

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