

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

LLUIS ERIKSSON

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

1. INTRODUCTION AND MAIN RESULT

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

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

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

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

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

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

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

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

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

The proof has three steps:

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

2. BALABAN'S CONSTRUCTIVE RG

We use the following result of Balaban.

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

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

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

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

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

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

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

3. SPECTRAL GAP AT THE TERMINAL SCALE

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

4. RG INVARIANCE OF THE MASS GAP

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

5. PROOF OF THE MAIN THEOREM

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

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

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

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

Therefore

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

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

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

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

6. DISCUSSION

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

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

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

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

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

6.3. The role of each paper.

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

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

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

Limitations:

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

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

NOTE ADDED IN v2

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

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

APPENDIX A. NUMERICAL VERIFICATION

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

REFERENCES

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

INDEPENDENT RESEARCHER

Email address: lluiseriksson@gmail.com