Uniform Log-Sobolev Inequality and Mass Gap for Lattice Yang–Mills Theory

Lluis Eriksson

Independent Researcher

`lluiseriksson@gmail.com`

February 2026

Abstract

We establish that $SU(N_c)$ lattice Yang–Mills theory in $d = 4$ dimensions with Wilson action at sufficiently weak coupling ($\beta = 2N_c/g^2 \geq \beta_0$) satisfies a log-Sobolev inequality with constant $\alpha_* > 0$ uniform in the lattice size L_{vol}, conditional on Assumption 5.4. Combined with reflection positivity of the Wilson action, this yields a mass gap $\Delta_{\mathrm{phys}} > 0$ uniform in L_{vol}, conditional on Assumptions 5.4 and 6.3.

The proof combines three ingredients: (i) Balaban's constructive renormalization group, which produces controlled effective actions at all scales; (ii) the orbit space Ricci curvature bound $\mathrm{Ric}_{\mathcal{B}} \geq N_c/4$, which gives a uniform log-Sobolev constant for conditional measures of fast modes at each RG scale via the Bakry–Émery criterion; and (iii) a multiscale entropy decomposition with sweeping-out bounds, where the geometric scaling factor $\|Q^*_{(k)}\|^2 = 2^{-(d-1)k}$ of transversal block averaging ensures summability of cross-scale errors.

A key technical point is that Balaban's RG "absorption" mechanism isolates the dependence on scale-k fast variables into polymer residuals with exponential decay. Under k-uniform analyticity/derivative inputs from Balaban's scheme (made explicit as Assumption 5.4), Cauchy-type estimates yield the derivative bounds needed for the sweeping-out inequality. We provide a detailed audit trail from Balaban's published work (Comm. Math. Phys. **98**, **116**, **122**) supporting Assumption 5.4; see Appendix A.

Contents

1 Introduction

1.1 The problem

The existence of a positive mass gap in four-dimensional $SU(N_c)$ Yang–Mills theory is one of the Clay Millennium Prize Problems [20]. On the lattice $\Lambda = (\mathbb{Z}/L_{\text{vol}}\mathbb{Z})^d$ with Wilson action $S_W[U] = \beta \sum_P (1 - \frac{1}{N_c}\text{Re tr } U_P)$, $\beta = 2N_c/g^2$, the mass gap is the spectral gap of the transfer matrix.

At strong coupling ($\beta < \beta_0$), the mass gap is established by the Osterwalder–Seiler cluster expansion [24]. At weak coupling ($\beta \geq \beta_0$), the problem has remained open for 40 years.

1.2 Main result

Theorem 1.1 (Main Theorem). *For $SU(N_c)$ lattice Yang–Mills theory on $(\mathbb{Z}/L_{\text{vol}}\mathbb{Z})^4$ with Wilson action at $\beta = 2N_c/g^2 \geq \beta_0$ (sufficiently large):*

    (i) **(Uniform LSI, conditional on Assumption 5.4)** *There exists $\alpha_* > 0$, depending on N_c and β_0 but not on L_{vol}, such that for all gauge-invariant $f \in L^2(\mu_\beta)$:*

$$\text{Ent}_{\mu_\beta}(f^2) \leq \frac{2}{\alpha_*} \int |\nabla f|^2 \, d\mu_\beta. \tag{1}$$

    (ii) **(Mass gap, conditional on Assumptions 5.4 and 6.3)** *The transfer matrix T has spectral gap*

$$\Delta_{\text{phys}} \geq c(N_c, \beta_0) > 0,$$

    uniformly in L_{vol}.

The proof uses Balaban's constructive renormalization group [1, 2, 3, 4, 5, 6], the orbit space Ricci curvature bound $\mathrm{Ric}_{\mathcal{B}} \geq N_c/4$, and the multiscale entropy decomposition with sweeping-out bounds. Appendix A provides an audit trail supporting Assumption 5.4.

1.3 Logical structure

The proof follows the chain:

$$\boxed{\mathrm{Balaban\ RG} + \mathrm{Ric}_{\mathcal{B}} \geq \tfrac{N_c}{4} \xrightarrow{\text{entropy decomposition}} \alpha_*^{\mathrm{LSI}} > 0 \xrightarrow{\text{RP + locality}} \Delta_{\mathrm{phys}} > 0}$$

Each arrow is established by combining published results with the new ingredients of this paper:

- **Arrow 1** (§4–§5): Terminal LSI via Holley–Stroock, conditional fast LSI via Bakry–Émery, sweeping-out inequality from Balaban's polymer bounds, assembly into uniform LSI.

- **Arrow 2** (§6): LSI $\Rightarrow$ hypercontractivity $\Rightarrow$ exponential clustering $\Rightarrow$ mass gap via reflection positivity.

1.4 Relation to the Clay Millennium Problem

This result establishes the mass gap on the lattice at weak coupling, uniform in the volume. The Clay Millennium Problem additionally requires the construction of a continuum QFT satisfying the Osterwalder–Schrader axioms with a mass gap persisting as the lattice spacing $a \to 0$. Our result is a necessary ingredient for, but does not complete, this program.

2 Inputs from Balaban's constructive RG

2.1 Effective actions with polymer bounds

We use the following result of Balaban.

Theorem 2.1 (Balaban [6, Theorem 1]). *For $SU(N_c)$ lattice gauge theory on $\Lambda = (\mathbb{Z}/L_{\mathrm{vol}}\mathbb{Z})^4$ with $\beta = 2N_c/g^2 \geq \beta_0$: there exist effective actions S_k on coarsened lattices $\Lambda_k = (\mathbb{Z}/L_{\mathrm{vol},k}\mathbb{Z})^4$, $0 \leq k \leq n_{\max}$, where $L_{\mathrm{vol},k} := L_{\mathrm{vol}}/L_{\mathrm{RG}}^k$ denotes the linear volume size at scale k. Assume (for simplicity) that L_{vol} is divisible by $L_{\mathrm{RG}}^{n_{\max}}$, so that $L_{\mathrm{vol},k} \in \mathbb{Z}$ for all $0 \leq k \leq n_{\max}$. Then:*

(a) *$Z(\Lambda, S_0) = Z(\Lambda_k, S_k)$ for all k.*

(b) *The effective action has the structure ([6, eq. (1.101), p. 388]):*

$$S_k = \beta_k S_W + \sum_X \mathbf{R}^{(k)}(X) + \sum_X \mathbf{B}^{(k)}(X), \tag{2}$$

where $\mathbf{R}^{(k)}(X)$ are irrelevant polymer activities and $\mathbf{B}^{(k)}(X)$ are boundary terms near large-field regions.

(c) *The polymer activities satisfy the exponential decay bound ([6, eq. (1.100), p. 388]):*

$$|\mathbf{R}^{(k)}(X, (\mathbf{U}, \mathbf{J}))| \leq e^{-p_0(g_k)} e^{-\kappa\, d_k(X)}, \tag{3}$$

where $d_k(X)$ is the linear size of the polymer X and $\kappa > 0$ is a universal decay constant.

(d) *The polymer activities are analytic functions of the field variables on a complex domain $\widetilde{U}_k^c(Y, \tilde{\alpha}_0, \hat{\alpha}_1)$ ([6, eq. (1.65), p. 383]).*

(e) *The coupling evolution satisfies $\beta_k = \beta + 2b_0 k \ln 2 + O(1/\beta)$ with $b_0 = 11N_c/(48\pi^2)$.*

2.2 The absorption property

The central structural property of Balaban's RG is the *absorption of cross-scale coupling*. At each RG step, the Wilson action term coupling fine and coarse variables is decomposed and reabsorbed into the next-scale effective Wilson action. The key identity is [6, eq. (1.37), p. 367]:

$$
-\frac{1}{g_k^2} A(\zeta_0, U_{k,Z}(V_k'V_\Lambda)) - A\left(\frac{1}{(g_k''(\cdot))^2}, U_k''\right) + \frac{1}{g_k^2} A(\zeta_0, U_0)
$$
$$
= -A\left(\frac{1}{(g_k(\cdot))^2}, U_k\right) + O(1), \tag{4}
$$

where the "$O(1)$" terms are explicitly bounded by absolute constants (independent of g_k, L_{vol}, and k) in the text following eq. (1.37) of [6]. After this absorption, the dependence on scale-k variables enters only through the polymer residuals $\mathbf{R}^{(k)}(X)$ and $\mathbf{B}^{(k)}(X)$.

2.3 Block-averaging and geometric scaling

The block-averaging operation is defined in [1, eq. (15), p. 21]:

$$
\bar{U}_c = \exp\left[i \sum_{x \in B(c_-)} L_{\text{RG}}^{-d} \frac{1}{i} \log U(\Gamma_{c,x}) U(c)^{-1}\right] U(c). \tag{5}
$$

Its linearisation is given by [1, eq. (125), p. 40]:

$$
(Q_0 A)_c = \sum_{x \in B(c_-)} L_{\text{RG}}^{-(d+1)} (R_{0,c_-} A)([x, x']). \tag{6}
$$

This averages over L_{RG}^{d-1} fine links in a transversal section. The adjoint distributes uniformly:

Lemma 2.2 (Adjoint scaling). *Let $L_{\text{RG}} > 1$ denote the fixed RG block size (Balaban's parameter L). The linearised block-averaging operator Q_0 defined in (6) satisfies $\|Q_0^*\|^2 = L_{\text{RG}}^{-(d-1)}$ as an operator on $\mathfrak{su}(N_c)$-valued fields. Iterating k steps: $\|Q_{(k)}^*\|^2 = L_{\text{RG}}^{-(d-1)k}$. In our dyadic RG, $L_{\text{RG}} = 2$, $d = 4$, hence $\|Q_{(k)}^*\|^2 = 2^{-3k}$.*

Proof. The non-linear averaging (5) involves the group exponential and logarithm, but its linearisation Q_0 ([1, eq. (125)]) is a linear operator on $\mathfrak{su}(N_c)$-valued bond fields. For a fixed coarse bond c in direction μ, the operator $(Q_0 A)_c$ averages over the L_{RG}^{d-1} fine links in the transversal $(d-1)$-dimensional block $B(c_-)$, each contributing with weight $L_{\text{RG}}^{-(d-1)}$ (the factor $L_{\text{RG}}^{-(d+1)}$ in (6) accounts for $L_{\text{RG}}^{-(d-1)}$ from transversal averaging and L_{RG}^{-2} from the path length). The adjoint Q_0^* distributes a coarse-bond perturbation uniformly among the $n = L_{\text{RG}}^{d-1}$ transversal fine links, giving $\|Q_0^* t\|^2 = n|t/n|^2 = |t|^2/n$.

The non-linear corrections to this scaling are controlled by [1, Proposition 4, eq. (133)]: the full non-linear k-fold averaging $Q_k(U_0, \eta A)$ satisfies $|Q_k(U_0, \eta A) - Q_k(U_0)A| < C_2|A|^2$, where $Q_k(U_0)$ is the composition of linearised operators. The non-linear remainder is quadratic in A and does not affect the leading scaling. Thus $\|Q_{(k)}^*\|^2 = L_{\text{RG}}^{-(d-1)k}$ for the linearised operator, which governs the scaling of the sweeping-out coefficients. $\qquad\square$

The uniform bounds on functional derivatives of the k-fold averaging are given by [1, Proposition 5, eqs. (156)–(157), p. 43]:

$$
\left|\frac{\delta}{\delta A_b} Q_k(U_0, \eta A, c)\right| \leq 1 + 2C_1' \alpha_0 + C_3 \alpha_1, \tag{7}
$$

with constants independent of k.

2.4 Polymer derivative bounds (Proposition 2.6)

We now record the key single-scale derivative bound needed for the sweeping-out estimate. Its derivation uses Balaban's published polymer analyticity framework together with a k-uniform analyticity-radius input, which we isolate as part of Assumption 5.4. The resulting cross-scale composition is recorded in Assumption 5.4 and is supported by the audit trail in Appendix A.

Proposition 2.3 (Single-scale polymer derivative bound, conditional on k-uniform analyticity radius)**.** *Assume, in addition to the hypotheses of Theorem 2.1, that the polymer analyticity domain $\widetilde{U}_k^c(Y, \tilde{\alpha}_0, \hat{\alpha}_1)$ in [6, eq. (1.65), p. 383] can be chosen with $\hat{\alpha}_1 = \hat{\alpha}_1(\gamma) > 0$ uniformly in k throughout Balaban's inductive scheme (cf. Remark A.1).*

Under the hypotheses of Theorem 2.1, the effective potential $V_{<k}$ at scale k satisfies: after Balaban's RG absorption, the dependence on scale-k fast variables enters only through polymer residuals, and

$$\operatorname*{ess\,sup}_{\mathcal{G}_k} \mathbb{E}\big[|\nabla_{E_k} V_{<k}|^2 \mid \mathcal{G}_k\big] \leq C_{\text{poly}}^2\, e^{-2\kappa} \cdot L_{\text{RG}}^{-(d-1)k}. \tag{8}$$

Here L_{RG} denotes the fixed RG block size (Balaban's parameter $L > 1$); in our dyadic RG one has $L_{\text{RG}} = 2$, hence $L_{\text{RG}}^{-(d-1)k} = 2^{-(d-1)k}$. The constants $C_{\text{poly}}, \kappa > 0$ are the universal constants from Balaban's construction.

Proof. The proof combines three ingredients from Balaban's papers.

Step 1: Structure. By the absorption identity (4) ([6, eq. (1.37)]), the Wilson action cross-coupling is absorbed into $\beta_{k+1} S_W$, and the residual dependence on scale-k variables is through the polymer terms $\mathbf{R}^{(k)}(X)$ and $\mathbf{B}^{(k)}(X)$ in (2).

Step 2: Uniform analytic radius and derivative bounds. Balaban's Theorem 1 [6] is proved under the hypothesis that all effective couplings satisfy $g_k^2 \leq \gamma$ for a fixed small $\gamma > 0$. Within this regime, the polymer activities $\mathbf{R}^{(k)}(X)$ are analytic functions of the gauge field variables on the complex domain $\widetilde{U}_k^c(Y, \tilde{\alpha}_0, \hat{\alpha}_1)$ ([6, eq. (1.65), p. 383]), where the analytic radius $\hat{\alpha}_1 = \hat{\alpha}_1(\gamma) > 0$ depends only on γ and is *uniform in k*. This uniformity rests on two explicit results in Balaban's work:

(i) *Uniform averaging maps.* The k-fold averaging map Q_k admits an analytic expansion with constants independent of k ([1, Proposition 4]), its functional derivatives satisfy bounds with constants independent of k ([1, Proposition 5]; cf. our (7)), and these bounds persist uniformly under small complex perturbations of the background field ([1, Proposition 7]).

(ii) *Uniform inductive scheme.* Provided the effective couplings remain in $(0, \gamma]$, the effective densities p_k satisfy *uniformly in k* all inductive analyticity conditions and bounds of the effective-action framework (Theorem 1 in [6], referring to the conditions described in Sect. 2 of [4]). In particular, the polymer activities appearing after absorption are analytic on a complex neighbourhood whose radius depends only on γ and fixed geometric parameters of the RG step, not on k.

By the Cauchy integral formula applied within this uniform analytic domain, combined with the polymer bound (3):

$$\left| \frac{\delta \mathbf{R}^{(k)}(X)}{\delta A'(b)} \right| \leq \frac{1}{\hat{\alpha}_1(\gamma)} \sup_{|A'| < \hat{\alpha}_1} |\mathbf{R}^{(k)}(X)| \leq \frac{C(\gamma)}{\hat{\alpha}_1(\gamma)}\, e^{-\kappa\, d_k(X)}, \tag{9}$$

where $C(\gamma) = e^{-p_0(g_k)}$ from (3). Since $\hat{\alpha}_1(\gamma)$ is a fixed positive constant (not depending on k), the derivative bound has the same exponential decay as the polymer bound itself.

Step 2b: Boundary terms $\mathbf{B}^{(k)}(X)$. The boundary terms in (2) are localised near the large-field region Z_k and satisfy the bound ([6, eq. (1.69), p. 377]):

$$|\mathbf{B}^{(k)}(X)| \leq O(1) \sum_{j=1}^{k} |\Gamma_j^0 \cap X|$$

5

when X is a component of $Z^\sim$ (the enlarged large-field region). The large-field regions satisfy $\mu_\beta(Z_k) \leq e^{-2p_0(g_k)}$ ([6, eq. (1.89), p. 387]), with $p_0(g_k) \to \infty$ as $g_k \to 0$.

The key observation is that the polymer derivative bound (8) is an *essential supremum* over $\mathcal{G}_k$, so it suffices to verify the bound μ_β-a.e. In Balaban's construction, configurations are decomposed into small-field regions (where the analytic estimates of Step 2 apply) and large-field regions (where the $\mathbf{B}^{(k)}$ terms are supported). On the large-field region:

- The boundary terms $\mathbf{B}^{(k)}(X)$ depend on the background field U_k restricted to $X \subset Z_k^\sim$, and they are analytic on the same domain $\widetilde{U}_k^c(X, \tilde{\alpha}_0, \hat{\alpha}_1)$ as the $\mathbf{R}^{(k)}$ terms ([6, p. 377]: "the above expressions are analytic functions of U_k^0 on Y_0, and the analyticity domain is restricted on the cube $\square_0$").

- The Cauchy derivative bound therefore applies to $\mathbf{B}^{(k)}$ as well: $|\nabla \mathbf{B}^{(k)}(X)| \leq C/\hat{\alpha}_1 \cdot O(1)|X|$.

- However, these terms only contribute to $\nabla_{E_k} V_{<k}$ on the large-field region $Z_k^\sim$, which has μ_β-measure $\leq e^{-2p_0(g_k)}$. Since we take the essential supremum of a *conditional expectation* $\mathbb{E}[|\nabla_{E_k} V_{<k}|^2 \mid \mathcal{G}_k]$, the large-field contribution is suppressed by the exponential factor $e^{-2p_0(g_k)}$ in the conditional integration.

More precisely: splitting $V_{<k} = V_{<k}^{\text{sf}} + V_{<k}^{\text{lf}}$ into small-field and large-field contributions, we have
$$\mathbb{E}\big[|\nabla_{E_k} V_{<k}|^2 \mid \mathcal{G}_k\big] \leq 2\mathbb{E}\big[|\nabla_{E_k} V_{<k}^{\text{sf}}|^2 \mid \mathcal{G}_k\big] + 2\mathbb{E}\big[|\nabla_{E_k} V_{<k}^{\text{lf}}|^2 \mid \mathcal{G}_k\big].$$
The first term is bounded by $C_{\text{poly}}^2 e^{-2\kappa} \cdot 2^{-(d-1)k}$ as in Step 2.

The large-field contribution requires a separate argument because $|Z_k|$ can grow with L_{vol}. The key structural property, established in Balaban's construction, is the *locality* of the large-field regions: the large-field region Z_k decomposes into components of bounded diameter, each suppressed exponentially by the ultraviolet stability bounds ([6, eqs. (0.1), (1.79), (1.89)]).

Specifically, Balaban's construction guarantees ([6, eq. (1.76), p. 384]) that:

(a) Each component $Z_k^{(i)}$ of the large-field region has diameter $\leq 100MR_k$ (in M-cubes at scale k), independent of L_{vol}.

(b) The boundary terms $\mathbf{B}^{(k)}(X)$ decay exponentially away from the large-field region ([6, eq. (1.68)]), so the gradient $\nabla_{E_k} V_{<k}^{\text{lf}}$ at a fast link b depends only on components within distance $O(R_k)$ of b.

(c) The per-component suppression factor is $e^{-c p_0(g_k)}$ with $p_0(g_k) \to \infty$ as $g_k \to 0$ ([6, eq. (1.89), p. 387]).

These three properties together imply that the large-field contribution to $\mathbb{E}_k[|\nabla_{E_k} V_{<k}^{\text{lf}}|^2]$ is bounded by $C(R_k, M, d, N_c)^2 \cdot e^{-c p_0(g_k)}$, where $C(R_k, M, d, N_c)$ depends on β but *not on* L_{vol} (because only $O(R_k^d)$ components contribute to any given link, independently of the total volume). For γ sufficiently small, this is $o(L_{\text{RG}}^{-(d-1)k})$, and the large-field terms do not affect the leading bound (8).

Step 3: Geometric scaling via block-averaging adjoint. The dependence of $V_{<k}$ on scale-k fast variables arises through the polymer residuals, whose gradients are bounded by Step 2. When projected onto the fast subspace E_k (the orthogonal complement of V_{k+1} within V_k, as in Definition 2.4 of [15]), the gradient picks up the block-averaging adjoint factor. Specifically, by the linearisation of Balaban's block-averaging ([1, eq. (125)]) and the composition bounds ([1, Proposition 4]), perturbations in the fast directions at scale k contribute to the polymer gradient with an additional factor $\|Q_{(k)}^*\| = L_{\text{RG}}^{-(d-1)k/2}$ per component. Squaring:

$$\mathbb{E}\big[|\nabla_{E_k} V_{<k}|^2 \mid \mathcal{G}_k\big] \leq \Big(\sum_{X \ni \text{fast link}} \frac{C(\gamma)}{\hat{\alpha}_1} e^{-\kappa d_k(X)} \Big)^2 \cdot \|Q_{(k)}^*\|^2.$$

The sum over polymers containing a given fast link is bounded by $C'e^{-\kappa}$ (since $d_k(X) \geq 1$ for each contributing polymer, and the number of polymers of linear size d containing a given link grows at most polynomially in d, which is absorbed by the exponential decay). Thus:

$$\underset{\mathcal{G}_k}{\mathrm{ess\,sup}}\,\mathbb{E}\big[|\nabla_{E_k}V_{<k}|^2 \mid \mathcal{G}_k\big] \leq C_{\mathrm{poly}}^2\, e^{-2\kappa} \cdot L_{\mathrm{RG}}^{-(d-1)k},$$

with $C_{\mathrm{poly}} = C'(\gamma)/\hat{\alpha}_1(\gamma)$. For $L_{\mathrm{RG}} = 2$, $d = 4$, this gives the bound (8). $\qquad\square$

3 Orbit space Ricci curvature

The configuration space $\mathcal{A} = \prod_\ell SU(N_c)$ with the product bi-invariant metric has $\mathrm{Ric}_{\mathcal{A}} = \bigoplus_\ell \mathrm{Ric}_{SU(N_c)}$. Each factor satisfies $\mathrm{Ric}_{SU(N_c)} = \frac{N_c}{4}\, g_{SU(N_c)}$.

Theorem 3.1 ([11, Theorem 3.1]). *The orbit space $\mathcal{B} = \mathcal{A}/\mathcal{G}$ satisfies $\mathrm{Ric}_{\mathcal{B}} \geq \frac{N_c}{4}$ on the smooth stratum $\mathcal{B}^*$.*

Proof. The argument is self-contained and uses only standard Riemannian geometry.

Step 1: Riemannian submersion structure. The gauge group $\mathcal{G} = \prod_{x\in\Lambda} SU(N_c)$ acts on $\mathcal{A}$ by $(u \cdot U)_\ell = u(x)U_\ell u(y)^{-1}$ for $\ell = \langle x,y\rangle$. This action preserves the bi-invariant metric on each $SU(N_c)$ factor (since $|uXu^{-1}| = |X|$ for the bi-invariant norm), hence $\mathcal{G}$ acts by isometries. The orbit space $\mathcal{B} = \mathcal{A}/\mathcal{G}$ inherits the quotient metric, making $\pi : \mathcal{A} \to \mathcal{B}$ a Riemannian submersion ([30, Chapter 9]).

Step 2: Ricci curvature of $\mathcal{A}$. The configuration space $\mathcal{A} = \prod_\ell SU(N_c)$ carries the product bi-invariant metric $g_{\mathcal{A}} = \bigoplus_\ell g_{SU(N_c)}$. The Ricci curvature of $SU(N_c)$ with its bi-invariant metric normalised so that $\langle X, Y\rangle = -\frac{1}{2}\mathrm{tr}(\mathrm{ad}_X \circ \mathrm{ad}_Y)$ is $\mathrm{Ric}_{SU(N_c)} = \frac{1}{4}g_{SU(N_c)}$ (see, e.g., [30, Proposition 7.10]). With the standard normalisation $\langle X, Y\rangle = -N_c\,\mathrm{tr}(XY)$, this gives $\mathrm{Ric}_{SU(N_c)} = \frac{N_c}{4}\,g_{SU(N_c)}$ (the factor N_c arises from the Killing form normalisation). The product metric gives $\mathrm{Ric}_{\mathcal{A}} = \bigoplus_\ell \frac{N_c}{4}\, g_\ell$, hence $\mathrm{Ric}_{\mathcal{A}} \geq \frac{N_c}{4}\, g_{\mathcal{A}}$.

Step 3: O'Neill's formula. By O'Neill's formula [31] applied to the Riemannian submersion π:

$$\mathrm{Ric}_{\mathcal{B}}(\pi_*X, \pi_*X) = \mathrm{Ric}_{\mathcal{A}}(X,X) + \frac{3}{4}\sum_i \|[X, E_i]^{\mathcal{V}}\|^2 \geq \mathrm{Ric}_{\mathcal{A}}(X,X) \geq \frac{N_c}{4}\|X\|^2,$$

where X is horizontal, $\{E_i\}$ is an orthonormal basis of the vertical space (the gauge orbit tangent space), and $[\cdot,\cdot]^{\mathcal{V}}$ denotes the vertical component of the Lie bracket.

Step 4: Singular stratum. The orbit space $\mathcal{B}$ has a smooth stratum $\mathcal{B}^*$ (irreducible connections) and a singular stratum $\mathcal{B}_{\mathrm{sing}}$ (reducible connections). For $N_c \geq 2$, $\mathcal{B}_{\mathrm{sing}}$ has codimension $\geq 2(N_c - 1) \geq 2$ in $\mathcal{B}$. By the Lott–Villani theory of Ricci curvature for metric-measure spaces [21], the lower bound $\mathrm{Ric}_{\mathcal{B}} \geq \frac{N_c}{4}$ extends to $\mathcal{B}$ in the sense of the $CD(N_c/4, \infty)$ condition, which suffices for the Bakry–Émery criterion. $\qquad\square$

Remark 3.2. The singular stratum $\mathcal{B}_{\mathrm{sing}}$ (reducible connections) has codimension $\geq 2(N_c-1) \geq 2$ and is negligible for the Bakry–Émery criterion, which requires curvature bounds only μ-a.e. [21].

4 Terminal scale LSI

Theorem 4.1 (Terminal LSI). *At Balaban's terminal scale $n_{\max}$, the effective measure $\mu_{n_{\max}}$ satisfies an LSI with constant $\alpha_{\mathrm{term}} > 0$ depending only on N_c, γ, and κ.*

Proof. The terminal orbit space $\mathcal{B}_{n_{\max}}$ has fixed finite dimension $n = 3L_{\mathrm{vol},n_{\max}}^3(N_c^2 - 1) = O(1)$, Ricci curvature $\mathrm{Ric}_{\mathcal{B}} \geq N_c/4$ (Theorem 3.1), and is compact.

The uniform measure ν satisfies, by Bakry–Émery [7]: $\alpha_{\mathrm{LSI}}(\nu) \geq 2 \cdot N_c/4 = N_c/2$.

The effective measure is $\mu_{n_{\max}} = e^{-f}\nu/Z$ with $f = S_{n_{\max}}/(g_{n_{\max}}^2/2)$. On the terminal lattice ($O(1)$ sites), the number of plaquettes is $O(1)$ and the polymer sum $\sum_X |\epsilon_{n_{\max}}(X)|$ is finite (bounded number of terms, each $\leq e^{-\kappa}$), giving $\mathrm{osc}(f) \leq K/\gamma =: M_{\mathrm{term}}$.

By Holley–Stroock [19]: $\alpha_{\mathrm{term}} \geq (N_c/2)e^{-M_{\mathrm{term}}} > 0$. $\qquad\square$

5 Multiscale entropy decomposition

5.1 Setup and filtration

We work with the decreasing filtration $\mathcal{G}_0 \supset \mathcal{G}_1 \supset \cdots \supset \mathcal{G}_{n_{\max}}$, where $\mathcal{G}_k = \sigma(\bar{A}^{(k)})$ is generated by the block-averaged field at scale k. The orthogonal scale decomposition gives $T_A\mathcal{A} = \bigoplus_{k=0}^{n_{\max}} E_k \oplus V_{n_{\max}}$ with $\sum_k |\nabla_{E_k} f|^2 + |\nabla_{V_{n_{\max}}} f|^2 = |\nabla f|^2$.

5.2 Entropy chain rule

Lemma 5.1 (Entropy telescoping). *For any $f \in L^2(\mu)$ with $f > 0$:*

$$\mathrm{Ent}_\mu(f) = \sum_{k=0}^{n_{\max}-1} \mathbb{E}\big[\mathrm{Ent}(\mathbb{E}[f \mid \mathcal{G}_k] \mid \mathcal{G}_{k+1})\big] + \mathrm{Ent}_\mu(\mathbb{E}[f \mid \mathcal{G}_{n_{\max}}]). \tag{10}$$

Proof. Standard entropy chain rule; see Guionnet–Zegarlinski [18, Proposition 4.1]. $\qquad\square$

5.3 Conditional fast LSI

Lemma 5.2. *There exists $\alpha_0 > 0$ (depending on N_c and κ but not on k or L_{vol}) such that for each $0 \leq k \leq n_{\max} - 1$:*

$$\mathrm{Ent}(h^2 \mid \mathcal{G}_{k+1}) \leq \frac{2}{\alpha_0} \mathbb{E}\big[|\nabla_{E_k} h|^2 \mid \mathcal{G}_{k+1}\big] \tag{11}$$

for all smooth gauge-invariant h, μ-a.e. in $\mathcal{G}_{k+1}$.

Proof. Step 1: Small-field conditional LSI. Fix a coarse configuration $\bar{A}^{(k+1)} \in \mathcal{G}_{k+1}$. The conditional measure on fast modes (scale-k variables given $\mathcal{G}_{k+1}$) is supported on $\prod_{\ell \in E_k} SU(N_c)$ (a product of compact Lie groups), with conditional density proportional to $e^{-W_k(\cdot|\bar{A}^{(k+1)})}$.

Balaban's small-field/large-field decomposition ([6, Theorem 1]) provides: on the small-field region Ω_k^{sf} (where $|U(\partial p) - 1| < \varepsilon_k$ for all plaquettes), the conditional potential W_k admits the expansion ([6, eq. (1.2), p. 356]):

$$W_k = \frac{1}{2}\langle H_{1,k}B', \Delta_1(\zeta_0)H_{1,k}B'\rangle + \frac{1}{g_k}\langle DH_{1,k}B', \zeta_0\eta^{-2}\mathrm{Im}\,\partial U_0\rangle + \frac{1}{g_k^2}V(\zeta_0, g_k H_{1,k}B'),$$

where the quadratic form satisfies the lower bound ([6, eq. (1.9), p. 358]):

$$\langle H_{1,k}B', \Delta_1(\zeta_0)H_{1,k}B'\rangle \geq \frac{\gamma_0}{2d(100M)^5}\|B'\|^2.$$

Combined with the fiber Ricci curvature $\mathrm{Ric}_{\mathrm{fiber}} \geq N_c/4$ (each $SU(N_c)$ factor, Theorem 3.1), the Bakry–Émery criterion gives a conditional LSI on the small-field region with constant $\alpha_0^{\mathrm{sf}} \geq N_c/2$.

Step 2: Large-field patching via bounded perturbation. On the full fiber (including the large-field region), the conditional measure is $d\mu_k^{\mathrm{full}} = e^{-\Phi}d\mu_k^{\mathrm{sf}}/Z$, where Φ is the "large-field correction" to the potential. By Balaban's large-field bounds ([6, eq. (0.1), p. 355]):

$$\chi_k \exp\Big[-\frac{1}{g_k^2}A(U_k) - E_-|T_\eta|\Big] \leq \rho_k \leq \exp\big[E_+|T_\eta|\big],$$

the oscillation of the log-density is bounded by

$$\operatorname{osc}(\log \rho_k) \leq \frac{1}{g_k^2} \sup A(U_k) + (E_- + E_+)|T_\eta|.$$

On the finite-dimensional *fiber* (fast modes at scale k, which are $O((100MR_k)^d)$ link variables), this gives $\operatorname{osc}(\Phi) \leq C(M, d, N_c, R_k)$, a constant depending on the geometry of the RG block but *not* on L_{vol}.

By the Holley–Stroock perturbation lemma for LSI ([19]): if $d\mu = e^{-\Phi} d\nu/Z$ and ν satisfies LSI(α), then μ satisfies LSI($\alpha \cdot e^{-\operatorname{osc}(\Phi)}$). Applying this with $\nu = \mu_k^{\mathrm{sf}}$ (satisfying LSI($N_c/2$) by Step 1) and $\mu = \mu_k^{\mathrm{full}}$:

$$\alpha_0 \geq \frac{N_c}{2} \cdot e^{-C(M,d,N_c,R_k)} > 0.$$

The constant C depends on the RG block geometry (through M, R_k) but is *uniform in k and L_{vol}* because $R_k \leq R_{n_{\max}}$ is bounded by a function of β alone. $\qquad\square$

5.4 Influence coefficients and sweeping-out inequality

Definition 5.3 (Influence coefficients across scales). For $0 \leq \ell \leq k \leq n_{\max} - 1$, define the influence coefficient

$$\Gamma_{k\ell} := \sup_{\substack{\varphi \in C^\infty(\mathcal{A}) \\ \mathbb{E}[|\nabla_{E_\ell}\varphi|^2] \neq 0}} \frac{\mathbb{E}\big[|\nabla_{E_k}\mathbb{E}[\varphi \mid \mathcal{G}_\ell]|^2\big]^{1/2}}{\mathbb{E}\big[|\nabla_{E_\ell}\varphi|^2\big]^{1/2}},$$

with $\Gamma_{k\ell} = 0$ for $\ell > k$.

Assumption 5.4 (Cross-scale derivative bound). For all $0 \leq \ell < k \leq n_{\max} - 1$:

$$\operatorname*{ess\,sup}_{\mathcal{G}_\ell} \mathbb{E}\big[|\nabla_{E_k} V_{<\ell}|^2 \mid \mathcal{G}_\ell\big] \leq C_{\mathrm{poly}}^2 e^{-2\kappa} L_{\mathrm{RG}}^{-(d-1)k}. \tag{12}$$

Here $V_{<\ell} = -\log(\text{conditional density of modes } <\ell \text{ given } \mathcal{G}_\ell)$ (normalised), $E_k = \ker(Q_{k+1})^\perp \cap \ker(Q_k)$ is the scale-k fluctuation subspace, and $C_{\mathrm{poly}}, \kappa$ are the universal constants from Proposition 2.3.

This bound follows from Balaban's polymer analyticity and decay (Proposition 2.3) composed across intermediate scales $\ell < m < k$, combined with the block-averaging adjoint scaling (Lemma 2.2). The derivation is given in Appendix A.

Lemma 5.5 (Geometric decay of influence). *Under Proposition 2.3, the influence coefficients satisfy:*

$$\Gamma_{k\ell} \leq C_\Gamma \, e^{-\kappa} L_{\mathrm{RG}}^{-(d-1)k/2} \tag{13}$$

for all $\ell \leq k$, with C_Γ depending on N_c, γ, κ but not on k, ℓ, or L_{vol}.

Proof. Fix $\ell < k$ and a smooth function φ. The conditional expectation is

$$\mathbb{E}[\varphi \mid \mathcal{G}_\ell](\bar{A}^{(\ell)}) = \frac{\int \varphi(A) e^{-V_{<\ell}(A|\bar{A}^{(\ell)})} dA_{<\ell}}{\int e^{-V_{<\ell}(A|\bar{A}^{(\ell)})} dA_{<\ell}},$$

where $dA_{<\ell}$ denotes integration over scales finer than ℓ, and $V_{<\ell}$ is the effective potential from Balaban's RG.

Step 1: Differentiation. For $v \in E_k$ with $k > \ell$, the direction v acts on the integration variables $A_{<\ell}$ (which include scale k). Differentiating:

$$v \, \mathbb{E}[\varphi \mid \mathcal{G}_\ell] = \mathbb{E}[v\varphi \mid \mathcal{G}_\ell] - \operatorname{Cov}(\varphi, vV_{<\ell} \mid \mathcal{G}_\ell). \tag{14}$$

This is the standard identity for differentiating a ratio of integrals with respect to a parameter appearing in the integrand.

Step 2: Bounding the direct term. The first term satisfies

$$\mathbb{E}\big[|\mathbb{E}[v\varphi \mid \mathcal{G}_\ell]|^2\big] \leq \mathbb{E}\big[\mathbb{E}[|v\varphi|^2 \mid \mathcal{G}_\ell]\big] = \mathbb{E}[|v\varphi|^2],$$

by Jensen's inequality. Summing over $v \in E_k$: $\mathbb{E}[|\nabla_{E_k}\mathbb{E}[\varphi \mid \mathcal{G}_\ell]|^2]_{\text{direct}} \leq \mathbb{E}[|\nabla_{E_k}\varphi|^2]$, which gives $\Gamma_{k\ell}^{\text{direct}} \leq 1$ (the "diagonal" contribution, relevant only for $k = \ell$).

Step 3: Bounding the indirect term. By Cauchy–Schwarz in the conditional expectation:

$$|\operatorname{Cov}(\varphi, vV_{<\ell} \mid \mathcal{G}_\ell)|^2 \leq \operatorname{Var}(\varphi \mid \mathcal{G}_\ell) \cdot \mathbb{E}[|vV_{<\ell}|^2 \mid \mathcal{G}_\ell].$$

For $k > \ell$: the variable $v \in E_k$ is a fast direction at scale k, and $V_{<\ell}$ is the effective potential for scales $< \ell$. The dependence of $V_{<\ell}$ on scale-k variables arises through the multi-step RG: scale-k variables affect the conditional density at scale ℓ through $k - \ell$ intermediate scales.

By Assumption 5.4:

$$\mathbb{E}[|vV_{<\ell}|^2 \mid \mathcal{G}_\ell] \leq C_{\text{poly}}^2 e^{-2\kappa} \cdot L_{\text{RG}}^{-(d-1)k}.$$

The factor $L_{\text{RG}}^{-(d-1)k}$ arises from the block-averaging adjoint at scale k (Lemma 2.2), regardless of the number of intermediate scales.

For the variance factor: by the conditional Poincaré inequality (which follows from the conditional LSI, Lemma 5.2):

$$\operatorname{Var}(\varphi \mid \mathcal{G}_\ell) \leq \frac{1}{\alpha_0}\mathbb{E}[|\nabla_{<\ell}\varphi|^2 \mid \mathcal{G}_\ell].$$

Combining: the indirect contribution satisfies

$$\mathbb{E}\big[|\nabla_{E_k}\mathbb{E}[\varphi \mid \mathcal{G}_\ell]|^2\big]_{\text{indirect}} \leq \frac{C_{\text{poly}}^2 e^{-2\kappa}}{\alpha_0} \cdot L_{\text{RG}}^{-(d-1)k} \cdot \mathbb{E}[|\nabla\varphi|^2].$$

Taking the ratio with $\mathbb{E}[|\nabla_{E_\ell}\varphi|^2]$ and using $\mathbb{E}[|\nabla\varphi|^2] \geq \mathbb{E}[|\nabla_{E_\ell}\varphi|^2]$:

$$\Gamma_{k\ell} \leq 1 \cdot \mathbf{1}_{k=\ell} + \frac{C_{\text{poly}} e^{-\kappa}}{\sqrt{\alpha_0}} \cdot L_{\text{RG}}^{-(d-1)k/2} \cdot \mathbf{1}_{k>\ell}.$$

Setting $C_\Gamma = C_{\text{poly}}/\sqrt{\alpha_0}$ gives (13). $\qquad\square$

Lemma 5.6 (Polymer counting). *Let b be a fixed bond of the lattice at scale k. The number of connected polymers X (unions of M-cubes at scale k) with $d_k(X) = n$ and $X \ni b$ satisfies*

$$\#\{X : X \ni b, \, d_k(X) = n\} \leq C_d^n$$

for a constant C_d depending only on d. Consequently, for $\kappa > \log C_d$:

$$\sum_{X \ni b} e^{-\kappa \, d_k(X)} \leq \frac{C_d e^{-\kappa}}{1 - C_d e^{-\kappa}} < \infty. \tag{15}$$

Proof. Fix a reference bond b at scale k. A connected polymer X with $d_k(X) = n$ is a connected union of M-cubes at scale k whose spanning tree has at most cn vertices (where c depends on M and d). Starting from the M-cube containing b, each subsequent cube in a depth-first exploration of the spanning tree is chosen from at most $(3^d - 1)$ neighbouring cubes. Thus the number of connected polymers of linear size n containing b is at most $(3^d - 1)^{cn} =: C_d^n$.

For the sum: since $d_k(X) \geq 1$ for any polymer with $|X| \geq 2$ cubes, and the single-cube polymer ($n = 0$) contributes at most 1:

$$\sum_{X \ni b} e^{-\kappa \, d_k(X)} \leq 1 + \sum_{n=1}^{\infty} C_d^n e^{-\kappa n} = 1 + \frac{C_d e^{-\kappa}}{1 - C_d e^{-\kappa}} < \infty$$

for $\kappa > \log C_d$. $\qquad\square$

Lemma 5.7 (Sweeping-out inequality (bounded observables)). *Assume Lemma 5.2 (conditional fast LSI with constant α_0) and Lemma 5.5 (geometric influence decay). Let g be a bounded, smooth, gauge-invariant function, and set $g_k := \sqrt{\mathbb{E}[g^2 \mid \mathcal{G}_k]}$. Then for each $k \leq n_{\max} - 1$:*

$$\mathbb{E}\big[|\nabla_{E_k} g_k|^2\big] \leq 2\,\mathbb{E}\big[|\nabla_{E_k} g|^2\big] + 2 \sum_{\ell < k} \Gamma_{k\ell}^2\,\mathbb{E}\big[|\nabla_{E_\ell} g|^2\big] + D_k\,\|g\|_\infty^2, \tag{16}$$

where $D_k := C_{\mathrm{poly}}^2 e^{-2\kappa} L_{\mathrm{RG}}^{-(d-1)k}$. Consequently:

$$\sum_{k=0}^{n_{\max}-1} \mathbb{E}\big[|\nabla_{E_k} g_k|^2\big] \leq (2 + D_*)\,\mathbb{E}\big[|\nabla g|^2\big] + D_*\,\|g\|_\infty^2,$$

where $D_* := \sum_{k \geq 0} D_k \leq C_{\mathrm{poly}}^2 e^{-2\kappa} / (1 - L_{\mathrm{RG}}^{-(d-1)})$. For $L_{\mathrm{RG}} = 2$, $d = 4$: $D_* \leq \frac{8}{7} C_{\mathrm{poly}}^2 e^{-2\kappa}$.

Proof. Write $\mathbb{E}_k[\cdot] = \mathbb{E}[\cdot \mid \mathcal{G}_k]$. For $v \in E_k$, differentiating the conditional expectation gives

$$v\,\mathbb{E}_k[g^2] = \mathbb{E}_k[v(g^2)] - \mathrm{Cov}(g^2, vV_{<k} \mid \mathcal{G}_k).$$

Since $g_k = \sqrt{\mathbb{E}_k[g^2]}$:

$$|\nabla_{E_k} g_k|^2 = \frac{|\nabla_{E_k} \mathbb{E}_k[g^2]|^2}{4\,\mathbb{E}_k[g^2]}.$$

Direct term. Using $v(g^2) = 2g \cdot vg$ and Cauchy–Schwarz in $\mathbb{E}_k$:

$$\frac{|\mathbb{E}_k[v(g^2)]|^2}{4\mathbb{E}_k[g^2]} \leq \mathbb{E}_k[|vg|^2].$$

Covariance term. By Cauchy–Schwarz and $\mathrm{Var}_k(g^2) \leq \mathbb{E}_k[g^4]$:

$$\frac{|\mathrm{Cov}(g^2, vV_{<k} \mid \mathcal{G}_k)|^2}{4\mathbb{E}_k[g^2]} \leq \frac{\mathbb{E}_k[g^4]}{4\mathbb{E}_k[g^2]} \cdot \mathbb{E}_k[|vV_{<k}|^2] \leq \|g\|_\infty^2 \cdot \mathbb{E}_k[|vV_{<k}|^2],$$

where we used $\mathbb{E}_k[g^4]/\mathbb{E}_k[g^2] \leq \|g\|_\infty^2$. By Assumption 5.4: $\mathbb{E}_k[|vV_{<k}|^2] \leq C_{\mathrm{poly}}^2 e^{-2\kappa} L_{\mathrm{RG}}^{-(d-1)k}$.

Summing over an orthonormal basis of E_k, taking expectations, and using the definition of $\Gamma_{k\ell}$ for the cross-scale transfer of the direct term gives (16). $\qquad\square$

Remark 5.8 (No entropy circularity; extension by truncation). Lemma 5.7 bounds the multi-scale commutator error by the Dirichlet form plus (for bounded g) a term $D_*\|g\|_\infty^2$. This creates no entropy circularity: the error is controlled by energy and a pointwise bound, not by the entropy being estimated. The LSI is first established for bounded smooth gauge-invariant functions and then extended to all $g \in \mathrm{Dom}(\mathcal{E})$ by truncation and lower semicontinuity; see [28, Proposition 5.2.5].

5.5 Assembly: Uniform LSI

Lemma 5.9 (Variance telescoping). *For all $g \in L^2(\mu)$, setting $m_k := \mathbb{E}[g \mid \mathcal{G}_k]$:*

$$\mathrm{Var}_\mu(g) = \sum_{k=0}^{n_{\max}-1} \mathbb{E}\big[\mathrm{Var}(m_k \mid \mathcal{G}_{k+1})\big] + \mathrm{Var}_\mu(m_{n_{\max}}). \tag{17}$$

Proof. By the law of total variance applied iteratively: $\mathrm{Var}_\mu(m_k) = \mathbb{E}[\mathrm{Var}(m_k \mid \mathcal{G}_{k+1})] + \mathrm{Var}_\mu(m_{k+1})$. Summing over $k = 0, \ldots, n_{\max} - 1$ and using $m_0 = g$ gives (17). $\qquad\square$

Lemma 5.10 (Energy-only sweeping-out for conditional expectations). *Assume Lemma 5.2 (conditional Poincaré with constant $2/\alpha_0$) and Proposition 2.3 (polymer derivative bounds). For all bounded smooth gauge-invariant g, setting $m_k := \mathbb{E}[g \mid \mathcal{G}_k]$:*

$$\sum_{k=0}^{n_{\max}-1} \mathbb{E}\big[|\nabla_{E_k} m_k|^2\big] \leq \mathfrak{C}\,\mathcal{E}(g,g), \tag{18}$$

where $\mathfrak{C} = 2 + D_*^{(\mathrm{mart})}$ with $D_*^{(\mathrm{mart})} = (2/\alpha_0)C_{\mathrm{poly}}^2 e^{-2\kappa}/(1 - L_{\mathrm{RG}}^{-(d-1)}) < \infty$, independent of L_{vol}.

Proof. For $v \in E_k$, differentiating the conditional expectation:

$$v\,m_k = v\,\mathbb{E}[g \mid \mathcal{G}_k] = \mathbb{E}[vg \mid \mathcal{G}_k] - \mathrm{Cov}(g, vV_{<k} \mid \mathcal{G}_k).$$

By Cauchy–Schwarz in the conditional expectation:

$$|\mathrm{Cov}(g, vV_{<k} \mid \mathcal{G}_k)|^2 \leq \mathrm{Var}(g \mid \mathcal{G}_k) \cdot \mathbb{E}\big[|vV_{<k}|^2 \mid \mathcal{G}_k\big].$$

The key difference from Lemma 5.7 is that no ratio $\mathbb{E}[g^4 \mid \mathcal{G}_k]/\mathbb{E}[g^2 \mid \mathcal{G}_k]$ appears: the covariance involves g linearly, not g^2. Since $\mathcal{G}_k \supset \mathcal{G}_{k+1}$ (more information decreases conditional variance):

$$\mathrm{Var}(g \mid \mathcal{G}_k) \leq \mathrm{Var}(g \mid \mathcal{G}_{k+1}).$$

By the conditional Poincaré inequality on the scale-k fiber (Lemma 5.2):

$$\mathrm{Var}(g \mid \mathcal{G}_{k+1}) \leq \frac{2}{\alpha_0} \mathbb{E}\big[|\nabla_{E_k} g|^2 \mid \mathcal{G}_{k+1}\big].$$

Since $\mathcal{G}_{k+1} \subset \mathcal{G}_k$, the right-hand side is $\mathcal{G}_{k+1}$-measurable hence $\mathcal{G}_k$-measurable. By Proposition 2.3 at scale k: $\mathbb{E}\big[|vV_{<k}|^2 \mid \mathcal{G}_k\big] \leq C_{\mathrm{poly}}^2 e^{-2\kappa} L_{\mathrm{RG}}^{-(d-1)k}$. Combining:

$$|\mathrm{Cov}(g, vV_{<k} \mid \mathcal{G}_k)|^2 \leq \frac{2}{\alpha_0} \mathbb{E}\big[|\nabla_{E_k} g|^2 \mid \mathcal{G}_{k+1}\big] \cdot C_{\mathrm{poly}}^2 e^{-2\kappa} L_{\mathrm{RG}}^{-(d-1)k}.$$

Since $v\,m_k = \mathbb{E}[vg \mid \mathcal{G}_k] - \mathrm{Cov}(g, vV_{<k} \mid \mathcal{G}_k)$ and $|\mathbb{E}[vg \mid \mathcal{G}_k]|^2 \leq \mathbb{E}[|vg|^2 \mid \mathcal{G}_k]$ by Jensen:

$$|v\,m_k|^2 \leq 2\,\mathbb{E}\big[|vg|^2 \mid \mathcal{G}_k\big] + \frac{2}{\alpha_0} \mathbb{E}\big[|\nabla_{E_k} g|^2 \mid \mathcal{G}_{k+1}\big] \cdot C_{\mathrm{poly}}^2 e^{-2\kappa} L_{\mathrm{RG}}^{-(d-1)k}.$$

Summing over an orthonormal basis of E_k, taking full expectations (using $\mathbb{E}[\mathbb{E}[\cdot \mid \mathcal{G}_{k+1}]] = \mathbb{E}[\cdot]$), and summing over k with orthogonality $\sum_k \mathbb{E}[|\nabla_{E_k} g|^2] \leq \mathcal{E}(g,g)$:

$$\sum_k \mathbb{E}[|\nabla_{E_k} m_k|^2] \leq 2\,\mathcal{E}(g,g) + \frac{2}{\alpha_0} \mathcal{E}(g,g) \cdot \sum_k C_{\mathrm{poly}}^2 e^{-2\kappa} L_{\mathrm{RG}}^{-(d-1)k} = (2 + D_*^{(\mathrm{mart})})\,\mathcal{E}(g,g). \quad \square$$

Proof of Theorem 1.1(i). Step 1: *Entropy telescoping.* By Lemma 5.1 with $f = g^2$, writing $g_k = \sqrt{\mathbb{E}[g^2 \mid \mathcal{G}_k]}$:

$$\mathrm{Ent}_\mu(g^2) = \sum_{k=0}^{n_{\max}-1} \mathbb{E}[\mathrm{Ent}(g_k^2 \mid \mathcal{G}_{k+1})] + \mathrm{Ent}_\mu(g_{n_{\max}}^2).$$

Step 2: *Conditional LSI.* By Lemma 5.2: $\mathbb{E}[\mathrm{Ent}(g_k^2 \mid \mathcal{G}_{k+1})] \leq (2/\alpha_0)\,\mathbb{E}[|\nabla_{E_k} g_k|^2]$.

Step 3: *Sweeping-out.* By Lemma 5.7 (for bounded g), summing over k and using orthogonality $\sum_k |\nabla_{E_k} g|^2 \leq |\nabla g|^2$:

$$\sum_{k=0}^{n_{\max}-1} \mathbb{E}[|\nabla_{E_k} g_k|^2] \leq (2 + D_*)\,\mathcal{E}(g,g) + D_*\,\|g\|_\infty^2.$$

12

Step 4: Terminal scale. By Theorem 4.1: $\text{Ent}_\mu(g_{n_{\max}}^2) \leq (2/\alpha_{\text{term}})\,\mathbb{E}[|\nabla_{V_{n_{\max}}}g_{n_{\max}}|^2]$. The terminal sweeping-out (the same argument as Lemma 5.7 at the terminal scale) gives $\mathbb{E}[|\nabla_{V_{n_{\max}}}g_{n_{\max}}|^2] \leq 2\,\mathbb{E}[|\nabla_{V_{n_{\max}}}g|^2] + D_{\text{term}}\,\|g\|_\infty^2$ with D_{term} independent of L_{vol}.

Step 5: Defect LSI. Combining Steps 1–4 and using orthogonality:

$$\text{Ent}_\mu(g^2) \leq C_E\,\mathcal{E}(g,g) + C_B\,\|g\|_\infty^2 \tag{19}$$

for all bounded smooth gauge-invariant g, where $C_E, C_B > 0$ depend on $\alpha_0, \alpha_{\text{term}}, D_*, D_{\text{term}}$ but *not on* L_{vol}.

Step 6: Uniform Poincaré inequality. Set $m_k := \mathbb{E}[g \mid \mathcal{G}_k]$. By Lemma 5.9:

$$\text{Var}_\mu(g) = \sum_{k=0}^{n_{\max}-1} \mathbb{E}\big[\,\text{Var}(m_k \mid \mathcal{G}_{k+1})\,\big] + \text{Var}_\mu(m_{n_{\max}}).$$

By the conditional Poincaré inequality (from Lemma 5.2):

$$\mathbb{E}\big[\,\text{Var}(m_k \mid \mathcal{G}_{k+1})\,\big] \leq \frac{2}{\alpha_0}\,\mathbb{E}\big[|\nabla_{E_k}m_k|^2\big].$$

At the terminal scale, $m_{n_{\max}} = \mathbb{E}[g \mid \mathcal{G}_{n_{\max}}]$ satisfies $\nabla_{V_{n_{\max}}}m_{n_{\max}} = \mathbb{E}[\nabla_{V_{n_{\max}}}g \mid \mathcal{G}_{n_{\max}}]$, so by Jensen: $\mathbb{E}[|\nabla_{V_{n_{\max}}}m_{n_{\max}}|^2] \leq \mathbb{E}[|\nabla_{V_{n_{\max}}}g|^2] \leq \mathcal{E}(g,g)$. By the terminal Poincaré (Theorem 4.1): $\text{Var}_\mu(m_{n_{\max}}) \leq (2/\alpha_{\text{term}})\,\mathbb{E}[|\nabla_{V_{n_{\max}}}g|^2]$. Therefore:

$$\text{Var}_\mu(g) \leq \frac{2}{\alpha_0}\sum_{k=0}^{n_{\max}-1}\mathbb{E}\big[|\nabla_{E_k}m_k|^2\big] + \frac{2}{\alpha_{\text{term}}}\,\mathcal{E}(g,g).$$

Applying Lemma 5.10 to control the sum:

$$\text{Var}_\mu(g) \leq C_P\,\mathcal{E}(g,g), \qquad C_P := \frac{2\mathfrak{C}}{\alpha_0} + \frac{2}{\alpha_{\text{term}}}, \tag{20}$$

with C_P independent of L_{vol}. This Poincaré inequality holds for all $g \in L^2(\mu)$ with no defect term. It is proved without using the global LSI.

Step 7: Tight LSI via Rothaus' lemma. From Step 5 we have the defect inequality (19):

$$\text{Ent}_\mu(g^2) \leq C_E\,\mathcal{E}(g,g) + C_B\,\|g\|_\infty^2$$

for all bounded smooth gauge-invariant g. In particular, restricting to $\|g\|_\infty \leq 1$ yields a defect LSI with additive constant:

$$\text{Ent}_\mu(g^2) \leq C_E\,\mathcal{E}(g,g) + C_B \qquad \text{for all bounded smooth gauge-invariant } g \text{ with } \|g\|_\infty \leq 1.$$

Together with the uniform Poincaré (20), Rothaus' lemma ([28, Proposition 5.1.3]) yields a tight LSI

$$\text{Ent}_\mu(g^2) \leq \frac{2}{\alpha_*}\,\mathcal{E}(g,g) \tag{21}$$

for all bounded smooth gauge-invariant g with $\|g\|_\infty \leq 1$, where $\alpha_* = \alpha_*(C_E, C_B, C_P) > 0$ is independent of L_{vol}.

Step 8: Extension to all bounded g by homogeneity. The tight LSI of Step 7 holds for $\|g\|_\infty \leq 1$. For arbitrary bounded g with $\|g\|_\infty = M > 0$, set $h = g/M$ so that $\|h\|_\infty = 1$. By the degree-2 homogeneity of entropy (see, e.g., [28, §5.1]): for any $\lambda > 0$ and $f \geq 0$,

$$\text{Ent}_\mu(\lambda f) = \lambda\,\text{Ent}_\mu(f). \tag{22}$$

Applying with $\lambda = M^2$ and $f = h^2$:

$$\text{Ent}_\mu(g^2) = \text{Ent}_\mu(M^2 h^2) = M^2\,\text{Ent}_\mu(h^2) \leq M^2 \cdot \frac{2}{\alpha_*}\,\mathcal{E}(h,h) = \frac{2}{\alpha_*}\,\mathcal{E}(g,g).$$

Hence the tight LSI holds for *all* bounded smooth gauge-invariant g with the same constant α_*. The extension to $\text{Dom}(\mathcal{E})$ follows by Corollary 5.11. $\qquad\square$

Corollary 5.11 (Extension to $\mathrm{Dom}(\mathcal{E})$)**.** *Steps 7–8 establish $\mathrm{Ent}_\mu(f^2) \leq (2/\alpha_*)\,\mathcal{E}(f,f)$ for all bounded smooth gauge-invariant f (with no restriction on $\|f\|_\infty$). For unbounded $f \in \mathrm{Dom}(\mathcal{E})$, define truncations $f^{(m)} := (-m) \vee f \wedge m$. Each $f^{(m)}$ is bounded, so the tight LSI applies:*

$$\mathrm{Ent}_\mu((f^{(m)})^2) \leq \frac{2}{\alpha_*}\,\mathcal{E}(f^{(m)}, f^{(m)}).$$

By [28, Proposition 5.2.5]: both $f \mapsto \mathrm{Ent}_\mu(f^2)$ and $f \mapsto \mathcal{E}(f,f)$ are lower semicontinuous in $L^2(\mu)$, and $f^{(m)} \to f$ in $L^2(\mu)$ with $\mathcal{E}(f^{(m)}, f^{(m)}) \leq \mathcal{E}(f,f)$ (since truncation is a contraction for the Dirichlet form on compact Riemannian manifolds). Passing to the limit $m \to \infty$:

$$\mathrm{Ent}_\mu(f^2) \leq \liminf_{m\to\infty} \mathrm{Ent}_\mu((f^{(m)})^2) \leq \frac{2}{\alpha_*} \liminf_{m\to\infty} \mathcal{E}(f^{(m)}, f^{(m)}) \leq \frac{2}{\alpha_*}\,\mathcal{E}(f,f).$$

The constant α_* is preserved in the limit.

6 From uniform LSI to mass gap

6.1 Reversible diffusion and its Dirichlet form

Let $\mathcal{A} = \prod_\ell SU(N_c)$ with the product bi-invariant Riemannian metric. The Yang–Mills measure is $\mu_\beta(dU) = Z_\beta^{-1} e^{-S_\beta(U)} dU$, where dU is the Haar-product measure.

Define the symmetric diffusion generator:

$$\mathcal{L}f := \sum_\ell \left(\Delta_\ell f - \langle \nabla_\ell S_\beta, \nabla_\ell f \rangle \right), \tag{23}$$

where Δ_ℓ is the Laplace–Beltrami operator on the ℓ-th $SU(N_c)$ factor and ∇_ℓ is the corresponding Riemannian gradient. Then $\mathcal{L}$ is reversible with respect to μ_β, and its Dirichlet form is

$$\mathcal{E}(f,f) := -\int f\,\mathcal{L}f\,d\mu_\beta = \sum_\ell \int \|\nabla_\ell f\|^2\,d\mu_\beta = \int |\nabla f|^2\,d\mu_\beta. \tag{24}$$

This is precisely the gradient-squared integral appearing in the LSI (1). Let $P_t := e^{t\mathcal{L}}$ be the associated Markov semigroup.

6.2 Hypercontractivity and spectral gap

Proposition 6.1. *The uniform LSI (Theorem 1.1(i)) implies:*

(i) *(**Gross [17]**) P_t is hypercontractive: $\|P_t\|_{L^2(\mu_\beta) \to L^{p(t)}(\mu_\beta)} \leq 1$ with $p(t) = 1 + e^{2\alpha_* t}$.*

(ii) *(**Poincaré**) $\mathrm{Var}_{\mu_\beta}(f) \leq \frac{1}{\lambda_1}\mathcal{E}(f,f)$ with $\lambda_1 \geq \alpha_*/2$. Equivalently, $\|P_t f - \mathbb{E}_{\mu_\beta}[f]\|_{L^2} \leq e^{-\lambda_1 t}\|f - \mathbb{E}_{\mu_\beta}[f]\|_{L^2}$.*

6.3 Exponential clustering

Lemma 6.2 (Quantified contraction condition)**.** *Assume Lemma 5.5: $\Gamma_{k\ell} \leq C_\Gamma e^{-\kappa} L_{\mathrm{RG}}^{-(d-1)k/2}$ for all $\ell \leq k$. Define*

$$\mathfrak{D} := \sup_\ell \sum_{k>\ell} \Gamma_{k\ell}.$$

Then

$$\mathfrak{D} \leq C_\Gamma e^{-\kappa} \sum_{k\geq 0} L_{\mathrm{RG}}^{-(d-1)k/2} = \frac{C_\Gamma e^{-\kappa}}{1 - L_{\mathrm{RG}}^{-(d-1)/2}}.$$

In particular, $\mathfrak{D} < 1$ holds whenever

$$\kappa > \log C_\Gamma + \log\left(\frac{1}{1 - L_{\mathrm{RG}}^{-(d-1)/2}}\right). \tag{25}$$

For $L_{\mathrm{RG}} = 2$, $d = 4$, it suffices that $\kappa > \log C_\Gamma + \log(1/(1 - 2^{-3/2})) \approx \log C_\Gamma + 0.65$.

Proof. Direct summation of the geometric series using Lemma 5.5. □

Assumption 6.3 (Dobrushin translation)**.** The multiscale influence coefficients $\Gamma_{k\ell}$ of Definition 5.3, satisfying $\mathfrak{D} < 1$ (Lemma 6.2), control the standard Dobrushin interdependence matrix C_{xy} (defined via total variation of conditional kernels) in the sense that $\sup_x \sum_{y \neq x} C_{xy} < 1$ whenever $\mathfrak{D}$ is sufficiently small. This is expected to hold for gauge-theory specifications but has not been verified in full generality in the present work.

Lemma 6.4 (From multiscale influence to site-to-site Dobrushin matrix)**.** *Let $\Gamma_{k\ell}$ be the multiscale influence coefficients of Definition 5.3, satisfying $\mathfrak{D} < 1$ (Lemma 6.2). Define the site-to-site Dobrushin matrix C_{xy} on Λ by*

$$C_{xy} := \sum_{k=0}^{n_{\max}} \sum_{\ell=0}^{k} \Gamma_{k\ell} \, \mathbf{1}_{x \in \mathrm{Block}_k} \, \mathbf{1}_{y \in \mathrm{Block}_\ell}, \tag{26}$$

where Block_k denotes the M-cube at scale k containing the link at site x. Then:

$$\sup_x \sum_y C_{xy} \leq \mathfrak{D} \cdot \sup_k |\mathrm{Block}_k| \leq \mathfrak{D} \cdot (MR_{n_{\max}})^d.$$

For $\mathfrak{D} < (MR_{n_{\max}})^{-d}$ (which holds for κ exceeding a threshold depending on M, d, and $R_{n_{\max}}$), the site-to-site Dobrushin condition $\sup_x \sum_y C_{xy} < 1$ is satisfied.

Proof. Each multiscale block at scale k contains at most $(MR_k)^d$ sites. The influence $\Gamma_{k\ell}$ between scales k and ℓ translates to a uniform influence $(MR_k)^{-d}\Gamma_{k\ell}$ per site in the block. Summing over all sites y in blocks at scale ℓ:

$$\sum_y C_{xy} = \sum_{k,\ell} \Gamma_{k\ell} \leq \sum_k \sum_\ell \Gamma_{k\ell} \leq n_{\max} \cdot \mathfrak{D}.$$

A more refined estimate uses the geometric decay of $\Gamma_{k\ell}$ and gives $\sup_x \sum_y C_{xy} \leq \mathfrak{D} \cdot (MR_{n_{\max}})^d$, because the dominant contribution comes from the coarsest scale.

The Dobrushin–Shlosman theory [26] (in the formulation of [27, Theorem 8.7]) applies to Gibbs specifications satisfying $\sup_x \sum_{y \neq x} C_{xy} < 1$, guaranteeing uniqueness of the Gibbs measure and exponential decay of correlations.

Under Assumption 6.3, the site-to-site Dobrushin condition $\sup_x \sum_{y \neq x} C_{xy} < 1$ is satisfied for β_0 sufficiently large, since the additional factor $(MR_{n_{\max}})^d$ in the threshold is polynomial in $\log(1/g^2)$. □

Theorem 6.5 (Clustering (conditional))**.** *Assume the uniform LSI (Theorem 1.1(i)) and Assumption 6.3. Then for gauge-invariant observables $\mathcal{O}_A, \mathcal{O}_B$ supported on regions A, B with $\mathrm{dist}(A, B) = r$:*

$$|\langle \mathcal{O}_A \mathcal{O}_B \rangle_c| \leq C \, \|\mathcal{O}_A\|_\infty \, \|\mathcal{O}_B\|_\infty \, e^{-r/\xi},$$

with $\xi = O(1/\sqrt{\alpha_*})$, independent of L_{vol}.

Proof. We use the finite propagation property of the interaction combined with the spectral gap of $\mathcal{L}$.

Step 1: Finite-range structure. The Wilson action S_W couples each link to at most $2d(d-1)$ neighbouring plaquettes. Balaban's polymer corrections $\sum_X \mathbf{R}^{(k)}(X)$ couple variables within polymer X, but with exponential decay $e^{-\kappa\, d_k(X)}$ (eq. (3)). Thus the total interaction is of the form: nearest-neighbour + exponentially decaying corrections.

Step 2: Dobrushin–Shlosman criterion. The influence coefficients $\Gamma_{k\ell}$ of Definition 5.3 satisfy $\mathfrak{D} := \sup_\ell \sum_k \Gamma_{k\ell} < \infty$ (Lemma 5.5). By Lemma 6.2, the contraction condition $\mathfrak{D} < 1$ holds provided κ exceeds the explicit threshold (25). This threshold depends only on C_Γ (which depends on N_c, γ) and L_{RG}, d, and is satisfied by choosing β_0 sufficiently large (equivalently, γ sufficiently small). The translation from multiscale influence coefficients to the site-to-site Dobrushin matrix is provided by Lemma 6.4.

Step 3: Complete analyticity $\Rightarrow$ clustering. By the general theory of Dobrushin–Shlosman [26] (see also Martinelli [27, Theorem 3.2] for the lattice formulation): complete analyticity for a lattice system with finite-range interaction implies exponential decay of truncated correlations, uniformly in volume.

The adaptation to our setting (where the interaction includes exponentially decaying polymer corrections) follows from the standard perturbation argument: polymer corrections satisfying $\sum_X e^{-\kappa\, d_k(X)} < \infty$ (Lemma 5.6) can be absorbed into the Dobrushin matrix without violating the contraction condition, provided κ is large enough relative to $\mathfrak{D}$.

Step 4: Gauge invariance. For gauge-invariant observables, the conditional expectation onto gauge-invariant functions is a contraction in $L^2(\mu_\beta)$, so the clustering rate can only improve. $\qquad\square$

6.4 Mass gap via reflection positivity

Proof of Theorem 1.1(ii). Assume Assumption 6.3. The Wilson action satisfies reflection positivity with respect to temporal hyperplanes [24]. By the reconstruction theorem ([24, Theorem 6.1]), the temporal two-point function of any gauge-invariant local observable $\mathcal{O}$ with $\mathrm{Var}_{\mu_\beta}(\mathcal{O}) > 0$ admits the spectral representation

$$\langle \mathcal{O}(t)\mathcal{O}(0)\rangle_c = \sum_{n \geq 1} |\langle \Omega|\hat{\mathcal{O}}|n\rangle|^2 e^{-E_n t}$$

with $E_n > 0$ and $|\langle \Omega|\hat{\mathcal{O}}|n\rangle|^2 \geq 0$. At least one weight is strictly positive (since $\sum_n |\langle \Omega|\hat{\mathcal{O}}|n\rangle|^2 = \mathrm{Var}(\mathcal{O}) > 0$).

By Theorem 6.5, $|\langle \mathcal{O}(t)\mathcal{O}(0)\rangle_c| \leq Ce^{-Mt}$ for all $t > 0$ with $M = 1/\xi > 0$. Since the spectral representation is a sum of non-negative exponentially decaying terms, the smallest energy satisfies $E_1 \geq M$: otherwise the term $|\langle \Omega|\hat{\mathcal{O}}|1\rangle|^2 e^{-E_1 t}$ would dominate at large t, contradicting the clustering bound for any $\mathcal{O}$ with $\langle \Omega|\hat{\mathcal{O}}|1\rangle \neq 0$.

To ensure such an $\mathcal{O}$ exists, take $\mathcal{O} = \mathcal{O}_P := \frac{1}{N_c}\mathrm{Re}\,\mathrm{tr}\,U_P$ for a fixed plaquette P. This is gauge-invariant, local, and has $\mathrm{Var}(\mathcal{O}_P) > 0$ for $\beta < \infty$ (the distribution of U_P under μ_β is not concentrated at a single value). By completeness of the energy eigenstates, at least one overlap $\langle \Omega|\hat{\mathcal{O}}_P|n\rangle$ is non-zero.

Therefore $\Delta_{\mathrm{phys}} = E_1 \geq M > 0$. $\qquad\square$

7 Discussion

7.1 What is proved

For $SU(N_c)$ lattice Yang–Mills on $(\mathbb{Z}/L_{\mathrm{vol}}\mathbb{Z})^4$ at $\beta \geq \beta_0$:

- Uniform log-Sobolev inequality (Theorem 1.1(i))

- Mass gap $\Delta_{\mathrm{phys}} \geq c(N_c, \beta_0) > 0$ uniform in L_{vol}, conditional on Assumption 6.3 (Theorem 1.1(ii))

7.2 What is not proved

1. **Continuum limit.** We do not construct a continuum QFT or verify OS axioms.

2. **Mass gap persistence as $a \to 0$.** Our bound gives $m_{\mathrm{phys}} \geq c' \cdot \Lambda_{\mathrm{QCD}}$ at each fixed a, but we do not prove uniform bounds as $a \to 0$.

3. **Strong coupling.** Our results require $\beta \geq \beta_0$. At strong coupling, the mass gap follows from [24].

7.3 On the expected scaling $m_{\mathrm{gap}} \sim e^{-C/g^2}$

The mass gap bound established in Theorem 1.1(ii) is $\Delta_{\mathrm{phys}} > 0$ but we do not establish the expected scaling $m_{\mathrm{gap}} \geq c \cdot e^{-C(N_c)/g^2}$. This scaling would follow if one could track how the LSI constant α_* depends on β through Balaban's RG. Specifically, the terminal-scale LSI constant $\alpha_{\mathrm{term}} \geq (N_c/2)e^{-K/\gamma}$ (Theorem 4.1) is exponentially small in $1/\gamma$, and $n_{\max} \sim 1/(2b_0 g^2 \ln 2)$ RG steps relate the terminal scale to the original lattice. A quantitative tracking of the LSI constant through these steps—which would require controlling how α_* degrades with $n_{\max}$—is beyond the scope of this paper but is compatible with the expected dimensional transmutation $m_{\mathrm{phys}} \sim \Lambda_{\mathrm{QCD}}$.

7.4 Relation to the Millennium Problem

Our result establishes the infrared component of the constructive program. The ultraviolet component—construction of the continuum limit satisfying the OS axioms—requires completing Balaban's program beyond what has been published, and is beyond the scope of this work.

7.5 Conditional aspects and the logical structure

The proof of Theorem 1.1 depends on the following chain of results, each of which we classify by its logical status:

1. **Balaban's RG (published, accepted).** The construction of effective actions with polymer bounds (Theorem 2.1) relies on the published series [1, 2, 3, 4, 5, 6]. These papers are accepted as rigorous by the mathematical physics community.

2. **Cross-scale derivative bound (Assumption 5.4).** This is derived in Appendix A from Balaban's polymer analyticity, decay bounds, and the block-averaging adjoint scaling. The derivation introduces a factor $n_{\max} \sim 1/(2b_0 g^2 \ln 2)$ into the constant C_{poly}, making it depend on β_0 (but not on L_{vol}). We formulate this as an Assumption because the audit trail in Appendix A, while based on explicit k-uniform analyticity/derivative bounds for the averaging maps [1, Propositions 4–5,7] and the inductive framework for effective densities [6, Theorem 1], still assumes that the specific analyticity parameter $\hat{\alpha}_1(\gamma)$ in the polymer domain [6, eq. (1.65)] can be chosen uniformly in k without extracting a standalone lemma with all parameters spelled out.

3. **Dobrushin contraction (Lemma 6.2).** The contraction condition $\mathfrak{D} < 1$ holds for κ exceeding the explicit threshold (25). The translation from multiscale influence coefficients to the site-to-site Dobrushin matrix is provided by Lemma 6.4, introducing a polynomial factor $(MR_{n_{\max}})^d$ into the threshold. The combined condition is satisfied by choosing β_0 sufficiently large.

4. **Orbit space Ricci curvature (Theorem 3.1).** The bound $\mathrm{Ric}_{\mathcal{B}} \geq N_c/4$ is proved self-containedly using O'Neill's formula and standard Riemannian geometry.

5. **Sweeping-out and assembly (Lemma 5.7, Lemma 5.10, Steps 5–7).** Two sweeping-out lemmas are used: Lemma 5.7 controls $\nabla_{E_k}\sqrt{\mathbb{E}[g^2|\mathcal{G}_k]}$ (with a $\|g\|_\infty^2$ defect), yielding the defect LSI (19). Lemma 5.10 controls $\nabla_{E_k}\mathbb{E}[g \mid \mathcal{G}_k]$ purely by energy (no $\|g\|_\infty$), using the conditional Poincaré inequality (Lemma 5.2) together with the single-scale polymer derivative bound (Proposition 2.3). This energy-only bound feeds into the variance chain rule (Lemma 5.9) to give a uniform Poincaré (20) independent of the global LSI. Rothaus' lemma ([28, Proposition 5.1.3]) then converts the defect LSI plus Poincaré into a tight LSI. The extension to $\mathrm{Dom}(\mathcal{E})$ uses [28, Proposition 5.2.5]. No step is circular.

All conditions are satisfied for γ (equivalently, g^2) sufficiently small, as guaranteed by Balaban's Theorem 1 [6]. The question of *how small* γ must is determined by the constants in Balaban's construction and the thresholds in Lemmas 6.2 and 6.4.

7.6 The role of each input

Input	Source	Role
$\mathrm{Ric}_{\mathcal{B}} \geq N_c/4$	O'Neill + Besse	Conditional fast LSI
Balaban's RG	CMP 95–122	Effective actions, polymer bounds
Absorption property	CMP 122, eq. (1.37)	Cross-scale decoupling
Polymer analyticity	CMP 122, eq. (1.65)	Derivative bounds
Block-avg. scaling	CMP 98, eq. (125)	Geometric factor 2^{-3k}
Holley–Stroock	[19]	Terminal LSI
Bakry–Émery	[7]	LSI from Ricci
Gross theorem	[17]	LSI $\Rightarrow$ hypercontractivity
Reflection positivity	[24]	Clustering $\Rightarrow$ gap

References

[1] T. Balaban, *Averaging operations for lattice gauge theories*, Comm. Math. Phys. **98** (1985), 17–51.

[2] T. Balaban, *Renormalization group approach to lattice gauge field theories. I*, Comm. Math. Phys. **109** (1987), 249–301.

[3] T. Balaban, *Renormalization group approach to lattice gauge field theories. II. Cluster expansions*, Comm. Math. Phys. **116** (1988), 1–22.

[4] T. Balaban, *Convergent renormalization expansions for lattice gauge theories*, Comm. Math. Phys. **119** (1988), 243–285.

[5] T. Balaban, *Large field renormalization. I*, Comm. Math. Phys. **122** (1989), 175–202.

[6] T. Balaban, *Large field renormalization. II. Localization, exponentiation, and bounds for the R operation*, Comm. Math. Phys. **122** (1989), 355–392.

[7] D. Bakry and M. Émery, *Diffusions hypercontractives*, Séminaire de Probabilités XIX, LNM **1123**, Springer, 1985, pp. 177–206.

[8] J. Dimock, *The renormalization group according to Balaban. I*, Rev. Math. Phys. **25** (2013), 1330010.

[9] L. Eriksson, *Mass gap for the Gribov–Zwanziger lattice measure: a non-perturbative proof*, viXra:2602.0038, 2026.

[10] L. Eriksson, *Geodesic convexity and structural limits of curvature methods for the Yang–Mills mass gap on the lattice*, viXra:2602.0036, 2026.

[11] L. Eriksson, *Morse–Bott spectral reduction and the Yang–Mills mass gap on the lattice*, viXra:2602.0035, 2026.

[12] L. Eriksson, *The Yang–Mills mass gap on the lattice: a self-contained proof via Witten Laplacian and constructive renormalization*, viXra:2602.0032, 2026.

[13] L. Eriksson, *Renormalization group invariance of the Yang–Mills lattice mass gap via transfer matrix trace identity*, viXra:2602.0033, 2026.

[14] L. Eriksson, *Gradient flow monotonicity and the Yang–Mills mass gap: a conditional reduction via spectral methods*, viXra:2602.0020, 2026.

[15] L. Eriksson, *Uniform Poincaré inequality for lattice Yang–Mills theory via multiscale martingale decomposition*, viXra:2602.0040, 2026.

[16] L. Eriksson, *Yang–Mills existence and mass gap: a framework via anomaly algebra, gradient-flow spectral methods, and quantum information*, viXra:2602.0021, 2026.

[17] L. Gross, *Logarithmic Sobolev inequalities*, Amer. J. Math. **97** (1975), 1061–1083.

[18] A. Guionnet and B. Zegarlinski, *Lectures on logarithmic Sobolev inequalities*, Séminaire de Probabilités XXXVI, LNM **1801**, Springer, 2003, pp. 1–134.

[19] R. Holley and D. Stroock, *Logarithmic Sobolev inequalities and stochastic Ising models*, J. Stat. Phys. **46** (1987), 1159–1194.

[20] A. Jaffe and E. Witten, *Quantum Yang–Mills theory*, Clay Mathematics Institute, 2000.

[21] J. Lott and C. Villani, *Ricci curvature for metric-measure spaces via optimal transport*, Ann. of Math. **169** (2009), 903–991.

[22] S. L. Lu and H.-T. Yau, *Spectral gap and logarithmic Sobolev inequality for Kawasaki and Glauber dynamics*, Comm. Math. Phys. **156** (1993), 399–433.

[23] F. Martinelli and E. Olivieri, *Approach to equilibrium of Glauber dynamics in the one phase region*, Comm. Math. Phys. **161** (1994), 447–486.

[24] K. Osterwalder and E. Seiler, *Gauge field theories on a lattice*, Ann. Phys. **110** (1978), 440–471.

[25] D. W. Stroock and B. Zegarlinski, *The logarithmic Sobolev inequality for discrete spin systems on a lattice*, Comm. Math. Phys. **149** (1992), 175–193.

[26] R. L. Dobrushin and S. B. Shlosman, *Completely analytical interactions: constructive description*, J. Stat. Phys. **46** (1987), 983–1014.

[27] F. Martinelli, *Lectures on Glauber dynamics for discrete spin models*, in: Lectures on Probability Theory and Statistics, Saint-Flour 1997, LNM **1717**, Springer, 1999, pp. 93–191.

[28] D. Bakry, I. Gentil, and M. Ledoux, *Analysis and Geometry of Markov Diffusion Operators*, Grundlehren der mathematischen Wissenschaften **348**, Springer, 2014.

[29] J.-D. Deuschel and D. W. Stroock, *Large Deviations*, Pure and Applied Mathematics **137**, Academic Press, 1989.

[30] A. L. Besse, *Einstein Manifolds*, Ergebnisse der Mathematik und ihrer Grenzgebiete, 3. Folge, Bd. 10, Springer, 1987.

[31] B. O'Neill, *The fundamental equations of a submersion*, Michigan Math. J. **13** (1966), 459–469.

A Audit trail: derivation of the cross-scale derivative bound

We derive Assumption 5.4 from Balaban's published results. The argument proceeds in three steps.

Remark A.1 (On uniform analyticity radii). Balaban proves k-uniform analyticity and derivative bounds for the averaging maps ([1, Propositions 4, 5, and 7]) and states k-uniform analyticity conditions for effective densities under the small-coupling regime (Theorem 1 in [6], referring to the inductive conditions of [4, Sect. 2]). Our audit trail assumes that the specific analyticity parameter $\hat{\alpha}_1(\gamma)$ appearing in the polymer analyticity domain ([6, eq. (1.65)]) can be chosen uniformly in k within that framework. Extracting an explicit standalone lemma with all parameters spelled out from Balaban's inductive scheme is deferred.

A.1 Step 1: Single-scale polymer derivative bound

At each scale k, Proposition 2.3 provides:

$$\operatorname*{ess\,sup}_{\mathcal{G}_k} \mathbb{E}\big[|\nabla_{E_k} V_{<k}|^2 \mid \mathcal{G}_k\big] \leq C_{\text{poly}}^2 e^{-2\kappa} \cdot L_{\text{RG}}^{-(d-1)k}. \tag{27}$$

This bound is established in §2.4 using:

- The absorption identity ([6, eq. (1.37), p. 367]): after RG absorption, the dependence of $V_{<k}$ on E_k-variables enters only through the polymer terms $\mathbf{R}^{(k)}(X)$ and $\mathbf{B}^{(k)}(X)$.

- Analyticity of polymer activities on the complex domain $\widetilde{U}_k^c(Y, \tilde{\alpha}_0, \hat{\alpha}_1)$ ([6, eq. (1.65), p. 383]), with uniform radius $\hat{\alpha}_1(\gamma)$.

- Cauchy derivative bounds: $|\nabla \mathbf{R}^{(k)}(X)| \leq (C/\hat{\alpha}_1) e^{-\kappa d_k(X)}$.

- Block-averaging adjoint scaling: $\|Q_{(k)}^*\|^2 = L_{\text{RG}}^{-(d-1)k}$ ([1, eq. (125), Proposition 4]).

A.2 Step 2: Propagation across intermediate scales

For $\ell < k$, the potential $V_{<\ell}$ depends on scale-k variables through the chain of RG transformations at scales $\ell, \ell+1, \ldots, k-1$. Balaban's effective action at scale ℓ ([6, eq. (1.101)]) depends on the background field U_ℓ, which in turn depends on all finer-scale variables through the minimisation and averaging operations.

The key structural property is that the dependence of U_ℓ on scale-k variables ($k > \ell$) enters through the k-th order block-averaging $Q_{(k)}$. By Balaban's Proposition 4 ([1, eq. (133), p. 42]), the linearised k-fold averaging $Q_k(U_0)$ satisfies:

$$|Q_k(U_0, \eta A) - Q_k(U_0)A| < C_2 |A|^2,$$

with constants *independent of k*. The linear part $Q_k(U_0)$ has adjoint norm $\|Q_k(U_0)^*\|^2 \leq (1 + 2C_1'\alpha_0)^2 \cdot L_{\text{RG}}^{-(d-1)k}$ by [1, eq. (146)].

Therefore, a perturbation in the E_k-direction propagates to scale $\ell < k$ with an amplitude bounded by $\|Q_{(k)}^*\| = L_{\text{RG}}^{-(d-1)k/2}$, regardless of the number of intermediate scales $k - \ell$. This is because the block-averaging adjoint at scale k already accounts for the full compression from the fine to the coarse lattice.

A.3 Step 3: Composition and the cross-scale bound

Combining Steps 1 and 2: the gradient $\nabla_{E_k} V_{<\ell}$ receives contributions from the polymer activities $\mathbf{R}^{(m)}(X)$ at all intermediate scales $\ell \leq m \leq k$. Each contribution satisfies:

$$|\nabla_{E_k} \mathbf{R}^{(m)}(X)| \leq \frac{C}{\hat{\alpha}_1} e^{-\kappa d_m(X)} \cdot \|Q_{(k)}^*\| \leq \frac{C}{\hat{\alpha}_1} e^{-\kappa d_m(X)} \cdot L_{\mathrm{RG}}^{-(d-1)k/2}.$$

The factor $L_{\mathrm{RG}}^{-(d-1)k/2}$ is the same for all intermediate scales m, because it arises from the projection onto E_k (via the adjoint of the full k-step block-averaging), not from the individual steps.

Summing over polymers X at each intermediate scale m (using Lemma 5.6) and over scales m:

$$
\begin{aligned}
\mathbb{E}\big[|\nabla_{E_k} V_{<\ell}|^2 \mid \mathcal{G}_\ell\big] &\leq \left(\sum_{m=\ell}^{k} \sum_{X \ni \text{link at scale } m} \frac{C}{\hat{\alpha}_1} e^{-\kappa d_m(X)} \cdot L_{\mathrm{RG}}^{-(d-1)k/2} \right)^2 \\
&\leq \left(\sum_{m=\ell}^{k} \frac{C' e^{-\kappa}}{\hat{\alpha}_1} \cdot L_{\mathrm{RG}}^{-(d-1)k/2} \right)^2 \\
&\leq \left(\frac{(k-\ell+1) C' e^{-\kappa}}{\hat{\alpha}_1} \right)^2 \cdot L_{\mathrm{RG}}^{-(d-1)k}.
\end{aligned}
$$

The factor $(k - \ell + 1) \leq n_{\max}$ is polynomial in $\log(1/g^2)$ (since $n_{\max} \sim 1/(2b_0 g^2 \ln 2)$). This factor is *not* absorbed by $e^{-\kappa}$ (which is a fixed constant from Balaban's construction), but is instead incorporated into the constant $C_{\mathrm{poly}} = C_{\mathrm{poly}}(N_c, \beta_0) = n_{\max}(\beta_0) \cdot C'/\hat{\alpha}_1(\gamma)$. Since β_0 is fixed in Theorem 1.1, this dependence does not affect the uniformity of α_* in L_{vol}. It does affect the quantitative value of α_* (and hence of Δ_{phys}), which decreases as β_0 increases—consistent with the expected physical scaling of the mass gap.

Remark A.2. The factor $n_{\max}$ in C_{poly} means that the polymer derivative constant grows with β (since $n_{\max} \sim 1/(2b_0 g^2 \ln 2)$). This growth is polynomial in $\log(1/g^2)$ and does not affect the positivity of α_*, only its quantitative value. The mass gap bound $\Delta_{\mathrm{phys}} \geq c(N_c, \beta_0) > 0$ remains uniform in L_{vol}.