$ with free boundary; $\mu_\Lambda \propto \exp(\beta \sum_{xy \in E} s_x \cdot s_y) \prod_x d\sigma(s_x)$ with σ the uniform probability measure on S^2 and $\beta > 0$; $G_1(x, y) = \langle s_x^1 s_y^1 \rangle_\Lambda = \langle s_x \cdot s_y \rangle_\Lambda / 3$; d is graph distance; $A_y = y + [-1, 1]^2$. The pinned 7×7 box at y is the measure on the 25 spins of $y + [-2, 2]^2$ with the ring $B_y = \{x : |x - y|_\infty = 3\}$ fixed to e_1 and no other spins.

Theorem 6 (conditional exponential decay). *Assume Lemmas 2 and 3 below. Let $\beta \in [0, 0.7701]$, or $\beta \in \{0.78, 0.785\}$. Then there is $\varphi_{\text{up}}(\beta) < 1$ (Tables 6 and 14; $\varphi_{\text{up}}(0.77) = 0.915117$) such that on every torus $\mathbb{T}_N$, $N \geq 8$,*

$$0 \leq \langle s_0 \cdot s_x \rangle \leq \varphi_{\text{up}}(\beta)^{\lfloor d(0, x)/3 \rfloor} \quad \text{for all } x,$$

and the same bound holds in a free-boundary box whenever $\text{dist}_\infty(0, \Lambda^c) \geq 3\lfloor d(0, x)/3 \rfloor + 4$. Every periodic or free infinite-volume limit point inherits the bound; at $\beta = 0.77$ the mass per lattice step is at least 0.02956. Uniqueness of the Gibbs state is neither used nor claimed.

Lemma chain. Each lemma is stated with its hypotheses and its status.

Lemma 0 (Ginibre 1970 [33]). For ferromagnetic $O(N)$ models on a finite graph with free or periodic boundary, $0 \leq G_1 \leq 1/3$. *Classical.*

Lemma 1 (Aizenman–Simon 1980 [8]). For finite $A \subset \Lambda$, $i \in A$, $j \notin A$: $G_1(i, j) \leq \beta \sum_{l \in A, k \notin A, l \sim k} G_1(i, l) G_1(k, j)$, both factors in the same measure μ_Λ (not Lieb’s improved form). Hypotheses: pair ferromagnetic isotropic couplings, free or periodic boundary, $N = 3$ (the proof uses the $N = 3, 4$ inequality $\langle \sigma_i^2 \sigma_l^2 \sigma_j^1 \sigma_k^1 \rangle \leq \langle \sigma_i^2 \sigma_l^2 \rangle \langle \sigma_j^1 \sigma_k^1 \rangle$ of Kunz–Pfister–Vuillermot and Dunlop in place of GKS-II). With $A = A_i$: $\varphi = \beta[4G_1(i, i+e) + 8G_1(i, i+d)]$.

Published; checked against the reproduction in [13], App. A.2, Thm. A.3 at $\lambda = 1$; the original was not accessible.

Lemma 2 (Wolff embedding and amplitude association). Write $s = (r\tau, q \cos \theta, q \sin \theta)$, $r \in [0, 1]$, $q = \sqrt{1 - r^2}$, $\tau = \pm 1$. Under σ , $r \sim U[0, 1]$ and (τ, θ) are uniform and independent. Given r , τ is an Ising model with couplings $\beta r_x r_y \geq 0$, independent of θ ; let n be its Edwards–Sokal bond variables. Then $G_1(x, y) = E[r_x r_y \mathbf{1}\{x \leftrightarrow y\}]$ and the joint law of (r, n) is positively associated. *Ingredients from Campbell–Chayes [30]: the amplitude-times-connectivity representation of the correlation and the FKG property of the amplitude variables (the r -marginal is MTP_2 because $\partial^2 \log / \partial r_x \partial r_z \geq 0$ by GKS-I/II for the Ising factor and Ginibre-I/II for the XY factor). Assembled here: the joint association of (r, n) from MTP_2 conditionals plus FK monotonicity in r . Hypotheses: ferromagnetic nearest-neighbour couplings, free or periodic boundary, $N = 3$.*

Lemma 3 (pinning comparison). If $y + [-3, 3]^2$ embeds in Λ and $l \in A_y$, then $G_{1,\Lambda}(y, l) \leq \langle s_y^1 s_l^1 \rangle^{\text{pin}}$, the pinned 7×7 box at y . *Steps: (A) $E[F] \leq E[F \mid r_{B_y} = 1]$ for $F = r_y r_l \mathbf{1}\{y \leftrightarrow l\}$, increasing in (r, n) , because $E[F \mid r_B = \rho]$ is increasing in ρ (Lemma 2); (B) given $r_B = 1$, fixing $\tau_B = +$ multiplies the (r, n) density by $2^{-k_B(n)}$, increasing in n , so FKG gives $E^{\text{pin}}[F] \geq E[F \mid r_B = 1]$ — this is the wired-boundary comparison of [30]; (C) given $s_B = e_1$ the interior is the pinned box (Markov property) and $\langle s_y^1 s_l^1 \rangle^{\text{pin}} \geq E^{\text{pin}}[F]$. Hypotheses as in Lemma 2. Status: assembled here from [30] and FKG; this is the step a referee should check first.*

Lemma T (kernel truncation). With $k_L(x) = \sum_{l \leq L} (2l + 1) i_l(\beta) P_l(x)$, $T_L = e^\beta - \sum_{l \leq L} (2l + 1) i_l(\beta)$ and $\varepsilon = T_L e^\beta$: $|k - k_L| \leq T_L$, $k \geq e^{-\beta}$, hence the ratio of Boltzmann weights lies in $[(1 - \varepsilon)^{60}, (1 + \varepsilon)^{60}]$ (60 bonds), and for $|F| \leq 1$, $|\langle F \rangle - \langle F \rangle_L| \leq \Delta = \kappa(1 + c_L)/(1 - \kappa)$ with $\kappa = ((1 + \varepsilon)^{60} - (1 - \varepsilon)^{60})/(2(1 - \varepsilon)^{60})$. At $\beta = 0.77$, $L = 6$: $\varepsilon = 2.750 \times 10^{-6}$, $\kappa = 1.650 \times 10^{-4}$, $\Delta_e = 1.85 \times 10^{-4}$, $\Delta_d = 1.79 \times 10^{-4}$. *Elementary; $i_l(\beta)$ and T_L as 256-bit arb enclosures.*

Lemma 4 (certified computation). The truncated 7×7 pinned model is evaluated exactly by a site-by-site transfer in the y_{lm} basis (axis e_1): horizontal integration is diagonal (Funk–Hecke), vertical bonds and ring fields are Gaunt multiplications with coefficients exact as $\text{sign} \cdot \sqrt{\text{rational}}$ (Racah formula, integer arithmetic), degree projections exact by the Gaunt selection rule, $M = 0$ sector only (6.2×10^7 slots at $L = 6$). *Arithmetic error model: IEEE-754 round-to-nearest, classical (non-Strassen) BLAS; every product carries a componentwise a-priori bound $E \geq |\hat{x} - x|$ propagated as $E' = [E, |\hat{X}|] \cdot [B_{\text{maj}}, \gamma_n B_{\text{maj}} + e_B]^T (1 + 3\gamma_n) + \text{guard}$ (γ_n majorant, valid for any summation order and with FMA); all bound terms are non-negative so their float evaluation is itself controlled; the resulting relative float error is $\approx 3 \times 10^{-9}$ on Z and 10^{-11} on correlators, negligible against Δ . Certified intervals: $c_e \in [0.122867, 0.123238]$, $c_d \in [0.086580, 0.086939]$, $\varphi_{\text{up}} = 0.77(4c_e^{\text{up}} + 8c_d^{\text{up}}) = 0.915117$ (outward rounding).*

Lemma 5 (iteration; Simon 1980 [34]). On $\mathbb{T}_N$ let $f(n) = \max\{G_1(y, x) : d(y, x) \geq n\}$. For $d(y, x) \geq 3$, $x \notin A_y$ and every outer neighbour k of A_y has $d(k, x) \geq d(y, x) - 3$; Lemmas 1 and 3 give $f(n) \leq \varphi_{\text{up}} f(n - 3)$, hence $G_1(0, x) \leq \frac{1}{3} \varphi_{\text{up}}^{\lfloor d/3 \rfloor}$ uniformly in N , and $\langle s_0 \cdot s_x \rangle = 3G_1 \leq \varphi_{\text{up}}^{\lfloor d/3 \rfloor}$. For free boxes the same holds under the distance condition of Theorem 6.

Lemma S (change of measure). $\mu_{\beta+\delta}$ has density $\propto e^{\delta H}$ relative to μ_β with $H \in [-60, 60]$ on the pinned box, so for $|F| \leq 1$: $|\langle F \rangle_{\beta+\delta} - \langle F \rangle_\beta| \leq \frac{1}{2}(e^{120|\delta|} - 1)\sqrt{\langle F^2 \rangle_\beta}$. Each grid point certifies $c_e, c_d, \langle F_e^2 \rangle, \langle F_d^2 \rangle$ at $L = 6$; δ_k is the largest half-width with $(\beta_k + \delta_k)[4(c_e + s_e) + 8(c_d + s_d)] < 1$. The 18 points of Table 14 cover $[0.7491374, 0.7701917]$ with overlapping intervals.

Table 14: The 18 certified grid points of the β -interval covering (kernel truncation $L = 6$), from `coverage_summary.json`: exact grid value β_k , certified $\varphi_{\text{up}}(\beta_k)$, half-width δ_k from Lemma S, and covered interval $[\beta_k - \delta_k, \beta_k + \delta_k]$. Consecutive intervals overlap (no gaps); the union is $[0.7491374, 0.7701917]$.

β_k (exact)	β_k	$\varphi_{\text{up}}(\beta_k)$	δ_k	covered interval
3/4	0.75000	0.82542	0.0008626	[0.7491374, 0.7508626]
75159/100000	0.75159	0.83226	0.0008282	[0.7507618, 0.7524182]
4707/6250	0.75312	0.83889	0.0007950	[0.7523250, 0.7539150]
75459/100000	0.75459	0.84531	0.0007629	[0.7538271, 0.7553529]
189/250	0.75600	0.85150	0.0007320	[0.7552680, 0.7567320]
15147/20000	0.75735	0.85747	0.0007022	[0.7566478, 0.7580522]
9483/12500	0.75864	0.86320	0.0006737	[0.7579663, 0.7593137]
18997/25000	0.75988	0.86874	0.0006461	[0.7592339, 0.7605261]
76107/100000	0.76107	0.87409	0.0006195	[0.7604505, 0.7616895]
76221/100000	0.76221	0.87924	0.0005940	[0.7616160, 0.7628040]
7633/10000	0.76330	0.88419	0.0005694	[0.7627306, 0.7638694]
15287/20000	0.76435	0.88898	0.0005457	[0.7638043, 0.7648957]
15307/20000	0.76535	0.89356	0.0005230	[0.7648270, 0.7658730]
76631/100000	0.76631	0.89797	0.0005011	[0.7658089, 0.7668111]
76723/100000	0.76723	0.90222	0.0004801	[0.7667499, 0.7677101]
76811/100000	0.76811	0.90630	0.0004599	[0.7676501, 0.7685699]
2403/3125	0.76896	0.91026	0.0004404	[0.7685196, 0.7694004]
76977/100000	0.76977	0.91404	0.0004217	[0.7693483, 0.7701917]

Proof of Theorem 6. Lemma 0 gives the lower bound. For $\beta < 3/4$ the Aizenman–Simon single-site criterion applies. For $\beta \in [0.7491374, 0.7701917]$ Lemma S and Table 14 give $\varphi_{\text{up}}(\beta) < 1$; at $\beta = 0.77, 0.78, 0.785$ Table 6 does. Lemmas 1, 3, 4 (with T) and 5 give the bound. $\square$

What the independent replay recomputes. `decay077_replay.py` (mpmath interval arithmetic, no python-flint, no transfer code) recomputes from the stored run files: the character coefficients $i_l(\beta)$ and tail T_L at every grid $\beta, \varepsilon, \kappa$, the truncation increments Δ_e, Δ_d , the outward rounding and the final φ_{up} from the stored truncated-model correlators and their stored float-error bounds, the mass bounds, and the covering half-widths δ_k and interval union of Lemma S. It does *not* re-run the transfer contraction of Lemma 4; it is an independent check of the error budget and of the assembly, not of the transfer arithmetic. The transfer arithmetic is checked instead by the independent heat-bath Monte Carlo comparisons on 2×2 , 3×3 and 5×5 pinned boxes, by D_4 symmetry between code paths to 12 digits, by L -convergence and by the small- β series limits (Section 4.7). A full second implementation of the contraction is not available and is listed as an open item.

Archive. The files named in Section 15 (`decay077_certificate.py`, `decay077_run.py`, `decay077_cover.py`, `decay077_replay.py`, `decay077_results.json`, `coverage_summary.json`, `decay077_preregistration.json`) are to be deposited with this version under a versioned DOI; the referee’s point that version 1 named files without a retrievable location is accepted.

B Response to the review of version 2.0

An informal review of version 2.0 (19 October 2026, file `review_v2_0.txt`, deposited) accepted the errata and the new sections and listed five “missing angles that could lead to a proof”. Each is answered below with its status in this version.

1. **Close the first unproved step of Section 10.** *Tested numerically; torus-native route closed or conditional.* The step was first corrected (factor 2, erratum E12) and restated in Θ -form, Eq. (12). It was then tested on exact circle transfer matrices for $L = 2\text{--}8$ at four

couplings (Section 10.1): the required margin is present on all 12 $O(3)$ doublings, grows with L , is at or above one loop for $L \geq 4$, and the $O(2)$ and λ -deformed controls fail as the step demands; the $4 \rightarrow 8$ values are provisional. The “torus-native version of the Cauchy–Schwarz step” suggested by the reviewer was attempted and is closed structurally (Section 10.2): the rotation-score band is empty for a single doubling, the method loses $C \log B$ per step, and conditioning a torus on cell rings produces pinned cells, so no induction closes in the torus class. What survives is Lemma T, an exact transfer from pinned cells to tori, which makes the imported pinned iteration a conditional, averaged, β -uniform stiffness inequality at half the one-loop rate (Corollary T1), conditional on the unrefereed 215-c estimates, the first unverified one being the weighted current estimate. The reviewer’s alternative, “confirmation that 215-a already supplies it with effective constants”, does not apply: 215-a gives fixed- β decay with doubly exponentially weak constants and no per-doubling statement. Separately, the proof chain was re-organized so that the step enters through the doubling defect, and piece (ii) of Section 6 was supplied (Theorem 1). *Update in v2.2:* the additive defect inequality C1’ of v2.1 is false at large β (E14) and the remaining statement is the reach form C1’’ (Section 8.4), shown sharp by the dichotomy Theorem 2; the Θ -form step of this item is its running-coupling form, whose weak-coupling constant is now fixed (Section 9), and the DMRG rates of Section 10.1 add a third, independent numerical check.

2. **Quantify the $6j$ /seam-flux recombination beyond the two-rotor model.** *Lemma R and the mode sum.* Lemma R (Section 10.2) decomposes the N -rotor ring exactly as $L_{\text{tot}}^2/(2N\Theta) + H'$ with H' $SO(3)$ -invariant, proves by Hörmander hypoellipticity and Perron–Frobenius that the composite inertia satisfies $\Theta_N < N\Theta$ strictly for $SO(3)$ and that the ring is rigid up to $\mathbb{Z}_N$ tunnelling for $O(2)$, and locates the whole gap in H' . The semiclassical sum over all $N - 1$ relative-orientation modes converts the single-seam $N^{-1/2}$ deficit into an N -independent drop per doubling that tends to $(N_c - 2) \ln 2 / (2\pi)$, and exact diagonalization at $N \leq 6$ shows N -independent drops (0.253/0.249 at $\beta_{\text{eff}} = 1$) with the $O(2)$ and λ controls quenched (provisional: $N \leq 6$, one coupling). Lemma R is qualitative with respect to the λ control; a quantitative lower bound (Temple/Kato) was attempted and not obtained. A “controlled expansion around the ideal-rotor fixed point” remains to be done.
3. **Turn the exact dynamic bound into a proof tool.** *Not pursued in this version.* The identity $-C'_M(0) = 2/\chi$ of Section 12 stands; a dynamic finite-size criterion or a disagreement-percolation bound on the generator was not attempted, because the finite criterion of Section 7 supplied the missing piece (ii) by a static route first. The dynamic route remains open and is listed as such.
4. **N -uniform strong-coupling certificate.** *Closed, negative.* Section 11, paragraph “No N -uniform $\ell_{\text{max}} = 1$ certificate”: for every g the certified $\ell = 1$ gap goes negative at a finite $N^*(g)$, for the rigorous floors ($N^* \approx 12$ at $g = 2$, $\lesssim 3.3g^2$ for $g \geq 3$) and for the best conceivable scalar floor ($N^* \approx 27g^4$); the loss is the extensive vacuum Q -dressing $Nq_{\text{vac}}/(d - E_0)$ with $q_{\text{vac}} \rightarrow 0.0975/g^2$, plus floor collapse for the rigorous floors; momentum resolution does not help. The reviewer’s suggested two-shell complement would move the cost wall but not the N -dependence; what is needed is a relative certificate for $E_1 - E_0$.
5. **Direct multi-angle sector weights.** *Addressed.* Section 4.6, paragraph “Direct sector weights by seam twist”, Table 5: $Z(\theta)/Z(0)$ measured at 13 angles by seam-twist thermodynamic integration at $\beta = 2.0, 1.8$ and $L = 16, 32, 64$, validated against the exact FK $Z(\pi)/Z(0)$ and the Casimir identity; the direct p_0 lie inside the rigorous brackets and 0.002–0.014 above the $\ell \leq 2$ lower bound; the singlet log-odds rise at every doubling ($\geq +0.169(16)$, $> 10\sigma$) and the increments lie on one curve in p_0 . This converts the provisional Λ increments of Section 10 into a primary observable at these six points; it confirms rather than kills the sector-contraction conjecture there, with the caveat that the two β do not overlap

in p_0 .

C Data and code for the version-2.1 and 2.2 results

All files are deposited with this version under the versioned DOI (Section 15). Each entry names the file, the section it supports and what it contains.

- `proof_chain_closure_2026-10-19.md` (v3; v2 superseded) — Sections 6, 7, 8: the chain with $C1''$ as the open link, the $C2$ correction, status of every link.
- `seed_push.md`, `seed_push_results.json`, `seed_push.py`, `seed_push_code.tar.gz` (modules incl. `sp_rho.py`), `preregistration.md` — Section 7 (v2.2): Lemmas U and $T1'$, the ring-5/6 exact traces with enclosures, MPS Rayleigh bounds, rows 3–9 of Table 7, Table 8, controls, pre-registration scorecard.
- `aspect_pocket_ising_2026-10-20.md`, `ising_aspect_scan.csv`, `ising_aspect_scan_fine.csv` — Section 7: seed aspect window on exact Ising.
- `liminf_formulation.md`, `liminf_formulation.py`, `liminf_results.csv`, `PREREG_liminf.md` — Section 8: Proposition 1, Theorems 2, 3, 4, thresholds, universal constant, Table 9.
- `laplace_betainf.md`, `laplace_betainf.py` (with `o3mc.c`), `laplace_betainf_results.csv`, `preregistered_predictions.csv` — Section 9: Theorem 5, $\gamma(L)$, $d(L)$, controls, Monte Carlo test.
- `rotor_semiclassical.md`, `rotor_semiclassical.py`, `run_scan.py`, `analyze_semiclassical.py`, `fb_check_N2.py`, `rotor_semiclassical_results.csv`, `rotor_semiclassical_fits.csv`, `fb_check_N2.csv`, `prereg_semiclassical.md` — Section 9.2.
- `rotor_rate.md`, `prereg.md`, `rotor_rate.py`, `run2.py`, `analyze.py`, `ed_obc_check.py`, `partB.py`, `rotor_rate_results.csv`, `rotor_rate_fits.csv`, `dmrg_energies.csv`, `rotor_rate_partB.csv` — Sections 10.1, 11: DMRG rates and the relative-certificate test.
- `purity_bootstrap_03.md` — Section 7: Lemmas P, B, C with proofs; $T0$, $T1$, $T2/T2'$, R , F ; results table; Ising and $O(2)$ controls; $\beta = 1.6$ Monte Carlo.
- `purity_results.json` — Section 7: every number of Theorem 1 and Table 7 (seed bounds, truncated values, Δ , Z_{trunc} with error, ring-8 Rayleigh sequence and lower bound, Proposition and chain outputs, truncation control, FK check, Ising validation per K , $O(2)$ control, $\beta = 1.6$ runs and purities).
- `purity_bootstrap_03.py` — Section 7: bootstrap algebra (`f_sq`, `G_boot`, `coupled_update`, `bootstrap_gap`, `iterated_certificate`), exact $\ell \leq 2$ character expansion and Gaunt tensor, ring transfer matrices (dense to ring 4, chunked matvec for ring 8), interval truncation constants (`delta_bound`, `delta_interp`), rounding bounds, Ising Kaufman Z , Monte Carlo purities.
- `purity_ising.py` — Section 7, controls: exact Kaufman torus partition function (mpmath), the purity criterion and $G(u)$; verified against brute-force enumeration to 10^{-14} for $L \leq 4$.
- `tasaki_verification_2026-10-19.md` — Section 7.6: what was fetched and read from [36], with the quoted statements.
- `c1_attack.md`, `c1_attack.py`, `lam_job.py`, `c1_tables.py`, `h0_check.py`, `ed_small.csv`, `ed_big.csv`, `ed_lam.csv`, `c1_doubling_table.csv`, `c1_rotor_ring_03.csv` — Section 10.2: Lemma T and its link to sector weights, Lemma R, the mode sum, the

rotor-ring exact diagonalization with controls, the estimate-by-estimate status table E0–E8 of the imported proof, Lemma 6.

- `gap_step_test.md` (pre-registration in §0), `gap_step_test.py`, `run_all.py`, `run_rest.py`, `run_lam.py`, `analyze_part1.py`, `make_tables_part1.py`, `gap_step_test.csv`, `gap_step_test_best.csv`, `gap_step_test_convergence.csv`, `gap_step_test_all.csv`, `v1_mc_check.json` — Section 10.1: transfer matrices, sectors, convergence tables, Monte Carlo validation, Table 10.
- `nuniform_cert.md`, `nuniform_cert.py`, `part2_rigorous.csv`, `part2_oracle.csv`, `part2_qweights.csv`, `part2_Wnorm.json`, `part2_nstar.json` — Section 11: the N -uniformity analysis and Table 12.
- `second_action_methods.md` (pre-registration in §1), `second_action.py`, `second_action_analysis.py`, `seam_analysis.py`, `second_action_results.csv` (blocks `map`, `map_replica`, `pair`, `seam`, `sector_weights`, `Lambda`), `second_action_map.png` — Section 4.6: Fig. 8, Table 5.
- `review_v2_0.txt` — Appendix B: the review being answered.
- `CHANGELOG_v2_1.md`, `CHANGELOG_v2_2.md` — the lists of changes from version 2.0 to 2.1 and from 2.1 to 2.2.

Numbers in this paper marked **[provisional]** rest on one of: an $\ell_{\max}$ truncation extrapolation ($L = 8$ rows of Section 10.1; $N = 8$ rotor rows), exact diagonalization at small N (Sections 10.2, 9.2), DMRG on open chains with $N \leq 16$ and $\ell_{\max} = 3$ (Section 10.1), a single Monte Carlo series ($\beta = 1.6$ seeds of Section 7; the defect values of Table 9), or the Langevin rows declared in Section 12. The exact statements (Theorems 1, 2, 3, 5, Proposition 1, Lemmas P, B, C, U, T1') do not depend on any provisional number; the constants of Theorem 5 are evaluated in double precision without interval arithmetic (declared).

References

- [1] A. M. Polyakov, Interaction of Goldstone particles in two dimensions. Applications to ferromagnets and massive Yang–Mills fields, *Phys. Lett. B* 59 (1975) 79.
- [2] P. Hasenfratz, M. Maggiore, F. Niedermayer, The exact mass gap of the $O(3)$ and $O(4)$ non-linear σ -models in $d = 2$, *Phys. Lett. B* 245 (1990) 522; P. Hasenfratz, F. Niedermayer, *Phys. Lett. B* 245 (1990) 529.
- [3] S. Caracciolo, R. G. Edwards, A. Pelissetto, A. D. Sokal, Asymptotic scaling in the two-dimensional $O(3)$ σ model at correlation length 10^5 , *Phys. Rev. Lett.* 75 (1995) 1891; reply, *Phys. Rev. Lett.* 76 (1996) 1179.
- [4] U. Wolff, Collective Monte Carlo updating for spin systems, *Phys. Rev. Lett.* 62 (1989) 361.
- [5] U. Wolff, Asymptotic freedom and mass generation in the $O(3)$ nonlinear σ -model, *Nucl. Phys. B* 334 (1990) 581.
- [6] B. Berg, M. Lüscher, Definition and statistical distributions of a topological number in the lattice $O(3)$ σ -model, *Nucl. Phys. B* 190 (1981) 412. [details unverified]
- [7] M. Lüscher, Does the topological susceptibility in lattice σ -models scale?, *Nucl. Phys. B* 200 (1982) 61. [details unverified]
- [8] M. Aizenman, B. Simon, Local Ward identities and the decay of correlations in ferromagnets, *Commun. Math. Phys.* 77 (1980) 137. (Original not accessed; statement checked against the abstract and against Appendix A of [13].)

- [9] O. A. McBryan, T. Spencer, On the decay of correlations in $SO(n)$ -symmetric ferromagnets, *Commun. Math. Phys.* 53 (1977) 299.
- [10] J. Fröhlich, T. Spencer, The Kosterlitz–Thouless transition in two-dimensional abelian spin systems and the Coulomb gas, *Commun. Math. Phys.* 81 (1981) 527.
- [11] L. Chayes, Discontinuity of the spin-wave stiffness in the two-dimensional XY model, *Commun. Math. Phys.* 197 (1998) 623.
- [12] J. Aru, C. Garban, A. Sepúlveda, Percolation for 2D classical Heisenberg model and exit sets of vector valued GFF, *Commun. Math. Phys.* 406 (2025) 37; arXiv:2212.06767.
- [13] C. Garban, N. de Montgolfier, Small perturbations of the $2d$ classical Heisenberg model, arXiv:2610.16824 (12 Oct 2026).
- [14] A. Patrascioiu, E. Seiler, Superinstantons and the reliability of perturbation theory in non-abelian models, *Phys. Rev. Lett.* 74 (1995) 1920.
- [15] A. Jaffe, E. Witten, Quantum Yang–Mills theory, Clay Mathematics Institute Millennium Prize problem description (2000).
- [16] OpenAI, OpenAI Math Release, <https://github.com/openai/math>, public commit `adc7f1241b` (6 Oct 2026). Family 215: *Exponential decay in two-dimensional classical $O(n)$ models* (23 Sep 2026); *Sharp mass bounds for the two-dimensional $O(4)$ model* (23 Sep 2026); *The canonical massive continuum limit of the two-dimensional $O(3)$ model* (4 Oct 2026); *An Isolated Particle Pole for the Two-Dimensional $O(3)$ Spin Field* (4 Oct 2026); *Exact mass asymptotics for the two-dimensional $O(4)$ lattice model* (5 Oct 2026). Unrefereed, AI-generated.
- [17] P. Hasenfratz, F. Niedermayer, Finite size and temperature effects in the AF Heisenberg model, *Z. Phys. B* 92 (1993) 91 (the rotator spectrum $E_\ell = \ell(\ell+1)/(2\Theta)$; see also P. Hasenfratz, H. Leutwyler, *Nucl. Phys. B* 343 (1990) 241). [details unverified]
- [18] H. Tasaki, Ground state of the $S = 1$ antiferromagnetic Heisenberg chain is topologically nontrivial if gapped, *Phys. Rev. Lett.* 134 (2025) 076602.
- [19] M. Lüscher, P. Weisz, U. Wolff, A numerical method to compute the running coupling in asymptotically free theories, *Nucl. Phys. B* 359 (1991) 221. [details unverified]
- [20] J. Kogut, L. Susskind, Hamiltonian formulation of Wilson’s lattice gauge theories, *Phys. Rev. D* 11 (1975) 395.
- [21] A. Ursachi, Local Weil positivity and arithmetic criteria for the Riemann hypothesis, preprint (2026), Proposition 7.3 (Schur-complement certificate with complementary floor). [details unverified] [DOI to be inserted by the author.]
- [22] M. Reed, B. Simon, *Methods of Modern Mathematical Physics IV: Analysis of Operators*, Academic Press (1978), Theorems XIII.44–45.
- [23] N. J. Higham, *Accuracy and Stability of Numerical Algorithms*, 2nd ed., SIAM (2002), Theorem 10.3.
- [24] S. Knabe, Energy gaps and elementary excitations for certain VBS-quantum antiferromagnets, *J. Stat. Phys.* 52 (1988) 627.
- [25] D. Gosset, E. Mozgunov, Local gap threshold for frustration-free spin systems, *J. Math. Phys.* 57 (2016) 091901.
- [26] B. Nachtergaele, The spectral gap for some spin chains with discrete symmetry breaking, *Commun. Math. Phys.* 175 (1996) 565. [details unverified]
- [27] M. Lemm, E. Mozgunov, Spectral gaps of frustration-free spin systems with boundary, *J. Math. Phys.* 60 (2019) 051901. [details unverified]
- [28] D. Bakry, M. Émery, *Diffusions hypercontractives*, Séminaire de Probabilités XIX, Lecture Notes in Math. 1123, Springer (1985) 177.

- [29] P. C. Hohenberg, B. I. Halperin, Theory of dynamic critical phenomena, *Rev. Mod. Phys.* 49 (1977) 435.
- [30] M. Campbell, L. Chayes, The isotropic $O(3)$ model and the Wolff representation, *J. Phys. A* 31 (1998) L255 (amplitude FK G, ordering under axial fields and couplings, wired-boundary comparison, amplitude-times-connectivity representation).
- [31] H. Molke, U. Wolff, Helicity modulus as renormalized coupling in the $O(3)$ σ -model, *Nucl. Phys. B* 603 (2001) 180; hep-lat/0012027. (Equation numbers and the matching coefficient 0.1635 are quoted from the referee report; the paper was not accessible from the sandbox.) [details unverified]
- [32] W. Bietenholz, P. de Forcrand, U. Gerber, H. Mejía-Díaz, I. O. Sandoval, Topological susceptibility of the 2d $O(3)$ model under gradient flow, *Phys. Rev. D* 98 (2018) 114501; arXiv:1808.08129. [details unverified]
- [33] J. Ginibre, General formulation of Griffiths’ inequalities, *Commun. Math. Phys.* 16 (1970) 310.
- [34] B. Simon, Correlation inequalities and the decay of correlations in ferromagnets, *Commun. Math. Phys.* 77 (1980) 111.
- [35] OpenAI, OpenAI Math Release, family 268, *The periodic spin-one Haldane gap* (24 Sep 2026), file `build/bootstrap.tex` at commit `adc7f1241b`: Lemma “Squaring probabilities”, Lemma “Coupled update”, Proposition “Uniform gap from two initial purities”. Unrefereed, AI-generated; only the analytic part is used here (Section 7).
- [36] H. Tasaki, expository notes on the OpenAI family-268 Haldane-gap argument, continuous-time and discrete-time versions (34 pp. each), https://haltasaki.github.io/writings/AI/OpenAI_Haldane_CT.pdf and [.../OpenAI_Haldane_DT.pdf](https://haltasaki.github.io/writings/AI/OpenAI_Haldane_DT.pdf), published 18 October 2026, revised 19 October 2026 (fetched 19 October 2026). Expository notes, not a research article; the author states that the computer-assisted inputs of the family-268 manuscript are not checked there. Cited here for the independent re-derivation of the analytic bootstrap, eq. (4.15), footnote 17 and Appendix D, Proposition D.1.
- [37] L. Hörmander, Hypoelliptic second order differential equations, *Acta Math.* 119 (1967) 147. (Used qualitatively in Lemma R together with the Chow–Rashevskii connectivity theorem and Bony’s strong maximum principle; standard statements, [details unverified] for the exact theorem numbers.)
- [38] D. A. Yarotsky, Ground states in relatively bounded quantum perturbations of classical lattice systems, *Commun. Math. Phys.* 261 (2006) 799. [details unverified]
- [39] B. Simon, Semiclassical analysis of low lying eigenvalues. I. Non-degenerate minima: asymptotic expansions, *Ann. Inst. H. Poincaré A* 38 (1983) 295. [details unverified] (Located, not read in full; cited for the type of statement needed in Lemma G1.)
- [40] B. Helffer, J. Sjöstrand, Puits de potentiel généralisés et asymptotique semi-classique (Puits multiples VI), *Ann. Inst. H. Poincaré Phys. Théor.* 46 (1987) 353 (submanifold wells). [details unverified] (Located, not read in full.)
- [41] B. Kaufman, Crystal statistics. II. Partition function evaluated by spinor analysis, *Phys. Rev.* 76 (1949) 1232 (exact torus partition function used in `purity_ising.py`).