

Erratum (v3): Theorems 1.1, 3.1 and 4.1 are not established

The Yang–Mills Mass Gap on the Lattice:
A Conditional Synthesis

Lluis Eriksson • ai.viXra:2602.0033 • July 2026

Neither of the two pillars of this paper is established: the terminal gap (Theorem 3.1) nor the doubling cascade that transports it (Theorem 4.1). Theorem 1.1 is therefore **not** established at several distinct points audited below. What the counterexamples of this note refute are general implications used by the proofs, one level below the numbered results. There are *two* such implications, with different premises and different conclusions, and each family satisfies the abstract hypotheses relevant to the implication it targets — not the union of both hypothesis sets.

The **exact-trace family targets the exact/additive extraction** (§2): positivity, self-adjointness, trace class, and the trace identity holding *exactly* for infinitely many M , hence $\bar{\delta} = 0$ and the printed window $M_0 \leq M$, $e^{-m_n M} \gg \bar{\delta}$ of the displayed estimate — the manuscript prints $\gg$ there, and with $\bar{\delta} = 0$ it holds however it is read. It does **not** satisfy the cap $m_n M \leq \kappa'/2$ of part (i) — its gap is a fixed $a > 0$, so $m_n M \rightarrow \infty$ — and it does not need to: what it answers is the uncapped exact/additive step. **The finite-error family targets the finite-window relative extraction** (§3): the approximate identity within the admitted error floor, the printed lower bound on the window, and the printed one-sided cap, for every $\kappa' \geq 1$. Each contradicts the conclusion of the implication it is aimed at, so **both abstract extraction implications are false**. *Singular* — “the abstract extraction implication” — *survived into the round that separated the two, because it had been made a required phrase the round before; a guard can preserve an error as faithfully as it preserves a correction. An earlier draft of this note said both families met every printed hypothesis. They do not, and they need not: a counterexample answers the implication it is aimed at, and stating more than that overstates it in the direction this whole correction campaign exists to avoid.*

It does not follow that a numbered result is false: Theorem 4.1 and Proposition 4.4 are stated about the operators T_n built in Step 1 from Balaban’s effective actions, neither is stated for an arbitrary positive trace-class family, and **none is shown to arise from the Balaban effective actions** or to coincide with the intended operators. *Nor is the converse claimed: that no such family can so arise is not established here either.* Accordingly, **the conclusion for the intended Yang–Mills transfer operators is not established, not refuted here.** *A previous round of this note declared the conclusion of Proposition 4.4(i) refuted outright, on this page and in the submission metadata. That went one level too far and is withdrawn; what changed is the classification, not the mathematics, and the families themselves stand.* The exact status, piece by piece:

- **the exact/additive extraction implication** used in Steps 3–4 of Theorem 4.1 and again in the display of Proposition 4.4 — **false**, by the exact-trace family;
- **the finite-window relative extraction implication** used in Proposition 4.4(i) — **false** under the standard uniform reading, by the finite-error family. *These two are defects of one architecture, but they are not one implication: their premises differ and so do their conclusions;*
- the additive estimate displayed in Proposition 4.4 — not established by the printed proof, and no survivor either: **the exact-trace family also contradicts the displayed additive estimate** read as a general implication;
- Proposition 4.4(i), the relative cascade — not established; as a general implication it fails quantitatively and not merely as an exact equality, the second family keeping $\frac{m_{n+1}}{2m_n} - 1 = 1$ rather than $O(e^{-\kappa'/2})$;
- Proposition 4.4(ii), the extensive regime — the two families do not reach its volume regime; it is *likewise not established by the printed proof*, since it rests on the same additive extraction

whose derivation is invalid through the volume-dependent operator. Its separate volume-cap algebra is not audited here;

- **(H-FACT) as consumed by Proposition 4.4** — the printed prefactor C_{fact} is normalised to one; for $C_{\text{fact}} > 1$ that needs $\kappa' > \log C_{\text{fact}}$ and $\kappa' - \log C_{\text{fact}}$ everywhere downstream, and no such renormalisation is stated. **Nor does the printed hypothesis specify coupling uniformity** (§3);
- Lemma 4.3(a) — a conditional algebraic identity;
- Lemma 4.3(b) — an invalid inference, using a lower bound as an upper bound.

Nothing here claims the physical gap differs from what the paper expected. For Theorems 1.1, 3.1, 4.1 and Proposition 4.4, what is withdrawn here is the printed derivation; no conclusion of theirs about the intended operators is refuted. What is false here, classified rather than counted: the exact/additive extraction implication; the finite-window relative extraction implication; and the printed orbit-space dimension corrected in §1.3. None of the three is a conclusion of a numbered result about the intended operators, and this note does not claim the list is closed — a closing count is the failure this chain has now recorded twice. This audit does not refute (H-BAL) or (H-FACT), but those hypotheses do not suffice for the printed conclusions without the additional operator and identification inputs stated below.

An earlier draft of this note stated that “the doubling cascade of §4 is unaffected and is the part of the paper that survives intact”. That is withdrawn: §2 below shows it is not. The provenance page appended here fixes the dictionary that §1.1 and §4.1 turn on.

Notation, before anything else. In this erratum $\nu_{\text{push}} := \pi_{\#} \text{Vol}_{\mathcal{A}}$, and every ν below means ν_{push} . The restored manuscript uses the same letter for a different object: Step 3 of Theorem 3.1 calls $\nu_{\text{Vol}} := d\text{Vol}_B/\text{Vol}_B(B)$ “the uniform measure ν ”. **The two are not identified**, and identifying them is what (I2b) of the appended page forbids. *Read one for the other and a printed derivation looks as though it already used the corrected measure.*

And four different constants are called C in this chain. The manuscript writes $C(N_c) = 24\pi^2/(11N_c)$ for the exponent of Theorem 1.1, C for the prefactor of (H-FACT), C again for the lower end of the window in Proposition 4.4 (“ $C \leq M$ ”), and C_0 for the volume constant of Remark 2.2. **This erratum writes C_{fact} for the (H-FACT) prefactor, M_0 for the window’s lower end, and C_β for the terminal-coupling constant it introduces in §4.2, keeps $C(N_c)$ and C_0 as printed, and writes a bare C only here and when quoting the manuscript’s own “ $C \leq M$ ”.** *An earlier draft of this note promised never to write a bare C and then wrote one, in §4.2, in the very round that put up the barrier.*

1 Theorem 3.1: the terminal gap

1.1 The printed Euclidean measure is not identified with the ground-state measure

Step 2 introduces $d\mu_{\text{print}} = e^{-f} d\text{Vol}_B/Z$ with $f = \beta_{n_{\text{max}}} V_W + \sum_X \epsilon_{n_{\text{max}}}(X)$ — the *effective action* — applies Holley–Stroock in Step 3, and writes $m_{n_{\text{max}}} = \frac{g_{n_{\text{max}}}^2}{2} \lambda_1$. That relation belongs to the Hamiltonian framework, where the ground-state transformation requires the measure to carry $f_0 = -\log \psi_0^2$, i.e. $\mu_0 = \psi_0^2 \nu_{\text{push}}$.

The Gibbs measure of a Euclidean effective action is *a priori* a distinct object, and its weighted Dirichlet-form gap *a priori* a distinct spectral quantity. **The printed derivation applies the Hamiltonian gap identity to an effective-action Gibbs measure without identifying that measure with the ground-state measure:** no identification of μ_{print} with $\mu_0 = \psi_0^2 \nu_{\text{push}}$ has been proved. In special models they may coincide, but in general the passage requires a transfer-operator or Osterwalder–Seiler identification theorem, not a change of notation. This is (I3) of the appended page.

Note the reference measure. What Step 2 prints is against Vol_B , not against ν_{push} , and it is a *regular-stratum* formula: no global extension, and no closed Dirichlet form for it, is specified. On B_{reg} , $d\nu_{\text{push}} = w d\text{Vol}_B$ with w the orbit-volume density, so

$$d\mu_{\text{print}} = Z_{\text{print}}^{-1} e^{-S_{\text{eff}}} d\text{Vol}_B = Z_{\text{print}}^{-1} e^{-(S_{\text{eff}} + \log w)} d\nu_{\text{push}} \quad \text{on } B_{\text{reg}},$$

which is an exact rewriting of one measure, with the same normaliser. Relative to ν_{push} the printed measure therefore carries the exponent $S_{\text{eff}} + \log w$, while the candidate ν_{push} -based measure $d\tilde{\mu} = \tilde{Z}^{-1} e^{-S_{\text{eff}}} d\nu_{\text{push}}$ carries S_{eff} ; they coincide only under an additional constancy or identification statement for w , none of which is established here. This matters for the repair, and is (I2b) of the appended page.

The programme already knew. The v2 of `ai.viXra:2602.0035` corrected exactly this substitution in its own v1 — the action-based measure being “*dimensionally inconsistent*” and, per its suite, failing to reproduce the gap at “> 50%, *growing*”. That correction did not reach this paper, though both are July 2026 documents of the same series. *What that measurement is, and is not:* it concerns the literal potential-based measure of that paper’s v1, not the ϵ -corrected effective-action measure printed here. **It supplies strong numerical counterevidence to a universal action-measure substitution, and reinforces that no identification may be made by notation; it is neither a proof that no general identification exists nor an analytic counterexample to Theorem 3.1 of the present paper.** *An earlier draft said it “confirms” the unavailability of a general identification — a non-rigorous one-plaquette computation cannot confirm a mathematical non-existence. The reason to withdraw the printed derivation is that no identification theorem has been proved, not that a number came out wrong.*

1.2 The geometric input is no longer available

Steps 1 and 3 use $\text{Ric}_B \geq N_c/4$ from `ai.viXra:2602.0036`, Theorem 3.1. That theorem is stated *at regular points* while Bakry–Émery is applied here globally, and the corollary that globalised it has been withdrawn. **Decisively, that theorem is itself not established by its printed proof:** its O’Neill trace omits the mixed horizontal–vertical sectional terms, which for a bi-invariant metric on a compact group are non-negative and subtract. **No pointwise Ricci bound on the orbit space may presently be cited.**

What is available is `ai.viXra:2602.0046 v3`: a global curvature-dimension bound, hence a Poincaré/log-Sobolev inequality, for (B, d_B, ν) with $\nu = \pi_{\#} \text{Vol}_A$ — a bound for K_{ν} , not for a realisation of the unweighted $\mathcal{K}_B^{\text{reg}}$. It cannot serve as the pointwise input used here.

A warning about the symbol ν , and it is load-bearing. **The ν of the appended page is not the ν of the restored pages.** That page writes $\nu_{\text{push}} = \pi_{\#} \text{Vol}_A$; Step 3 of Theorem 3.1 calls $\nu_{\text{Vol}} = d\text{Vol}_B / \text{Vol}_B(B)$ “the uniform measure ν ”. Reading one for the other makes the printed derivation look as though it already used the corrected measure — which is exactly the identification (I2b) forbids. *Nor does it repair Step 3 by a change of reference measure:* by (I2b), **replacing $d\text{Vol}_B$ by $d\nu_{\text{push}}$ under a fixed exponent is not a free rewriting.** Two routes are open and they are different theorems. **(A)** Keep μ_{print} , whose exponent relative to ν_{push} is $S_{\text{eff}} + \log w$: one must then extend the regular-stratum formula globally and bound $\text{osc}(S_{\text{eff}} + \log w)$, for which the bound on $\text{osc}(S_{\text{eff}})$ does not suffice — if w degenerates at strata of enlarged stabiliser, $\log w$ is exactly the term that destroys a bounded-perturbation argument, and **no control of $\log w$ is supplied** in this chain. **(B)** Reformulate over (B, ν_{push}) with $d\tilde{\mu} = \tilde{Z}^{-1} e^{-S_{\text{eff}}} d\nu_{\text{push}}$, where 2602.0046 does supply the base inequality — but that is not the measure the paper printed, and a theorem is then owed showing it is the correct Euclidean object.

1.3 A four-dimensional action on a three-dimensional orbit space

Step 1 fixes the space as $B_{n_{\text{max}}}$ and prints its dimension as $3L_{n_{\text{max}}}^3 (N_c^2 - 1)$. **That is the dimension of the spatial configuration manifold $\mathcal{A}_{\text{spat}} = SU(N_c)^{3L^3}$, not of the orbit space:** the site

gauge group $SU(N_c)^{L^3}$ has dimension $L^3(N_c^2 - 1)$ and acts with discrete generic stabiliser, so the **principal-stratum dimension is $2L_{n_{\max}}^3(N_c^2 - 1)$** , and at reducible configurations the quotient is stratified with no single manifold dimension. *An earlier draft of this note repeated the printed figure as though it were correct.* The $O(L^3)$ versus $O(L^4)$ mismatch below is unchanged by the correction. Step 1 also says “*the spatial slice is three-dimensional*”; Step 2 bounds the oscillation with $O(L_{n_{\max}}^4)$ plaquettes. Both are $O(1)$ terminally, so no numerical bound breaks — but the mismatch shows f is the four-dimensional Euclidean effective action while μ is treated as living on the three-dimensional spatial orbit space. (H-FACT) approximately controls the factorisation of the *partition function*; it does not build a Hamiltonian on B^{spat} nor identify its ground state.

2 Theorem 4.1: the limit is taken over operators that change with the volume

This is the defect that removes the last standing pillar, and it is independent of everything in §1.

Step 1 defines T_n on $\mathcal{H}_n = L^2(B_n^{\text{spat}})$ for $\Lambda_n = (\mathbb{Z}/L_n\mathbb{Z})^4$ with $L_n = L/2^n$, using — in the paper’s own words — “*the convention of equal spatial and temporal extent*”. **Both the Hilbert space and the operator therefore depend on L_n .**

Step 2 sets $M := L_n$ and obtains $\text{Tr}(T_{n+1}^{M/2}) = \text{Tr}(T_n^M)$, then justifies $M \rightarrow \infty$ by *varying the original isotropic size L* . But varying L varies $L_n = M$, hence varies the spatial slice, hence varies $\mathcal{H}_n$, T_n , the eigenvalues $\lambda_j^{(n)}$ and the gap m_n — all of which Step 3 then treats as constants in M when it writes

$$\text{Tr}(T_n^M) = (\lambda_0^{(n)})^M (1 + e^{-m_n M} + O(e^{-m'_n M})).$$

That expansion extracts the spectrum only when a *fixed* operator is raised to a growing power. Here the power and the operator grow together.

The trace identities do not imply the conclusion. An abstract example suffices to show that **the doubling and the displayed additive estimate do not follow from the family of identities alone** — not that the identities carry no spectral information at all, which would be a stronger claim and is not made here. For *fixed* $a, b > 0$ with $b \neq 2a$, and for each even M , put

$$T_{n,M} = \begin{pmatrix} 1 & 0 \\ 0 & e^{-a} \end{pmatrix}, \quad T_{n+1,M/2} = c_M \begin{pmatrix} 1 & 0 \\ 0 & e^{-b} \end{pmatrix}, \quad c_M = \left(\frac{1 + e^{-aM}}{1 + e^{-bM/2}} \right)^{2/M}.$$

Both matrices are positive, self-adjoint and trace class; *positivity of a and b is what makes the top eigenvalue the first entry*, so the gaps are exactly a and b . Then $\text{Tr}(T_{n+1,M/2}^{M/2}) = \text{Tr}(T_{n,M}^M)$ for every even M , while b ranges over all positive values other than $2a$. This is not offered as a model of anything physical, and **it is not shown to arise from any lattice gauge theory** — nor is the converse shown; what it does show is that the inference used in Steps 3–4 is not valid for an M -dependent family.

With a third eigenvalue, and against the additive estimate as well. The two-by-two family leaves the printed remainder $O(M^{-1}e^{-(m'_n - m_n)M})$ with no third eigenvalue to refer to, which invites the objection that it is testing a degenerate case. It is not needed. Fix $a, b, r, s > 0$ with $b \neq 2a$, $r > a$ and $s > b$, and for each even M put

$$T_{n,M} = \text{diag}(1, e^{-a}, e^{-r}), \quad T_{n+1,M/2} = c_M \text{diag}(1, e^{-b}, e^{-s}),$$

$$c_M = \left(\frac{1 + e^{-aM} + e^{-rM}}{1 + e^{-bM/2} + e^{-sM/2}} \right)^{2/M}.$$

Then $\text{Tr}(T_{n+1,M/2}^{M/2}) = \text{Tr}(T_{n,M}^M)$ *exactly*, for every even M , so one may take $\delta_n = \delta_{n+1} = 0$; the orderings $1 > e^{-a} > e^{-r}$ and $1 > e^{-b} > e^{-s}$ hold, and the gaps and next distinct gaps are $m_n = a$,

$m_{n+1} = b$, $m'_n = r$, $m'_{n+1} = s$. The displayed additive estimate of Proposition 4.4 would give

$$b = 2a + O(M^{-1}e^{-(r-a)M}),$$

whose right-hand side tends to $2a$ along the infinitely many admissible M , contradicting $b \neq 2a$. **So the exact-trace family also contradicts the displayed additive estimate**, and not only the relative cascade built on it. *Verified at $a = 1/3$, $b = 1$, $r = 2$, $s = 5$ for $M = 4, 8, 20, 60, 200, 1000$: the two traces agree to sixty digits while $|b - 2a| = 1/3$ stays fixed and the printed remainder falls below 10^{-700} . An earlier draft of this note recorded the additive estimate as untouched by these families; §3 said so too.*

What these families do and do not refute. They are matrices, not the transfer operators of Step 1: **none is shown to arise from the Balaban effective actions**, and this note does not claim they arise from any lattice gauge theory — nor does it claim they cannot. **Each family satisfies the abstract hypotheses relevant to the implication it targets**, and only those: the exact-trace family has a fixed gap $a > 0$, so $m_n M \rightarrow \infty$ and the one-sided cap of part (i) is not among its hypotheses, while the finite-error family of §3 does satisfy that cap. What they refute are two *general* implications. **The exact-trace family targets the exact/additive extraction:** positivity, self-adjointness, trace class and the exact trace identity for infinitely many M , therefore the doubling and the displayed additive estimate — which is what Steps 3–4 use. **The finite-error family targets the finite-window relative extraction:** the approximate identity, the error floor and the one-sided window, therefore a relative cascade with error $O(e^{-\kappa'/2})$ — which is what Proposition 4.4(i) uses. Theorem 4.1 and Proposition 4.4 are stated about the operators built in Step 1, not about that abstract class, so their conclusions are *not established* for the intended operators rather than refuted. *This distinction is the one the note previously collapsed, calling the conclusion of Proposition 4.4(i) false outright.*

A second, local defect of Step 3: multiplicities. Even for a fixed operator the printed expansion is not correct as written. It sets the coefficient of $e^{-m_n M}$ to 1 and takes $m'_n = -\log(\lambda_2/\lambda_0) > m_n$, then adds that if λ_1 is degenerate the coefficient should be its multiplicity. But with the printed enumeration a degenerate λ_1 gives $\lambda_2 = \lambda_1$ and hence $m'_n = m_n$, not $m'_n > m_n$: the remainder is then of the same order as the term it is supposed to be smaller than. Nor does “positive, self-adjoint and trace class” make λ_0 simple — that needs a positivity-improving or Perron–Frobenius input. The correct form groups *distinct* eigenvalues $\lambda_0 > \lambda_1 > \lambda_2$ with multiplicities d_j :

$$\mathrm{Tr}(T^M) = d_0 \lambda_0^M \left[1 + \frac{d_1}{d_0} e^{-m M} + O(e^{-m' M}) \right], \quad m' > m.$$

This does not change the verdict — Theorem 4.1 is already not established by the M -dependence above — but it is recorded because this note claims to have audited the cascade, and because the corrected form is what any repair must use.

Why the suite does not see it. The verification miniature is a one-dimensional $\mathbb{Z}_N$ clock chain under decimation. In one dimension, varying the length varies only the temporal direction: there is no growing spatial slice, so the operator stays fixed and the extraction is legitimate. The miniature therefore tests a situation from which the defect is absent by construction.

What would repair it. A fixed spatial slice by itself is not sufficient. A family of *rectangular* lattices with the spatial slice held fixed and only the temporal extent varied removes the M -dependence of $\mathcal{H}_n$ and T_n , which is the defect above; but the extraction still needs an identity to feed it, and the printed one is $\mathrm{Tr}(T_{n+1}^{M/2}) = \mathrm{Tr}(T_n^M)$ obtained by varying the *isotropic* size. *Nor is a fixed slice necessary: the last two routes below do not presuppose one, and an earlier draft of this note asserted a necessity it had not shown.* Three routes, then:

- a fixed-spatial-slice family *together with* an RG/transfer-matrix identity valid for all temporal extents over that fixed space — both, not either;
- or an operator relation such as T_{n+1} conjugate to T_n^2 , which gives the doubling directly and needs no limit;
- or a **quantitative, normalised spectral-convergence theorem** for the volume-dependent family. **Ordinary spectral convergence is insufficient**, and the exact-trace family above already refutes it: there $\log c_M \rightarrow 0$, so $T_{n+1,M/2} \rightarrow \text{diag}(1, e^{-b}, e^{-s})$ in operator norm while $T_{n,M}$ does not move at all; every eigenvalue converges, the limiting gaps are a and b , and still $b \neq 2a$. The reason is that $c_M - 1 \sim (2/M)(e^{-aM} - e^{-bM/2})$ is of the *same* exponential order as the first excited contribution, and absorbs it exactly. What a repair needs is **leading-eigenvalue normalisation relative to the actual finite- M first-excited contributions**, and it must carry the multiplicities this note has just insisted on. Write $d_{0,n}(M)$ for the multiplicity of the top eigenvalue, so that the leading trace term is $d_{0,n}(M)\lambda_{0,n}(M)^M$, and set

$$Q_M := \frac{d_{0,n+1}(M/2) \lambda_{0,n+1}(M/2)^{M/2}}{d_{0,n}(M) \lambda_{0,n}(M)^M},$$

$$A_n(M) := \frac{d_{1,n}(M)}{d_{0,n}(M)}, \quad A_{n+1}(M/2) := \frac{d_{1,n+1}(M/2)}{d_{0,n+1}(M/2)}.$$

Suppose $m_n(M) \rightarrow m_n^\infty > 0$ and $m_{n+1}(M/2) \rightarrow m_{n+1}^\infty > 0$. The normalisation condition required by the trace comparison is

$$|Q_M - 1| = o\left(A_n(M)e^{-m_n(M)M} + A_{n+1}(M/2)e^{-m_{n+1}(M/2)M/2}\right).$$

The earlier condition $Q_M = 1 + o(e^{-\mu_\infty M})$, with $\mu_\infty = \min\{m_n^\infty, m_{n+1}^\infty/2\}$, is **necessary-looking but not sufficient**: it compares with a limiting exponential scale rather than the finite- M corrections the trace must resolve. The boxed condition is required together with convergence of the ratios of *distinct* eigenvalues, **subexponential control of the excited multiplicities**

$$\frac{1}{M}|\log(d_{1,n}(M)/d_{0,n}(M))| \rightarrow 0, \quad \frac{1}{M}|\log(d_{1,n+1}(M/2)/d_{0,n+1}(M/2))| \rightarrow 0,$$

and **uniform control of the total excited trace tail**

$$\sum_{j \geq 2} \frac{d_{j,n}(M)}{d_{1,n}(M)} e^{-[m_{j,n}(M) - m_{1,n}(M)]M} \rightarrow 0,$$

and the same at the next scale with $M/2$ in place of M throughout. For an exact trace identity these conditions do force the limiting doubling. Indeed, put

$$X_M = A_n(M)e^{-m_n(M)M}, \quad Y_M = A_{n+1}(M/2)e^{-m_{n+1}(M/2)M/2}.$$

The two tail conditions reduce the normalised trace identity to $1 + X_M(1 + o(1)) = Q_M[1 + Y_M(1 + o(1))]$. Here both $o(1)$ terms are **multiplicative relative-tail errors**; an additive absolute $o(1)$ would be too weak because it could absorb either exponentially small first-excited term. The boxed condition then gives $X_M - Y_M = o(X_M + Y_M)$, hence $X_M/Y_M \rightarrow 1$. Taking logarithms, dividing by M , and using the two subexponential multiplicity conditions yields $m_n^\infty = m_{n+1}^\infty/2$. Before passing to the limit the same argument gives the stronger exponent matching

$$m_n(M)M - \frac{1}{2}m_{n+1}(M/2)M - \log \frac{A_n(M)}{A_{n+1}(M/2)} \rightarrow 0.$$

This is the missing sufficiency argument; without the boxed finite- M comparison, the conclusion does not follow. **Level-by-level multiplicity control is not sufficient**: it bounds each row of an enumeration, not how many rows there are.

Positivity improving removes only the top-eigenvalue multiplicity; it says nothing about $d_1, d_2, \dots$, nor about the spectral counting function.

The multiplicity and tail conditions are both necessary, and each is forced by its own counterexample; they are not sufficient without the boxed normalisation. The first family below has one excited level of exponentially large multiplicity; the second has exponentially many simple excited levels. **Neither the multiplicity condition nor the tail condition excludes both.**

First: multiplicity. For even M put $D_M = \lfloor e^M \rfloor$, $q = e^{-2}$, and let $P_N = \frac{1}{N} \mathbf{1} \mathbf{1}^*$. Take

$$T_{n,M} = q I_{D_M+1} + (1-q) P_{D_M+1}, \quad T_{n+1,M/2} = c_M (q I_2 + (1-q) P_2),$$

with $c_M = [(1 + D_M e^{-2M}) / (1 + e^{-M})]^{2/M}$. Every entry is strictly positive, so both are *positivity improving*; the spectra are 1 simple and e^{-2} with multiplicity D_M , respectively 1 and e^{-2} simple. The traces agree exactly, both equal to $1 + D_M e^{-2M}$, so $\delta_n = \delta_{n+1} = 0$; both gaps are 2, so $m_{n+1}^\infty \neq 2m_n^\infty$; $d_{0,n} = d_{0,n+1} = 1$; the distinct-eigenvalue ratios are constant; and $\mu_\infty = \min\{2, 1\} = 1$ while the normalisation quotient is $c_M^{M/2} = (1 + D_M e^{-2M}) / (1 + e^{-M}) = 1 + O(e^{-2M}) = 1 + o(e^{-\mu_\infty M})$. **The normalisation condition and the positivity-improving alternative are both satisfied, and the doubling still fails.** The mechanism is the excited multiplicity: $D_M e^{-2M} \sim e^{-M}$, so the trace sees the effective exponent $2 - \frac{1}{M} \log D_M \rightarrow 1$, not the spectral gap 2. Here $\frac{1}{M} \log(d_1/d_0) \rightarrow 1$; there is no level $j \geq 2$, so the tail sum is empty and does not see it.

Second: counting. Keep $D_M = \lfloor e^M \rfloor$ and choose distinct $\varepsilon_{j,M} \in [-\frac{1}{4}, \frac{1}{4}]$ with $\sum_{j \leq D_M} \varepsilon_{j,M} = 0$ — equispaced and symmetric will do — and order them so that $\varepsilon_{1,M} = \max_j \varepsilon_{j,M}$. Put $\lambda_{j,M} = e^{-2}(1 + \varepsilon_{j,M})^{1/M}$ and

$$T_{n,M} = \text{diag}(1, \lambda_{1,M}, \dots, \lambda_{D_M,M}), \quad T_{n+1,M/2} = c_M \text{diag}(1, e^{-2}),$$

with the same c_M . Then $\text{Tr}(T_{n,M}^M) = 1 + e^{-2M} \sum_j (1 + \varepsilon_j) = 1 + D_M e^{-2M}$, the traces agree exactly again, **every multiplicity is one**, the excited ratios converge to e^{-2} uniformly, both limiting gaps are 2, and the normalisation quotient is the same $1 + o(e^{-\mu_\infty M})$. **So the whole multiplicity condition is satisfied trivially — and the doubling still fails.** What fails now is the tail: each term of $\sum_{j \geq 2}$ lies in $[\frac{3}{5}, 1]$ and there are $D_M - 1$ of them, so the sum grows like e^M . *Checked at $M = 4, 8, 16, 24$: the traces agree, the normalisation error $|Q_M - 1|/e^{-\mu_\infty M}$ falls to 3.2×10^{-11} , and the tail sum exceeds 1.6×10^{10} . An earlier draft of this note wrote “and likewise for the levels making up the remainder”, which is not a verifiable hypothesis; the defect it was gesturing at is the counting function, and this family is what makes that concrete.* **Note the arguments: the operator compared at the next scale is $T_{n+1,M/2}$, so its leading eigenvalue is evaluated at $M/2$, and the gaps are the limits, not constants in M .** Or else an identity already normalised so that the leading eigenvalue cannot absorb the excited term.

Third: the limiting scale is not the finite- M correction. Let $P_2 = \frac{1}{2} \mathbf{1} \mathbf{1}^*$ and $A(t) = e^{-t} I_2 + (1 - e^{-t}) P_2$, whose entries are strictly positive and whose eigenvalues are $1, e^{-t}$. Fix $a = 1/3$, $b = 1$, and for even M put

$$a_M = a + M^{-1/2}, \quad b_{M/2} = b + (M/2)^{-1/2},$$

$$T_{n,M} = A(a_M), \quad T_{n+1,M/2} = c_M A(b_{M/2}), \quad c_M = \left(\frac{1 + e^{-a_M M}}{1 + e^{-b_{M/2} M/2}} \right)^{2/M}.$$

The traces agree exactly; both operators are positivity improving; all top and first excited multiplicities are one; the tail is empty; and the gaps converge to a and b . With $\mu_\infty = 1/3$, the old condition still holds: $Q_M - 1 = O(e^{-M/3 - \sqrt{M}}) = o(e^{-\mu_\infty M})$. Nevertheless $b \neq 2a$.

Here $Q_M - 1$ is of the same order as the actual dominant finite- M excited correction, so the boxed condition fails exactly where it should. Thus positivity improving does not rescue the R29 condition.

The original fixed-gap exact-trace family of Section 2 has $d_0 = 1$ at both scales and gives $1 + \Theta(e^{-\mu_\infty M})$ in that quotient, so the condition is not vacuous: it excludes it. An earlier draft of this note wrote the numerator at M rather than $M/2$, and the gaps as constants — in a note whose central defect is confusing what moves with M — and normalised only the leading eigenvalues, ignoring the very multiplicity correction it had just declared indispensable. An earlier draft of this note offered plain uniform spectral convergence as a repair, which its own counterexample refutes.

None appears. An earlier draft of this note listed the first item as one alternative among four, which understates what it owes.

Theorem 4.1 is not established by its printed proof.

3 What this costs

Proposition 4.4 silently normalises the prefactor of (H-FACT) to one. (H-FACT) is printed as

$$|\delta_n| \leq C_{\text{fact}} e^{-\kappa'}, \quad \kappa' > 0 \text{ independent of } L_n.$$

Proposition 4.4(i) invokes “if $\bar{\delta} \leq e^{-\kappa'}$ uniformly in L_n (Hypothesis (H-FACT))”. **The prefactor has been set to one.** The implication from $C_{\text{fact}} e^{-\kappa'}$ to $e^{-\kappa'}$ is automatic when $C_{\text{fact}} \leq 1$; for $C_{\text{fact}} > 1$ it needs the renormalised exponent

$$\tilde{\kappa} = \kappa' - \log C_{\text{fact}},$$

hence $\kappa' > \log C_{\text{fact}}$ to leave a positive exponent at all, and then $\tilde{\kappa}$ in place of κ' everywhere downstream: the cap of Lemma 4.3 becomes $m_n M \leq \tilde{\kappa}/2$ and the cascade factor becomes $1 + O(n_{\max} e^{-\tilde{\kappa}/2}) = 1 + O(n_{\max} \sqrt{C_{\text{fact}}} e^{-\kappa'/2})$, so C_{fact} enters the accumulation as $\sqrt{C_{\text{fact}}}$. **No such renormalisation is stated anywhere**, and the manuscript nowhere bounds C_{fact} . *An earlier draft of this note called the substitution “a strictly stronger bound”; that is not universal — it is an equality of content when $C_{\text{fact}} = 1$ and a weakening when $C_{\text{fact}} < 1$. This is independent of the accumulation and of the resolution defects below: before any iteration, the proof does not literally consume the hypothesis its own theorem declares.* The counterexample below is stated against the proposition’s own, stronger, printed bound $\bar{\delta} \leq e^{-\kappa'}$, so it is untouched by this.

Proposition 4.4 — the window controls the wrong side. Its additive estimate has the shape $m_{n+1} = 2m_n + O(\bar{\delta} e^{m_n M}/M) + O(e^{-(m'_n - m_n)M}/M)$. Writing $x = m_n M$ and dividing by $m_n = x/M$ to get the relative estimate the cascade then uses, what is needed is

$$\bar{\delta} e^{m_n M} \ll m_n M \quad \text{and} \quad e^{-(m'_n - m_n)M} \ll m_n M.$$

The paper instead requires only $e^{-x} \gg \bar{\delta}$ and then caps $x \leq \kappa'/2$. **Neither of the two conditions follows:** as $x \rightarrow 0$ the printed condition becomes automatically easy, since $e^{-x} \rightarrow 1$, while $\bar{\delta} e^x/x$ grows without bound. **A fixed error floor can hide a positive gap**, and a very small gap is not an advantage against a fixed error — it falls below the resolution.

A finite positive counterexample, as a family parametrised by κ' . Put $\bar{\delta} = e^{-\kappa'}$ and $x = \bar{\delta}/2$, and choose any even $M \geq M_0$, the printed lower bound on the window (the manuscript’s C in “ $C \leq M$ ”). Take $T_n = \text{diag}(1, e^{-x/M})$ and $T_{n+1} = \text{diag}(1, e^{-4x/M})$, with gaps $m_n = x/M$ and $m_{n+1} = 4x/M$, so $m_{n+1} \neq 2m_n$. Put $\delta_n = 0$ and $\delta_{n+1} = (1 + e^{-x})/(1 + e^{-2x}) - 1$, which satisfies $0 < \delta_{n+1} < x < \bar{\delta}$. Then $\text{Tr}(T_{n+1}^{M/2})(1 + \delta_{n+1}) = \text{Tr}(T_n^M)(1 + \delta_n)$ exactly: the approximate partition identity holds within the admitted errors, both gaps are positive, $e^{-x} \gg \bar{\delta}$ holds with margin $e^{\kappa'}$ up

to a factor tending to 1, the window $x \leq \kappa'/2$ holds **for every** $\kappa' \geq 1$, and the doubling fails. Not for every positive value of the parameter: that window is $e^{-\kappa'} \leq \kappa'$, first true at $\kappa' = W(1) = 0.5671\dots$, and the same condition one scale up, $2x \leq \kappa'/2$, is $2e^{-\kappa'} \leq \kappa'$, first true at $W(2) = 0.8526\dots$; an earlier draft asserted it for all positive values, which is false below those two thresholds. Taking $\kappa' \geq 1$ clears both, and clears $e^{-2x} \geq \bar{\delta}$ at the next scale as well.

And it fails relatively, not merely exactly. Proposition 4.4(i) does not promise exact doubling: it promises $1 + O(e^{-\kappa'/2})$ per step. Here $m_n = x/M$ and $m_{n+1} = 4x/M$, so

$$\frac{m_{n+1}}{2m_n} - 1 = 1 \quad \text{for every } \kappa',$$

which is not $O(e^{-\kappa'/2})$ as $\kappa' \rightarrow \infty$. The flat-ASCII transcriptions that earlier drafts carried in this prose have been removed; they belong in the all-ASCII submission sheet, where they are, and not in the typeset page. Equivalently, at $n_{\max} = 1$ the downward form gives $m_n/(m_{n+1}/2) - 1 = -\frac{1}{2}$, again a fixed nonzero constant. Checked numerically for $\kappa' = 10, 20, 40, 80$: the relative error stays at 1 while $e^{-\kappa'/2}$ falls to 4×10^{-18} . **The separation condition and the one-sided cap are met with growing margin** — $e^{-x}/\bar{\delta} \sim e^{\kappa'}$ and $\kappa'/2 - x \rightarrow \infty$ — **while the admitted-error bound is met uniformly and not with growing margin**, since $\delta_{n+1}/\bar{\delta} \rightarrow \frac{1}{4}$ as $\kappa' \rightarrow \infty$; and $M \geq M_0$ is met by a fixed M . An earlier draft of this note wrote that every hypothesis was met with growing margin, which is false for those last two. **If the implicit O -constant is allowed to depend arbitrarily on κ' , the printed estimate is not quantitative enough to support its subsequent “harmless error” use; under the standard uniform reading, the family above refutes it as a general implication, which is not the same as refuting the conclusion for the intended operators.** What this family reaches is the step that turns the additive estimate into a relative cascade and a uniform positive lower bound; **the additive estimate itself is reached by the exact-trace family of §2**, where $\bar{\delta} = 0$ and the remainder is genuinely present. An earlier draft of this note said the additive estimate survived both. An earlier draft stopped at “the doubling fails”, which by itself only denies exact equality — not what the proposition claims.

The finite-error spectral window is two-sided, roughly $\bar{\delta}e^{m_n M} \ll m_n M \lesssim \log(1/\bar{\delta})$, together with quantitative isolation from the next *distinct* spectral value — the third-eigenvalue term $O(e^{-(m'_n - m_n)M}/M)$ is dropped with no hypothesis at all. For $\bar{\delta} \leq e^{-\kappa'}$ and $m_n M \leq \kappa'/2$ the lower end requires $m_n M \gg e^{-\kappa'/2}$, and **nothing in (H-FACT), Theorem 3.1 or Lemma 4.3 supplies that resolution floor. Theorem 3.1 supplies the lower bound $m_{n_{\max}} \geq c_0$, and no hypothesis compares $c_0 L_{n_{\max}}$ with that floor.** If the certified lower bound lies below the floor, the available information does not guarantee that the actual gap can be resolved — *the actual gap may still be larger*. An earlier draft wrote that the smallness of c_0 “can put the gap below” the floor, which repeats, three lines after denouncing it, the very lower-bound-as-value slip attributed to Lemma 4.3. This defect precedes the $n_{\max}e^{-\kappa'/2}$ accumulation: it bites at a single step. And the proof inserts the error factors “and repeats Step 4”, so it inherits the M -dependence of §2 before the error floor $\bar{\delta}$ is even reached.

Lemma 4.3 — a conditional identity, and an inequality used backwards. That $m_{n+1} = 2m_n$ and $L_{n+1} = L_n/2$ imply $m_{n+1}L_{n+1} = m_nL_n$ is valid algebra, but its premise is exactly the unestablished conclusion of Theorem 4.1. **Lemma 4.3(a) stands only as a conditional algebraic implication**, not as an effective invariance of the cascade.

Part (b) is worse, and it is an independent failure. Theorem 3.1 supplies only $m_{n_{\max}} \geq c_0$. Lemma 4.3 writes “ $m_{n_{\max}} = c_0 = (\gamma N_c/8)e^{-K/\gamma}$ is exponentially small in $1/\gamma$ ” — **it replaces a lower bound by an equality and then uses the smallness of c_0 as an upper bound on the gap.** A small lower bound says nothing about how large the gap is: $m_{n_{\max}} = 1$ with $c_0 = 10^{-19}$ satisfies $m_{n_{\max}} \geq c_0$ and violates the window. **The claim that the window is automatic is not established**, and it would not be established even if Theorem 3.1’s own proof were repaired: what the window needs is an *upper* bound on $m_{n_{\max}}$, or the terminal-window condition assumed outright. The suite does not catch this either — checking that $c_0 C_0$ is tiny checks the size of the

lower bound, not the size of the gap. *An earlier draft of this note misquoted that constant: it dropped the γ prefactor and wrote κ for γ in the exponent. Those are two different parameters — the terminal coupling and the polymer decay constant of Theorem 2.1(b) — and a note whose object is the direction of an inequality cannot make its case with a formula the paper never wrote. An earlier draft of this note called part (b) merely a second consumption of Theorem 3.1’s numerical value. It is not: Theorem 3.1 never supplied that value, and this is a direction error of its own. It is the third inequality-direction slip recorded in this chain.*

The printed control on the accumulated factorisation error ceases to be small as $g^2 \rightarrow 0$. Proposition 4.4 and the main proof produce $m_0 = (m_{n_{\max}}/2^{n_{\max}})(1 + O(n_{\max}e^{-\kappa'/2}))$, and the proof then replaces that factor by $1 + o(1)$, noting that it is harmless for $g^2 \geq 1/(b_0e^{\kappa'/2})$ and that “for smaller g^2 the constant degrades explicitly but remains positive” for $n_{\max}e^{-\kappa'/2} \leq \frac{1}{2}$, “which we fold into g_0^2 ”. **That folding runs the wrong way.** Since $n_{\max} = \frac{1}{2b_0 \log 2}(1/g^2 - 1/\gamma) + O(1)$, at fixed κ' one has $n_{\max}e^{-\kappa'/2} \rightarrow \infty$ as $g^2 \rightarrow 0$, and the smallness condition is a lower bound $g^2 \gtrsim ce^{-\kappa'/2}$, which cannot be absorbed into a hypothesis of the form $0 < g^2 \leq g_0^2$ — shrinking g_0 makes it worse.

What that shows, exactly. Write the relative remainder as $R(g)$, so that the printed estimate is an upper bound

$$|R(g)| \leq A_{\text{rel}}(g) n_{\max}(g) e^{-\tilde{\kappa}(g)/2}$$

for some explicit A_{rel} . **The printed upper bound ceases to be small** and becomes unbounded as $g \rightarrow 0$. **This does not prove that the actual error grows:** $R \equiv 0$ satisfies any such bound. What follows is only that *the printed estimate can no longer justify the replacement by $1 + o(1)$ — information insufficient, not conclusion false. Earlier drafts of this note, and its submission sheet, said that the accumulated error itself grows and diverges. That reads an upper bound as a value, which is the same slip this note attributes to Lemma 4.3(b), with the inequality turned round.*

And the printed hypothesis does not specify coupling uniformity. (H-FACT) says only that $\kappa' > 0$ is independent of L_n . It does not say that κ' or C_{fact} is independent of g , nor that either is uniform on $0 < g^2 \leq g_0^2$. **The printed (H-FACT) does not suffice under either reading:** under the first **the printed upper bound becomes unbounded** and therefore ceases to justify $R(g) = o(1)$ — **this does not prove that the actual accumulated error diverges**; under the second the hypothesis is under-specified and must be supplemented. *An earlier draft of this note wrote that the accumulated error diverges here, four lines after distinguishing a bound from a value. Two readings, and the paper fixes neither:*

- if C_{fact} , κ' and A_{rel} are uniform in g , the printed control parameter becomes unbounded and cannot be folded into a hypothesis $0 < g^2 \leq g_0^2$;
- if they may depend on g , write $\tilde{\kappa}(g) = \kappa'(g) - \log C_{\text{fact}}(g)$, and **choose an explicit $A_{\text{rel}}(g) \geq 1$** with $|R(g)| \leq A_{\text{rel}}(g) n_{\max}(g) e^{-\tilde{\kappa}(g)/2}$ — a specified function, not “the” implicit constant of an $O(\cdot)$, which is not unique and can be scaled up at will; $A_{\text{rel}} \geq 1$ costs nothing, since one may replace any bound by its maximum with 1, and it keeps $\log A_{\text{rel}}$ non-negative. The growth condition is then imposed on that specified bound:

$$\tilde{\kappa}(g) - 2 \log n_{\max}(g) - 2 \log A_{\text{rel}}(g) \rightarrow +\infty,$$

If A_{rel} is only known to be bounded above, **a sufficient condition** — not an equivalent one — is $\tilde{\kappa}(g) = 2 \log n_{\max}(g) + \omega(1)$, which since $n_{\max} \asymp g^{-2}$ reads $\tilde{\kappa}(g) \geq 4 \log(1/g) + \omega(1)$; **the chosen explicit bound A_{rel} must be uniformly bounded in g** for that reduction to be available at all. For a factor merely bounded by some $\eta < 1$ rather than $o(1)$, the requirement relaxes to $\tilde{\kappa} \geq 2 \log n_{\max} + 2 \log(A_{\text{rel}}/\eta)$. **None of this is an $O(1)$ statement:** a bounded difference keeps the factor below a constant, it does not send it to zero. **That growth condition is an additional missing hypothesis.** *An earlier draft of this note wrote $+O(1)$ where $\omega(1)$ is needed, and left A_{rel} out of the condition while admitting three lines later that it too needs uniform control.*

*An earlier draft of this note said the flat verdict was “the first reading only”. That was itself wrong: the printed hypothesis fails under both. The fault of the flat verdict was that it did not say which reading it was arguing from, not that its verdict was too strong. **This propagates to Theorem 1.1’s own constant.*** The paper declares $c(N_c) > 0$ to depend only on N_c and the universal constants of Balaban’s construction; that cannot be maintained unless C_{fact} , κ' and the implicit O -constants are themselves controlled uniformly over the admitted data. What would suffice: a two-sided coupling window rather than the half-line; per-scale errors summable uniformly in n_{max} ; a bound improving with g or with the scale; or a proof that the accumulation renormalises without the n_{max} factor.

Remark 4.5 — half stands. That the temporal nonlocality of Balaban’s effective actions “*is not an artifact*” stands. That (H-FACT) is “*the sharply isolated open point of this paper*” does **not**: there are now at least three separate unbuilt bridges — the Euclidean-action/ground-state measure identification (§1.1), the four-dimensional-action/three-dimensional-operator construction (§1.3), and the fixed-operator gap extraction (§2).

Theorem 1.1 — not established at several distinct points audited below, and its printed quantifier is not justified either. The terminal-gap derivation and the automatic-window claim are separate failures, not one input consumed twice; it now also loses the transport of the gap from the terminal scale back to the original one. **The automatic-window failure is independent of the terminal-gap proof:** it reverses the role of a lower bound and would survive any repair of Theorem 3.1 (§3). The failure of the transport step is a third, independent defect. *An earlier draft of this note called the first two “two consumptions of the same broken input”. That was still true of the terminal gap and false of the window, and it survived into the same round that proved it false.*

And a fourth debt, of domain rather than of proof. The blocking sequence $\Lambda_k = (\mathbb{Z}/(L/2^k)\mathbb{Z})^4$ requires **admissible volumes**: Remark 2.2 and Step 2 restrict to $2^{n_{\text{max}}+1} \leq L \leq C_0 2^{n_{\text{max}}}$, equivalently $L = 2^{n_{\text{max}}}q$ with $2 \leq q \leq C_0$ — so $L \in 2^{n_{\text{max}}}\mathbb{N}$ alone is not enough: it admits $L = 2^{n_{\text{max}}}$, a terminal slice of extent one, which Remark 2.2 itself excludes. **And the lower endpoint depends on g through n_{max} ,** so it cannot be replaced by a fixed L_0 as $g^2 \rightarrow 0$. Theorem 1.1 nevertheless quantifies over every L in an interval, and no rounding, interpolation or irregular-blocking argument is given for the rest. **The statement is not literally established with its printed quantifier over all L .** *This is the same correction the companion erratum applies to Theorem 6.3 of [ai.viXra:2602.0035](#); it had not been propagated back to the paper the cascade comes from.*

4 Two residual pieces of structure, with their exact limitations

Neither is a printed statement that survives as printed. Both are pieces of structure that a reformulated route could reuse, and each is stated here with what it still owes.

4.1 The transfer operator is on the intended sector, but the measure is unstated

Theorem 4.1 writes $\mathcal{H}_n = L^2(B_n^{\text{spat}})$ without saying whether the measure is $\nu_{\text{spat}} = \pi_{\#}\text{Haar}$, the Riemannian volume of the regular stratum, or another. *An earlier draft of this note called this “the right object, by (I1)”;* that was premature and is withdrawn, and it contradicted the appended page, which records that this paper’s spatial domain is not yet fixed. The accurate statement is: **Theorem 4.1 intends a transfer operator on the gauge-invariant spatial sector, but the printed notation does not specify the measure; identification (I1) applies only after $\nu_{\text{spat}} = \pi_{\#}\text{Haar}$ is taken explicitly.**

4.2 The terminal oscillation bound needs a two-sided coupling hypothesis

Step 2 bounds $\text{osc}(V_W) = O(1)$ by the plaquette count and $\sum_X |\epsilon_{n_{\max}}(X)| \leq C_\epsilon(\kappa, d) < \infty$, both correctly. It then substitutes $g_{n_{\max}}^2 = \gamma$ — equivalently $\beta_{n_{\max}} = 2N_c/\gamma$ — and concludes $\text{osc}(f) \leq K/\gamma$.

But Theorem 2.1(c) supplies only $g_{n_{\max}}^2 \leq \gamma$, i.e. $\beta_{n_{\max}} \geq 2N_c/\gamma$ — a *lower* bound on $\beta_{n_{\max}}$. An upper bound on $\beta_{n_{\max}} \text{osc}(V_W)$ requires an *upper* bound on $\beta_{n_{\max}}$, and the printed inequality runs the other way. *An earlier draft of this note asserted that this bound “does establish what it claims uniformly”; that is withdrawn.*

And it is used twice, not once. The last line of the same proof concludes

$$m_{n_{\max}} = \frac{1}{2}g_{n_{\max}}^2 \lambda_1 \geq \frac{1}{8}\gamma N_c e^{-K/\gamma},$$

replacing the prefactor $g_{n_{\max}}^2$ by γ inside a *lower* bound. The hypothesis $g_{n_{\max}}^2 \leq \gamma$ does not license that either: it makes the prefactor smaller, not larger. So the unsupported replacement $g_{n_{\max}}^2 = \gamma$ occurs at both ends of the proof — once to bound $\beta_{n_{\max}}$ from above in the oscillation estimate, once to bound the kinetic prefactor from below — and the printed one-sided hypothesis supplies neither inference. **A two-sided terminal-coupling bound is required** — $c\gamma \leq g_{n_{\max}}^2 \leq \gamma$ for some $c > 0$, or directly $\beta_{n_{\max}} \leq C_\beta/\gamma$ — and it repairs both places at once, with constants degraded by c . (The direction of this slip is the same as one corrected in [ai.viXra:2602.0052](#): substituting an endpoint into a bound that moves the wrong way.)

What survives is only the finite-dimensional oscillation estimate for the explicit Euclidean effective action, namely

$$\text{osc}(f) \leq \beta_{n_{\max}} C_W + 2C_\epsilon, \quad C_W := \text{osc}(V_W), \quad C_\epsilon := C_\epsilon(\kappa, d).$$

This is an explicit finite-dimensional inequality. It becomes uniform in the claimed parameters only after a two-sided terminal-coupling bound is supplied — note that the right-hand side still contains $\beta_{n_{\max}}$, which is exactly the quantity not yet bounded above. And **it does not control the ground-state exponent** $f_0 = -\log \psi_0^2$; it therefore neither supplies an (A-osc)-type hypothesis about the ground state nor removes the separate identification required by (I3).

5 The citation loop, and what a clean repair would need

Theorem 3.1 proves the terminal gap by “*applying the argument of* [3, Theorem 5.3]”, where [3] is [ai.viXra:2602.0035](#); that paper cites this one as discharging its (H1). The loop is not logically vicious — this paper re-runs the method with its own inputs — but **this paper is not an independent confirmation of that one**, and neither discharges a hypothesis of the other.

A standalone terminal lemma stated on (B, ν, K_ν) would *not* by itself suffice: for the Euclidean measure $e^{-S_{\text{eff}}}\nu$ it would remain a statement about a Euclidean diffusion. To discharge a Hamiltonian gap the lemma must be stated either for $\mu_0 = \psi_0^2\nu$, or together with an explicit transfer-operator construction identifying the two quantities.

6 Scope

Not examined here: Balaban’s package, the numerical suite beyond the observation of §2, Theorem 2.1 and the coupling-evolution erratum of Remark 2.3, and the existence of a positive self-adjoint trace-class transfer matrix at each scale, which Theorem 4.1 assumes as part of its hypothesis. *This audit does not refute* (H-BAL) *or* (H-FACT) — both are, as printed, unusually candid — but those hypotheses do not suffice for the printed conclusions without further inputs, **including** all of the following:

1. a *globally specified* Euclidean measure with a closed form, along one of the two branches of (I2b) — and hence either control of $\text{osc}(\log w)$ on route (A), or a theorem identifying the ν -based measure as the correct object on route (B) (§1.1–1.2);
2. a four-dimensional-action to three-dimensional-space transfer-operator or Hamiltonian construction (§1.3);
3. the ground-state-measure identification (I3);
4. the renormalisation $\kappa' \mapsto \kappa' - \log C_{\text{fact}}$ that Proposition 4.4 needs in order to consume (H-FACT) as printed when $C_{\text{fact}} > 1$, together with $\kappa' > \log C_{\text{fact}}$ (§3);
5. *either* a restriction of the theorem to a coupling window chosen so that $A_{\text{rel}} n_{\text{max}}(g) e^{-\tilde{\kappa}/2} \leq \eta < 1$ on it, if C_{fact} , κ' and A_{rel} are uniform in g — a compact window alone only makes n_{max} finite, not small against the error floor, and in any case this changes the statement rather than repairing its printed quantifier over all sufficiently small couplings — *or*, if they may depend on g , the growth condition $\tilde{\kappa}(g) - 2 \log n_{\text{max}}(g) - 2 \log A_{\text{rel}}(g) \rightarrow +\infty$ (§3);
6. a two-sided terminal-coupling bound (§4.2);
7. the admissible-volume condition $2^{n_{\text{max}}+1} \leq L \leq C_0 2^{n_{\text{max}}}$ for the blocking sequence, whose lower endpoint moves with g (§3);
8. an *upper* bound on $m_{n_{\text{max}}}$, or the terminal-window condition as a hypothesis, for Lemma 4.3(b) (§3);
9. per-scale factorisation errors that stay summable as $g^2 \rightarrow 0$, which the printed (H-FACT) does not guarantee under either reading of its coupling-uniformity (§3);
10. a two-sided spectral-resolution window for the approximate trace extraction, including $\bar{\delta} e^{m_n L_n} \ll m_n L_n$ and quantitative isolation from the next distinct spectral value (§3);
11. for the cascade (§2), *either* a fixed-spatial-slice family *together with* an identity valid for all temporal extents over that fixed space, *or* an operator relation, *or* a *normalised* spectral-convergence theorem in the sense of §2, with normalisation relative to the actual finite- M first-excited contributions, *subexponential control of the excited multiplicities, and uniform control of the total excited trace tail*. The displayed multiplicity and tail conditions are necessary but not sufficient without that finite- M normalisation; ordinary spectral convergence and the old limiting-scale normalisation are insufficient, positivity improving removes only the top-eigenvalue multiplicity, and the fixed slice alone does not suffice.

An earlier draft of this section listed only items 3–5. Items 1 and 2 are defects this same note had already separated, and the closing list dropped them; it is written with “including” now because a closing list of obligations that reads as exhaustive is the same failure as a survivors list closed too early.

Provenance. §1 was found by asking which measure each paper of this chain uses, and by auditing the cited Ricci theorem rather than the citing sentence. **§2, §4.1 and §4.2 came from external audits of the earlier drafts of this very note**, which had certified the doubling cascade as surviving intact, the spatial sector as the right object, and the oscillation bound as uniform. All three were premature. A second audit round supplied the rest: the *second* inverted use of the coupling hypothesis in the last line of the same proof, the positivity condition $a, b > 0$ without which the abstract example has the wrong top eigenvalue, the multiplicity defect of Step 3, and the distinction between three proof failures and three independent routes. A third round found the drift *on the appended page itself*, which had recorded this paper’s printed measure against ν instead of Vol_B — the very substitution its new (I2b) forbids — and then found that first correction overshooting, asserting the two measures *are* different when that needs w non-constant and is not proved here. A fourth round supplied the three-eigenvalue family, the thresholds $W(1)$ and $W(2)$ for the window, and the separation between a false general implication and a refuted printed conclusion — the note had crossed it in the direction that overstates. The pattern — a survivors list closed by checking what a statement depends on rather than by reading its proof — is the same one this correction campaign has now recorded five times, and is the reason the appended page exists. That page is reproduced verbatim: paraphrase is how the columns drift.

Object Provenance for the Orbit-Space Gap Chain

A common page, reproduced verbatim in ai.viXra:2602.0033, 2602.0035 and 2602.0036

Lluis Eriksson • July 2026

Why this page exists. Four papers of this series state, prove, or consume spectral claims about objects that are routinely called “the gap on the orbit space”. They are not the same object. A pointwise curvature *claim* stated in one of them was consumed, in three places, as if it were established and as if it applied to the same object. It is neither: its printed proof is not presently established, and the sound synthetic result available elsewhere concerns a measure and an operator that *have not been identified* with the ones the consuming steps need. The error was invisible because every paper used the same word. This page fixes the dictionary. **A citation discharges a hypothesis only when all four columns agree.**

Notation barrier

This page writes ν for $\nu_{\text{push}} := \pi_{\#} \text{Vol}_{\mathcal{A}}$, and never for anything else. The host papers of this chain use the same letter for a different object: 2602.0033 (Theorem 3.1, Step 3) and 2602.0035 (Theorem 5.3, Step 2) both write “the uniform measure ν ” for the normalised Riemannian volume

$$\nu_{\text{Vol}} := d\text{Vol}_B / \text{Vol}_B(B).$$

ν_{push} and ν_{Vol} are not identified, and identifying them is exactly what (I2b) below forbids without a statement about the orbit-volume density. A reader who carries the host’s ν into this page, or this page’s ν into the host’s Step 3, will read a printed derivation as though it already used the corrected measure. *Where the distinction bites, the subscripts are written out.*

The dictionary

Source	Space, measure	Operator	What it can conclude
2602.0036	$B_{\text{reg}}, d\text{Vol}_B$	$\mathcal{K}_B^{\text{reg}}$: unweighted Laplace–Beltrami <i>expression</i> ; global realisation unspecified	<i>targets</i> a pointwise Ricci bound on the regular stratum; printed proof not presently established
2602.0046	$B, \nu = \pi_{\#} \text{Vol}_{\mathcal{A}}$	$K_{\nu} \geq 0$, operator of the form on $L^2(B, \nu)$	global Poincaré/LSI, singular locus included
2602.0033, as printed	spatial orbit space; printed <i>regular-stratum formula</i> $d\mu_{\text{print}} = Z_{\text{print}}^{-1} e^{-S_{\text{eff}}} d\text{Vol}_B$ on B_{reg} , against Vol_B , not ν ; <i>no global extension specified</i>	weighted diffusion of μ_{print} , <i>if</i> a global closed form is constructed	neither a Poincaré/LSI statement nor a Hamiltonian gap is presently established
route (A): keep μ_{print}	(B, ν) ; a candidate <i>global extension</i> of μ_{print} , exponent $S_{\text{eff}} + \log w$	weighted form over K_{ν}	needs w globally, and $\text{osc}(S_{\text{eff}} + \log w)$
route (B): new measure	(B, ν) ; a <i>separately defined candidate</i> measure $d\tilde{\mu} = \tilde{Z}^{-1} e^{-S_{\text{eff}}} d\nu$	weighted form over K_{ν}	needs a theorem that this is the right object
2602.0035, as printed	$L^2(B, d\text{Vol}_B)$, global realisation not fixed; $\mu_{0,\text{print}} = \psi_{0,\text{print}}^2 d\text{Vol}_B$	Hamiltonian on $L^2(B, d\text{Vol}_B)$; diffusion of $\mu_{0,\text{print}}$	parts (i) and (ii) are not established by the printed proofs
candidate physical restatement	$L^2(\mathcal{A})^{\mathcal{G}} \cong L^2(B, \nu)$; $\mu_{0,\nu} = \psi_{0,\nu}^2 \nu$	$H_{\nu} = \frac{g^2}{2} K_{\nu} + V_{\text{pot}}$; then diffusion of $\mu_{0,\nu}$	qualitative fixed-volume positivity; a rate only under (A-osc) $_{\nu}$

The four identifications that must never be silent

(I1) $L^2(B, \nu) \cong L^2(\mathcal{A}, \text{Vol}_{\mathcal{A}})^{\mathcal{G}}$ — **exact, and free.** Pullback $f \mapsto f \circ \pi$ gives $\int_B |f|^2 d\nu = \int_{\mathcal{A}} |f \circ \pi|^2 d\text{Vol}_{\mathcal{A}}$, which is the definition of $\nu = \pi_{\#} \text{Vol}_{\mathcal{A}}$. This is an isometry, not an approximation, and it holds across the singular locus. It is the reason the metric-measure quotient is the natural physical object: gauge-invariant functions on the smooth compact $\mathcal{A}$ are $L^2(B, \nu)$.

(I2) K_{ν} versus $\mathcal{K}_B^{\text{reg}}$ — **not free, and not even the same kind of object.** Fix the convention here, so that nothing depends on the host paper's. On the regular stratum write $d\nu = w d\text{Vol}_B$, with w the orbit-volume density, and define the *formal differential expressions*

$$\mathcal{L}_B^{\text{reg}} := \text{div}_B \nabla, \quad \mathcal{L}_{\nu}^{\text{reg}} := w^{-1} \text{div}_B(w \nabla), \quad \mathcal{K}_{\bullet}^{\text{reg}} := -\mathcal{L}_{\bullet}^{\text{reg}},$$

so that

$$\mathcal{L}_{\nu}^{\text{reg}} f = \mathcal{L}_B^{\text{reg}} f + \langle \nabla \log w, \nabla f \rangle.$$

Let $K_{\nu} \geq 0$ denote the *global* self-adjoint operator associated with the closed Dirichlet form on $L^2(B, \nu)$.

No global self-adjoint realisation of $\mathcal{K}_B^{\text{reg}}$ is fixed by this page. B_{reg} is the regular stratum with the singular locus removed and is in general an incomplete manifold; $\mathcal{K}_B^{\text{reg}}$ on $C_c^{\infty}(B_{\text{reg}})$ is symmetric and non-negative, but it has not been shown to be essentially self-adjoint, and it may admit several self-adjoint extensions. Choosing the Friedrichs extension is already a decision about behaviour across $B \setminus B_{\text{reg}}$, and must be specified and justified separately — this is one of the gaps that forces the withdrawal of the global corollary in 2602.0036.

The two *regular-stratum differential expressions* coincide only if w is constant on each relevant component. *Constant stabiliser type does not by itself imply constant orbit volume* — on the principal stratum all stabilisers are already conjugate, and the induced metric on the orbits, hence their volume, may still vary with the configuration; reducible connections make a global identification less tenable still. **A spectral bound for the global K_{ν} does not by itself transfer to any chosen self-adjoint realisation of $\mathcal{K}_B^{\text{reg}}$, or conversely.** Comparison theorems under further hypotheses are not excluded; what does not exist is a free identification — and on the $\mathcal{K}_B^{\text{reg}}$ side, there is not yet a single global operator for such a theorem to be about.

(I2b) $d\text{Vol}_B$ versus ν *at the level of the perturbed measure* — **also not free.** This is a separate step from (I2), and it is the one most easily skipped, because both measures are written with the same exponent. On B_{reg} , $d\nu = w d\text{Vol}_B$ with w the orbit-volume density, so the measure a paper prints against Vol_B is, written against ν ,

$$d\mu_{\text{print}} = Z_{\text{print}}^{-1} e^{-S_{\text{eff}}} d\text{Vol}_B = Z_{\text{print}}^{-1} e^{-(S_{\text{eff}} + \log w)} d\nu \quad \text{on } B_{\text{reg}},$$

an *exact* rewriting of one and the same measure, with the same normaliser. The candidate ν -based measure $d\tilde{\mu} = \tilde{Z}^{-1} e^{-S_{\text{eff}}} d\nu$ carries the exponent S_{eff} and, in general, a *different* normaliser. **On B_{reg} the printed measure has exponent $S_{\text{eff}} + \log w$ relative to ν , whereas the candidate ν -based measure has exponent S_{eff} . They coincide only under an additional constancy or identification statement for w , none of which is established here. Replacing $d\text{Vol}_B$ by $d\nu$ under a fixed exponent is therefore not a free rewriting.**

And the converse is not asserted either. This page does not claim the two measures are *proved* distinct: that would need w non-constant, up to normalisation, on the relevant domain. Reducible connections make a gratuitous global identification implausible, but implausible is not proved, and the discipline this page enforces cuts both ways — “no identification has been proved” must not become “an inequality has been proved”.

Two routes are open, and they are not the same route. **(A) Keep the printed measure.** Put $\Phi_{\text{print}} = S_{\text{eff}} + \log w$, so that $d\mu_{\text{print}} = Z_{\text{print}}^{-1} e^{-\Phi_{\text{print}}} d\nu$. This needs a global construction of the

measure from the regular-stratum formula, control of w approaching the singular locus, and a bound on $\text{osc}(S_{\text{eff}} + \log w)$ — for which a bound on $\text{osc}(S_{\text{eff}})$ alone does not suffice: if w degenerates at strata of enlarged stabiliser, $\log w$ is exactly the term that destroys a bounded-perturbation argument. **No such control is supplied anywhere in this chain.** (B) **Change to the ν -natural measure $\tilde{\mu}$,** which is *a separately defined candidate measure* and not the one the paper printed. Here 2602.0046 can supply the base functional inequality for K_ν , but a new theorem is required to show that $\tilde{\mu}$, and not μ_{print} , is the correct Euclidean or transfer-theoretic object. Note that route (A) is not free either: “keep the printed measure” still means *constructing a global extension of a regular-stratum formula*, which is itself one of the open obligations. **A ν -based base inequality is therefore an input to one of these two routes, not a repair of a printed Vol_B -based derivation by change of reference measure.** Neither route by itself supplies the four-dimensional-action to three-dimensional-space step, the identification with μ_0 , or a transfer-operator construction.

(I3) μ_{print} **has not been identified with $\mu_0 = \psi_0^2 \nu$ — not free, and the load-bearing gap.** With $K_\nu \geq 0$ as in (I2) — equivalently, by (I1), the operator of the kinetic form $\mathcal{E}_\nu(f, f) = \int_B |\nabla f|^2 d\nu$ on $L^2(\mathcal{A}, \text{Vol}_\mathcal{A})^G$ — let V_{mult} be a specified real multiplication potential and set $H_{\text{mult}} = \frac{g^2}{2} K_\nu + V_{\text{mult}}$, with a normalised, strictly positive ground state ψ_0 , and put $d\mu_0 = \psi_0^2 d\nu$. *The subscripted name is deliberate. This page is reproduced verbatim inside papers that have already bound V to a gauge-fixed potential and W to the Weyl group, so no bare letter is safe here: a page written to prevent symbol collisions must not cause one.* The ground-state transformation then gives

$$E_1 - E_0 = \frac{g^2}{2} \lambda_1(\mu_0), \quad f_0 = -\log \psi_0^2,$$

where E_1 denotes the first spectral value strictly above the simple ground-state energy E_0 , and, *defined explicitly to avoid the reciprocal ambiguity in the term “Poincaré constant”,*

$$\lambda_1(\mu_0) := \inf \left\{ \mathcal{E}_{\mu_0}(f, f) : f \in \text{Dom } \mathcal{E}_{\mu_0}, \int f d\mu_0 = 0, \|f\|_{L^2(\mu_0)} = 1 \right\}.$$

That is, $\lambda_1(\mu_0)$ is the Poincaré *spectral gap* — the largest $\rho \geq 0$ with $\text{Var}_{\mu_0}(f) \leq \rho^{-1} \mathcal{E}_{\mu_0}(f, f)$ for all admissible f — and *not* the constant $C_P = \rho^{-1}$ of the variance inequality, which some authors also call the Poincaré constant.

The Gibbs measure of a Euclidean effective action is *a priori* a distinct mathematical object, and its weighted Dirichlet-form gap is *a priori* a distinct spectral quantity. In the symbols fixed above,

No globally defined printed or corrected Euclidean measure has been identified with $\mu_0 = \psi_0^2 \nu$.

That covers μ_{print} , route (A)’s global extension of it, and route (B)’s $\tilde{\mu}$ alike — stated this way so that (I3) neither singles out one candidate nor silently presupposes (I2b). No such identification has been proved here; in special models the objects may coincide, but passing from one to the other in general requires a transfer-operator or Osterwalder–Seiler identification theorem, not a change of notation. **An oscillation bound for S_{eff} does not by itself bound $\text{osc}(-\log \psi_0^2)$.**

Constants travel with their metric normalisation

$\text{Ric} \geq N_c/2$ for $\langle X, Y \rangle = -\text{tr}(XY)$ and $\text{Ric} \geq N_c/4$ for $-2\text{tr}(XY)$ are the *same* bound in two normalisations, not a sharpening of one by the other. Any statement quoting $N_c/2$ or $N_c/4$ must name its inner product.

How to use this page

Before writing that a companion result discharges a hypothesis, fill the four columns for the result and for the hypothesis. If any column differs, the citation does not discharge; at most it supplies an input to a further identification, which must then be stated as its own step. In particular:

- 2602.0046 supplies global geometry on the metric-measure object identified by (I1): a bound for K_ν on $L^2(B, \nu)$. It does not supply (I1) — that is free — and it does *not* supply a bound for any realisation of the unweighted $\mathcal{K}_B^{\text{reg}}$; nor does it attempt pointwise bounds on B_{reg} , as its own text says. Note that it reaches its conclusion by synthetic quotient stability and *explicitly bypasses the O’Neill computation*. That is now the only sound geometric input in this chain: the O’Neill trace printed in 2602.0036 omits the mixed horizontal–vertical sectional terms and does not establish its own conclusion (see that record’s erratum). **No pointwise Ricci bound on B_{reg} may be cited as available.**
- 2602.0033, as printed, targets a weighted diffusion for $d\mu_{\text{print}} = Z_{\text{print}}^{-1} e^{-S_{\text{eff}}} d\text{Vol}_B$ on B_{reg} — against Vol_B , *not* against ν , and with no global extension or associated closed Dirichlet form specified. It does not presently establish a Euclidean Poincaré/LSI statement, and it does not establish a Hamiltonian or transfer-operator gap without (I3). **Citing 2602.0046 does not repair it by writing the same exponent against ν :** by (I2b) that is not a free rewriting. Either route (A) is taken and $\text{osc}(S_{\text{eff}} + \log w)$ is controlled, or route (B) is taken and the new measure is shown to be the right object.
- 2602.0035 v2 already uses a ground-state measure, so (I3) is not the missing bridge in its printed quantitative proof — that was the correction its v2 made to its own v1. Its printed-to-physical transition **crosses (I2) and (I2b)**: the unweighted realisation against K_ν , and $d\text{Vol}_B$ against ν . The candidate physical restatement then requires $(A\text{-osc})_\nu$, a hypothesis about *its own* ground state. **Identification (I3) re-enters only if that Hamiltonian construction is used to identify a Euclidean or transfer-operator object**, as in the route towards (H1) of 2602.0033. Among the four identifications on this page, its qualitative fixed-volume positivity needs only (I1); it additionally uses the standard compact-elliptic and positivity-improving inputs — self-adjoint realisation, compact resolvent, connectedness — and it does not require any treatment of the quotient singularities. *An earlier version of this page said flatly that the paper needs (I3) for its quantitative bound. That was wrong, and it is the kind of error this page exists to catch: it named the wrong hypothesis for a citation to discharge.*

Provenance. This page was written after an audit found the same curvature *claim* consumed in three places under three different pairings of measure and operator, and after a proposed repair — citing the metric-measure result in place of the regular-stratum one — was found to change the object rather than fix the proof. A later pass found that the regular-stratum claim is not established by its own printed proof either, so the defect is now twofold: an invalid derivation in one paper, and a correct theorem in another about a different object. A third pass found the drift *on this page*: the row for 2602.0033 had recorded the printed measure against ν when the paper writes it against Vol_B , which is precisely the substitution (I2b) now forbids. It was introduced by writing the row from the intended repair rather than from the printed text. Hence the rule below, and hence (I2b) as a numbered step rather than a remark. The same pass then caught the correction overshooting: (I2b) had asserted the two measures *are* different, which needs w non-constant and is not proved here. Both directions are now stated at the strength they have. A fourth pass found that (I3) wrote its Hamiltonian as $\frac{g^2}{2} K_\nu + V$ — and V is already bound, in 2602.0035, to the gauge-fixed $S_{\text{YM}} - \log \det M$. A page reproduced verbatim inside that paper cannot borrow a symbol the host has spent. The first repair reached for a bare W , which that same paper binds to the Weyl group S_{N_c} : **no bare letter is safe in a page meant to be embedded, only a subscripted name**. The same pass found the collision that matters most, because it is in the *measure* column: both host papers write “the uniform measure ν ” for the normalised $d\text{Vol}_B$, which is what this page’s own (I2b) forbids identifying with $\pi_\# \text{Vol}_A$. Hence the notation barrier above. A fifth pass found the row for 2602.0035 built from its intended repair rather than its printed text — the identical error this page had already corrected once, for 2602.0033 — and found the claim that that paper needs (I3) to be simply wrong: its v2 already uses a ground-state measure. It records a failure mode that is not covered by the usual distinction between a false statement and an invalid proof: **a correct theorem about the wrong object for the consuming dependency**. Reproduce it verbatim; do not paraphrase it, because paraphrase is how the columns drift.

THE YANG–MILLS MASS GAP ON THE LATTICE: A CONDITIONAL SYNTHESIS

LLUIS ERIKSSON

ABSTRACT. This is the synthesis paper of a series studying the Yang–Mills mass gap via the geometry of the gauge orbit space. We show, conditionally on two explicitly stated hypotheses, that $SU(N_c)$ lattice Yang–Mills theory in $d = 4$ dimensions with Wilson action at sufficiently weak coupling has a positive mass gap $m_{\text{gap}} \geq c(N_c) e^{-C(N_c)/g^2} > 0$ in lattice units, uniformly in $L_0 \leq L \leq C_0 e^{C/g^2}$, combining Balaban’s constructive renormalization group, a Holley–Stroock/Bakry–Émery spectral gap bound at the terminal scale, and a transfer-matrix trace identity. Version 1 framed the result as conditional on a single informal compatibility assumption. Version 2 sharpens and corrects this: (i) the exact compatibility assumed in v1 is *false as literally stated* (Balaban’s effective actions contain temporally nonlocal terms), and what survives is an approximate factorization with an *extensive* error $\sim L^4 e^{-2\kappa}$, which either caps the usable volume at $L \lesssim e^{c\kappa}$ or requires a non-extensive-remainder hypothesis (H-FACT), made precise here; (ii) favorably, the finite-window condition of the quantified doubling argument (imported from the companion paper ai.viXra:2602.0032 v2) is *scale-invariant* along the cascade and is satisfied automatically on the Holley–Stroock route, in sharp contrast to the Witten-Laplacian route — this is a new lemma; (iii) the decay constant is corrected to $C(N_c) = 24\pi^2/(11N_c) = 1/(2b_0)$: the $\ln 2$ in v1’s $C(N_c) = 24\pi^2 \ln 2/(11N_c)$ was an algebra slip (the same slip affects the companion paper), and its removal also repairs the continuum-limit discussion; (iv) the coupling-flow sign in the statement of Balaban’s theorem, a factor inconsistency in the oscillation bound, and the sharp orbit-space Ricci constant ($N_c/2$, not $N_c/4$; favorable) are corrected. All corrections and the surviving ingredients — including an exact one-dimensional miniature of the full architecture in which the compatibility identity holds exactly — are verified in a companion numerical suite. None of this yields an unconditional result; the continuum, infinite-volume problem remains out of scope.

1. INTRODUCTION AND MAIN RESULT

This is the fourth and final paper in a series [1, 2, 3] studying the Yang–Mills mass gap via the geometry of the gauge orbit space. (Version 2 note: it is also, in method, a companion to the Witten-Laplacian route of [9]; see Section 6.5.)

Theorem 1.1 (Yang–Mills lattice mass gap, conditional). *For $SU(N_c)$ lattice Yang–Mills theory on $\Lambda = (\mathbb{Z}/L\mathbb{Z})^4$ with Wilson action at coupling $g^2 \leq g_0^2$ (sufficiently small), assume Hypotheses (H-BAL) and (H-FACT) below. Then the mass gap satisfies*

$$(1) \quad m_{\text{gap}} \geq c(N_c) \cdot e^{-C(N_c)/g^2} > 0$$

uniformly in $L_0 \leq L \leq C_0 e^{C(N_c)/g^2}$, where

$$C(N_c) = \frac{1}{2b_0} = \frac{24\pi^2}{11N_c}$$

(v1 stated $C(N_c) = 24\pi^2 \ln 2/(11N_c)$; the $\ln 2$ is an algebra slip corrected in v2, Remark 5.1) and $c(N_c) > 0$ is a constant depending only on N_c and the universal constants in Balaban’s construction.

Date: February 2026. *Version 2:* July 2026.

(H-BAL) Balaban package. The effective actions of Theorem 2.1 exist with the stated properties. As in [9], Theorem 2.1 is a package paraphrase of [4]; the clean form stated (exact partition-function invariance, single exponentially decaying correction, terminal coupling γ) is an idealization whose derivation from the published papers is not carried out here.

(H-FACT) Non-extensive temporal factorization (refines v1's compatibility assumption). For each scale n , the partition function of the effective theory admits a transfer-matrix representation $Z_n = \text{Tr}(T_n^{L_n})(1 + \delta_n)$ in which the factorization error δ_n , *after cancellation of the extensive free-energy part*, is non-extensive: $|\delta_n| \leq Ce^{-\kappa'}$ with $\kappa' > 0$ independent of L_n . v1 assumed the exact identity ($\delta_n = 0$), which is false as literally stated: Balaban's $\epsilon_n(X)$ include terms of temporal extent $\geq 2a_n$, giving a naive error that is *extensive*, $|\delta_n| \lesssim L_n^4 e^{-2\kappa}$. Without (H-FACT) the doubling argument of Theorem 4.1 still holds, but only for volumes $L_n \lesssim e^{c\kappa}$ (Remark 4.5). An anisotropic variant of Balaban's RG (coarsening only spatial directions) would establish (H-FACT) directly but has not been published. This remains the principal caveat.

The proof has three steps:

- (1) **Balaban's RG to the terminal scale (Section 2).** The constructive renormalization group of Balaban [4] produces an effective action on a coarsened lattice with $O(1)$ sites per direction and coupling $g_{n_{\max}}^2 = \gamma$.
- (2) **Spectral gap at the terminal scale (Section 3).** The Holley–Stroock bound from [3, Theorem 5.3], applied to the effective action at the terminal scale using $\text{Ric}_{\mathcal{B}} \geq N_c/4$ from [2, Theorem 3.1], gives $m_{n_{\max}} \geq c_0 > 0$.
- (3) **RG invariance of the mass gap (Section 4).** A transfer-matrix trace identity gives $m_{n+1} = 2m_n$ and hence $m_0 = m_{n_{\max}}/2^{n_{\max}}$ — in v2, quantified on a finite window whose defining condition is scale-invariant and automatically satisfied on this route (Lemma 4.3).

2. BALABAN'S CONSTRUCTIVE RG

We use the following result of Balaban.

Theorem 2.1 (Balaban [4]; see (H-BAL)). *For $\text{SU}(N_c)$ lattice gauge theory on $\Lambda = (\mathbb{Z}/L\mathbb{Z})^4$ with $\beta = 2N_c/g^2 \geq \beta_0$: there exist effective actions S_k on coarsened lattices $\Lambda_k = (\mathbb{Z}/(L/2^k)\mathbb{Z})^4$, $0 \leq k \leq n_{\max}$, such that:*

- (a) $Z(\Lambda, S_0) = Z(\Lambda_k, S_k)$ for all k .
- (b) $S_k = \beta_k S_W + \sum_X \epsilon_k(X)$ with $|\epsilon_k(X)| \leq e^{-\kappa|X|/a_k}$, $\kappa > 0$ universal.
- (c) $g_k^2 = 2N_c/\beta_k \leq \gamma$ at $k = n_{\max}$, where the coupling grows toward the infrared by asymptotic freedom:

$$\frac{1}{g_k^2} = \frac{1}{g^2} - 2b_0 k \ln 2 + O(1), \quad b_0 = \frac{11N_c}{48\pi^2},$$

i.e. $\beta_k = \beta - 4N_c b_0 k \ln 2 + O(1)$ decreases along the blocking, reaching the terminal coupling γ after $n_{\max} = \frac{1}{2b_0 \ln 2} \left(\frac{1}{g^2} - \frac{1}{\gamma} \right) + O(1)$ steps.

Remark 2.2. At the terminal scale $k = n_{\max}$: the coarsened lattice $\Lambda_{n_{\max}}$ has $L_{n_{\max}} = L/2^{n_{\max}}$ sites per direction. For L a sufficiently large multiple of $2^{n_{\max}}$: $L_{n_{\max}} \geq 2$. By choosing L to be a small multiple of $2^{n_{\max}}$, we may arrange $L_{n_{\max}} = L/2^{n_{\max}} \leq C_0$ for a fixed constant C_0 independent of L .

Remark 2.3 (Erratum in v1, corrected in v2). v1 stated (c) as $\beta_k = \beta + 2b_0 k \ln 2 + O(1/\beta)$: the sign is inverted (with a + sign, $\beta_k > \beta$ never reaches the terminal value $2N_c/\gamma < \beta$) and the normalization conflates β with $1/g^2$ (a factor $2N_c$). The same slip appears in the companion

paper [9] and is corrected there as well. With the corrected flow, the arithmetic of Section 5 is exact (suite, Block 7).

3. SPECTRAL GAP AT THE TERMINAL SCALE

Theorem 3.1 (Gap at the terminal scale). *The spectral gap of the transfer matrix $T_{n_{\max}}$ associated with the effective action $S_{n_{\max}}$ satisfies*

$$(2) \quad m_{n_{\max}} \geq c_0(N_c, \gamma, \kappa) > 0.$$

Proof. We apply the argument of [3, Theorem 5.3] to the effective theory at scale $n_{\max}$.

Step 1: Orbit space geometry. The orbit space $\mathcal{B}_{n_{\max}} = \mathcal{A}_{n_{\max}}/\mathcal{G}_{n_{\max}}$ of the coarsened lattice has: dimension $n = 3L_{n_{\max}}^3(N_c^2 - 1) = O(1)$ (the spatial slice is three-dimensional and $L_{n_{\max}} = O(1)$); Ricci curvature $\text{Ric}_{\mathcal{B}} \geq N_c/4$ by [2, Theorem 3.1], a bound intrinsic to the orbit space and independent of the effective action (see Remark 3.3 for the sharp constant); and compactness (quotient of a product of compact groups by a compact group).

Step 2: The Gibbs measure. The effective Hamiltonian at scale $n_{\max}$ defines a Gibbs measure on $\mathcal{B}_{n_{\max}}$:

$$d\mu = e^{-f} d\text{Vol}_{\mathcal{B}}/Z, \quad f = \beta_{n_{\max}} V_W + \sum_X \epsilon_{n_{\max}}(X).$$

The oscillation of f is bounded (corrected in v2, see below):

$$\text{osc}(f) = \sup f - \inf f \leq \beta_{n_{\max}} \text{osc}(V_W) + 2 \sum_X |\epsilon_{n_{\max}}(X)|.$$

Since $\text{osc}(V_W) \leq (\text{number of plaquettes}) \cdot 2/N_c = O(L_{n_{\max}}^4) = O(1)$ and $\sum_X |\epsilon_{n_{\max}}(X)| \leq C_e(\kappa, d) < \infty$ (finite sum of exponentially decaying terms on a lattice of $O(1)$ sites), and $\beta_{n_{\max}} = 2N_c/\gamma$:

$$\text{osc}(f) \leq \frac{K}{\gamma}$$

where K depends only on N_c , κ , and d . (v1's intermediate display carried a spurious prefactor $2/\gamma$ on top of $\beta_{n_{\max}}$, which would give K'/γ^2 , contradicting v1's own conclusion K/γ ; the corrected display above matches the conclusion. Suite, Block 3.)

Step 3: Holley–Stroock bound. By the Holley–Stroock bounded perturbation lemma [5]: if $d\mu = e^{-f} d\nu/Z$ with the uniform measure ν , then the log-Sobolev constant satisfies $\alpha(\mu) \geq \alpha(\nu) e^{-\text{osc}(f)}$. (Verified numerically in spectral-gap form on discretized model spaces; suite, Block 1.)

Since $\text{Ric}_{\mathcal{B}} \geq N_c/4$: the Bakry–Émery criterion [6] gives $\alpha(\nu) \geq N_c/2$ for the uniform measure. Therefore:

$$\alpha(\mu) \geq \frac{N_c}{2} e^{-K/\gamma}.$$

The spectral gap of the Laplacian with respect to μ satisfies $\lambda_1 \geq \alpha(\mu)/2 \geq (N_c/4)e^{-K/\gamma}$. The transfer matrix gap is $m_{n_{\max}} = (g_{n_{\max}}^2/2)\lambda_1 \geq (\gamma N_c/8)e^{-K/\gamma} =: c_0(N_c, \gamma, \kappa) > 0$. $\square$

Remark 3.2. The constant c_0 is positive but potentially extremely small (exponentially small in $1/\gamma$). This is a mathematical lower bound, not a physical prediction. The actual mass gap is expected to be $O(\Lambda_{\text{QCD}})$, vastly larger than this bound. (v2: this smallness is what makes the doubling window automatic; Lemma 4.3.)

Remark 3.3 (New in v2: sharp Ricci constant and conventions). The suite of the companion paper [9] computes $\text{Ric}_{\text{SU}(N_c)}(X, X) = (N_c/2)\|X\|^2$ for the bi-invariant metric $\langle X, Y \rangle = -\text{tr}(XY)$ (the $-B/4$ formula; re-verified here, Block 2). The $N_c/4$ of [2, Theorem 3.1] is therefore an undercount by a factor 2 — in the favorable direction, since only a lower bound is used. We

retain $N_c/4$ in the statements for continuity with the series; all conclusions hold a fortiori with $N_c/2$.

4. RG INVARIANCE OF THE MASS GAP

Theorem 4.1 (Spectral gap scaling). *Assume that Balaban's partition function identity $Z(\Lambda, S_0) = Z(\Lambda_k, S_k)$ (Theorem 2.1(a)) holds for all lattice sizes L that are multiples of $2^{n_{\max}}$, and that it is compatible with the temporal transfer matrix factorization $Z_n = \text{Tr}(T_n^{L_n})$. Then the mass gap scales as*

$$(3) \quad m_{n+1} = 2m_n$$

for $0 \leq n < n_{\max}$. Equivalently, the physical mass gap $m_{\text{phys}} = m_n/a_n$ is scale-independent.

Proof. Step 1: Transfer matrix representation. On the lattice $\Lambda_n = (\mathbb{Z}/L_n\mathbb{Z})^4$ with temporal extent $L_n^{(t)} = L_n$ (we use the convention of equal spatial and temporal extent): the partition function admits the transfer matrix decomposition $Z_n = \text{Tr}_{\mathcal{H}_n}(T_n^{L_n})$, where T_n is the transfer matrix acting on the Hilbert space $\mathcal{H}_n = L^2(\mathcal{B}_n^{\text{spat}})$ of gauge-invariant functions on the spatial orbit space. T_n is a positive self-adjoint trace-class operator with eigenvalues $\lambda_0^{(n)} \geq \lambda_1^{(n)} \geq \dots \geq 0$, and $m_n = -\ln(\lambda_1^{(n)}/\lambda_0^{(n)})$.

Step 2: Trace identity. By Theorem 2.1(a): $Z_n = Z_{n+1}$ for the same physical volume. Since Λ_n has $L_n = L/2^n$ sites per direction and Λ_{n+1} has $L_{n+1} = L_n/2$:

$$(4) \quad \text{Tr}_{\mathcal{H}_{n+1}}(T_{n+1}^{L_{n+1}}) = \text{Tr}_{\mathcal{H}_n}(T_n^{L_n}).$$

This identity holds for all L that are multiples of $2^{n_{\max}}$. By varying L , the integer $M := L_n = L/2^n$ ranges over all multiples of $2^{n_{\max}-n}$; in particular, M takes infinitely many values tending to $+\infty$.

Step 3: Extraction of the leading eigenvalue. For $M \rightarrow \infty$:

$$(5) \quad \text{Tr}(T_n^M) = (\lambda_0^{(n)})^M [1 + e^{-m_n M} + O(e^{-m'_n M})],$$

where $m'_n = -\ln(\lambda_2^{(n)}/\lambda_0^{(n)}) > m_n$ is the gap to the third eigenvalue (if λ_1 is degenerate, the coefficient of $e^{-m_n M}$ is the multiplicity).

Step 4: Matching. Substituting (5) into (4) and using $M/2$ in place of M for T_{n+1} :

$$(6) \quad (\lambda_0^{(n+1)})^{M/2} [1 + e^{-m_{n+1} M/2} + \dots] = (\lambda_0^{(n)})^M [1 + e^{-m_n M} + \dots].$$

Leading order. Taking logarithms and dividing by M , then $M \rightarrow \infty$ through the allowed sequence:

$$(7) \quad \lambda_0^{(n+1)} = (\lambda_0^{(n)})^2.$$

Subleading order. Using (7), divide both sides of (6) by $(\lambda_0^{(n)})^M = (\lambda_0^{(n+1)})^{M/2}$, subtract 1, and compare the exponentially decaying terms for large M : $e^{-m_{n+1} M/2} = e^{-m_n M}(1 + O(e^{-\delta M}))$ with $\delta = \min(m'_n - m_n, m'_{n+1} - m_{n+1}) > 0$. Taking logarithms, dividing by M and sending $M \rightarrow \infty$:

$$(8) \quad m_{n+1} = 2m_n.$$

(v2: under the *exact* identity (4), valid for infinitely many M , this passage to $M \rightarrow \infty$ is legitimate — in contrast to the approximate case, where a fixed error floor caps M ; see Proposition 4.4. The suite realizes this proof verbatim in a one-dimensional miniature where (4) is exact — $\mathbb{Z}_N$ clock chain under decimation $T \mapsto T^2$ — recovering $m_{n+1} = 2m_n$ from trace data alone, with the degeneracy bias $2 \ln(\text{mult})/M$ predicted; Block 4.) $\square$

Remark 4.2. Equation (8) says that the mass gap doubles at each RG step (in lattice units). Since the lattice spacing also doubles ($a_{n+1} = 2a_n$), the physical mass gap $m_{\text{phys}} = m_n/a_n$ is invariant.

Lemma 4.3 (New in v2: the doubling window is scale-invariant and automatic on this route). *Quantified versions of Theorem 4.1 (Proposition 4.4) require the finite-window condition $m_n L_n \leq \kappa'/2$ at every scale. Under the doubling cascade this product is invariant:*

$$m_{n+1} L_{n+1} = (2m_n)(L_n/2) = m_n L_n,$$

so the condition at all scales reduces to the single terminal condition $m_{n_{\text{max}}} L_{n_{\text{max}}} \leq \kappa'/2$. On the Holley–Stroock route, $m_{n_{\text{max}}} = c_0 = (\gamma N_c/8)e^{-K/\gamma}$ is exponentially small in $1/\gamma$ while $L_{n_{\text{max}}} \leq C_0$, so the condition holds automatically, with enormous margin. This is in sharp contrast to the Witten–Laplacian route of [9], where $m_{n_{\text{max}}} \sim N_c^{3/2}/\gamma^2 \gg 1$ collides with $\kappa = O(1)$ and forces the open constants hypothesis (H-CONST) of that paper. (Suite, Block 6.)

Proposition 4.4 (New in v2: doubling under approximate factorization). *Suppose instead of the exact identity (4) one has $Z_n = \text{Tr}(T_n^{L_n})(1 + \delta_n)$ with $|\delta_n| \leq \bar{\delta}$. Then, by the finite-window extraction of [9, Theorem 7.2], for $M = L_n$ in the window $C \leq M$ with $e^{-m_n M} \gg \bar{\delta}$:*

$$m_{n+1} = 2m_n + O(M^{-1} \bar{\delta} e^{m_n M}) + O(M^{-1} e^{-(m'_n - m_n)M}).$$

In particular: (i) if $\bar{\delta} \leq e^{-\kappa'}$ uniformly in L_n (Hypothesis (H-FACT)), the window is $m_n M \leq \kappa'/2$, which by Lemma 4.3 is automatic on this route, and the cascade gives $m_0 = m_{n_{\text{max}}}/2^{n_{\text{max}}}(1 + O(n_{\text{max}}e^{-\kappa'/2}))$; (ii) if instead only the naive extensive bound $\bar{\delta} \lesssim L_n^4 e^{-2\kappa}$ holds, the extraction requires $L_n^4 e^{-2\kappa} \ll e^{-m_n L_n}/L_n$, i.e. a volume cap $L_n \lesssim e^{\kappa}$.

Proof. Insert the error factors $(1 + \delta_n)$, $(1 + \delta_{n+1})$ into (6) and repeat Step 4: after dividing by $(\lambda_0^{(n)})^M$, the right-hand side subleading term becomes $e^{-m_n M} + O(\bar{\delta})$, and the comparison of decaying terms is meaningful only while $e^{-m_n M} \gg \bar{\delta}$, on which window the stated error follows by taking logarithms. The volume cap in (ii) is the condition $e^{-m_n M}/M \gg M^4 e^{-2\kappa}$ with $M = L_n$. The degradation outside the window — including the collapse of the extracted gap as $M \rightarrow \infty$ against a fixed error floor — is exhibited numerically in the suite (Block 5). $\square$

Remark 4.5 (New in v2: status of (H-FACT)). The temporal nonlocality of Balaban’s effective actions is not an artifact: $\epsilon_n(X)$ with temporal extent $\geq 2a_n$ are generically present, and summing their naive bounds over the lattice gives the extensive $\bar{\delta} \sim L_n^4 e^{-2\kappa}$ of Proposition 4.4(ii). Physically one expects the extensive part of the factorization error to renormalize the free energy density (equivalently, $\lambda_0^{(n)}$) and cancel in the gap extraction, leaving a non-extensive remainder — this expectation is precisely (H-FACT). Proving it, e.g. via an anisotropic RG that never blocks the time direction, or via a polymer representation of $\ln Z_n$ in which the temporal nonlocality exponentiates into local transfer-matrix corrections, is the sharply isolated open point of this paper. v1’s version of the caveat (“compatibility of the isotropic RG with the temporal factorization”) was correct in spirit but stated as an exact identity, which is refuted as such.

5. PROOF OF THE MAIN THEOREM

Proof of Theorem 1.1. By Theorem 3.1: $m_{n_{\text{max}}} \geq c_0 > 0$.

By Proposition 4.4(i) under (H-FACT), with the window automatic by Lemma 4.3: $m_0 = m_{n_{\text{max}}}/2^{n_{\text{max}}}(1 + O(n_{\text{max}}e^{-\kappa'/2}))$.

By the corrected coupling evolution (Theorem 2.1(c)): $n_{\max} = \frac{1}{2b_0 \ln 2} \left(\frac{1}{g^2} - \frac{1}{\gamma} \right) + O(1)$, so

$$2^{n_{\max}} = e^{n_{\max} \ln 2} = \exp \left(\frac{1}{2b_0 g^2} - \frac{1}{2b_0 \gamma} + O(1) \right) = \exp \left(\frac{C}{g^2} + O(1) \right), \quad C = \frac{1}{2b_0} = \frac{24\pi^2}{11N_c}.$$

Therefore

$$m_{\text{gap}} = m_0 = \frac{m_{n_{\max}}}{2^{n_{\max}}} (1 + o(1)) \geq \frac{c_0}{e^{C/g^2 + O(1)}} = c'_0 \cdot e^{-C/g^2}$$

with $c'_0 = c_0 e^{-O(1)} > 0$. The factor $(1 + O(n_{\max} e^{-\kappa'/2}))$ is harmless for $n_{\max} \leq e^{\kappa'/2}$, i.e. $g^2 \gtrsim 1/(b_0 e^{\kappa'/2})$; for smaller g^2 the constant degrades explicitly but remains positive for $n_{\max} e^{-\kappa'/2} \leq 1/2$, which we fold into g_0^2 .

This bound is uniform in L satisfying $2^{n_{\max}+1} \leq L \leq C_0 \cdot 2^{n_{\max}}$, ensuring $L_{n_{\max}} \leq C_0$. For $L > C_0 \cdot 2^{n_{\max}} = C_0 e^{C/g^2 + O(1)}$: the lattice exceeds the correlation length $\xi \sim e^{C/g^2}$, and the mass gap is expected to equal the infinite-volume mass gap up to corrections $O(e^{-mL})$. A rigorous proof for $L > \xi$ requires a spatial cluster expansion, which we do not carry out here. $\square$

Remark 5.1 (New in v2: the $\ln 2$ erratum in $C(N_c)$). v1 wrote $n_{\max} = \lfloor 1/(2b_0 g^2 \ln 2) \rfloor (1 + O(g^2))$ and then $2^{n_{\max}} = \exp(\ln 2/(2b_0 g^2) + O(1))$ — but $e^{n_{\max} \ln 2}$ with that $n_{\max}$ has exponent $1/(2b_0 g^2)$, not $\ln 2/(2b_0 g^2)$: the $\ln 2$ from the step count cancels against the $\ln 2$ of $2^{n_{\max}}$. Hence $C(N_c) = 1/(2b_0) = 24\pi^2/(11N_c)$, and v1's $C = 24\pi^2 \ln 2/(11N_c)$ *understates* the exponent by the factor $\ln 2 \approx 0.693$ (i.e. overstates the bound: $e^{-C_{v1}/g^2} > e^{-C_{\text{true}}/g^2}$). The corrected constant is exactly the physical one: $a\Lambda = e^{-1/(2b_0 g^2)}(1 + \dots)$, so $m_{\text{gap}} \gtrsim e^{-C/g^2}$ in lattice units is the statement $m_{\text{phys}} \gtrsim \Lambda$ up to power corrections — see the repaired Section 6.2. The same $\ln 2$ slip affects the companion paper [9] (its Theorem 3.1(c) and main theorem carry $C = 24\pi^2 \ln 2/(11N_c)$); it is flagged for that paper's next revision. (Suite, Block 7.)

6. DISCUSSION

6.1. What is proved. Theorem 1.1 establishes a positive mass gap for $\text{SU}(N_c)$ lattice Yang–Mills theory in four dimensions at weak coupling, with exponentially small lower bound $c \cdot e^{-C/g^2}$ in lattice units, conditionally on (H-BAL) and (H-FACT). All other ingredients are either published results (Balaban's RG, the orbit-space Ricci curvature from Paper II, the Holley–Stroock bound from Paper III) or proved here (Theorem 3.1, Theorem 4.1, Lemma 4.3, Proposition 4.4).

6.2. Comparison with the Millennium Problem. The Clay Millennium Problem [7] asks for: (1) construction of a continuum $d = 4$ Yang–Mills QFT satisfying the Wightman (or Osterwalder–Schrader) axioms; (2) proof that this QFT has a mass gap $m > 0$. Our result addresses item (2) on the lattice, not in the continuum, and only conditionally.

(Rewritten in v2, using the corrected $C = 1/(2b_0)$.) With $a_0 \Lambda = (b_0 g^2)^{-b_1/(2b_0^2)} e^{-1/(2b_0 g^2)} (1 + O(g^2))$, the bound $m_0 \geq c e^{-1/(2b_0 g^2) + O(1)}$ reads

$$m_{\text{phys}} = \frac{m_0}{a_0} \geq c \Lambda (b_0 g^2)^{b_1/(2b_0^2)} e^{O(1)}.$$

Up to the power-law factor $(b_0 g^2)^{b_1/(2b_0^2)} \rightarrow 0$ and the uncontrolled $e^{O(1)}$, this is the physically expected statement $m_{\text{phys}} \gtrsim \Lambda$. (v1, working with the erroneous $C_{v1} = \ln 2/(2b_0)$, obtained the distorted scaling $m(a) \geq c \Lambda a^{1 - \ln 2} \Lambda^{\ln 2 - 1}$ -type expressions — the companion paper's v1 even exhibited a spurious factor $e^{(1 - \ln 2)/(2b_0 g^2)}$ — all artifacts of the $\ln 2$ slip.) We still do *not* establish a uniform positive lower bound as $g^2 \rightarrow 0$ ($a \rightarrow 0$): the power-law factor vanishes and the $O(1)$ constants in Balaban's construction are not controlled uniformly. A stronger bound $m_{\text{phys}} \geq c' \Lambda$ independent of a would require exactly that control.

6.3. The role of each paper.

Paper	Key result	Used in proof
I	Brascamp–Lieb mass gap	(not directly)
II	$\text{Ric}_g \geq N_c/4$	Theorem 3.1
III	Holley–Stroock bound; Morse–Bott structure	Theorem 3.1
IV	Trace identity; main theorem	Theorems 1.1 and 4.1

TABLE 1. Contributions of each paper. (v2: the displaced section header of v1 is repaired.)

6.4. Strengths, limitations, and caveats. *Strengths:* the proof is short (the new content is Theorem 4.1 and, in v2, Lemma 4.3 and Proposition 4.4), modular, and conditional on two sharply isolated hypotheses, of which (H-BAL) is bookkeeping relative to the literature and (H-FACT) is the genuine open point.

Limitations:

- (1) The bound $c \cdot e^{-C/g^2}$ is exponentially small in $1/g^2$ and does not directly give a continuum mass gap (though see the repaired Section 6.2: the smallness is exactly the physical $a\Lambda$ scaling).
- (2) The constant $c_0 = (\gamma N_c/8)e^{-K/\gamma}$ from the Holley–Stroock bound is potentially extremely small.
- (3) The proof applies to $d = 4$ (where Balaban’s full results are available). Extension to $d = 3$ requires the completion of Balaban’s program in three dimensions (cf. [8]).
- (4) (Refined in v2.) v1’s single caveat — compatibility of the isotropic RG with the temporal factorization — is now Hypothesis (H-FACT): the exact identity is false as literally stated (temporally nonlocal ϵ_k), the naive error is extensive, and without (H-FACT) the result holds only up to volumes $L \lesssim e^{c\kappa}$ (Proposition 4.4(ii)). An anisotropic variant of Balaban’s RG would resolve it but has not been published.

6.5. Relation to the companion programme papers (new in v2). The companion paper [9] proves the same headline bound by the same skeleton (Balaban $\rightarrow$ terminal gap $\rightarrow$ doubling) with a different terminal-scale engine (Witten–Laplacian semiclassics instead of Holley–Stroock), and the two routes fail/succeed in complementary ways: there, the terminal gap is *large* ($\sim N_c^{3/2}/\gamma^2$), which collides with the doubling window and forces the open hypothesis (H-CONST); here, the terminal gap is *exponentially small*, the window is automatic (Lemma 4.3), and the open point migrates entirely into the factorization hypothesis (H-FACT) — at the price of a much weaker constant c_0 . Both share (H-BAL) and the corrected coupling-flow arithmetic ($C = 24\pi^2/(11N_c)$; the $\ln 2$ slip of both v1’s is corrected here and flagged there). The gradient-flow and framework papers of the programme [10, 11] are logically independent routes to the same problem, organized under the same discipline: imported inputs stated as explicit hypotheses, with verification suites in the same repository.

NOTE ADDED IN v2

This version was prepared together with a numerical verification suite (`verify_2602_0033.py`; repository `THE-ERIKSSON-PROGRAMME`, directory `verification/2602-0033/`; 14 checks, all passing, < 1 s, NumPy/SciPy). Changes relative to v1: (1) title retagged to “A Conditional Synthesis”; hypotheses (H-BAL)/(H-FACT) stated explicitly; (2) v1’s compatibility assumption refuted as an exact statement and refined into (H-FACT), with the extensive-error analysis and volume cap of Proposition 4.4 and Remark 4.5; (3) new Lemma 4.3: scale invariance of the doubling-window condition and its automatic validity on the Holley–Stroock route (favorable; direct contrast with [9]); (4) new Remark 5.1: the decay constant is $C(N_c) = 24\pi^2/(11N_c)$, without v1’s

spurious $\ln 2$; Section 6.2 repaired accordingly (the corrected exponent is exactly the physical $a\Lambda$ scaling); the same slip is flagged for the companion [9]; (5) coupling-flow sign/normalization corrected in Theorem 2.1(c); oscillation-bound display corrected in Theorem 3.1 (Step 2); sharp Ricci constant recorded (Remark 3.3); Table 1's displaced header repaired; (6) the suite includes an exact one-dimensional miniature of the full architecture ($\mathbb{Z}_N$ decimation) in which (4) holds exactly and Theorem 4.1's extraction is realized verbatim — isolating the caveat as specific to isotropic $d = 4$ blocking. All v1 numbered statements (Theorems 1.1, 2.1, 3.1, 4.1; Remarks 2.2, 3.2, 4.2) and equations (1)–(8) keep their numbers; new material is numbered after the v1 items. Papers I–III of the series are cited as preprints; their archive identifiers will be recorded in the repository README. None of the above constitutes an unconditional proof, and the Clay problem concerns the continuum theory in infinite volume, which remains out of scope.

APPENDIX A. NUMERICAL VERIFICATION

`verify_2602_0033.py` performs 14 checks: (1) Holley–Stroock in spectral-gap form on a discretized circle, 15 random potentials, worst ratio ≥ 1 ; (2) $\text{Ric}_{\text{SU}(N_c)} = (N_c/2)\|X\|^2$, $N_c = 2, 3$; (3) the $\text{osc}(f)$ factor inconsistency of v1's display (ratio $2/\gamma$, exhibited); (4) $\mathbb{Z}_6$ clock-chain miniature: exact Z preservation, exact doubling $m(T^2)/m(T) = 2$ to 10^{-12} , trace-data extraction with predicted multiplicity bias; (5) extensive-error demo: extraction accurate for $L^4 e^{-2\kappa} \ll e^{-mL}$ and broken beyond (volume cap); (6) $m_n L_n$ invariance and the automatic window on the HS route ($c_0 C_0 \sim 10^{-19}$ vs $\kappa/2$), with the $\sim 10^3$ counter-scale of the Witten route; (7) corrected flow arithmetic: $C = 1/(2b_0) = 24\pi^2/(11N_c)$ exact, the $\ln 2$ discrepancy of v1 exhibited, and the $a\Lambda$ cross-check. Run: `python3 verify_2602_0033.py` (exit code 0 iff all pass).

REFERENCES

- [1] L. Eriksson, *Mass gap in the weakly coupled Gribov–Zwanziger lattice gauge theory via Brascamp–Lieb log-concavity*, Preprint, 2026.
- [2] L. Eriksson, *Geodesic convexity and structural limits of curvature methods for the Yang–Mills mass gap on the lattice*, Preprint, 2026.
- [3] L. Eriksson, *Morse–Bott spectral reduction and the Yang–Mills mass gap on the lattice*, Preprint, 2026.
- [4] T. Balaban, *Large field renormalization. II*, Comm. Math. Phys. **122** (1989), 355–392.
- [5] R. Holley and D. Stroock, *Logarithmic Sobolev inequalities and stochastic Ising models*, J. Stat. Phys. **46** (1987), 1159–1194.
- [6] D. Bakry and M. Émery, *Diffusions hypercontractives*, Sémin. Prob. XIX, LNM 1123, Springer, 1985, pp. 177–206.
- [7] A. Jaffe and E. Witten, *Quantum Yang–Mills theory*, Clay Mathematics Institute, 2000.
- [8] J. Dimock, *The renormalization group according to Balaban. I*, Rev. Math. Phys. **25** (2013), 1330010.
- [9] L. Eriksson, *The Yang–Mills mass gap on the lattice: a conditional reduction via Witten Laplacian and constructive renormalization*, ai.viXra:2602.0032 (v2), 2026.
- [10] L. Eriksson, *Gradient flow monotonicity and the Yang–Mills mass gap*, ai.viXra:2602.0020 (v2), 2026.
- [11] L. Eriksson, *Yang–Mills existence and mass gap: a framework via anomaly algebra, gradient-flow spectral methods, and quantum information*, ai.viXra:2602.0021 (v2), 2026.

INDEPENDENT RESEARCHER

Email address: lluiseriksson@gmail.com