Exact-Kernel Fourier Families Obstruct Every Penalty Factored Through a Block Map in a Flat Lattice Gauge Form

Lluis Eriksson

August 4, 2026 (version 1.4)

Abstract

Let a fine periodic lattice have side LN' and let Q_L be the L^{-d}-normalized block average of length-L line integrals. In four dimensions, the rescaling $Q_L \mapsto LQ_L$ repairs the elementary constant-field scaling obstruction to a Poincaré estimate. More generally, for a power rescaling $L^\sigma Q_L$ in dimension d, that sector only forces $\sigma \geq (d-2)/2$; equality is the unique scale-neutral choice, not a sufficient coercivity theorem. We prove a stronger factor-through-kernel no-go: no scalar functional $\Phi_L(Q_L A)$ with $\Phi_L(0) = 0$ can yield full-space coercivity with a constant uniform in the block side. The functional may depend on L and may be nonlinear, discontinuous, anisotropic, or unbounded; no positivity or growth assumption is used. For every $L \geq 2$, fixed $N' \geq 1$, and $N_c \geq 2$, we embed within-block Fourier phases in a real two-plane of the internal coordinate space and construct transverse one-cochains. The first mode A_L satisfies

$$Q_L A_L = 0, \qquad \operatorname{div} A_L = 0,$$

and, for every dimension $d \geq 2$,

$$\|A_L\|^2 = (LN')^d, \qquad \langle A_L, K_0 A_L \rangle = (LN')^d \lambda_L, \qquad \lambda_L = 4\sin^2\left(\frac{\pi}{L}\right) \leq \frac{4\pi^2}{L^2}.$$

Thus every factorized penalty vanishes on the witness. In the linear case, the Rayleigh quotient of $K_0 + (T_L Q_L)^*(T_L Q_L)$ is exactly λ_L for every scale-dependent continuous linear coarse post-conditioner T_L, with no bound on $\|T_L\|$. Higher frequencies give an orthogonal $2R$-dimensional real subspace on which the maximal Rayleigh quotient is $4\sin^2(\pi R/L)$. Hence the ordered spectrum of the linear operator has at least $2R$ eigenvalues below $4\pi^2 R^2/L^2$; choosing $R = \lfloor L^\alpha \rfloor$, $0 < \alpha < 1$, produces a growing low-energy cluster rather than a single exceptional line. Every admissible full-space Poincaré constant obeys $C_{\mathrm{P}}(L) \geq 1/\lambda_L \geq L^2/(4\pi^2)$, so no constant chosen before L can work. We also give an exact repair budget. If the block map is perturbed to $\widetilde{Q}_L$, and a positive term V_L is added, the same witness has Rayleigh quotient

$$\lambda_L + \frac{\left\|T_L \widetilde{Q}_L A_L\right\|^2}{\|A_L\|^2} + \frac{\langle A_L, V_L A_L \rangle}{\|A_L\|^2}.$$

Uniform coercivity therefore requires a scale-independent positive lift of this channel; the condition is also sufficient on the witness line, but not on the full space. Here "volume-uniform" means uniform while L varies at an arbitrary fixed positive coarse side N'; it is not a claim about $N' \to \infty$ at fixed L.

The construction, kernel identities, exact Hodge energy, factorized-penalty no-go, repair Rayleigh identity, and necessary repair budget are formalized in Lean 4. The focused axiom audit reports only `propext`, `Classical.choice`, and `Quot.sound`. The higher-frequency spectral

multiplicity theorem is proved analytically here and is not claimed as part of the formalized surface. The theorem concerns the stated flat full-domain form; it does not contradict gauge-restricted propagator constructions and makes no infinite-volume, continuum, or Yang–Mills mass-gap claim.

Keywords. Lattice gauge theory; block averaging; Poincaré inequality; Fourier mode; postconditioning; formal verification; Lean 4.

1 Introduction

Finite-range positive operators admit Combes–Thomas localization estimates whose constants depend on the lower spectral edge [5]. In a block-spin treatment of lattice gauge fields, a natural finite-dimensional precision is a flat Hodge operator plus a positive block penalty. The question addressed here is sharply limited: can any penalty factored through the original block variable, in the form

$$\langle A, K_0 A\rangle + \Phi_L(Q_L A), \qquad \Phi_L(0) = 0, \tag{1}$$

control the full fine one-cochain norm with a constant independent of the block side L? The familiar linear postconditioned form $\Phi_L(B) = \|T_L B\|^2$ is a special case.

The factor L is not arbitrary. With the line-integral block map used below, it exactly balances direction-wise constant fields in dimension four. The corresponding general-dimensional arithmetic is often compressed into the phrase "critical exponent $(d-2)/2$." We separate two statements that this phrase can obscure: exponents at or above that threshold avoid growth forced by constant fields, while equality is selected only by exact scale neutrality. Neither statement is sufficient for full-space coercivity. The block map also has mean-zero transverse modes in its exact kernel. Because a map applied *after* Q_L cannot distinguish zero from zero, no factorized functional that vanishes at the coarse zero field can lift them. This includes every scalar, anisotropic, mixing, or unbounded scale-dependent linear postconditioner. The within-block Fourier family makes the obstruction local to every block and supplies a growing cluster of low modes, not merely one adversarial vector.

The main result has eight components.

(i) The constant-sector calculation isolates exactly what the exponent $(d-2)/2$ does and does not imply.

(ii) The witness is defined on the exact finite periodic lattice and works for every $L \geq 2$, with no parity restriction.

(iii) Its cancellation under Q_L and under the flat divergence is exact in every block and at every positive coarse side N'.

(iv) Its Hodge Rayleigh quotient is the first-cycle eigenvalue $4\sin^2(\pi/L)$, hence is $O(L^{-2})$.

(v) Frequencies $1 \leq r \leq R$ produce an orthogonal real subspace of dimension $2R$ and an explicit min–max bound on the $2R$-th eigenvalue.

(vi) The no-go is uniform over arbitrary factorized scalar functionals Φ_L satisfying only $\Phi_L(0) = 0$; linear postconditioners are a strict subclass.

(vii) The quantifier order is explicit and machine checked: the family T_L or Φ_L, and one proposed constant C_{P}, are chosen before all block sides L.

(viii) An exact repair identity quantifies how much a changed block map or an added positive Hessian must contribute on the missed channel.

An earlier square-wave witness gives the weaker even-scale quotient $8/L$. It already disproves the same gate, but the Fourier construction identifies the actual low-frequency mechanism and strengthens the quantitative lower bound on Poincaré constants from linear to quadratic growth.

The conclusion is a domain diagnosis, not a general no-go for lattice gauge renormalization. Classical constructions use gauge restrictions, quotient structures, or background-dependent operators. The present theorem instead tests the unrestricted flat form in (1). A repaired theorem may still hold on a physically justified sector, after changing the information captured before postconditioning, or after adding a positive interacting Hessian that lifts the explicit transverse family.

2 Finite lattice, block map, and quantifiers

Fix integers $d \geq 2$, $L \geq 1$, $N' \geq 1$, and $N_c \geq 2$. The fine site set is

$$\Lambda_{L,N'} = \left(\mathbb{Z}/(LN')\mathbb{Z}\right)^d,$$

and a positive bond is (x,μ) with $x \in \Lambda_{L,N'}$ and $\mu \in \{0,\ldots,d-1\}$. The real internal coordinate space is $E_{N_c} = \mathbb{R}^{N_c^2-1}$. For a one-cochain $A(x,\mu) \in E_{N_c}$, set

$$\|A\|^2 = \sum_{x\in\Lambda_{L,N'}} \sum_{\mu=0}^{d-1} \|A(x,\mu)\|^2 .$$

At the trivial background, let D_1 be the discrete exterior derivative and δ the discrete divergence. The flat Hodge operator K_0 satisfies

$$\langle A, K_0 A\rangle = \|D_1 A\|^2 + \|\delta A\|^2 . \tag{2}$$

In the formal development this is derived from the adjoint-defined operator; it is not introduced as an axiom.

For a coarse site $y \in (\mathbb{Z}/N'\mathbb{Z})^d$, write

$$B_y = \{Ly + r : 0 \leq r_\nu < L \text{ for every } \nu\}.$$

For a coarse positive bond (y,μ), define the linearized line-integral average

$$(Q_L A)(y,\mu) = L^{-d} \sum_{x\in B_y} \sum_{k=0}^{L-1} A(x + ke_\mu, \mu). \tag{3}$$

For a continuous linear endomorphism T_L of the coarse one-cochain space, the tested precision is

$$H_{L,T} = K_0 + (T_L Q_L)^*(T_L Q_L). \tag{4}$$

This class includes scalar rescaling $T_L = s_L I$, direction-dependent weights, and mixing of coarse sites, directions, or internal coordinates. No uniform operator-norm bound on T_L is imposed.

Proposition 2.1 (what the constant-sector wall fixes)**.** *Let $A_v(x,\mu) = v_\mu$ be a direction-wise constant field and put $S_v = \sum_\mu \|v_\mu\|^2 > 0$. Then*

$$\|A_v\|^2 = L^d (N')^d S_v, \qquad \|L^\sigma Q_L A_v\|^2 = L^{2\sigma+2}(N')^d S_v. \tag{5}$$

Consequently, a Poincaré estimate with a constant independent of unbounded L can pass this sector only if

$$\sigma \geq \frac{d-2}{2}. \tag{6}$$

Equality is the unique exponent for which the two norms in (5) have identical scale dependence; larger exponents overweight the constant sector and also pass this necessary test.

Proof. There are $(LN')^d$ fine sites. For a constant field, the inner line sum in (3) supplies L, while the L^{-d} normalization cancels the L^d source sites in a block, so $Q_L A_v = Lv$ on each coarse bond. This gives (5). Substitution in a Poincaré estimate and cancellation of the positive factor $(N')^d S_v$ yields $L^{d-2-2\sigma} \leq C_{\mathrm{P}}$. Uniformity for arbitrarily large L forces (6); equality is exactly the zero exponent in this last power. □

Remark 2.2 (necessary is not sufficient)**.** *The threshold in Proposition 2.1 is sufficient only for preventing the* constant-sector *lower bound from diverging. Calling $\sigma = (d-2)/2$ necessary and sufficient for the full gate would silently add scale neutrality as a hypothesis and would ignore the kernel exhibited below.*

Definition 2.3 (critical full-space gate)**.** In $d = 4$, fix $N' \geq 1$, $N_c \geq 2$, and the adjoint-model parameter defining the flat maps. The gate asserts that there is a number $C_{\mathrm{P}} > 0$ such that, for every integer $L \geq 1$ and every fine one-cochain A,

$$\|A\|^2 \leq C_{\mathrm{P}} \left(\langle A, K_0 A\rangle + \|LQ_L A\|^2 \right). \tag{7}$$

Definition 2.4 (postconditioned full-space gate)**.** Fix the same parameters and an arbitrary family $(T_L)_{L\geq 1}$ of continuous linear coarse-space endomorphisms. The postconditioned gate asserts that there is one $C_{\mathrm{P}} > 0$, chosen after the family but before the scale, such that for every $L \geq 1$ and every fine one-cochain A,

$$\|A\|^2 \leq C_{\mathrm{P}} \left(\langle A, K_0 A\rangle + \|T_L Q_L A\|^2 \right). \tag{8}$$

The critical scalar gate is the special case $T_L = LI$ in $d = 4$.

Definition 2.5 (factorized full-space gate)**.** Fix the same parameters and an arbitrary family of scalar functionals

$$\Phi_L : \{\text{coarse one-cochains at side } N'\} \longrightarrow \mathbb{R}, \qquad \Phi_L(0) = 0.$$

The factorized gate asserts that there is one $C_{\mathrm{P}} > 0$, chosen after the family but before the scale, such that for every $L \geq 1$ and every fine one-cochain A,

$$\|A\|^2 \leq C_{\mathrm{P}} \left(\langle A, K_0 A\rangle + \Phi_L(Q_L A) \right). \tag{9}$$

No linearity, continuity, positivity, or growth hypothesis on Φ_L is needed for the no-go below. If "penalty" is reserved for nonnegative functionals, that narrower class is of course included. Definition 2.4 is recovered by $\Phi_L(B) = \|T_L B\|^2$.

Remark 2.6 (meaning of volume-uniform)**.** *The coarse side N' is arbitrary but fixed inside one instance of the gate. The family (T_L) or (Φ_L) and C_{P} may depend on N', N_c, and the fixed adjoint model, but not on the subsequently quantified L. The fine volume $(LN')^4$ then varies and the same C_{P} must work for all of it. This paper does not call a bound "volume-uniform" merely because it is uniform as $N' \to \infty$ at fixed L; that is a different quantifier statement.*

3 The block-periodic Fourier witness

Let

$$\zeta_L = \exp(2\pi i/L), \qquad L \geq 2.$$

Then $\zeta_L^L = 1$, $\zeta_L \neq 1$, and the finite geometric sum gives

$$\sum_{r=0}^{L-1} \zeta_L^r = 0. \tag{10}$$

Choose orthonormal vectors $u_0, u_1 \in E_{N_c}$, possible because $N_c^2 - 1 \geq 3$, and define the real-linear isometry

$$\iota : \mathbb{C} \longrightarrow E_{N_c}, \qquad \iota(a + ib) = au_0 + bu_1.$$

For $m \in \mathbb{Z}/(LN')\mathbb{Z}$, define the block-periodic internal profile

$$\phi_L(m) = \iota\Big(\zeta_L^{m \bmod L}\Big). \tag{11}$$

It has unit norm and exact zero mean in every length-L block.

Choose distinct directions $i \neq j$ and set

$$A_L(x, \mu) = \begin{cases} \phi_L(x_j), & \mu = i, \\ 0, & \mu \neq i. \end{cases} \tag{12}$$

The field points in direction i and varies only in the transverse direction j.

Lemma 3.1 (norm and block kernel)**.** *For every $L \geq 2$ and $N' \geq 1$,*

$$\|A_L\|^2 = (LN')^d, \qquad Q_L A_L = 0. \tag{13}$$

Consequently $(T_L Q_L)A_L = 0$ for every linear postconditioner T_L. More generally, $\Phi_L(Q_L A_L) = 0$ for every scalar functional with $\Phi_L(0) = 0$.

Proof. Only direction-i bonds contribute to the norm, one at each fine site, and $\|\phi_L(m)\| = 1$. For the block map, a coarse bond outside direction i sees only zero components. In direction i, line shifts preserve x_j. After the line sum is factored out, the sum over the within-block j-offset is (10). This cancellation occurs separately in every source block, hence for every N'. □

Lemma 3.2 (transverse divergence)**.** *At the trivial background,*

$$\delta A_L = 0.$$

Proof. Only the direction-i component enters. That component depends only on x_j and is invariant under forward and backward shifts in direction i. □

The witness therefore lies simultaneously in the exact kernel of the block map and the flat gauge-constraint kernel. Its remaining energy is purely curvature energy.

4 Exact first-mode energy

Define

$$\lambda_L = |\zeta_L - 1|^2. \tag{14}$$

For every r, multiplication by the unit complex number ζ_L^r gives

$$\left\|\iota(\zeta_L^r) - \iota(\zeta_L^{r+1})\right\|^2 = |1 - \zeta_L|^2 = \lambda_L, \tag{15}$$

including the periodic jump $r = L - 1$ because $\zeta_L^L = 1$. Elementary trigonometry yields

$$\lambda_L = 2 - 2\cos(2\pi/L) = 4\sin^2(\pi/L) \le \frac{4\pi^2}{L^2}. \tag{16}$$

Proposition 4.1 (exact flat Hodge energy)**.** *For the transverse witness* (12),

$$\langle A_L, K_0 A_L\rangle = (LN')^d \lambda_L. \tag{17}$$

Proof. Lemma 3.2 removes the divergence term in (2). In the plaquette curl, every direction pair other than the pair containing i and j vanishes: either the field component is zero or its profile is unchanged by the shift. Each (i, j) plaquette contributes the common value λ_L from (15). There is one contributing geometric plaquette at each of the $(LN')^d$ fine sites under the convention used by D_1. Summation gives (17). □

Theorem 4.2 (postconditioner-invariant Rayleigh quotient)**.** *For every continuous linear coarse-space endomorphism T_L,*

$$\frac{\langle A_L, H_{L,T} A_L\rangle}{\|A_L\|^2} = \lambda_L = 4\sin^2(\pi/L). \tag{18}$$

Thus the best full-space coercivity constant c_L of $H_{L,T}$ obeys

$$c_L \le \lambda_L \le \frac{4\pi^2}{L^2}. \tag{19}$$

Proof. Linearity gives $T_L 0 = 0$, so Lemma 3.1 removes the block term from (4), irrespective of $\|T_L\|$. Divide Proposition 4.1 by the positive norm in Lemma 3.1; the entire volume factor cancels. □

5 A spectral ladder, not an isolated witness

The same cancellation is available at every nonzero within-block frequency. For $1 \le r < L$ and $c \in \mathbb{C}$, define the real one-cochain

$$A_{r,c}(x, \mu) = \begin{cases} \iota(c\exp(2\pi i r x_j/L)), & \mu = i, \\ 0, & \mu \ne i, \end{cases} \tag{20}$$

and let $E_r = \{A_{r,c} : c \in \mathbb{C}\}$. Each E_r is a real two-dimensional subspace. The root-of-unity sum, transverse-divergence calculation, and plaquette calculation used above give

$$Q_L A_{r,c} = 0, \qquad \delta A_{r,c} = 0, \qquad \langle A_{r,c}, K_0 A_{r,c}\rangle = \lambda_{r,L}\|A_{r,c}\|^2, \quad \lambda_{r,L} = 4\sin^2(\pi r/L). \tag{21}$$

Theorem 5.1 (growing low-energy multiplicity)**.** *Let $L \geq 3$ and $1 \leq R \leq \lfloor (L-1)/2 \rfloor$. For every continuous linear postconditioner T_L, the spaces $E_1, \dots, E_R$ are mutually orthogonal, and*

$$S_R := E_1 \oplus \cdots \oplus E_R \subseteq \ker Q_L, \qquad \dim_{\mathbb{R}} S_R = 2R.$$

Every nonzero $A \in S_R$ satisfies

$$\frac{\langle A, H_{L,T} A \rangle}{\|A\|^2} \leq \lambda_{R,L} = 4 \sin^2(\pi R/L) \leq \frac{4\pi^2 R^2}{L^2}. \tag{22}$$

Consequently, if $\mu_1(H_{L,T}) \leq \mu_2(H_{L,T}) \leq \cdots$ are the eigenvalues counted with multiplicity, then

$$\mu_{2R}(H_{L,T}) \leq \frac{4\pi^2 R^2}{L^2}. \tag{23}$$

In particular, for every $0 < \alpha < 1$ and all sufficiently large L, the choice $R = \lfloor L^\alpha \rfloor$ gives at least $2\lfloor L^\alpha \rfloor$ eigenvalues bounded by $4\pi^2 L^{-2(1-\alpha)}$.

Proof. The internal embedding satisfies $\langle \iota(z), \iota(w) \rangle = \mathrm{Re}(z\overline{w})$. Summing the inner product of an r-profile and an s-profile over x_j therefore reduces, for $r \neq s$, to the real part of a nontrivial root-of-unity sum and vanishes. This proves orthogonality and the dimension count. Equations (21) show that the block term vanishes and that the quadratic form is diagonal on the orthogonal sum. Since $r \mapsto \sin(\pi r/L)$ is increasing on the stated range, the Rayleigh quotient of any vector in S_R is at most $\lambda_{R,L}$. The last inequality in (22) is $\sin t \leq t$. The finite dimensional min–max principle applied to the $2R$-dimensional test subspace gives (23). Finally, $\lfloor L^\alpha \rfloor \leq (L-1)/2$ for all sufficiently large L, and $R^2/L^2 \leq L^{-2(1-\alpha)}$. □

Theorem 5.1 concerns the linear self-adjoint operator $H_{L,T}$. For a nonlinear factorized functional Φ_L, an eigenvalue statement need not make sense; nevertheless, every vector in every E_r is still erased by Q_L, so the domination inequality fails already on E_1. The multiplicity theorem is analytic and is not part of the Lean surface claimed below.

6 No uniform critical coercivity

Theorem 6.1 (factor-through-kernel all-scales no-go)**.** *Fix $N' \geq 1$, $N_c \geq 2$, the adjoint-model parameter defining the flat operators, and any family $(\Phi_L)_{L \geq 1}$ of scalar functionals on the coarse field space with $\Phi_L(0) = 0$. The factorized full-space gate in Definition 2.5 is false. More quantitatively, at every $L \geq 2$ any constant valid at that scale must satisfy*

$$C_{\mathrm{P}}(L) \geq \frac{1}{\lambda_L} \geq \frac{L^2}{4\pi^2}. \tag{24}$$

Proof. Insert A_L into (9). By Lemma 3.1 and Proposition 4.1,

$$(LN')^4 \leq C_{\mathrm{P}} (LN')^4 \lambda_L.$$

The volume factor is positive, so $1 \leq C_{\mathrm{P}} \lambda_L$, proving the first inequality in (24); the second follows from (16). If one C_{P} worked for all L, it would force $L^2 \leq 4\pi^2 C_{\mathrm{P}}$ for every L, a contradiction. □

Corollary 6.2 (arbitrary linear postconditioning)**.** *The postconditioned gate in Definition 2.4 is false for every family of continuous linear maps T_L, even if their norms diverge with L.*

Proof. Apply Theorem 6.1 to $\Phi_L(B) = \|T_L B\|^2$. $\square$

Corollary 6.3 (critical and supercritical scalar normalizations)**.** *No real scalar sequence s_L can restore full-space uniform coercivity for $K_0 + (s_L Q_L)^*(s_L Q_L)$. In particular, neither the critical choice $s_L = L^{(d-2)/2}$ nor any faster scalar growth repairs the full-space gate.*

Proof. Take $T_L = s_L I$ in Corollary 6.2. $\square$

The point is stronger than "the critical exponent is insufficient." Even an arbitrary rule depending only on the retained coarse variable acts after the witness has already been erased. The proof uses only $\Phi_L(0) = 0$; linearity was never the essential hypothesis.

The theorem isolates the lower spectral edge as the obstruction. Separate finite-range, kernel-amplitude, ball-count, and block-metric localization estimates are not invalidated. What fails is the positive scale-independent coercivity input required to turn those estimates into a uniform inverse or inverse-square-root bound on the full domain.

7 Exact repair budget on the missed channel

The kernel theorem says precisely what postconditioning cannot do. It also allows a quantitative statement about changes made *before* that kernel or outside the block penalty. Let $\widetilde{Q}_L$ be any linear replacement for Q_L, let T_L be any continuous linear postconditioner, and let V_L be a positive semidefinite self-adjoint operator on fine one-cochains. Set

$$\widetilde{H}_L = K_0 + (T_L \widetilde{Q}_L)^*(T_L \widetilde{Q}_L) + V_L \tag{25}$$

and define the two dimensionless lifts

$$\varepsilon_L = \frac{\left\|T_L \widetilde{Q}_L A_L\right\|}{\|A_L\|}, \qquad \nu_L = \frac{\langle A_L, V_L A_L\rangle}{\|A_L\|^2} \geq 0. \tag{26}$$

Theorem 7.1 (sharp witness-channel repair criterion)**.** *For every $L \geq 2$,*

$$\frac{\left\langle A_L, \widetilde{H}_L A_L \right\rangle}{\|A_L\|^2} = \lambda_L + \varepsilon_L^2 + \nu_L. \tag{27}$$

Consequently, if one constant $C_{\mathrm{P}} > 0$ gives full-space coercivity for $\widetilde{H}_L$ at every scale, then

$$\varepsilon_L^2 + \nu_L \geq C_{\mathrm{P}}^{-1} - \lambda_L \quad (L \geq 2), \qquad \liminf_{L\to\infty}(\varepsilon_L^2 + \nu_L) \geq C_{\mathrm{P}}^{-1}. \tag{28}$$

Conversely, uniform coercivity restricted to the one-dimensional spaces $\operatorname{span}\{A_L\}$ *holds if and only if*

$$\inf_{L\geq 2}(\lambda_L + \varepsilon_L^2 + \nu_L) > 0. \tag{29}$$

This equivalence is only a witness-sector statement; it is not sufficient for coercivity on the full one-cochain space.

Proof. Expand the three quadratic terms in (25), use Proposition 4.1, and divide by $\|A_L\|^2$ to obtain (27). Testing a full-space Poincaré estimate on A_L gives $1 \leq C_{\mathrm{P}}(\lambda_L + \varepsilon_L^2 + \nu_L)$, hence (28) because $\lambda_L \to 0$. On the line spanned by A_L, homogeneity makes the displayed Rayleigh quotient the unique quotient, so a uniform lower bound is equivalent to (29). $\square$

The exact identity (27) and the finite-scale necessary budget obtained by testing the full-space estimate are machine checked in Lean. The asymptotic liminf statement and the witness-line equivalence (29) are elementary analytic consequences; no full-space sufficiency claim is made or smuggled into the hypotheses.

The formula distinguishes three logically different repairs. A postconditioner after the original Q_L has $\varepsilon_L = 0$ exactly and is therefore powerless. A new pre-kernel measurement may give $\varepsilon_L > 0$. An additional Hessian contributes ν_L. Any proposed interacting continuation must prove, rather than assume, that their combined contribution stays uniformly positive on this explicit family.

There is a useful scaling corollary. Write $\widetilde{Q}_L = Q_L + R_L$, take the constant-sector critical scalar $T_L = L^{(d-2)/2} I$, and put

$$\eta_L = \frac{\|R_L A_L\|}{\|A_L\|}. \tag{30}$$

Because $Q_L A_L = 0$, the block contribution is exactly $\varepsilon_L^2 = L^{d-2}\eta_L^2$. In four dimensions, with $V_L = 0$, a uniform repair of even this one channel requires

$$\liminf_{L\to\infty} L^2 \eta_L^2 > 0; \tag{31}$$

leakage $o(L^{-1})$ cannot work. This exposes the scale at which a changed block observable must see the missing Fourier mode. It does not manufacture such an observable or establish its gauge covariance.

8 Relation to prior gauge-RG constructions

Bałaban's propagator papers study the Gaussian vector-field theory underlying lattice gauge renormalization and then extend the analysis to operators with restrictions at several scales and to iterated renormalization transformations [1, 2]. His separate treatment of averaging operations studies their regularity on selected classes of gauge fields and gauge transformations [3]. In that program, block averages are not used in isolation: gauge conditions, restricted domains, minimizers, and background-dependent propagators are part of the construction. Modern expositions of covariant axial gauge likewise emphasize an axial restriction adapted to block averaging [4].

There is therefore no conflict between classical positivity statements on a gauge-fixed or otherwise restricted subspace and Theorem 6.1. The claims have different domains. The exact delta established here is:

> For the explicit finite periodic line-integral map (3), the flat full-space form has block-periodic transverse Fourier families in the simultaneous kernels of Q_L and the flat divergence. Therefore every term $\Phi_L(Q_L A)$ with $\Phi_L(0) = 0$ misses them exactly. In the linear case, $K_0 + (T_L Q_L)^*(T_L Q_L)$ has at least $2R$ eigenvalues no larger than $4\pi^2 R^2/L^2$ for $1 \le R \le \lfloor (L-1)/2 \rfloor$, at every $N' \ge 1$.

This is a no-go for that unrestricted coercivity proposal, not a priority claim over the general observation that gauge fields require gauge fixing, and not a no-go for Bałaban's restricted propagators. Indeed, the result explains why a domain restriction or quotient is mathematically load-bearing rather than cosmetic.

9 Machine-checked realization

The Lean 4 development is an eight-module theorem chain, plus three focused oracle files, headed by

`YangMills/RG/PhysicalCriticalRescalingFourierFactorizedRepair.lean`.

Table 1 maps the analytic statements to selected declarations.

The constant-sector norm identity and lower bound used in Proposition 2.1 are formalized in `PhysicalPoincareCriticalRescaling.lean`; the specialization to a general real exponent σ is the elementary power comparison displayed in the proof. The Fourier chain then tests the sector that the constant-field calculation cannot see.

Mathematical content	Lean declaration
Root-of-unity block cancellation	`sum_blockFourierRoot_pow_eq_zero`
Lie-valued cancellation in every block	`sum_blockFourierProfile_block_eq_zero`
Eigenvalue bound $\lambda_L \le (2\pi/L)^2$	`blockFourierEigenvalue_le`
Exact block-map kernel	`flatBlockConstraintQCLM_blockFourierMode_eq_zero`
Zero transverse divergence	`gaugeConstraintQCLM_blockFourierModeCochain_eq_zero_of_ne`
Exact Hodge quadratic form	`flatGaugeHodgeK0_inner_blockFourierModeCochain_of_ne`
Exact Rayleigh quotient	`blockFourierMode_rayleigh_eq`
Necessary bound $1 \le C_{\mathrm{P}}\lambda_L$	`criticalRescaledFlatPoincare_fourier_lower_bound`
All-scales no-go	`volumeUniformCriticalRescaledFlatPoincareGate_fourier_false`
Arbitrary-postconditioner lower bound	`postconditionedFlatPoincare_fourier_lower_bound`
Postconditioned all-scales no-go	`volumeUniformPostconditionedFlatPoincareGate_fourier_false`
Arbitrary factorized-penalty lower bound	`factorizedFlatPoincare_fourier_lower_bound`
Factor-through-kernel all-scales no-go	`volumeUniformFactorizedFlatPoincareGate_fourier_false`
Exact repaired witness quotient	`repairedFlatPoincare_blockFourier_rayleigh_eq`
Necessary repaired full-space budget	`repairedFlatPoincare_blockFourier_budget`

Table 1: Theorem-to-artifact map.

The formal result uses the exact norm definition $\lambda_L = \|\zeta_L - 1\|^2$ and proves the bound $\lambda_L \le (2\pi/L)^2$. Equation (16) is the elementary closed-form identification of that norm. The original focused oracle

`YangMills/RG/PhysicalCriticalRescalingFourierNoGoOracle.lean`

reports exactly

[`propext`, `Classical.choice`, `Quot.sound`]

for all nine audited declarations, with no project axiom and no `sorryAx`. A second focused oracle,

`YangMills/RG/PhysicalCriticalRescalingFourierPostconditionedNoGoOracle.lean`,

reports the same list for the block-kernel theorem, the exact Hodge energy, the eigenvalue estimate, and both postconditioned theorems. The third oracle,

`YangMills/RG/PhysicalCriticalRescalingFourierFactorizedRepairOracle.lean`,

reports the same list for six declarations: the reused exact kernel and Hodge energy, the factorized lower bound and all-scales no-go, and the repaired Rayleigh identity and necessary budget.

Immutable source and reproduction

The original seven-file theorem source is fixed at repository commit

`f21539ed0bb880a04078de369bf5cbf063f7b101`.

Its former head theorem file has LF-normalized SHA-256

`5F0890D14EA6981CBE6459605A634386ED2C844F6D03A6947E33B34250E2F5AE`.

The v1.4 factorized-and-repair extension has 12,064 LF bytes and SHA-256

`D7586D24D7B7E9050001AC5A8E4A8061A3EE3F7D441C5373186EBAEB7B97DE38`,

and its 537-byte focused oracle has SHA-256

`6334EE25A6A0D8778A91CC49269BC92563C327DD1BD0E5CA9CF477F5E90738E1`.

The companion packet `critical_rescaling_fourier_no_go_sources_v1.4.zip` has SHA-256

`EF5501145398B5F672DCDFB66206C1F5F280F2C61B231D2C9AC0E1A4E693F82A`.

It contains the eight theorem modules, all three focused oracles, the v1.4 manifest and audit record, the Lean toolchain, and the pinned Lake files. The toolchain is Lean `v4.29.0-rc6`. The pinned Mathlib revision is

`07642720480157414db592fa85b626dafb71355b`.

From the repository root, run

```
lake build
YangMills.RG.PhysicalCriticalRescalingFourierFactorizedRepair
lake env lean
YangMills/RG/PhysicalCriticalRescalingFourierFactorizedRepairOracle.lean
```

The v1.4 reproduction on August 4, 2026 used a Colab Pro+ CPU/high-RAM runtime (50.99 GB visible RAM) with no GPU, under the visibly confirmed account `lluiseriksson@gmail.com`. A disconnected preparation reconstructed the public base `04f87347f3e4d46a05e77bc1c70855794e111477` plus 17 LF-normalized frozen bodies, checking every body hash before execution. Lake update took 124.412 seconds and cache retrieval 6.507 seconds. The first target pass materialized the 8,185 imported jobs and exposed a target-local arithmetic proof defect; the judge and criterion were not changed. After two source-only proof repairs, the final 12,064-byte body was frozen at the hash above. The final target invocation succeeded with 8,186 jobs in 18.835 seconds and materialized the target `.olean`. The focused oracle then exited successfully in 14.953 seconds and returned exactly [`propext`, `Classical.choice`, `Quot.sound`] for all six declarations. The final frozen run spanned 19:07:36.898702–19:08:10.693578 UTC; the full reconstruction audit spanned 18:56:52.878156–19:08:10.693578 UTC. The runtime was visibly open no later than 18:48:10.965571 UTC and was disconnected and deleted immediately after the audit. The only two warnings were pre-existing style lints in imported files. A local code-token scan found no `sorry`, `admit`, or project `axiom` in the new source. The notebook is archived at Colab audit notebook 1iXBNBeBEGr7O3pEbuRtac1VE1mMHX0Lh. This is reproduction evidence and a frozen object for external audit, not an external certification.

10 What remains open

Restricted or quotient sector. A repaired domain may remove or quotient the obstructive transverse family, but the restriction must be physically justified and stable under the renormalization operations. Coercivity on that new domain remains to be proved.

Additional positive Hessian. An interacting Wilson Hessian or another positive term V_L could lift the family. Theorem 7.1 gives its exact required contribution on this channel, but establishing a lower bound for the physical interacting Hessian remains open. Such a bound is not a consequence of the flat calculation and is not assumed in the no-go theorem.

Changed block observable. A map $\widetilde{Q}_L$ that measures information discarded by Q_L may lift the witness. Equation (31) gives the necessary four-dimensional leakage scale for the critical scalar normalization. Constructing a local, gauge-covariant, renormalization-compatible map with this property and proving full-space coercivity remain open.

No continuum conclusion. The theorem is finite-dimensional. It proves no positive coercivity for an interacting theory, no infinite-volume measure, no continuum limit, no Osterwalder–Schrader reconstruction, and no Yang–Mills mass gap.

11 Conclusion

Critical scalar rescaling can correct a normalization mismatch but cannot remove an exact kernel; neither can any functional of the block variable performed after that kernel has formed, provided it vanishes at the coarse zero field. The constant-sector wall selects $\sigma = (d-2)/2$ only when exact scale neutrality is required; uniformity of that sector alone allows every larger exponent. None of these scalar choices acts on $\ker Q_L$. The first within-block Fourier phase produces, at every $L \geq 2$ and every positive coarse side, a transverse divergence-free cochain annihilated by Q_L. Its exact Rayleigh quotient is $4\sin^2(\pi/L)$, so the best full-space spectral lower bound is at most $4\pi^2/L^2$. Higher frequencies give a $2R$-dimensional orthogonal test space below $4\pi^2R^2/L^2$, hence a growing low-energy spectral cluster in the linear problem. The current flat full-domain gate is therefore impossible for every factorized penalty $\Phi_L(Q_LA)$ with $\Phi_L(0) = 0$; continuity, linearity, positivity, and bounded growth play no role.

The negative result is operational: it identifies the precise input that must change. A successful continuation must justify a different domain, measure the missing channel before postconditioning, or add an operator that supplies uniform energy to the explicitly exhibited Fourier sector. The machine-checked repair identity and necessary budget quantify this last requirement without assuming full-space sufficiency.

A The square-wave precursor

For even $L = 2M$, the real profile that equals $+1$ on the first half of each block and -1 on the second half also has zero block mean. Its two jumps per block have squared magnitude four, giving total one-dimensional energy $8N'$ and Rayleigh quotient $8/L$. This witness is useful because it requires only a single internal vector and elementary interface counting. The Fourier witness is stronger in three ways: it works for every $L \geq 2$, it identifies the first periodic eigenmode, and it improves the quantitative obstruction from $8/L$ to $4\sin^2(\pi/L) = O(L^{-2})$.

References

[1] T. Bałaban, *Propagators and renormalization transformations for lattice gauge theories. I*, Communications in Mathematical Physics **95** (1984), 17–40. doi:10.1007/BF01215753.

[2] T. Bałaban, *Propagators and renormalization transformations for lattice gauge theories. II*, Communications in Mathematical Physics **96** (1984), 223–250. doi:10.1007/BF01240221.

[3] T. Bałaban, *Averaging operations for lattice gauge theories*, Communications in Mathematical Physics **98** (1985), 17–51. doi:10.1007/BF01211042.

[4] J. Dimock, *Covariant axial gauge*, Letters in Mathematical Physics **105** (2015), 959–987. doi:10.1007/s11005-015-0763-0.

[5] J. M. Combes and L. Thomas, *Asymptotic behaviour of eigenfunctions for multiparticle Schrödinger operators*, Communications in Mathematical Physics **34** (1973), 251–270. doi:10.1007/BF01646473.

[6] L. de Moura and S. Ullrich, *The Lean 4 theorem prover and programming language*, Automated Deduction – CADE 28, LNCS 12699 (2021), 625–635. doi:10.1007/978-3-030-79876-5_37.

[7] The mathlib Community, *The Lean mathematical library*, Proceedings of CPP 2020 (2020), 367–381. doi:10.1145/3372885.3373824.