

Erratum and restoration (v2): one symbol carries two angles, several printed statements are false, (H1) is not discharged, and (H-FACT) does not replace (H2)

Morse–Bott Spectral Reduction and the
Yang–Mills Mass Gap on the Lattice

Lluis Eriksson • ai.viXra:2602.0035 • July 2026

This version does two things at once. It *restores* the revision prepared for this record in July 2026, which was mis-filed onto `ai.viXra:2602.0052` and so never appeared here; and it *corrects* that revision. **It refutes several printed statements**, and they do not all fail for the same reason. **The angle ambiguity accounts for some, but not all, of the defects:** the symbol θ carries both the loop holonomy angle and the per-link background angle, which differ by a factor L (§2), and that invalidates equation (5), the shift used in equation (6), the Fourier labels of Proposition 3.3 and the frequencies (9). **Proposition 2.2(b), the additive split in equation (6), the centraliser count and the normal-bundle classification fail for independent reasons.** Two status claims — that (H1) is discharged, and that (H2) is refined into (H-FACT) — are withdrawn in §4.

1 What happened to this record

The file prepared as v2 of this paper was uploaded, during the replacement campaign of 2026-07-06, onto the record `ai.viXra:2602.0052`, whose own paper (*Interface Lemmas*) it displaced. This record therefore still showed v1 while three companion papers cited “2602.0035 v2” — a version that existed only as the mis-filed file. The pages after this note are that revision, restored; the corrections below apply to it.

2 One symbol, two angles

§2.1 defines the total holonomy $g_\mu = \prod_{t=0}^{L-1} U_\mu(t\hat{\mu}) = e^{i\Theta_\mu}$ — writing Θ for what the manuscript calls θ . The flat, translation-invariant representative that the Fourier analysis then uses must have

$$U_\mu(x) = e^{i\Theta_\mu/L},$$

because there are L links around the cycle. **Any corrected proof must distinguish the loop-holonomy angle from the per-link angle.** Two consequences follow immediately, and both are printed the other way.

The covariant symbol. For a fluctuation in root direction α the eigenvalue of the covariant difference is

$$4 \sin^2\left(\frac{k_\mu + \alpha(\Theta_\mu)/L}{2}\right) \quad \text{— in flat notation, } k + \text{alpha}(\text{Theta})/L \text{ —}$$

and not $4 \sin^2((k_\mu + \alpha(\Theta_\mu))/2)$.

The metric. In total-holonomy coordinates the per-link variation is $\delta\Theta_\mu/L$, so $\sum_x |\delta\Theta_\mu/L|^2 = L^d L^{-2} |\delta\Theta_\mu|^2 = L^{(d-2)} |\delta\Theta_\mu|^2$, whereas equation (5) prints $L^d |\delta\Theta_\mu|^2$.

Only two readings are available, and each demands a correction. If θ is the total-holonomy angle, as Proposition 2.2 says, then the factors $1/L$ are missing from the metric and from the shifted momentum. If θ is the per-link angle, it is no longer the coordinate of Proposition 2.2 and its period is $2\pi/L$, with the corresponding gauge identifications. **The manuscript uses one symbol for both without declaring a change of variable**, so equation (5), the mass term $m_a^2(\theta)$, the proof of (6) and the frequencies (9) cannot all be correct as printed.

3 What is false as printed

3.1 Proposition 2.2(b): the $SU(2)$ quotient is misidentified

The proposition states $\mathcal{M}_{\text{flat}} \cong [0, \pi]^d / \mathbb{Z}_2$ with $\theta_\mu \mapsto \pi - \theta_\mu$. But the $SU(2)$ Weyl action is simultaneous inversion: $(\theta_1, \dots, \theta_d) \mapsto (-\theta_1, \dots, -\theta_d)$ on $(S^1)^d$. The printed proof passes from $\theta \mapsto 2\pi - \theta$ to $\theta \mapsto \pi - \theta$ with no change of coordinates; these are not the same map. The paper contradicts itself on it: under inversion the fixed points are the 2^d tuples with $\theta_\mu \in \{0, \pi\}$, while under the printed map the only interior fixed point is $(\pi/2, \dots, \pi/2)$ — and Proposition 3.3 then places $\theta = 0$ on the orbifold locus, which the printed map does not fix. **Proposition 2.2(b) is false as printed.** The correct statement is $\mathcal{M}_{\text{flat}}^{SU(2)} = (S^1)^d / \{\theta \sim -\theta\}$; a cube presentation is possible but needs its boundary identifications fixed with care.

3.2 Equation (5): the metric factor

As computed in §2, the induced metric in total-holonomy coordinates carries L^{d-2} , not L^d . **Equation (5) is false as printed** under the reading Proposition 2.2 fixes.

3.3 Equation (6): two independent failures

The theorem asserts $\text{Hess}(V_{\text{pot}})^\perp = \frac{2}{g^2} \sum_{k \neq 0, a} [\widehat{k}^2 + m_a^2(\theta)] |\delta A_\perp^a(k)|^2$.

First failure — the split. Its own proof gives the covariant-difference eigenvalue as $4 \sum_\mu \sin^2((k_\mu + \phi_\mu)/2)$, and

$$4 \sin^2\left(\frac{k + \phi}{2}\right) \neq 4 \sin^2\left(\frac{k}{2}\right) + 4 \sin^2\left(\frac{\phi}{2}\right).$$

A test inside the smooth stratum: $SU(2)$, $L = 6$, $\Theta = \pi/2$, $\alpha = 2$, $k = \pi/3$. Then $\phi = \alpha\Theta/L = \pi/6$, the left side is $4 \sin^2(\pi/4) = 2$ and the right side is $3 - \sqrt{3} \approx 1.268$. $\Theta = \pi/2$ is a regular holonomy, so the test lies where Theorem 3.1 is restricted. *An earlier draft used $k = \phi = \pi/3$, which under the corrected convention forces $\Theta = \pi$ — central holonomy, outside the smooth stratum, and so an avoidable objection about the domain.* **Equation (6) is not valid as a general equality away from trivial holonomy** — stated that way because special pairs (k, ϕ) may agree by accident. The proof concedes it without saying so: it obtains (6) *at trivial holonomy* and, elsewhere, replaces the equality by an inequality with a cross term. That inequality may serve an argument for positivity; it does not establish the displayed formula.

Second failure, independent of the first — the shift. By §2 the shift is $\phi_\mu = \alpha(\Theta_\mu)/L$, not $\alpha(\Theta_\mu)$. So even the inequality, and the mass term m_a^2 , are written with the wrong argument.

Why the suite does not catch it. The Hessian is validated at trivial holonomy, where the shift vanishes and the defect is therefore absent. This is the same pattern as the one-dimensional miniature in [ai.vixra:2602.0033](#): **the test does not examine the regime in which the defect lives**. Two papers of this chain now have a check with that property: a fact about how test cases are chosen, not a coincidence.

3.4 Proposition 3.3: the count, and the Fourier labels

The proposition counts $d(N_c^2 - 1 - r)$ massless non-Cartan $k = 0$ constant modes at any point of enhanced stabiliser. Two things are wrong.

The count. Its own mass formula makes a root direction α massless exactly when its pairing with every holonomy angle is an integer multiple of 2π — that is, exactly on the common centraliser. **The correct count is $d(\dim \mathfrak{z}_\Theta - r)$** , in flat notation d times $(\dim \mathfrak{z}_\Theta - r)$, with $\mathfrak{z}_\Theta$ the algebra of the common centraliser; the printed count holds only when that centraliser is the whole group. For $SU(3)$ at a point with centraliser $S(U(2) \times U(1))$, $\dim \mathfrak{z}_\Theta = 4$ and $r = 2$, giving $2d$, not $6d$.

The labels. **Covariantly constant modes need not have Fourier label $\mathbf{k}=0$.** Take $SU(2)$ with the central holonomy $g_\mu = -I$, $\Theta_\mu = \pi$, and the adjoint root $\alpha = 2$: the per-link shift is $2\pi/L$, so $\lambda_\alpha(k = 0) = 4d \sin^2(\pi/L) > 0$, while the zero sits at the allowed momentum $k_\mu = -2\pi/L$. **The centraliser is everything here, so this is precisely a case the proposition claims**, and the modes it claims at $k = 0$ are not at $k = 0$. The accurate statement is: at an enhanced-stabiliser holonomy there

are $d(\dim \mathfrak{z}_\Theta - r)$ additional covariantly constant centraliser directions, and their Fourier labels depend on the chosen flat representative.

And “normal to $\mathcal{M}_{\text{flat}}$ ” is not a valid classification at a singular point, where there is no ordinary normal bundle. At the trivial point the quartic form is $\propto \sum_{\mu < \nu} |X_\mu, X_\nu|^2$, which *vanishes on commuting tuples*: taking all X_μ proportional to one non-Cartan generator keeps the connection flat, with no quartic growth. The accurate classification is that all constant directions lie in the kernel of the Hessian at the trivial point, commuting tuples lie in the tangent cone to the flat connections, and non-commuting tuples exhibit the positive quartic term.

What the numerics support is that for $SU(2)$ on a 2^3 spatial lattice the Hessian kernel has dimension 30 and a non-commuting direction has quartic ratio 15.99. **Verified, not proved**, and for that instance.

3.5 Proposition 4.1: the cutoff does not do what it is asked to do

Correcting the Fourier labels refutes two further printed claims, and this propagates into the Born–Oppenheimer reduction rather than stopping at Proposition 3.3.

The cutoff. Proposition 4.1 says the sums (8) run over $k \neq 0$ and thereby omit the toron modes. The mode of §3.4 — $g_\mu = -I$, $\alpha = 2$, $k_\mu = -2\pi/L$ — has $k \neq 0$, lies *inside* those sums, and has zero frequency. **The $k \neq 0$ cutoff does not exclude all covariantly constant modes.**

The uniform gap. The printed normal-mode bound is $\omega_{\perp, \min} = 2\sqrt{2}\sin(\pi/L) > 0$, uniform over the smooth stratum. Put $\Theta_\mu = \pi - \epsilon$ with $\epsilon \neq 0$, so the holonomy is regular, and keep $k_\mu = -2\pi/L$: then $k_\mu + \alpha\Theta_\mu/L = -2\epsilon/L$, the eigenvalue is $\lambda_\epsilon = 4d\sin^2(\epsilon/L)$ and the corresponding frequency of (9) is $\omega_\epsilon = \sqrt{2\lambda_\epsilon} = \sqrt{8d}|\sin(\epsilon/L)| \rightarrow 0$ as $\epsilon \rightarrow 0$, at *fixed* L — written as a frequency so that the comparison is with like for like — while the printed bound does not move. **The normal-mode gap closes on approach to the enhanced-stabiliser locus.** So there is no such uniform bound over the smooth stratum.

Both claims of Proposition 4.1 are false as printed. The printed harmonic Born–Oppenheimer reduction, and any argument relying on a uniformly gapped normal sector, can at most hold on a region uniformly separated from the enhanced-stabiliser locus, with a holonomy-dependent normal gap. *A more general global theorem is not excluded*: it would additionally require the quartic-sector construction listed in §7(C2). Equations (8)–(10), the bound on the remaining eigenvalues, and the proof’s explanatory remark all inherit this.

3.6 What this leaves of Morse–Bott

The smooth-stratum conclusion is not refuted but is not established by the printed proof. Equation (6) is false as an equality, and the shift it uses is wrong by a factor L , so the printed non-degeneracy argument is gone. A corrected proof must distinguish the two angles and use the actual shifted spectrum; the conclusion is plausible and may well survive, but this manuscript does not establish it. The frequencies (9) inherit both defects, in a proposition the text already labels *formal*.

4 (H1), (H2), and what Theorem 6.3 is

4.1 (H1) is not discharged

The restored text asserts in four places — Proposition 6.4, Remark 6.6, item (4) of the Note added, and §7.3 — that (H1) is discharged by `ai.viXra:2602.0033 v2`, Theorem 3.1. **All four are withdrawn.** That companion’s erratum records that its Theorem 3.1 is not established by its printed proof, and that its Theorem 4.1 — the doubling cascade — is independently not established, its transfer operator changing with the isotropic volume over which the limit is taken.

Only the negative status assessment of Proposition 6.4 remains accurate: at the relevant coupling the terminal gap is open, Balaban’s regime being intermediate rather than strong. Its *positive* appeal to Corollary 5.5 is not available either, since that corollary rests on Theorem 5.3 against Vol_B and on the withdrawn Ricci bound; it would need the corrected physical-measure construction and $(A\text{-osc})_\nu$ of §5.4.

Moreover (H1) as printed — “*the spectral gap of the transfer matrix associated with $S_{n_{\max}}$ on the orbit space*” — **does not by itself fix a space, a measure or an operator**, as the audit of 2602.0033 showed. A usable (H1) must name the physical transfer operator, its fixed spatial slice and its Hilbert space.

4.2 (H-FACT) does not replace (H2)

Proposition 6.5, Remark 6.6, §7.3 and Note added (4) present (H2) — *the physical gap is RG-invariant* — as *refined* into (H-FACT), the non-extensive temporal factorisation hypothesis of the companion. **That is withdrawn too, and it is not a matter of wording.** Getting from (H-FACT) to (H2) was the work of Theorem 4.1 and Proposition 4.4 of 2602.0033, and that extraction is exactly what fails. **H-FACT does not by itself imply H2** with the printed proofs; it is at most one input to a future proof of (H2). §7.3, which called (H-FACT) the remaining analytic content of the programme, is withdrawn with it.

4.3 Theorem 6.3

The algebraic implication survives on admissible volumes, under *Balaban’s Theorem 6.1* — (H-BAL) in the coordinated vocabulary, and the theorem’s own first line assumes it — together with the original (H1) and (H2), once those hypotheses are attached to precisely specified operators. *An earlier draft of this note omitted (H-BAL) from that list.*

And not with the same prefactor. **Equation (18), not the final proof line, retains the c_0 of (H1):** it states $m_{\text{gap}} \geq c_0 e^{-C/g^2}$ with the very c_0 of that hypothesis, whereas **the proof already introduces c'_0** , writing $m_0 \geq c_0 e^{-C/g^2 - O(1)} = c'_0 e^{-C/g^2}$ after absorbing the bounded remainder. **The exponential constant C is unchanged** — $C = 1/(2b_0) = 24\pi^2/(11N_c)$ — but the displayed theorem should carry the **possibly reduced positive prefactor c'_0** . This is not an obstruction to the e^{-C/g^2} behaviour; it corrects the displayed statement, and the defect is in (18) alone. *An earlier draft of this note said the proof absorbed the remainder into the same symbol. It does not: that reading came from a text extraction, in which the prime on c'_0 does not survive — a failure mode this programme has recorded before, and one to check against the rendered page.* **Theorem 6.3 retains the original H2 as a hypothesis;** a hypothesis that is not discharged is still a hypothesis, and it may *not* be presented as a theorem under (H1) plus (H-FACT). But the statement is not literally established with its printed quantifier over all L : the proof fixes $n_{\max}$ by the coupling and works on lattices of extent $L/2^{n_{\max}}$, which is a lattice only for $L \in 2^{n_{\max}}\mathbb{N}$, or for other L under an argument not given. *An earlier draft of this note called the statement untouched while conceding the missing condition; those cannot both hold, and the second is the correct one.*

A citation loop is declared. This paper’s (H1) was said to be discharged by 2602.0033, whose Theorem 3.1 proves its terminal gap by “applying the argument of” this paper’s Theorem 5.3. Neither discharges a hypothesis of the other.

5 Theorem 5.3: on a measure not identified with the physical one

5.1 Part (i): recoverable, but not by the printed proof

Part (i) takes H to be “(the Friedrichs extension of) an elliptic operator with smooth potential on $L^2(B)$ ”. $L^2(B)$ is written without a measure; and $B = \mathcal{A}/\mathcal{G}$ is a *stratified space*, the reducible connections having stabilisers of positive dimension, so “the Friedrichs extension” presupposes a closed form the paper does not fix. This is the gap that forces the withdrawal of Corollaries 3.2–3.3 in [ai.viXra:2602.0036](#). **The repair does not need the quotient at all:** by (I1) of the appended page, pullback is an isometry $L^2(B, \nu) \cong L^2(\mathcal{A}, \text{Vol}_{\mathcal{A}})^{\mathcal{G}}$ with $\nu = \pi_{\#} \text{Vol}_{\mathcal{A}}$, and $\mathcal{A} = SU(N_c)^E$ is a smooth compact connected manifold, where compact resolvent, positivity improvement and uniqueness of the ground state are standard. It gives $\Delta > 0$ at fixed volume, qualitative only.

5.2 Part (ii): the reference measure

Step 1 puts $d\mu_0 = \psi_0^2 d\text{Vol}_B$. The ground-state identity is correct and is the right correction of v1, whose measure e^{-2V_{pot}/g^2} was dimensionally inconsistent and failed numerically by more than fifty per cent. **But**

$L^2(B, d\text{Vol}_B)$ is **not identified with the physical measure**: the physical sector is $L^2(B, \nu) \cong L^2(\mathcal{A})^{\mathcal{G}}$ by (I1). On the regular stratum $d\nu = w d\text{Vol}_B$, so the two measures are related by that density, and **the associated generators differ by a drift** $\langle \nabla \log w, \nabla \cdot \rangle$. They coincide only under an additional constancy or identification statement for w , none of which is established here. *No inequality between the two is asserted*; what is asserted is that no identification is available.

Two warnings about symbols, both load-bearing. First, **the appended page records this paper in two rows**: an *as-printed* row carrying $\mu_{0,\text{print}} = \psi_{0,\text{print}}^2 d\text{Vol}_B$, and a *candidate physical-restatement* row carrying $\mu_{0,\nu} = \psi_{0,\nu}^2 d\nu$. An earlier single-row version exposed the intended physical object but failed to record the printed one, and is superseded. Second, **the ν of the appended page is not the ν of the restored pages**: the page writes $\nu_{\text{push}} = \pi_{\#} \text{Vol}_{\mathcal{A}}$, while Step 2 of Theorem 5.3 calls $\nu_{\text{Vol}} = d\text{Vol}_B / \text{Vol}_B(B)$ “the uniform measure ν ”. Identifying the two is precisely what (I2b) forbids.

5.3 Part (ii): the base inequality

Step 2 claims $\alpha(d\text{Vol}_B) \geq N_c/2$ from $\text{Ric}_B \geq N_c/4$ by Bakry–Émery, hence $\lambda_1(d\text{Vol}_B) \geq \alpha/2 \geq N_c/4$. That pointwise Ricci bound is Theorem 3.1 of [ai.viXra:2602.0036](#) and is **not established**: its O’Neill trace omits mixed horizontal–vertical sectional terms which for a bi-invariant metric are non-negative and subtract. **No pointwise Ricci bound on the orbit space may presently be cited.**

5.4 One repair for both — and what it costs

Work on $\mathcal{H}_{\text{phys}} = L^2(\mathcal{A}, \text{Vol}_{\mathcal{A}})^{\mathcal{G}} \cong L^2(B, \nu)$; take $H_{\nu} = \frac{g^2}{2} K_{\nu} + V_{\text{pot}}$ — in flat notation, take $H_{\text{nu}} = (g^2/2) K_{\text{nu}} + V_{\text{pot}}$, and **not** the gauge-fixed $V = S_{\text{YM}} - \log \det M$, whose symbol this manuscript has already fixed — with K_{ν} the operator of the closed form on $L^2(B, \nu)$; put $d\mu_0 = \psi_{0,\nu}^2 d\nu$. For the base inequality use [ai.viXra:2602.0046 v3](#), which gives $\text{RCD}^*(N_c/4, \dim \mathcal{A})$ for (B, d_B, ν) , is valid across the singular locus, and *explicitly bypasses* the O’Neill computation. Holley–Stroock against ν then gives the rate.

This is not free. H_{ν} is a separately defined operator that **has not been identified with the printed H** ; correspondingly its ground state $\psi_{0,\nu}$ has not been identified with the printed one. So the printed oscillation hypothesis does not supply what the corrected route needs, and that route requires a *new* hypothesis, in flat notation (A-osc)_{nu}:

$$(\text{A-osc})_{\nu} : \quad \text{osc}(-\log \psi_{0,\nu}^2) \leq C L^d / g^2.$$

The route is worth recording — it replaces two unavailable inputs with one available one, and takes the quotient singularities out of the base inequality — but it also *moves* the semiclassical hypothesis to a new object. It is not carried out here and gives no uniformity in L . **The present fixed-volume argument does not by itself establish (H1)**, and it has no bearing on (H2). *An earlier draft of this note said it bore in no way on (H1) either, while §7 called it a route towards (H1); §7(B) now states what would have to be built for the two to meet.*

6 What this paper does establish, stated exactly

1. **The Hamiltonian/path-integral distinction** of §7.2, in its narrow form: the multiplication potential $V_{\text{pot}} = S_{\text{YM}}$ does not contain the gauge-fixed term $-\log \det M$, which establishes a distinction between the two formulations. Its third listed ingredient, $\text{Ric}_B > 0$, is not currently available, and this narrow statement does not depend on it. **It does not establish that all orbit-volume or Faddeev–Popov geometry is absent from the physical Hamiltonian**, whose kinetic operator is K_{ν} . **Orbit-volume geometry is present in the physical measure and in the drift** $\langle \nabla \log w, \nabla \cdot \rangle$ **of K_{ν}** ; the Faddeev–Popov term is present in the gauge-fixed formulation; and this note does not identify w with $\det M$, so it does not claim that the same term has merely moved from one formulation to the other. *Two earlier drafts erred in opposite directions: the first certified the distinction unqualified, the second said the geometry “has moved from” the potential — itself an identification, asserted in the sentence denying one.* **And the measure formula printed alongside it is itself false: Remark 1.1 double-counts the**

Faddeev-Popov determinant. It writes the gauge-fixed measure as $e^{-V} \det M dA$ while setting $V = S_{\text{YM}} - \log \det M$, which gives $e^{-S_{\text{YM}}} (\det M)^2 dA$. The two correct forms, each with one determinant, are

$$e^{-S_{\text{YM}}} \det M dA \quad \text{or} \quad e^{-(S_{\text{YM}} - \log \det M)} dA,$$

in flat notation $\exp(-S_{\text{YM}}) \det M dA$ and $\exp(-(S_{\text{YM}} - \log \det M)) dA$, and the printed combination is neither. The narrow distinction between V_{pot} and the gauge-fixed effective potential survives that correction. Relatedly, Lemma 5.1 should not be read as an ordinary Hessian on B “at every point of $\mathcal{M}_{\text{flat}}$ ”: on the stratified quotient no such Hessian is fixed. What holds is that **the lifted second variation on A is non-negative** — every flat connection is a global minimum of the Wilson action on the smooth $\mathcal{A}$ — together with the corresponding statement on B_{reg} .

2. **Fixed-volume positivity**, once restated on $L^2(\mathcal{A}, \text{Vol}_{\mathcal{A}})^{\mathcal{G}}$. Qualitative only.
3. **The exact ground-state-measure identity**, stated against the reference measure it should carry. It is what gave this revision its reason to exist, and it **invalidates the v1 derivation unless the required identification is proved; the one-plaquette suite supplies strong numerical counterevidence, not an analytic refutation** — the same rule this note applies to every other number the suite produces. *An earlier draft called it a refutation.*
4. **The constant $C = 24\pi^2/(11N_c) = 1/(2b_0)$** , resolving v1’s internal inconsistency between its introduction and its Theorem 6.3.
5. **Remark 5.4** as a description of the conditional bound: not uniform in L .
6. **Remark 5.6** in part: the one-plaquette evidence and the conclusion of non-sharpness stand. Its inference that the gap on B is $O(g)$ comes from Proposition 4.1, whose frequencies inherit the defects of §2–§3, and is *not* certified here.
7. **The $SU(2)$, 2^3 , trivial-holonomy numerics** as evidence for that case.

7 Three lists, and they are not the same list

(A) **The conditional implication of Theorem 6.3** needs: (H-BAL), that is Balaban’s Theorem 6.1, which its own statement assumes; a precise (H1), naming the physical transfer operator, its fixed spatial slice and its Hilbert space; the original (H2), or a valid derivation of it — for which (H-FACT) is at most one input; and the admissible-volume condition. **Theorem 6.3, with (H-BAL), (H1) and (H2) assumed directly, uses nothing from list (C).**

(B) **A candidate route towards (H1)**, and only that, would have to **construct from S_{nmax} a transfer operator on a fixed spatial slice**, specify its Hilbert space and measure, identify its generator or its gap with the corrected physical Hamiltonian H_{ν} of §5.4, and control the effect of the ϵ_X terms and their temporal nonlocality — and only then would $(A\text{-osc})_{\nu}$ and terminal-scale uniformity become relevant. **(A-osc)_{nu} is specific to the candidate route to (H1)**; it is not a hypothesis of Theorem 6.3, which assumes (H1) outright. **The present fixed-volume route does not discharge (H1)**: it supplies none of those identifications, it yields a bound at fixed volume where (H1) is a statement at the terminal scale, and H_{ν} is a Wilson Hamiltonian on a spatial configuration space rather than the transfer operator of a four-dimensional effective action. *An earlier draft listed $(A\text{-osc})_{\nu}$ among the requirements of (A), and described (B) as though the two objects were already the same one.*

(C1) **The printed formulas that must be corrected**: Proposition 2.2(b), equations (5) and (6), Proposition 3.3 and its Fourier labels, and equations (8)–(10). This list is independent of (A). *An earlier draft put (A) and (C) into one sentence, which overstated what the conditional chain depends on.*

(C2) **What a Born–Oppenheimer theorem needs beyond those corrections**: a region uniformly separated from the singular strata; a uniform normal gap *on that region*, which §3.5 shows does not exist on the smooth stratum as a whole; control of the reduction errors; and a construction that splices in or absorbs the quartic sector on approach to the singular locus. **(C1) is necessary but not sufficient for a rigorous Born–Oppenheimer theorem**, and the restored text concedes as much in asking for a normal-gap condition it does not verify.

8 Two blocks of the restored text are superseded wholesale

Rather than chase each sentence: **Table 1 of Section 7.1 is superseded** by §§3–7 of this note, and **the Note added in v2 is superseded** likewise. Neither may be read as a current status summary. Table 1 still lists as “gap proved” the Paper II row resting on $\text{Ric}_B \geq N_c/4$ and the Paper III row resting on ground-state measure plus Holley–Stroock, both of which §5 withdraws. The Note added still carries item (1) on Morse–Bott and Proposition 3.3, item (4) on (H1) discharged and (H2) refined, and **item (5), conclusion unaffected, is withdrawn** together with them.

What survives from those two blocks: the constant $C = 24\pi^2/(11N_c)$ and the resolution of v1’s internal inconsistency about it; the repairs of broken citations, of the non-existent “Theorem 5.5”, and of the displaced headers; and **the record that the v1 derivation is invalid without the missing identification, and that the suite gives strong numerical counterevidence to its literal measure e^{-2V_{pot}/g^2}** . Everything else in Table 1 and in the Note added is to be read through §§3–7 above.

9 Scope

Not audited here: the Balaban package (H-BAL); **the remainder of Proposition 4.1 beyond the two claims refuted in §3.5**; and the numerical suite beyond the points named. *An earlier draft listed Proposition 4.1 as unaudited outright, which contradicted §3.5 of the same note* — a partial audit of a proposition is not a certification of the rest of it, and it is not a refusal to look either.

Provenance. §4 follows from the errata attached to 2602.0036 and 2602.0033. §5.2 was found by asking which reference measure each paper writes its ground-state measure against — the question the appended page exists to force — and that page’s single row for this paper exposed the intended physical object while failing to record the printed one — itself a provenance error, now split in two. **§2 came from an external audit of the second draft of this note**, which had found three independent defects and still missed the angle ambiguity creating several more; the first draft had opened by claiming nothing was false at all. *Each round has found the previous round’s survivors list too short* — the pattern this campaign keeps measuring.

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.

MORSE–BOTT SPECTRAL REDUCTION AND THE YANG–MILLS MASS GAP ON THE LATTICE

LLUIS ERIKSSON

ABSTRACT. We establish results toward the $SU(N_c)$ lattice Yang–Mills mass gap. First, the Wilson potential on the gauge orbit space $\mathcal{B} = \mathcal{A}/\mathcal{G}$ is Morse–Bott *along the smooth stratum* of its critical manifold $\mathcal{M}_{\text{flat}}$ (the flat connections), and we derive the Born–Oppenheimer effective Hamiltonian on $\mathcal{M}_{\text{flat}}$ (Theorem 3.1, Proposition 4.1); version 2 records that the Morse–Bott non-degeneracy *fails* at the orbifold locus of $\mathcal{M}_{\text{flat}}$ — including $\theta = 0$ — where the $k = 0$ non-Cartan “toron” modes are normal to $\mathcal{M}_{\text{flat}}$ with quartic potential, as exhibited numerically on a real lattice (new Proposition 3.3, following ai.viXra:2602.0032 v2). Second, the Faddeev–Popov obstruction identified in Paper II applies to the path integral but not to the Hamiltonian on $\mathcal{B}$: the spectral gap is strictly positive for each fixed lattice size L (unconditionally, by compactness and Perron–Frobenius), and a quantitative rate follows from the exact ground-state-measure identity plus Holley–Stroock, conditionally on a semiclassical oscillation bound (Theorem 5.3, reformulated in v2: version 1 derived the rate from the measure e^{-2V_{pot}/g^2} , whose oscillation is $O(L^d/g^4)$, dimensionally inconsistent with the $O(L^d/g^2)$ it used; the corrected exponent is $-\ln \psi_0^2$, with the identity and scalings verified numerically). Third, the physical mass gap $m \sim e^{-C/g^2}$ follows if the spectral gap at Balaban’s terminal renormalization scale is bounded below (Theorem 6.3); version 2 corrects the decay constant to $C(N_c) = 24\pi^2/(11N_c) = 1/(2b_0)$, resolving an internal inconsistency of v1 (whose introduction stated the correct constant while Theorem 6.3 carried a spurious $\ln 2$), and records that hypothesis (H1) is discharged by the Holley–Stroock bound of the companion synthesis paper ai.viXra:2602.0033 (v2), while (H2) is refined there into the non-extensive factorization hypothesis (H-FACT). All corrections are verified in a companion numerical suite.

1. INTRODUCTION

1.1. Context and motivation. This is the third paper in a series [1, 2] studying the Yang–Mills mass gap on the lattice via the geometry of the gauge orbit space. Paper I [1] established a mass gap for the Gribov–Zwanziger lattice measure via Brascamp–Lieb convexity. Paper II [2] proved three results: (i) geodesic convexity of $-\text{Re Tr } U$ on $B_{\pi/2}(I) \subset SU(N_c)$; (ii) $\text{Ric}_{\mathcal{B}} \geq N_c/4$ for the orbit space; and (iii) a universal obstruction showing that the Yang–Mills–Faddeev–Popov potential $V = S_{\text{YM}} - \ln \det \mathcal{M}$ is not geodesically convex at the trivial vacuum.

The present paper makes three advances beyond Paper II. First, we perform a Morse–Bott analysis of the Wilson potential on $\mathcal{B}$, deriving the Born–Oppenheimer effective Hamiltonian on the flat connection moduli space $\mathcal{M}_{\text{flat}}$ (Sections 3 and 4). Second, we observe that the Paper II obstruction applies to the path integral potential $V = S_{\text{YM}} - \ln \det \mathcal{M}$ but not to the physical Hamiltonian $H = \frac{g^2}{2}(-\Delta_{\mathcal{B}}) + V_{\text{pot}}$, where $V_{\text{pot}} = S_{\text{YM}} \geq 0$. This yields a mass gap for each fixed lattice size via the ground-state-measure/Holley–Stroock argument (Section 5). Third, we formulate a conditional mass gap theorem: if the spectral gap at Balaban’s terminal renormalization scale is $O(1)$, the physical mass gap is $O(e^{-C/g^2})$ (Section 6). (v2: the status of each step is sharpened; see the Note added.)

1.2. The Kogut–Susskind Hamiltonian on the orbit space. We work with $\mathrm{SU}(N_c)$ lattice Yang–Mills on $\Lambda_{\mathrm{lat}} = (\mathbb{Z}/L\mathbb{Z})^d$, $d \geq 3$. The orbit space $\mathcal{B} = \mathcal{A}/\mathcal{G}$ carries the Riemannian metric induced from the bi-invariant product metric on $\mathcal{A} = \prod_{\ell} \mathrm{SU}(N_c)$ (see [2, Section 3]). The Kogut–Susskind Hamiltonian on $\mathcal{B}$ is

$$(1) \quad H = \frac{g^2}{2}(-\Delta_{\mathcal{B}}) + V_{\mathrm{pot}},$$

where

$$(2) \quad V_{\mathrm{pot}}([U]) = \frac{2}{g^2} \sum_{\square} \left(1 - \frac{1}{N_c} \mathrm{Re} \mathrm{Tr} U_{\square}\right) \geq 0$$

is the Wilson potential on $\mathcal{B}$. The mass gap is $m_{\mathrm{gap}} = E_1 - E_0 = \lambda_1(H)$.

Remark 1.1 (Hamiltonian vs. path integral). The Hamiltonian (1) acts on $L^2(\mathcal{B}, d\mathrm{Vol}_{\mathcal{B}})$ and involves only the Wilson potential $V_{\mathrm{pot}} \geq 0$. In contrast, the Faddeev–Popov path integral uses the gauge-fixed measure $e^{-V} \det \mathcal{M} d\mathcal{A}$ on the Gribov region, where $V = S_{\mathrm{YM}} - \ln \det \mathcal{M}$ includes the FP determinant. The convexity obstruction of [2, Theorem 4.1] applies to V (which has negative Hessian in zero-mode directions) but does not apply to V_{pot} (which has non-negative Hessian at its minimum). This distinction is central to the results of this paper.

1.3. Main results. (i) **Morse–Bott structure (Theorem 3.1, amended in v2).** The Wilson potential V_{pot} on $\mathcal{B}$ is Morse–Bott *along the smooth stratum of $\mathcal{M}_{\mathrm{flat}}$* , with normal Hessian $(2/g^2)[\hat{k}^2 + m_a^2(\theta)]$ for $k \neq 0$; at the orbifold locus (including $\theta = 0$) the $k = 0$ non-Cartan modes are normal with vanishing Hessian and quartic potential (Proposition 3.3). (ii) **Born–Oppenheimer effective Hamiltonian (Proposition 4.1).** The low-lying spectrum of H is formally determined by $H_{\mathrm{eff}} = \frac{g^2}{2}(-\Delta_{\mathcal{M}_{\mathrm{flat}}}) + V_{\mathrm{BO}}(\theta) + O(g^2)$. (iii) **Fixed-volume mass gap (Theorem 5.3).** For each fixed L , $m_{\mathrm{gap}} > 0$ unconditionally; under a semiclassical oscillation bound, $m_{\mathrm{gap}} \geq c(L, N_c, d) g^2 > 0$ quantitatively. (iv) **Conditional mass gap (Theorem 6.3).** If the spectral gap at Balaban’s terminal RG scale satisfies $m_{n^*} \geq c_0 > 0$, then $m_{\mathrm{gap}} \geq c_0 e^{-C/g^2}$ with

$$C = \frac{24\pi^2}{11N_c}$$

(as correctly stated in v1’s introduction; v1’s Theorem 6.3 carried a spurious $\ln 2$, corrected here).

1.4. Organization. Section 2 describes the flat connection moduli space. Section 3 proves the Morse–Bott structure and its v2 amendment. Section 4 derives the Born–Oppenheimer effective Hamiltonian. Section 5 proves the fixed-volume mass gap. Section 6 formulates the conditional mass gap. Section 7 discusses the remaining steps.

2. THE FLAT CONNECTION MODULI SPACE

2.1. Flat connections on the torus. A lattice connection $U = \{U_{\mu}(x)\} \in \mathcal{A}$ is flat if $U_{\square} = I$ for every plaquette. On the periodic lattice, a flat connection is determined (up to gauge equivalence) by its holonomies $g_{\mu} = \prod_{t=0}^{L-1} U_{\mu}(t\hat{\mu}) \in \mathrm{SU}(N_c)$ around the d non-contractible loops, satisfying $[g_{\mu}, g_{\nu}] = I$.

Definition 2.1. *The flat connection moduli space is*

$$(3) \quad \mathcal{M}_{\mathrm{flat}} = \{(g_1, \dots, g_d) \in \mathrm{SU}(N_c)^d : [g_{\mu}, g_{\nu}] = I \ \forall \mu, \nu\} / \text{conjugation}.$$

Since commuting elements of $SU(N_c)$ lie in a common maximal torus $T \cong U(1)^{N_c-1}$, we may simultaneously diagonalize: $g_\mu = \exp(i\theta_\mu)$ with θ_μ diagonal traceless. The residual gauge freedom is the Weyl group $W = S_{N_c}$.

Proposition 2.2. *For $SU(N_c)$ on $(\mathbb{Z}/L\mathbb{Z})^d$:*

$$(4) \quad \mathcal{M}_{\text{flat}} \cong (T^d)^{N_c-1}/W$$

where $T^d = (\mathbb{R}/2\pi\mathbb{Z})^d$ and $W = S_{N_c}$ acts by simultaneous permutation of eigenvalues. In particular: (a) $\dim \mathcal{M}_{\text{flat}} = d(N_c - 1)$; (b) for $SU(2)$: $\mathcal{M}_{\text{flat}} \cong [0, \pi]^d/\mathbb{Z}_2$ with $\theta_\mu \mapsto \pi - \theta_\mu$; (c) $\mathcal{M}_{\text{flat}}$ is compact.

Proof. (a): each θ_μ has $N_c - 1$ independent components. (b): for $SU(2)$ the maximal torus is $U(1)$ with $\theta \in [0, 2\pi)$ and the Weyl group acts by $\theta \mapsto 2\pi - \theta$, which on $[0, \pi]$ becomes $\theta \mapsto \pi - \theta$. (c): quotient of a compact set by a finite group. $\square$

2.2. The metric on $\mathcal{M}_{\text{flat}}$. The tangent space to $\mathcal{M}_{\text{flat}}$ at $[\theta]$ consists of constant Cartan variations $\delta\theta_\mu \in \mathfrak{t}$ (zero-momentum modes). The induced metric from $\mathcal{A}$ is

$$(5) \quad \langle \delta\theta, \delta\theta' \rangle_{\mathcal{M}_{\text{flat}}} = N \sum_{\mu} \langle \delta\theta_\mu, \delta\theta'_\mu \rangle_{\mathfrak{t}},$$

where $N = L^d$ is the number of lattice sites (the constant mode is replicated at each site).

3. MORSE–BOTT STRUCTURE OF THE WILSON POTENTIAL

Theorem 3.1 (Morse–Bott structure; amended in v2). *The Wilson potential $V_{\text{pot}} : \mathcal{B} \rightarrow [0, \infty)$ of (2) is Morse–Bott along the smooth stratum of its critical manifold $\mathcal{M}_{\text{flat}} = V_{\text{pot}}^{-1}(0)$ (in v1 this theorem asserted the Morse–Bott property without restriction; see Proposition 3.3). Specifically:*

- (a) $V_{\text{pot}}([U]) = 0$ if and only if $[U] \in \mathcal{M}_{\text{flat}}$.
- (b) $\text{Hess}(V_{\text{pot}})|_{T\mathcal{M}_{\text{flat}}} = 0$.
- (c) For normal fluctuations with Fourier mode $k \neq 0$,

$$(6) \quad \text{Hess}(V_{\text{pot}})^\perp(\delta A, \delta A) = \frac{2}{g^2} \sum_{k \neq 0} \sum_a [\hat{k}^2 + m_a^2(\theta)] |\delta A_\perp^a(k)|^2,$$

where $m_a^2(\theta) = 4 \sum_\mu \sin^2(\alpha_a \cdot \theta_\mu/2)$ and α_a are the roots of $\mathfrak{su}(N_c)$. Since $\hat{k}^2 > 0$ for $k \neq 0$: this part of the normal Hessian is strictly positive definite. (v2: the $k = 0$ non-Cartan fluctuations are also normal to $\mathcal{M}_{\text{flat}}$; their Hessian is $m_a^2(\theta)$, which vanishes on the orbifold locus — Proposition 3.3.)

Proof. Part (a): $V_{\text{pot}} \geq 0$ with equality iff $U_\square = I$ for all $\square$. Part (b): $V_{\text{pot}} = 0$ on all of $\mathcal{M}_{\text{flat}}$. Part (c): at a flat connection with constant holonomy, the covariant finite difference acting on a fluctuation in the root space of α_a gives, in Fourier space, $|\nabla_\mu \hat{A}_a(k)|^2 = 4 \sin^2((k_\mu + \alpha_a \cdot \theta_\mu)/2) |\hat{A}_a(k)|^2$; the plaquette Hessian (lattice curl squared, restricted to transverse modes) then yields (6) at $\theta = 0$ exactly and for general θ with the stated positivity for $k \neq 0$, since the shifted momenta satisfy $\sum_\mu \hat{k}_\mu(\theta)^2 \geq \hat{k}^2 + m_a^2(\theta) - C|\theta||\hat{k}| > 0$. The verification suite validates (6) non-perturbatively at $\theta = 0$: the finite-difference Hessian of S_W on the full configuration space of a 2^3 spatial $SU(2)$ lattice equals $(1/N_c)(d^T d) \otimes I_3$ entry by entry (lattice Maxwell per adjoint color; the overall metric normalization $1/N_c$ is absorbed into constants). $\square$

Remark 3.2. For the “neutral” modes ($\alpha_a = 0$, i.e., Cartan modes with $k \neq 0$): $m_a^2 = 0$ and the normal Hessian reduces to $(2/g^2)\hat{k}^2 > 0$. These remain massless at non-trivial holonomy and are the softest *quadratic* normal modes.

Proposition 3.3 (New in v2: failure of Morse–Bott non-degeneracy at the orbifold locus). *At points $\theta \in \mathcal{M}_{\text{flat}}$ with enhanced stabilizer — the orbifold locus of $(T^d)^{N_c-1}/W$, which includes $\theta = 0$ — the $k = 0$ non-Cartan constant fluctuations $((d)(N_c^2 - 1 - r)$ of them on the spatial slice, $r = N_c - 1$) are normal to $\mathcal{M}_{\text{flat}}$ and have vanishing Hessian: along them V_{pot} is quartic ($V_{\text{pot}} \sim \frac{1}{2N_c g^2} \sum \| [X_\mu, X_\nu] \|^2$). These are the toron modes of the small-volume analyses [12, 13]. Numerically (suite; the computation of ai.viXra:2602.0032 v2, re-run here): for SU(2) on a 2^3 spatial lattice, the kernel of the finite-difference Hessian of S_W at $\theta = 0$ has dimension $30 = 21$ (gauge) + 3 (Cartan moduli) + 6 (non-Cartan $k = 0$), exceeding $\dim T(\mathcal{G} \cdot \mathcal{M}_{\text{flat}}) = 24$, with quartic ratio $f(2t)/f(t) = 15.99 \approx 2^4$; at generic θ the kernel drops to 26 and Morse–Bott is restored. Consequently the Morse–Bott property holds exactly on the smooth stratum and fails on the orbifold locus.*

4. BORN–OPPENHEIMER EFFECTIVE HAMILTONIAN

Proposition 4.1 (Born–Oppenheimer reduction, formal). *In the semiclassical regime $g \rightarrow 0$, the low-lying spectrum of H (1) on $\mathcal{B}$ is formally determined (up to corrections of order $O(g^2)$) by the effective Hamiltonian on $\mathcal{M}_{\text{flat}}$:*

$$(7) \quad H_{\text{eff}} = \frac{g^2}{2} (-\Delta_{\mathcal{M}_{\text{flat}}}) + V_{\text{BO}}(\theta) + O(g^2),$$

where the Born–Oppenheimer potential is the sum of zero-point energies of normal-mode oscillators:

$$(8) \quad V_{\text{BO}}(\theta) = \frac{1}{2} \sum_{k \neq 0} \sum_a \omega_{k,a}(\theta) - \frac{1}{2} \sum_{k \neq 0} \sum_a \omega_{k,a}(0),$$

with normal-mode frequencies

$$(9) \quad \omega_{k,a}(\theta) = \sqrt{g^2 \cdot \frac{2}{g^2} [\hat{k}^2 + m_a^2(\theta)]} = \sqrt{2} [\hat{k}^2 + m_a^2(\theta)]^{1/2}.$$

More precisely: the first $\dim H^0(\mathcal{M}_{\text{flat}}, \mathbb{R})$ eigenvalues of H lie in $[E_0, E_0 + Cg^2]$ and agree with those of H_{eff} up to $O(g^4)$; the remaining eigenvalues satisfy $E_n \geq E_0 + \omega_{\perp, \min} - O(g)$, $\omega_{\perp, \min} = \sqrt{2} \hat{k}_{\min} = \sqrt{2} \cdot 2 \sin(\pi/L)$.

Proof. As in v1: adapted coordinates (y, z) near $\mathcal{M}_{\text{flat}}$, harmonic quantization of the normal modes with mass $1/g^2$ and frequencies (9), projection onto the z -ground state, giving $V_{\text{BO}} = E_{\perp}(\theta) - E_{\perp}(0)$. Higher corrections are formally $O(g^2)$. A rigorous justification (e.g. via [14]) requires a gap condition for the normal-mode Hamiltonian that we do not verify; in addition (v2), at the orbifold locus of $\mathcal{M}_{\text{flat}}$ the quartic toron modes of Proposition 3.3 are outside the harmonic scheme: the sums in (8) run over $k \neq 0$ and omit them, so (7) is at best valid away from the orbifold locus (or must be supplemented by the small-volume quartic sector of [12, 13]). The Born–Oppenheimer potential is not used in the proofs of Theorems 5.3 and 6.3. $\square$

Remark 4.2 (BO potential vs. GPY potential). V_{BO} involves $\sum \sqrt{\hat{k}^2 + m^2}$, while the Gross–Pisarski–Yaffe effective potential [3] involves $\frac{1}{2} \sum \ln(\hat{k}^2 + m^2)$. These are different mathematical objects: V_{BO} governs the Hamiltonian spectrum, V_{GPY} the Euclidean path integral. Both have the same qualitative behavior and are proportional to $\sum_a m_a^2(\theta)$ at leading order.

Remark 4.3 (New in v2: the metric factor in Section 4.1). The kinetic term in (7) carries the metric (5), i.e. $-\Delta_{\mathcal{M}_{\text{flat}}} = -N^{-1} \partial_\theta^2$ with $N = L^d$. v1’s Section 4.1 quoted the oscillator frequency as $\omega_\theta = g(4\sqrt{2}J_1)^{1/2}$, omitting N : the correct expression is $\omega_\theta = g(4\sqrt{2}J_1/N)^{1/2}$. Since $J_1 = \sum_{k \neq 0} 1/|\hat{k}|$ grows like L^d in $d = 3$ spatial dimensions ($J_1/L^3 \rightarrow 0.45$ numerically),

$J_1/N = O(1)$ and the advertised conclusion — spectral gap $O(g)$ on $\mathcal{M}_{\text{flat}}$ — is unaffected. The suite verifies the Hessian $\text{Hess } V_{\text{BO}}(0) = 4\sqrt{2}J_1 I$ (v1's expansion $V_{\text{BO}} \approx 2\sqrt{2}J_1 \sum \theta_\mu^2$, exact), the positivity of (10) on a grid, and the convergence of J_1/L^3 . Note also that (8)–(9) use the dispersion $\sqrt{\hat{k}^2 + m_a^2}$, which is the internally consistent one (the companion ai.viXra:2602.0032 v1 defined its BO potential with a different, twisted-momentum dispersion and proved properties of this one; its v2 adopts the present convention).

4.1. Explicit computation for $\text{SU}(2)$. For $\text{SU}(2)$ on $(\mathbb{Z}/L\mathbb{Z})^3$: the roots are $\alpha = \pm 2$, $\mathcal{M}_{\text{flat}} \cong [0, \pi]^3/\mathbb{Z}_2$, and $m^2(\theta) = 4 \sum_{\mu=1}^3 \sin^2 \theta_\mu$. The BO potential:

$$(10) \quad V_{\text{BO}}(\theta) = \sqrt{2} \sum_{k \neq 0} \left[\left(\hat{k}^2 + 4 \sum_{\mu} \sin^2 \theta_\mu \right)^{1/2} - |\hat{k}| \right].$$

For θ small: $V_{\text{BO}}(\theta) \approx 2\sqrt{2}J_1 \sum_{\mu} \theta_\mu^2$, where $J_1 = \sum_{k \neq 0} 1/|\hat{k}|$ is a convergent lattice sum for $d = 3$. The effective Hamiltonian is a 3D Schrödinger operator on $[0, \pi]^3/\mathbb{Z}_2$ with confining quadratic potential and spectral gap $O(g)$ (harmonic frequency $\omega_\theta = g(4\sqrt{2}J_1/N)^{1/2}$; Remark 4.3).

5. FIXED-VOLUME MASS GAP VIA THE GROUND-STATE MEASURE

5.1. The key observation: no FP obstruction. The Hamiltonian (1) uses the potential $V_{\text{pot}} = S_{\text{YM}} \geq 0$. The Bakry–Émery Ricci curvature of a Gibbs-type measure $d\mu = e^{-f} d\text{Vol}_{\mathcal{B}}/Z$ is

$$(11) \quad \text{Ric}_\mu = \text{Ric}_{\mathcal{B}} + \text{Hess}(f).$$

By [2, Theorem 3.1]: $\text{Ric}_{\mathcal{B}} \geq N_c/4$ (the sharp constant is $N_c/2$; suite, and ai.viXra:2602.0032 v2). The crucial question is the sign of $\text{Hess}(f)$ near the minimum.

Lemma 5.1. *At every point of $\mathcal{M}_{\text{flat}}$, the Hessian of V_{pot} is non-negative in all directions:*

$$(12) \quad \text{Hess}(V_{\text{pot}})|_{\mathcal{M}_{\text{flat}}} \geq 0.$$

Proof. $V_{\text{pot}} \geq 0$ on all of $\mathcal{B}$ and $V_{\text{pot}} = 0$ on $\mathcal{M}_{\text{flat}}$, so $\mathcal{M}_{\text{flat}}$ is a global minimum and the Hessian there is non-negative. (v2: non-negativity is of course compatible with the additional zero directions of Proposition 3.3; positivity in *all* normal directions holds only on the smooth stratum.) $\square$

Remark 5.2. Contrast with the path integral potential $V = S_{\text{YM}} - \ln \det \mathcal{M}$: [2, Theorem 4.1] shows $\text{Hess}(V)$ is negative in zero-mode directions. This negativity is entirely due to $-\ln \det \mathcal{M}$ and does not affect the Hamiltonian.

5.2. The mass gap theorem.

Theorem 5.3 (Fixed-volume Hamiltonian mass gap; reformulated in v2). *For $\text{SU}(N_c)$ lattice Yang–Mills on $(\mathbb{Z}/L\mathbb{Z})^d$, $d \geq 3$, $g^2 \leq g_0^2$:*

- (i) (Unconditional.) *For each fixed $L \geq 2$, $m_{\text{gap}} > 0$.*
- (ii) (Quantitative, under (A-osc) below.)

$$(13) \quad m_{\text{gap}} \geq \frac{g^2 N_c}{16} e^{-C_2 L^d / g^2} > 0,$$

where $C_2 > 0$ depends only on N_c and d . The bound degrades exponentially with the volume and is not uniform in L .

(A-osc) *The ground state ψ_0 of H satisfies $\text{osc}(-\ln \psi_0^2) \leq C'_2 L^d / g^2$.*

Proof. (i) $\mathcal{B}$ is compact, H is (the Friedrichs extension of) an elliptic operator with smooth potential on $L^2(\mathcal{B})$: the spectrum is discrete, and the ground state is simple and strictly positive (Perron–Frobenius for the positivity-improving semigroup e^{-tH}). Hence $E_1 > E_0$.

(ii) *Step 1: exact ground-state-measure identity (corrected formulation).* Let $\psi_0 > 0$ be the ground state and $d\mu = \psi_0^2 d\text{Vol}_{\mathcal{B}}$. The ground-state transformation $f \mapsto \psi_0 f$ is a unitary from $L^2(\mu)$ to $L^2(\text{Vol})$ conjugating $H - E_0$ into the μ -reversible Dirichlet-form operator $\frac{g^2}{2} L_\mu$, $\langle f, L_\mu f \rangle_\mu = \int |\nabla f|^2 d\mu$. Hence *exactly*

$$(14) \quad m_{\text{gap}} = \frac{g^2}{2} \lambda_1(\mu).$$

(v1’s display (14), the near-flat region Ω_δ , belonged to the superseded derivation; the equation number now carries the identity (14) — see the Note added. The suite verifies this identity to 10^{-9} relative precision on a one-plaquette model, where it holds by matrix similarity.) v1 instead applied the argument to the measure $e^{-2V_{\text{pot}}/g^2} d\text{Vol}$; since V_{pot} already contains the prefactor $2/g^2$ in (2), that exponent is $O(L^d/g^4)$ — dimensionally inconsistent with the $O(L^d/g^2)$ oscillation v1’s own estimate used, and numerically that measure does not reproduce the gap (suite: errors $> 50\%$, growing as $g \rightarrow 0$). The correct exponent is $f_0 := -\ln \psi_0^2$.

Step 2: Holley–Stroock. By the Holley–Stroock bounded perturbation lemma [7] applied to $d\mu = e^{-f_0} d\text{Vol}/Z$ against the uniform measure ν :

$$(15) \quad \alpha(\mu) \geq \alpha(\nu) \cdot e^{-\text{osc}(f_0)}.$$

Since $\text{Ric}_{\mathcal{B}} \geq N_c/4$, Bakry–Émery [10] gives $\alpha(\nu) \geq N_c/2$, so $\lambda_1(\mu) \geq \alpha(\mu)/2 \geq (N_c/4)e^{-\text{osc}(f_0)}$.

Step 3: the oscillation input. Under (A-osc): $\text{osc}(f_0) \leq C'_2 L^d/g^2$, whence (13) with $C_2 = C'_2$ after absorbing constants. (A-osc) is the semiclassical WKB/Agmon estimate: $f_0 \approx (2/g^2) \times (\text{Agmon distance to } \mathcal{M}_{\text{flat}})$, which is $O(L^d/g^2)$ since the action density is bounded; the suite verifies the scaling $\text{osc}(f_0) \cdot g^2 \approx \text{const}$ on the one-plaquette model (10.3, 10.7, 10.9 for $g^2 = 0.5, 0.3, 0.2$). We do not prove (A-osc) on $\mathcal{B}$ here; it replaces v1’s inconsistent bookkeeping by an explicit, numerically supported hypothesis. $\square$

Remark 5.4. The bound (13) is not uniform in L due to the oscillation factor $e^{-O(L^d/g^2)}$. The uniform-in- L bound requires either (a) a Bakry–Émery argument avoiding the global oscillation bound, or (b) a localization technique (Balaban’s RG) replacing the global oscillation by a local one. We pursue (b) in Section 6.

Corollary 5.5. *For each fixed $L \geq 2$ and $g^2 \leq g_0^2$, under (A-osc):*

$$(16) \quad m_{\text{gap}}(L, g) \geq c(L, N_c, d) \cdot g^2 > 0.$$

Remark 5.6 (New in v2: sharpness). The bound (16) is far from sharp: on the one-plaquette model the true gap tends to $\sqrt{2} = O(1)$ as $g \rightarrow 0$ (suite), and Proposition 4.1 suggests the gap on $\mathcal{B}$ is $O(g)$ (soft modes on $\mathcal{M}_{\text{flat}}$), not $O(g^2)$. The content of Theorem 5.3 is positivity with an explicit (if pessimistic) rate, not sharp asymptotics.

6. CONDITIONAL MASS GAP VIA MULTI-SCALE ANALYSIS

6.1. Balaban’s constructive renormalization group.

Theorem 6.1 (Balaban [6, Theorem 1]; package paraphrase). *For $\text{SU}(N_c)$ lattice Yang–Mills on $(\mathbb{Z}/L\mathbb{Z})^d$ with $d = 4$ and $\beta \geq \beta_0$: there exist effective actions $S_0 = S_{\text{Wilson}}, S_1, \dots, S_n$ on successively coarsened lattices with $|\Lambda_k| = L^d/2^{kd}$, such that:*

- (a) $Z_0 = Z_k$ for all $k \leq n$.

(b) *The effective coupling flows toward the infrared by asymptotic freedom:*

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

(equivalently $\beta_k = \beta - 4N_c b_0 k \ln 2 + O(1)$, decreasing; *v1* stated $\beta_k = \beta + 2b_0 \ln 2^k$, with inverted sign and β conflated with $1/g^2$ — the erratum shared with *ai.viXra:2602.0032/0033 v1*, corrected in all three *v2*'s).

(c) *The effective action has the form*

$$(17) \quad S_k(U) = \beta_k S_W(U) + \sum_X \epsilon_k(X),$$

with $|\epsilon_k(X)| \leq e^{-\kappa|X|/a_k}$, $\kappa > 0$ universal.

(d) *The construction is valid for $k \leq n_{\max}$, the largest integer with $g_k^2 = 2N_c/\beta_k \leq \gamma$, $\gamma > 0$ a small universal constant.*

Remark 6.2. Balaban's published results cover $d = 4$. The extension to $d = 3$ is expected to be simpler (super-renormalizable) but has not been published in the same generality; the $d = 3$ case has been partially carried out by Dimock [8, 9].

6.2. The conditional mass gap.

Theorem 6.3 (Conditional mass gap). *Assume Balaban's Theorem 6.1 (for $d = 4$; see Remark 6.2 for $d = 3$). Assume further:*

- (H1) *The spectral gap of the transfer matrix associated with $S_{n_{\max}}$ on the orbit space $\mathcal{B}_{n_{\max}}$ satisfies $m_{n_{\max}} \geq c_0 > 0$, with c_0 depending only on N_c , d , γ .*
- (H2) *Balaban's RG preserves the exponential decay of connected correlations of gauge-invariant observables (equivalently, the physical mass gap is RG-invariant).*

Then:

$$(18) \quad m_{\text{gap}} \geq c_0 \cdot e^{-C/g^2}, \quad C = \frac{1}{2b_0} = \frac{24\pi^2}{11N_c},$$

uniformly in L , for $g^2 \leq g_0^2$. (*v1*'s Theorem 6.3 stated $C = \ln 2/(2b_0) = 24\pi^2 \ln 2/(11N_c)$, in contradiction with *v1*'s own introduction (1.3.iv), which had the correct $24\pi^2/(11N_c)$: the $\ln 2$ from the step count cancels against the $\ln 2$ of $2^{n_{\max}}$, cf. *ai.viXra:2602.0033 v2*, Remark 5.1. Resolved in favor of the introduction.)

Proof. By (H2), the physical gap is RG-invariant, so in lattice units at the original scale $m_0 = m_{n_{\max}}/2^{n_{\max}}$. By (H1): $m_{n_{\max}} \geq c_0$. By the corrected coupling evolution: $n_{\max} = \frac{1}{2b_0 \ln 2} \left(\frac{1}{g^2} - \frac{1}{\gamma} \right) + O(1)$, so

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

Therefore $m_0 \geq c_0 e^{-C/g^2 - O(1)} = c'_0 e^{-C/g^2}$. □

6.3. Evidence for hypotheses (H1) and (H2).

Proposition 6.4. *Hypothesis (H1) holds if γ is large enough that the terminal theory is in the strong-coupling regime (Osterwalder–Seiler cluster expansion [4]). Balaban's γ is a small constant, so the terminal theory is at intermediate coupling, where the gap is not established by classical methods; Corollary 5.5 (*v1* cited a non-existent “Theorem 5.5” here) gives $m_{n_{\max}} \geq c(\gamma)g_{n_{\max}}^2 > 0$ for each fixed coarsened lattice size, under (A-osc). (*v2*: (H1) is now discharged by the companion synthesis paper: *ai.viXra:2602.0033 v2*, Theorem 3.1, proves $m_{n_{\max}} \geq$*

$(\gamma N_c/8)e^{-K/\gamma} > 0$ at the terminal scale via Holley–Stroock, using $L_{n_{\max}} = O(1)$ so that the oscillation is manifestly $O(1/\gamma)$ — no (A-osc)-type input needed there.)

Proposition 6.5. *Hypothesis (H2) is expected to follow from Balaban’s preservation of the partition function (Theorem 6.1(a)) together with the exponential decay of the block-averaging kernel; the precise extraction requires detailed analysis of his cluster expansion. (v2: this is now made precise in ai.viXra:2602.0033 v2: the exact-compatibility reading is false as literally stated (temporally nonlocal ϵ_k), the naive factorization error is extensive, and (H2) is refined into the non-extensive factorization hypothesis (H-FACT) there, with a finite-window doubling theorem and a scale-invariant window condition that is automatic when c_0 is small. (H2) should therefore be read as (H-FACT).)*

Remark 6.6 (New in v2: status within the series). Combining the companions: (H1) is discharged by ai.viXra:2602.0033 v2 (Theorem 3.1); (H2) is refined to (H-FACT) (ibid., Proposition 4.4); the toron amendment of Proposition 3.3 comes from ai.viXra:2602.0032 v2 (Proposition 5.3); the constant C is common to all three v2’s. The net conditional statement of the series is: *Balaban package (H-BAL) + non-extensive temporal factorization (H-FACT) $\Rightarrow$ volume-uniform lattice mass gap ce^{-C/g^2} at weak coupling*, with the present paper supplying the Hamiltonian-side geometry (no FP obstruction, fixed-volume positivity) and the companions the terminal gap and the RG cascade.

7. DISCUSSION

7.1. Summary of the three-paper program.

Paper	Result	Gap proved	Uniform in L ?
I	Brascamp–Lieb on Gribov region	$O(g^2)$	Yes
II	$\text{Ric}_{\mathcal{B}} \geq N_c/4$; convexity obstruction	$O(g^2)$ (kinetic)	Yes
III	Ground-state measure + Holley–Stroock	$c(L) \cdot g^2$	No (Remark 5.4)
III	Conditional (Balaban)	$O(e^{-C/g^2})$	Yes

TABLE 1. Mass gap results across the three papers. (v2: v1’s displaced subsection header and orphaned footnote are repaired.)

7.2. The Hamiltonian vs. path integral distinction. The central insight of this paper is the distinction between $V_{\text{pot}} = S_{\text{YM}} \geq 0$ and $V = S_{\text{YM}} - \ln \det \mathcal{M}$. The FP obstruction of [2, Theorem 4.1] applies only to the latter. The Hamiltonian approach avoids it because: (1) the Hamiltonian acts on gauge-invariant functions on $\mathcal{B}$ without gauge fixing; (2) $V_{\text{pot}} \geq 0$ has non-negative Hessian at its minimum (Lemma 5.1); (3) $\text{Ric}_{\mathcal{B}} > 0$ and the ground-state measure concentrates near $\mathcal{M}_{\text{flat}}$. Methods based on the physical Hamiltonian may succeed where path-integral methods fail.

7.3. The remaining step. In v1 this read: “bound the spectral gap of the transfer matrix associated with Balaban’s effective action at the terminal scale”. In v2, per Proposition 6.4, that step is discharged by the companion synthesis paper; the remaining analytic content of the programme is **(H-FACT)** (non-extensive temporal factorization of Balaban’s effective actions, or an anisotropic RG variant), together with the derivation of the idealized Balaban package (H-BAL) from the primary sources.

7.4. Toward the Millennium Problem. The Clay problem [11] requires a continuum theory satisfying the Wightman axioms with a mass gap. Our lattice results (including the conditional Theorem 6.3) concern the lattice theory only. The remaining steps — continuum limit, OS/Wightman axioms, mass gap persistence — are beyond the scope of this paper. None of the results here is an unconditional solution of the Millennium Problem, nor close to one: the conditional chain still rests on (H-BAL) and (H-FACT), and the continuum questions are untouched.

NOTE ADDED IN V2

Prepared with the companion suite `verify_2602_0035.py` (repository THE-ERIKSSON-PROGRAMME, `verification/2602-0035/`; 14 checks, all passing, ~ 10 s, NumPy/SciPy). Changes: (1) Theorem 3.1 restricted to the smooth stratum; new Proposition 3.3 (toron modes; kernel 30 vs 24, quartic ratio 16, numerically re-verified); Proposition 4.1 amended accordingly; (2) Theorem 5.3 reformulated: positivity unconditional (compactness/Perron–Frobenius); rate via the *exact* ground-state-measure identity plus Holley–Stroock, under the explicit oscillation hypothesis (A-osc); v1’s measure e^{-2V_{pot}/g^2} has oscillation $O(L^d/g^4)$ and does not reproduce the gap (both exhibited numerically); new Remark 5.6 on non-sharpness (true toy gap $\rightarrow \sqrt{2}$); (3) Theorem 6.3: $C = 24\pi^2/(11N_c)$, resolving v1’s internal inconsistency (introduction correct, theorem carrying the $\ln 2$ slip shared with the companions); coupling flow direction corrected in Theorem 6.1(b); (4) series status: (H1) discharged and (H2) refined to (H-FACT) via ai.viXra:2602.0033 v2 (Propositions 6.4–6.5, Remark 6.6); (5) new Remark 4.3 (metric factor N in Section 4.1; conclusion unaffected); (6) repairs: broken citations “[?, ?]” in v1’s Section 6.1 (now [5, 6]), the non-existent “Theorem 5.5” in Proposition 6.4 (now Corollary 5.5), duplicated “Step 4” header, displaced Section 7.1 header and orphaned table footnote; equation numbers (1)–(18) are preserved in count and position, with the caveat that the rewritten proof of Theorem 5.3 retires v1’s display (14) (the region Ω_δ) and its number now carries the ground-state identity (14); (7) references to Papers I–III fixed with their archive identifiers: Paper II = ai.viXra:2602.0036, Paper III = this paper (= ai.viXra:2602.0035); Paper I is identified as ai.viXra:2602.0038 (submitted title differs from the working title cited in v1). All v1 numbered statements (1.1–6.5) and equations (1)–(18) keep their numbers.

APPENDIX A. NUMERICAL VERIFICATION

`verify_2602_0035.py`: (1) one-plaquette model: exact ground-state-measure identity (10^{-9}), failure of the literal e^{-2V_{pot}/g^2} measure ($> 50\%$, growing), $\text{osc}(-\ln \psi_0^2)g^2 \approx \text{const}$ vs $\text{osc}(2V_{\text{pot}}/g^2)g^4 = 8.00$ exact, true gap $\rightarrow \sqrt{2}$; (2) SU(2) 2^3 finite-difference Hessian: Maxwell match, kernel 30/26, quartic ratio 15.99; (3) Hess $V_{\text{BO}}(0) = 4\sqrt{2}J_1I$, positivity, $J_1/L^3 \rightarrow 0.45$; (4) $C = 1/(2b_0) = 24\pi^2/(11N_c)$ and flow direction; (5) $\text{Ric}_{\text{SU}(N_c)} = (N_c/2)\|X\|^2$. Run: `python3 verify_2602_0035.py` (exit 0 iff all pass).

REFERENCES

- [1] L. Eriksson, “Mass gap for the Gribov–Zwanziger lattice measure: a non-perturbative proof,” ai.viXra:2602.0038, 2026 (cited in v1 under the working title “Mass gap in the weakly coupled Gribov–Zwanziger lattice gauge theory via Brascamp–Lieb log-concavity”).
- [2] L. Eriksson, “Geodesic convexity and structural limits of curvature methods for the Yang–Mills mass gap on the lattice,” ai.viXra:2602.0036, 2026.
- [3] D. J. Gross, R. D. Pisarski, and L. G. Yaffe, “QCD and instantons at finite temperature,” *Rev. Mod. Phys.* **53**, 43–80 (1981).
- [4] K. Osterwalder and E. Seiler, “Gauge field theories on a lattice,” *Ann. Phys.* **110**, 440–471 (1978).
- [5] T. Balaban, “Large field renormalization. I,” *Comm. Math. Phys.* **122**, 175–202 (1989).

- [6] T. Balaban, “Large field renormalization. II. Localization, exponentiation, and bounds for the R operation,” *Comm. Math. Phys.* **122**, 355–392 (1989).
- [7] R. Holley and D. Stroock, “Logarithmic Sobolev inequalities and stochastic Ising models,” *J. Stat. Phys.* **46**, 1159–1194 (1987).
- [8] J. Dimock, “The renormalization group according to Balaban. I. Small fields,” *Rev. Math. Phys.* **25**, 1330010 (2013).
- [9] J. Dimock, “The renormalization group according to Balaban. III. Convergence,” *Ann. Henri Poincaré* **15**, 2133–2175 (2014).
- [10] D. Bakry and M. Émery, “Diffusions hypercontractives,” *Sém. Prob. XIX, LNM 1123*, Springer, 1985, pp. 177–206.
- [11] A. Jaffe and E. Witten, “Quantum Yang–Mills theory,” Clay Mathematics Institute, 2000.
- [12] M. Lüscher, “Some analytic results concerning the mass spectrum of Yang–Mills gauge theories on a torus,” *Nucl. Phys. B* **219**, 233–261 (1983).
- [13] J. Koller and P. van Baal, “A non-perturbative analysis in finite volume gauge theory,” *Nucl. Phys. B* **302**, 1–64 (1988).
- [14] J.-M. Combes, P. Duclos, and R. Seiler, “The Born–Oppenheimer approximation,” in *Rigorous Atomic and Molecular Physics*, Plenum, 1981, pp. 185–213.
- [15] 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.
- [16] L. Eriksson, “The Yang–Mills mass gap on the lattice: a conditional synthesis,” *ai.viXra:2602.0033* (v2), 2026.

INDEPENDENT RESEARCHER

Email address: lluiseriksson@gmail.com