Coarse information and finite rank obstructions for flat lattice Hodge forms

Lluis Eriksson

9 November 2026 — integrated geometry revision candidate

Abstract

We study how finite observations and finite-rank repairs can restore coercivity for unit-spacing flat lattice Hodge forms. An explicitly constructed real Fourier band gives necessary output-dimension budgets, including tangent obstructions for differentiable nonlinear observations. Singular Schur compatibility explains why kernel nonnegativity alone is insufficient. Positive results include scalar and matrix critical inverse-moment bounds for global spectral observations, exact finite prescribed-mean penalty thresholds, and constructive spatial means. For coefficients depending on any $k < D$ coordinate directions, we construct actual disjoint rectangular averages with penalty coefficient $K = c$ and the critical target-order observation budget. The construction derives local energy control from the coefficients and charges a balanced base tree through its invariant directions; it handles odd sizes and singleton cells. At least one exactly invariant coordinate remains required by that argument. An explicit corollary of the existing threshold construction handles arbitrary arrangements of scalar coefficient values 1 and H: the actual critical moment gives target-order rectangular means with the same global penalty coefficient $K = c$. Complementary examples distinguish successful observations from minimum-coefficient, full-tree, point-support and strong Green-moment certificates. The latter have matching all-rectangle certificate-cost regimes. A weaker detour exclusion and an explicit coexact parity example delimit an additional local comparison route. Further exact boundary-energy calculations rule out uniformly controlled exact trace extension based only on internal energy, including near-one critical moments, and, in $D \geq 4$, a size-independent sum of separately minimized harmonic costs over disjoint rectangles. The discrepancy also occupies a growing-dimensional space, ruling out uniformly bounded-rank repairs of that particular comparison. Nevertheless, the underlying binary family admits the constructive partition just described. Local shell estimates also lower the inverse-moment exponent from $D/2$ to $(D-1)/2$ with a necessary thickness cost; the thickness-uniform $D/2$ estimate and sharp local counterexamples are grouped with them. A coefficient-selected unit shell supplies actual disjoint face means. Two matched quantitative counterexamples distinguish discarded canonical interfaces from an actual off-center annular slowdown; successful full-cube means on both families keep the unrestricted question separate. Complete analytical proofs and separately pinned formal endpoints and finite diagnostic records are distinguished throughout. Classical methods are attributed, and originality remains subject to further comparison. The unrestricted critical rectangular problem is open. No interacting or continuum Yang–Mills mass gap is asserted.

Theorem, scope and evidence guide

The immediate public predecessor is the existing AIRR version 11. Historical versions, reviews and their exact hashes remain separate. This is a changed integrated candidate, not an assigned public version or acceptance.

Result family	Mathematical scope	Proof and evidence location
Fourier and tangent budgets	Actual finite flat spaces; necessary budgets	Sections 2–13; historical declaration map and source anchors in Section 14.
Critical spectral sufficiency	Scalar critical inverse moments; global observations	Sections 16–18; analytical finite proof, separate formal endpoints.
Spatial construction adapters	Supplied cubes, fixed gaps and actual coexact counterexamples	Sections 19–24 and 29; full finite Sobolev and partition proofs retained.
Matrix and exact penalty results	Matrix-weighted global criteria; specified mean maps	Section 28; paired finite records are distinct from all-size proofs.
Laminate and Green costs	One-coordinate profiles; costs of a sufficient certificate	Section 30 and following sections; source units 184–185.
Partial coordinate invariance	Every $1 \leq k < D$; actual means, $K = c$	Section 35; source188, full source audit and paired finite tree diagnostics.
Detour geometry boundary	Distance two suffices in $D \geq 3$; nearest-neighbour independence fails	Section 36; source190 and full involved audit, finite replay pending.
Boundary transfer and binary means	Growing-rank comparison obstruction; arbitrary scalar binary arrangements	Sections 37–38; comparison proofs and binary corollary of Section 21; source audits192–199, inventory correction201.
Local shell input	Actual internal energy; two inverse-moment endpoints and sharp powers	Section 39; grouped analytical sources205/216, finite checks separately scoped.
Selected observations and limits	Disjoint shell means; sharp internal repair and actual annular costs	Section 40; full grouped proofs218–220, original source reports and finite transcripts.
Unrestricted question	Fully variable coefficients; still open	Section 27 and final scope paragraph; no physical continuum conclusion.

The immutable programme source/evidence root is `https://github.com/lluiseriksson/THE-ERIKSSON-PROGRAMME`. At commit `d8d2c66ce76103e302e08cb74f878d46b7a60262`, the unit188 proof and complete evidence are under `papers/flat-coarse-rank/research/`: `2026-11-07_partial-invariance188.txt` and `evidence188/`. The latter retains the complete source audit, every limitation, the Linux driver, raw diagnostic archive and resource record. Revision221's own build manifest must pin this changed manuscript. No source-audit verdict or older PDF score transfers as its numerical exact-PDF assessment.

1 Purpose and relation to the existing record

This text consolidates a continuation of the exact-kernel obstruction in [4]. That paper exhibits fields in the kernel of a particular line-integral block map. Changing the map can remove a particular witness. Here the obstruction depends instead on the amount of coarse information: a whole low-energy subspace is too large to be seen by a fixed-dimensional measurement together with a bounded-rank repair. The map need not factor through the original block variable, be local, or obey a norm bound; it is real-linear in the unrestricted-penalty results. The penalty applied to its output may be nonlinear. Later sections treat differentiable observations with a locally upper-quadratic penalty and exact nonlinear regular sectors. The restricted-sector result is proved separately and does not follow merely from failure of an estimate on the full space.

Fourier diagonalization, rank-nullity and finite-dimensional Schur complements are classical. For the singular Schur complement and pseudoinverse background, see Albert and Carlson–Haynsworth–Markham [6, 7]. Finite-difference Hodge theory has a long history; Dodziuk's work [5] provides historical context, rather than this paper's specific lattice normalization. We make no claim to a new general spectral principle or an exhaustive priority search. The contribution is a concrete, reproducibly checked constraint on proposed repair routes in this programme. This consolidation is intended for the existing paper dossier, rather than an additional deposit for each lemma.

Gauge-fixing and block averaging in renormalization-group constructions require their own domain and scaling analysis. Dimock [8] develops covariant axial gauge for an abelian model and discusses its relation to Balaban's programme. His introduction distinguishes a finite unit lattice from a torus with small spacing, equivalently a large unit-spacing volume. Our unrestricted, unit-spacing flat form is not identified with every constrained fluctuation form in those constructions. This negative diagnostic does not refute their gauge-fixed estimates or establish an interacting bound.

2 Spaces and the actual flat form

Fix integers $d \geq 1$, $n \geq 1$ and $N_c \geq 2$. Set $X_n = (\mathbb{Z}/n\mathbb{Z})^d$ and $m = N_c^2 - 1$. The real one-cochain space is

$$E_n = \{A : X_n \times \{1, \dots, d\} \longrightarrow \mathbb{R}^m\}, \qquad \|A\|^2 = \sum_{x \in X_n} \sum_{i=1}^{d} \|A_i(x)\|^2.$$

The internal coordinates are those of the repository's Euclidean Lie-coordinate model. No identification with a matrix Lie algebra is required here. For positive unordered plaquettes $i < k$, define

$$\begin{aligned} (\mathrm{d}A)_{ik}(x) &= A_k(x + e_i) - A_k(x) - A_i(x + e_k) + A_i(x), \\ (\delta A)(x) &= \sum_i (A_i(x) - A_i(x - e_i)). \end{aligned}$$

The flat operator $K_{0,n}$ is the actual repository Hodge operator, with

$$\langle A, K_{0,n} A \rangle = \|\mathrm{d}A\|^2 + \|\delta A\|^2. \tag{1}$$

The plaquette norm sums each pair once. All differences have unit spacing; there is no hidden n^2 or block-side squared factor.

Choose a direction j and two distinct internal unit coordinates u, v. They exist since $m \geq 3$. For $f : X_n \to \mathbb{C}$, put

$$(E_j f)_i(x) = \begin{cases} (\Re f(x))u + (\Im f(x))v, & i = j, \\ 0, & i \neq j. \end{cases}$$

This is a real-linear injection and $\|E_j f\|^2 = \sum_x |f(x)|^2$. Writing $D_i f(x) = f(x + e_i) - f(x)$, the actual curl and divergence give

$$\langle E_j f, K_{0,n} E_j f \rangle = \sum_{i \neq j} \sum_x |D_i f(x)|^2 + \sum_x |f(x) - f(x - e_j)|^2 = \sum_i \sum_x |D_i f(x)|^2. \tag{2}$$

The last equality is finite translation invariance. The exclusive two plaquette incidences and the divergence term account for every direction exactly once, with no extra factor two.

3 The original block kernel and its repair budget

We retain the principal results of the historical paper [4] before explaining the new information-budget obstruction. The original Appendix A (an even-block square-wave witness with quotient $8/\ell$) remains in the preserved v1; it is not reproduced here. Use positive integers ℓ, N', set $n = \ell N'$, and for a coarse site y put $B_y = \{\ell y + r : 0 \leq r_i < \ell\}$. The original line-integral map is

$$(Q_\ell A)_i(y) = \ell^{-d} \sum_{x \in B_y} \sum_{k=0}^{\ell-1} A_i(x + k e_i).$$

For a constant field $A_i(x) = v_i$ and $S = \sum_i \|v_i\|^2 > 0$,

$$\|A\|^2 = \ell^d (N')^d S, \qquad \|\ell^\sigma Q_\ell A\|^2 = \ell^{2\sigma+2} (N')^d S.$$

A uniform Poincare estimate must therefore have $\sigma \geq (d-2)/2$ if it is to survive this sector. Equality alone makes the two powers identical. Larger exponents pass this necessary test too; neither equality nor the inequality is a sufficient full-space condition. This follows by cancelling $(N')^d S$ to obtain $\ell^{d-2-2\sigma} \leq C_P$. The real-exponent comparison is analytic here.

For $d \geq 2$ choose distinct directions p, q and the real isometry $\iota(z) = (\Re z)u + (\Im z)v$ from Section 2. For $\ell \geq 2$ define

$$A_i^{(\ell)}(x) = \mathbf{1}_{i=p}\, \iota(e^{2\pi \mathrm{i} x_q/\ell}), \qquad \lambda_\ell = |e^{2\pi \mathrm{i}/\ell} - 1|^2 = 4\sin^2(\pi/\ell).$$

The character is well defined on $X_{\ell N'}$. The sum over the q-offset in every block is a nontrivial root-of-unity sum and vanishes. Shifts in the line direction p do not change the profile. Hence

$$Q_\ell A^{(\ell)} = 0, \quad \delta A^{(\ell)} = 0, \quad \|A^{(\ell)}\|^2 = (\ell N')^d, \quad \langle A^{(\ell)}, K_{0,n} A^{(\ell)} \rangle = (\ell N')^d \lambda_\ell. \tag{3}$$

Only the plaquette pair containing p, q contributes to the curl energy; its difference has constant squared norm λ_ℓ at every site. This also explains the exact energy in (3), including periodic boundary jumps.

Consequently any scale-wise estimate

$$\|A\|^2 \leq C_P(\ell)(\langle A, K_{0,n} A \rangle + \Psi_\ell(Q_\ell A)), \qquad \Psi_\ell(0) = 0,$$

forces $C_P(\ell) \geq \lambda_\ell^{-1} \geq \ell^2/(4\pi^2)$. No fixed $C_P > 0$ before all block sides can work. This includes $\Psi_\ell(y) = \|T_\ell y\|^2$ for arbitrary linear postconditioners, and all critical or faster scalar rescalings. There is no regularity or sign assumption on the scalar function Ψ_ℓ.

The original spectral strengthening is retained analytically. For $1 \leq r \leq R \leq \lfloor(\ell-1)/2\rfloor$, replace the profile by $\iota(ce^{2\pi \mathrm{i} r x_q/\ell})$, $c \in \mathbb{C}$. These are mutually orthogonal real two-dimensional spaces in $\ker Q_\ell$. Derivative orthogonality must be checked too: every forward difference multiplies a character by its frequency multiplier, so the same root-of-unity Gram sum annihilates cross terms between distinct frequencies. Thus for $H_\ell = K_{0,n} + (T_\ell Q_\ell)^*(T_\ell Q_\ell)$ the sum of these spaces has dimension $2R$ and maximum Rayleigh quotient at most $4\sin^2(\pi R/\ell) \leq 4\pi^2 R^2/\ell^2$. Finite-dimensional min–max gives $\mu_{2R}(H_\ell) \leq 4\pi^2 R^2/\ell^2$. Taking $R = \lfloor \ell^\alpha \rfloor$, $0 < \alpha < 1$, for sufficiently large ℓ, yields a growing cluster below $4\pi^2 \ell^{-2(1-\alpha)}$. Eigenvalue language is restricted to this linear self-adjoint operator, not applied to a general nonlinear penalty.

Finally, for any linear changed measurement $\widetilde{Q}_\ell$, linear postconditioner T_ℓ and positive self-adjoint fine term V_ℓ, put

$$\varepsilon_\ell = \frac{\|T_\ell \widetilde{Q}_\ell A^{(\ell)}\|}{\|A^{(\ell)}\|}, \qquad \nu_\ell = \frac{\langle A^{(\ell)}, V_\ell A^{(\ell)}\rangle}{\|A^{(\ell)}\|^2}.$$

Set $\widetilde{H}_\ell = K_{0,n} + (T_\ell \widetilde{Q}_\ell)^*(T_\ell \widetilde{Q}_\ell) + V_\ell$. Expansion gives the exact repaired quotient

$$\frac{\langle A^{(\ell)}, \widetilde{H}_\ell A^{(\ell)}\rangle}{\|A^{(\ell)}\|^2} = \lambda_\ell + \varepsilon_\ell^2 + \nu_\ell.$$

A full-space constant $c > 0$ therefore requires $\varepsilon_\ell^2 + \nu_\ell \geq c - \lambda_\ell$ at every scale, hence $\liminf_{\ell\to\infty}(\varepsilon_\ell^2 + \nu_\ell) \geq c$. On the single witness line, uniform coercivity is equivalent to $\inf_{\ell\geq 2}(\lambda_\ell + \varepsilon_\ell^2 + \nu_\ell) > 0$; this is not full-space sufficiency. In particular if $\widetilde{Q}_\ell = Q_\ell + R_\ell$, $T_\ell = \ell^{(d-2)/2} I$ and $\eta_\ell = \|R_\ell A^{(\ell)}\|/\|A^{(\ell)}\|$, then $\varepsilon_\ell^2 = \ell^{d-2}\eta_\ell^2$. In dimension four and with $V_\ell = 0$, leakage $o(\ell^{-1})$ cannot even repair this line. The liminf, real-exponent, spectral multiplicity and witness-line equivalence are analytic statements. The older exact kernel, norm-form energy, finite budget and dimension-four factorized no-go have their own preserved formal source and reproductions, separately anchored below.

4 A constructed low band

For a nonnegative integer R with $2R < n$, let $\Lambda_R = \{-R, \ldots, R\}^d$ and $B = |\Lambda_R| = (2R+1)^d$. Define

$$\chi_k(x) = \exp(2\pi \mathrm{i}\, k \cdot x/n), \qquad f_a = \sum_{k\in\Lambda_R} a_k \chi_k, \qquad S_R = \{E_j f_a : a \in \mathbb{C}^{\Lambda_R}\} \subset E_n.$$

Lemma 4.1. *The real subspace S_R has dimension $2B$, and for every $A \in S_R$,*

$$\langle A, K_{0,n} A\rangle \leq b_R(n)\|A\|^2, \qquad b_R(n) = \frac{4\pi^2 d R^2}{n^2}. \tag{4}$$

Proof. For $k, \ell \in \Lambda_R$, the product root-of-unity sum is $\sum_x \overline{\chi_k(x)}\chi_\ell(x) = n^d \mathbf{1}_{k=\ell}$. The no-alias condition ensures that coordinate differences, of absolute value at most $2R < n$, cannot be nonzero multiples of n. Thus synthesis is injective and its complex image has dimension B. As a real space that image has dimension $2B$; applying the real-linear injection E_j preserves

dimension. The factor two represents independent real and imaginary coefficients stored in two real internal coordinates, not two copies of a real scalar band.

For completeness, energy orthogonality is needed as well as mass orthogonality. For each i,

$$D_i f_a = \sum_k a_k (e^{2\pi \mathrm{i} k_i/n} - 1)\chi_k.$$

Applying the same Gram sum to these derivative amplitudes gives

$$\|E_j f_a\|^2 = n^d \sum_k |a_k|^2,$$
$$\langle E_j f_a, K_{0,n} E_j f_a \rangle = n^d \sum_k |a_k|^2 \sum_i |e^{2\pi \mathrm{i} k_i/n} - 1|^2.$$

Since $|e^{\mathrm{i}t} - 1| \le |t|$ and $|k_i| \le R$, each sum of multipliers is at most $4\pi^2 dR^2/n^2$. This proves the bound for every member of the constructed image, not only its individual generators. □

5 The finite information budget

Let F be a finite-dimensional real vector space, $C : E_n \to F$ a linear measurement and $V : E_n \to E_n$ a linear operator. No symmetry or positivity of V is needed. Let $\Phi : F \to \mathbb{R}$ be any function with $\Phi(0) = 0$.

Theorem 5.1. *If $c > b_R(n)$ and*

$$c\|A\|^2 \le \langle A, K_{0,n} A\rangle + \Phi(CA) + \langle A, VA\rangle \quad \textit{for all } A \in E_n, \tag{5}$$

then $2(2R+1)^d \le \dim F + \operatorname{rank} V$.

Proof. If the budget is smaller than $\dim S_R$, restrict the map (C, V) to S_R. Its image has dimension at most $\dim F + \operatorname{rank} V$. Rank-nullity yields a nonzero $A \in S_R$ with $CA = VA = 0$. Equation (5) and Lemma 4.1 imply $c\|A\|^2 \le b_R(n)\|A\|^2$, a contradiction. □

The codomain dimension is a conservative measurement budget. We do not label this formal endpoint as a sharper theorem with $\operatorname{rank} C$ in its place. No regularity, sign, polynomial form or growth restriction on Φ enters.

Proposition 5.2. *If $P : E_n \to E_n$ is linear, $c > b_R(n)$ and $\dim F + \operatorname{rank} P < 2(2R+1)^d$, the inequality*

$$c\|A\|^2 \le \langle A, K_{0,n} A\rangle + \Phi(CA) \qquad (A \in \ker P)$$

is false.

Proof. Apply the same rank-nullity argument to $(C, P)|_{S_R}$. Its nonzero kernel vector belongs to the stipulated sector as well as the unobserved band. □

For example, $d = 4, n = 16, R = 1$ gives dimension 162 and $b_R(n) = \pi^2/16 < 1$. Twelve measurement coordinates and a repair of rank at most 149 cannot yield (5) with $c = 1$. This illustrates a necessary budget, not the programme's chosen target constant.

6 Uniform refinement and sector exclusion

Fix a positive integer N' and put $n_L = (L+1)N'$ for nonnegative integers L. The finite-dimensional space F is fixed. Allow arbitrary families of linear maps $C_L : E_{n_L} \to F$, arbitrary functions $\Phi_L : F \to \mathbb{R}$ with $\Phi_L(0) = 0$, and linear maps $V_L : E_{n_L} \to E_{n_L}$ with $\operatorname{rank} V_L \le r$ for one fixed nonnegative integer r.

Theorem 6.1. *There is no $c > 0$ for which (5) holds for every L and every $A \in E_{n_L}$ with these families. Likewise, for linear P_L of rank at most r, there is no positive uniform c for the sector estimates of Proposition 5.2 on all* $\ker P_L$.

Proof. Set $M = \dim F$ and $R = M + r + 1$. Since $d \ge 1$, $2(2R+1)^d > M + r$. Choose L so large that $2R < n_L$ and $b_R(n_L) < c$. Here is a fully finite choice argument, without an assumed limit theorem: with $a = 4\pi^2 dR^2$, take an integer $L > \max\{2R, a/c, 1\}$. Since $N' \ge 1$, $n_L \ge L + 1 > L > 1$, and $a < cL \le cn_L^2$. Theorem 5.1 now contradicts the rank bound. Use Proposition 5.2 for the separate sector conclusion. □

The constants c, r precede the quantifier over L. Arbitrary rescalings of C_L preserve its fixed codomain and do not escape the theorem. Removing harmonic or constant directions is covered only when a concrete linear constraint describing the proposed sector has uniformly bounded rank. A quotient or its norm is not silently identified with such a sector. Linearity is essential: if nonlinear observations were allowed, the single real observation $C_L(A) = \|A\|^2$, with $\Phi_L(y) = y$ and $V_L = 0$, would give the uniform lower bound $c = 1$ since $K_{0,n_L} \succeq 0$. This counterexample explains why finite-dimensional output alone is not the hypothesis.

7 Constructed derivative repairs and tangent observations

The bounded-rank obstruction does not exhaust the available repairs. For the actual componentwise forward gradient $\mathcal{G}A = (D_k A_i)_{x,i,k}$, finite periodic summation and cancellation of the mixed differences give

$$\|\mathcal{G}A\|^2 = \langle A, K_{0,n} A\rangle.$$

This is an identity on the whole cochain space, with unordered plaquettes counted once; it is not an assumed comparison on a supplied band. Let W be the diagonal multiplication by real weights $w_{x,i,k}$, and construct $V_w = \mathcal{G}^* W \mathcal{G}$. If $w_{x,i,k} \le M$ and $M \ge 0$, then

$$\langle A, V_w A\rangle = \sum_{x,i,k} w_{x,i,k} \|(D_k A_i)(x)\|^2 \le M \langle A, K_{0,n} A\rangle. \tag{6}$$

A lower bound on the weights is unnecessary for this upper estimate. Consequently the constructed Fourier band satisfies

$$q_w(A) := \langle A, (K_{0,n} + V_w)A\rangle \le (1+M) b_R(n) \|A\|^2 \qquad (A \in S_R).$$

These operator identities and inequalities are formally proved in the separately pinned source reported below.

Theorem 7.1 (Finite-dimensional tangent obstruction)**.** *Let E be a finite-dimensional real inner-product space, F a finite-dimensional real normed space, and $S \subseteq E$ a linear subspace*

with $\dim S > \dim F$. *Let* $q : E \to \mathbb{R}$ *satisfy* $q(x) \le b\|x\|^2$ *on* S. *Suppose* $C : E \to F$ *is Fréchet differentiable at zero,* $C(0) = 0$, *and for some* $\delta > 0$ *and finite* $K \in \mathbb{R}$,

$$\Phi(y) \le K\|y\|^2 \qquad (\|y\| < \delta). \tag{7}$$

If $c > b$, *the full-space estimate* $c\|x\|^2 \le q(x) + \Phi(C(x))$ *is impossible.*

Proof. Put $L = DC(0)$. Rank-nullity applied to the actual restriction $L|_S$ produces $v \in S \cap \ker L$, $v \neq 0$. Differentiability gives $\|C(tv)\|/\|tv\| \to 0$ as nonzero real $t \to 0$. For small such t, (7) applies and the proposed bound implies

$$c \le b + K\frac{\|C(tv)\|^2}{\|tv\|^2}.$$

Taking the limit gives $c \le b$, a contradiction. The kernel vector and the remainder limit are derived, not introduced as hypotheses of the theorem. □

The three focal Lean endpoints include this theorem, its remainder lemma and the actual physical specialization with threshold $(1 + M)b_R(n)$. The proof does not require an explicit assumption $\Phi(0) = 0$. A local-ball version is also valid by the same shrinking-ray argument, while the original formal endpoint above is stated on the full space. The later local kernel and average-weight endpoint in Section 9 has its own separately reproduced proof.

This result trades hypotheses. The global linearity of C is replaced by differentiability at zero, while an arbitrary zero-vanishing penalty is replaced by genuine quadratic upper control. Merely asking C and Φ to be C^1 with zero derivatives at zero is insufficient: for $1 < \alpha < 2$,

$$C(x) = \|x\|^\alpha, \qquad \Phi(y) = |y|^{2/\alpha} \quad \Longrightarrow \quad \Phi(C(x)) = \|x\|^2.$$

Both derivatives at zero vanish. Off zero, $\|DC(x)\| = \alpha\|x\|^{\alpha-1}$ and $|\Phi'(y)| = (2/\alpha)|y|^{2/\alpha-1}$ tend to zero, so both maps are indeed C^1 there and at zero. But $\Phi(y)/|y|^2 = |y|^{2/\alpha-2}$ is unbounded above near zero. On the actual lattice, setting $w = 0$ gives a uniform lower bound with $c = 1$ because $K_{0,n} \succeq 0$. This is an exact counterexample, not a numerical test.

8 Exact nonlinear regular sectors

The next theorem is an analytic consequence of rank-nullity and the classical implicit function theorem. The actual curve construction and local-energy obstruction now have separately pinned formal proofs, identified in the evidence section. The pinned formal library supplies an implicit parametrization with derivative the inclusion of the derivative kernel [9]; we construct and use that parametrization explicitly below rather than assuming the curve needed for the conclusion.

Theorem 8.1 (Regular-sector tangent obstruction)**.** *Let* E *be a finite-dimensional real inner-product space and* F, G *finite-dimensional real normed spaces. Let* $T : E \to E$ *be linear and set* $q(x) = \langle x, Tx\rangle$. *Suppose a linear subspace* S *satisfies* $q(v) \le b\|v\|^2$ *on* S *and* $\dim S > \dim F + \dim G$. *Let* C *be differentiable at zero with* $C(0) = 0$ *and let* Φ *satisfy* (7). *Let* $P : E \to G$ *be* C^1 *on a neighborhood of zero with* $P(0) = 0$ *and* $A = DP(0)$ *surjective. For every* $c > b$ *and every* $r > 0$, *there is no estimate*

$$c\|x\|^2 \le q(x) + \Phi(C(x)) \qquad (P(x) = 0,\ \|x\| < r). \tag{8}$$

Proof. Put $L = DC(0)$. Apply rank-nullity to $(L, A)|_S : S \to F \times G$ to obtain $v \in S$, $v \neq 0$, with $Lv = Av = 0$. Choose a linear complement W of $H = \ker A$. Surjectivity implies $A|_W : W \to G$ is bijective. Apply the implicit function theorem to $(u, w) \mapsto P(u + w)$ at $(0, 0) \in H \times W$. It produces a C^1 map ψ near zero in H with

$$\psi(0) = 0, \quad P(u + \psi(u)) = 0, \quad D\psi(0) = -(A|_W)^{-1} A|_H = 0.$$

Thus $\gamma(t) = tv + \psi(tv)$ lies in the exact sector for small real t, $\gamma(0) = 0$, and $\gamma(t) = tv + e(t)$ with $\|e(t)\| = o(|t|)$. In particular $\|\gamma(t)\|/|t| \to \|v\| > 0$, so the nonzero points of this curve eventually lie in the stipulated ball.

The curve need not remain in S. This is why the quadratic form assumption is used: direct expansion gives

$$q(\gamma(t)) = t^2 q(v) + t\langle v, Te(t)\rangle + t\langle e(t), Tv\rangle + \langle e(t), Te(t)\rangle.$$

Since T is bounded in finite dimension, division by t^2 gives $q(\gamma(t))/t^2 \to q(v)$. Differentiability of C and $Lv = 0$ give $C(\gamma(t)) = Le(t) + o(|t|) = o(|t|)$. Apply the assumed bound along this exact curve and use (7). Division by t^2 and the limit yield $c\|v\|^2 \leq q(v) \leq b\|v\|^2$, contrary to $c > b$. □

Unlike the preceding full-space result, an arbitrary energy with an upper bound only on S would not justify this proof: the constructed exact-sector curve generally leaves S. The displayed expansion supplies the missing comparison for the actual quadratic energy.

A concrete curved example also shows why an exact vector in S is not the right object. Write $E = \mathbb{R}^4$ with coordinates (u_1, u_2, y, z), put $P(u, y, z) = z - \|u\|^2$, $C(u, y, z) = y$, $\Phi(y) = y^2$, $S = \{z = 0\}$ and $q(u, y, z) = z^2$. Then $DP(0)$ is surjective and $\dim S = 3 > 1 + 1$. But $S \cap \ker C \cap P^{-1}(0) = \{0\}$: the linear band contains no nonzero exact-sector unseen vector. Nevertheless $\gamma(t) = (t, 0, 0, t^2)$ lies in the exact sector and has repaired Rayleigh quotient $t^2/(1 + t^2) \to 0$. The implicit-curve argument is genuinely needed.

Nor can quadratic q in Theorem 8.1 be replaced merely by a C^1 energy stationary at zero with the same upper bound on S. With the same E, C, Φ, S, instead put

$$P(u, y, z) = z - \|u\|^{4/3}, \qquad e(u, y, z) = |z|^{3/2}.$$

Both maps are C^1, $DP(0)(u, y, z) = z$ is surjective, $De(0) = 0$, and $e = 0$ on S, so $b = 0$. On the exact sector,

$$e(x) + \Phi(C(x)) = \|u\|^2 + y^2, \qquad \|x\|^2 = \|u\|^2 + y^2 + \|u\|^{8/3}.$$

For $\|x\| < 1$, $\|u\|^{8/3} \leq \|u\|^2$, so the local sector bound with $c = 1/2 > b$ holds. Differentiability of the norm powers at zero follows from their exponents being greater than one; their gradients off zero tend to zero. This exact counterexample refutes a proposed C^1-energy weakening, while the actual quadratic Hodge/repair energy retains the required expansion.

The regularity condition cannot simply be dropped. Take $E = \mathbb{R}^3$, $F = G = \mathbb{R}$, $C = 0$, $\Phi = 0$, $T = 0$, $S = E$, and $P(x) = \|x\|^2$. Then P is smooth, $DP(0) = 0$, and the strict dimension budget holds, but $P^{-1}(0) = \{0\}$ and the estimate with $c = 1$ is vacuous on the sector. This refutes the proposal to replace submersion by bounded derivative rank. It does not refute Theorem 8.1, because $DP(0)$ is not surjective. Critical constraint maps require a separate analysis.

Proposition 8.2 (Uniform physical consequence)**.** *Let $n_L \to \infty$ with fixed d, N_c, and let F, G have fixed finite dimensions. For each L, allow its own C_L, P_L, Φ_L satisfying the hypotheses of Theorem 8.1. Let q_L be the actual flat Hodge form plus the constructed weighted derivative repair with weights $w_L \le M_L$, $M_L \ge 0$, and $M_L/n_L^2 \to 0$. Even on the exact sectors $P_L = 0$ and arbitrarily small balls of positive radius r_L, no positive uniform coercivity constant exists.*

Proof. Choose a fixed integer R with $2(2R+1)^d > \dim F + \dim G$. For all large L, $2R < n_L$ and the constructed band obeys $q_L \le (1 + M_L) b_R(n_L) \|\cdot\|^2$. Its coefficient tends to zero. For any proposed $c > 0$, choose one such L with this coefficient below c and apply Theorem 8.1 at that fixed lattice. The curve is shrunk only after choosing L: no uniform implicit-function radius, inverse-derivative norm, penalty constant or r_L is required. □

These are finite unit-spacing obstructions for regular nonlinear sectors, not identifications of a physical gauge-fixed domain or a quotient norm.

9 Directional mean weights and local nonlinear observations

Pointwise bounds on a derivative repair are stronger than the following actual spatial-mean condition. Fix $d \ge 2$, $N_c \ge 2$, and distinct directions i, j. For nonnegative componentwise weights $w_n(k, (x, l))$, put

$$Q_n(A) = \langle A, K_{0,n} A \rangle + \sum_{x,k,l} w_n(k, (x, l)) \|A_l(x + e_k) - A_l(x)\|^2, \qquad b_{ij}(n) = n^{-d} \sum_x w_n(i, (x, j)).$$

The repair is the actual $G^* W_n G$ of the componentwise forward gradient. In particular $Q_n \ge 0$; signed weights are not included in this section. Let $D_n = \ker \delta$ be the actual flat divergence-free space.

Theorem 9.1 (Local nonlinear average-weight obstruction)**.** *For each $n \ge 1$, let F_n be a finite-dimensional real normed space, $C_n : E_n \to F_n$ be Fréchet differentiable at zero with $C_n(0) = 0$, and let $\Phi_n : F_n \to \mathbb{R}$ satisfy*

$$\Phi_n(y) \le K_n \|y\|^2 \quad (\|y\| < \delta_n), \qquad \delta_n > 0.$$

Suppose one fixed $c > 0$ and positive radii r_n satisfy

$$c\|A\|^2 \le Q_n(A) + \Phi_n(C_n(A)) \quad (A \in D_n,\ \|A\| < r_n). \tag{9}$$

For every integer $0 \le R < n$,

$$c(2R - \dim F_n) \le 2(1 + b_{ij}(n)) \sum_{s=1}^{R} |e^{2\pi i s/n} - 1|^2 \le \frac{8\pi^2 (1 + b_{ij}(n)) R^3}{n^2}. \tag{10}$$

Consequently, if $(1 + b_{ij}(n))/n^2 \to 0$, then $\dim F_n \to \infty$. No uniform bound on K_n, δ_n^{-1} or r_n^{-1} is required.

Proof. Write $L_n = DC_n(0)$. If $v \in D_n \cap \ker L_n$, differentiability gives $C_n(tv) = o(|t|)$. For all sufficiently small nonzero t, both the local coercivity and penalty bound apply. Divide by t^2 and let $t \to 0$ to get

$$c\|v\|^2 \le Q_n(v). \tag{11}$$

This argument is performed separately at each n; it does not impose uniform neighborhoods or derivative remainder rates.

Use the two internal coordinates u, v of Section 2 to regard a complex number as a real internal vector. For $1 \le s \le R$ and $a \in \{1, \mathrm{i}\}$, define $\psi_{s,a}$ by taking only component j nonzero and setting it to $n^{-d/2}$ times the internal vector corresponding to $ae^{2\pi \mathrm{i} s x_i/n}$. The finite-character orthogonality identity and the real orthogonality of the two quadratures give an orthonormal family of $2R$ fields. This also covers even n; the two quadratures use two internal coordinates, not two allegedly independent scalar sine and cosine functions. Since $i \neq j$, these fields are independent of x_j, so $\delta\psi_{s,a} = 0$. Their span $S_{R,n}$ is thus an actual $2R$-dimensional subspace of D_n.

Every nonzero forward difference of $\psi_{s,a}$ is in direction i and has the same squared norm at every site, namely $n^{-d}|e^{2\pi \mathrm{i} s/n} - 1|^2$. The exact Hodge and repair energies are therefore

$$Q_n(\psi_{s,a}) = (1 + b_{ij}(n))|e^{2\pi \mathrm{i} s/n} - 1|^2.$$

In particular the trace of the compression of Q_n to $S_{R,n}$ is the middle expression in (10); no invariance of $S_{R,n}$ under the weighted operator is asserted or needed.

For completeness, the positive trace/rank argument is as follows. Set $U = S_{R,n} \cap \ker L_n$. Rank-nullity gives $\dim U \ge 2R - \dim F_n$. Extend an orthonormal basis of U to one of $S_{R,n}$. Equation (11) bounds the terms on U below by c; positivity of Q_n makes every remaining diagonal term nonnegative. Hence $c(2R - \dim F_n) \le \mathrm{tr}_{S_{R,n}} Q_n$. Finally $|e^{2\pi \mathrm{i} s/n} - 1| \le 2\pi s/n \le 2\pi R/n$, which proves the last inequality in (10).

Fix R. The right side tends to zero under the stated mean condition. Eventually it is less than c. If $\dim F_n < 2R$, integrality would make the left side at least c. Thus $\dim F_n \ge 2R$ eventually. Arbitrarily large fixed R prove the claimed divergence. □

The condition concerns one directional spatial average, not the maximum of all weights. For example, a weight of size n^{d+1} at one site in component (i, j) and zero elsewhere has $b_{ij}(n) = n$, while its maximum is not subquadratic. Theorem 9.1 still excludes bounded output dimension under its other hypotheses. This example is a direct calculation, not an additional formal endpoint. The nonnegative-weight condition is essential to the trace argument as used here: the complement of the derivative kernel must contribute nonnegative energy. This is not a universal extension of older statements allowing signed weights with a pointwise absolute cap.

9.1 A constructed quadratic-weight boundary example

The exponent 2 cannot be replaced by an unrestricted allowance of quadratic-size mean weights. Let $\mu_j(A) = n^{-d} \sum_x A_j(x)$, with output in the same fixed space $F = (\mathbb{R}^{N_c^2-1})^d$ for every n. A finite path comparison on the actual periodic lattice gives

$$\|A - \mathrm{lift}\, \mu(A)\|^2 \le dn^2 \|GA\|^2, \qquad \|\mathrm{lift}\, \mu(A)\|^2 = n^d \|\mu(A)\|^2.$$

One way to obtain the first inequality is to write each difference of site values as a path of at most dn forward edges, apply finite Cauchy–Schwarz, and average the pairs of sites; the displayed constant is deliberately nonoptimal. The exact modular-torus reindexing and its transport to the physical cochain norm are also proved in the companion formal construction. Using $\|x + y\|^2 \le 2\|x\|^2 + 2\|y\|^2$, choose actual weights $w_n(k, (x, l)) = 2dn^2$ and the quadratic penalty $\Phi_n(y) = 2n^d \|y\|^2$. Then

$$\|A\|^2 \le \langle A, G^* W_n G A \rangle + \Phi_n(\mu(A)) \le Q_n(A) + \Phi_n(\mu(A))$$

on the entire space, hence also on D_n. The mean observation is an actual continuous linear map with finite-dimensional output type independent of n, not a postulated observation certificate. Here $b_{ij}(n) = 2dn^2$ and the subquadratic hypothesis fails. The example allows $K_n = 2n^d$ to vary, exactly as the obstruction does. It proves neither an optimal boundary constant nor positive coercivity for an interacting gauge operator. The path/mean mechanism is classical; see the primary exposition in [11].

9.2 Why the penalty assumption remains

The C^1 counterexample after Theorem 7.1 also applies to the actual nonnegative Q_n and to D_n. A second useful counterexample retains a polynomial observation and a globally Lipschitz penalty:

$$C(A) = \|A\|^2 \in \mathbb{R}, \qquad \Phi(y) = c|y|, \qquad DC(0) = 0, \qquad \Phi(C(A)) = c\|A\|^2.$$

This example already gives the bound for the actual admissible choice $w = 0$, with $Q_n = K_{0,n} \succeq 0$. The zero-form interpretation is only an abstract illustration, not an instance of this nonnegative-weight lattice family. No finite local quadratic upper constant for Φ exists, since $c|y|/y^2 = c/|y|$. Likewise, $C(A) = \|A\|^{3/2}$ and $\Phi(y) = c|y|^{4/3}$ are both C^1 with zero derivatives at zero, yet their composition is exactly quadratic. These are complete analytic counterexamples, not new Lean claims.

There is a distinct analytic regularity tradeoff, not a hypothesis removal for merely differentiable observations. If, on each fixed derivative-kernel direction, $C(tv) = t^2 H_v + o(t^2)$, and $\Phi(0) = 0$, $D\Phi(0) = 0$, then $\Phi(C(tv)) = o(t^2)$. Indeed $\|C(tv)\|/t^2$ is bounded and $|\Phi(y)| \le \epsilon\|y\|$ near zero, including at zero; multiply the two estimates and let $\epsilon \downarrow 0$. A C^2 observation and a nonnegative differentiable penalty vanishing at zero suffice for this step, because a differentiable local minimum has zero derivative. No actual observation's C^2 regularity is proved here, and no formal composition of this alternate route is claimed.

The strengthened mean-weight result is an output-rank divergence theorem. Its fixed-R argument does not assert a positive density of observations per lattice volume. Fixed $c > 0$, nonnegative weights and the stated penalty condition remain explicit. In particular no interacting or continuum Yang–Mills hypothesis, ultraviolet activity estimate, or physical spectral gap is supplied by this section.

10 Signed weights and the limits of mean criteria

Write $R_w = G^* W G$ and $Q_w(A) = \langle A, K_{0,n} A\rangle + \langle A, R_w A\rangle$. For signed weights this form need not be nonnegative. Its spatial means alone no longer supply the positive trace comparison used above. The following sufficient condition is nevertheless available without a lower bound on any original weight. Define the *pointwise* positive part and its selected directional mean by

$$w^+(k,(x,l)) = \max\{w(k,(x,l)), 0\}, \qquad b^+_{ij}(n) = n^{-d} \sum_x w^+(i,(x,j)).$$

In general b^+_{ij} is not the positive part of the signed mean.

Theorem 10.1 (A sufficient signed-weight condition)**.** *Keep the observation, penalty, local divergence-free domain and fixed $c > 0$ of Theorem 9.1, but use arbitrary signed actual weights in*

the local coercivity inequality. For every $0 \le R < n$,

$$c(2R - \dim F_n) \le \frac{8\pi^2(1 + b_{ij}^+(n))R^3}{n^2}.$$

If $(1 + b_{ij}^+(n))/n^2 \to 0$, *then* $\dim F_n \to \infty$. *The original* $C_n, L_n, \Phi_n, c, K_n, \delta_n, r_n$ *are retained.*

Proof. The exact weighted-gradient energy gives, for every actual cochain,

$$\langle A, R_w A\rangle = \sum_{x,k,l} w(k,(x,l))\|D_k A_l(x)\|^2 \le \sum_{x,k,l} w^+(k,(x,l))\|D_k A_l(x)\|^2 = \langle A, R_{w^+} A\rangle.$$

The inequality uses only nonnegativity of the squared differences. Thus the original local coercivity inequality transfers upward to the constructed nonnegative repair R_{w^+}, and Theorem 9.1 applies. Nonnegativity of w^+ and the comparison are proved from their actual definitions, not supplied as fields of a comparison structure. □

This changes the explicit mean hypothesis for signed weights. It neither proves the original weights nonnegative nor removes the local regularity, quadratic penalty or fixed positive coercivity requirements.

10.1 Zero signed means with fixed output

There is no counterpart using only ordinary signed directional means, even when *every* such mean is zero. Fix a site $o \in X_n$, set $V = n^d$, $a = 2dn^2$, and construct actual weights

$$w_n(k,(x,l)) = a\big(1 - V\mathbf{1}_{x=o}\big).$$

Since $|X_n| = V$, their ordinary spatial means are all exactly zero. Let $\mu(A)$ be the directional mean above, and let $J_n(A) = (D_k A_l(o))_{k,l}$ collect all actual forward differences at o. Use the same finite output type for every side:

$$F = (\mathbb{R}^m)^d \times (\mathbb{R}^m)^{d\times d}, \quad C_n(A) = (\mu(A), J_n(A)), \quad \|(y,z)\|_F = \max\{\|y\|_2, \|z\|_2\}.$$

Its dimension is $(d + d^2)m$, independent of n. The companion code proves the fixed finite output type; this displayed numerical dimension is an analytic calculation, not an additional named oracle endpoint.

Proposition 10.2. *The actual continuous linear observation* C_n *and the nonnegative quadratic penalty*

$$\Phi_n(y,z) = 2V\|y\|_2^2 + aV\|z\|_2^2$$

satisfy $C_n(0) = 0$, $DC_n(0) = C_n$, $\Phi_n(y,z) \le V(2+a)\|(y,z)\|_F^2$, *and*

$$\|A\|^2 \le Q_{w_n}(A) + \Phi_n(C_n(A)) \qquad (A \in E_n).$$

Consequently all the original local observation and penalty hypotheses can hold with fixed output and $c = 1$, *although all signed means vanish.*

Proof. Linearity follows directly from the finite sums and differences. The penalty bound follows from each projection norm being at most the specified product max norm. The exact energy expansion is

$$\langle A, R_{w_n} A\rangle = a\|GA\|^2 - aV\|J_n(A)\|_2^2.$$

Adding the penalty cancels the negative term and leaves

$$Q_{w_n}(A) + \Phi_n(C_n(A)) = \langle A, K_{0,n} A\rangle + a\|GA\|^2 + 2V\|\mu(A)\|_2^2 \geq \|A\|^2.$$

The last step is the actual mean/path estimate of the preceding section, with $a = 2dn^2$, together with $K_{0,n} \succeq 0$. □

The coefficient $K_n = V(2 + a)$ may grow with the lattice side, exactly as allowed by the proposed ordinary-mean extension. This is not a transfer of the desired coercivity into an assumed observation certificate. The weights, observation, penalty and cancellation are constructed. For $n > 1$ their positive-part mean is $a(1 - V^{-1})$, of quadratic rather than subquadratic order, so they do not contradict Theorem 10.1.

As an analytic diagnostic, take $n \geq 3$, $i \neq j$, a unit internal vector u, and $A_j(x) = u\mathbf{1}_{x_i = o_i}$, with other components zero. Then $\delta A = 0$, $\|GA\|^2 = 2n^{d-1}$ and $\|J_n(A)\|^2 = 1$, giving

$$Q_{w_n}(A) = n^{d-1}\big(2 - a(n-2)\big) < 0.$$

This explains the failure of the original positive trace argument. The diagnostic field calculation is analytic; the formal fixed-output family theorem does not claim to check this additional field explicitly.

10.2 An attributed square cancellation

Fix distinct directions i, k, a component $j \neq i$ (for instance $j = k$), and $o \in X_n$. For an amplitude $A_0 \geq 0$, define the repair weights additively by

$$\begin{aligned} w(p) = {} & A_0\mathbf{1}_{p=(i,(o,j))} - 3A_0\mathbf{1}_{p=(k,(o,j))} \\ & - 3A_0\mathbf{1}_{p=(i,(o+e_k,j))} - 3A_0\mathbf{1}_{p=(k,(o+e_i,j))}. \end{aligned}$$

Writing $D_s A_j(x) = A_j(x + e_s) - A_j(x)$, commuting shifts give

$$D_i A_j(o) = D_k A_j(o) + D_i A_j(o + e_k) - D_k A_j(o + e_i).$$

The squared norm of this sum is at most three times the sum of the three squared norms. The actual finite weighted-energy sum therefore proves $\langle A, R_w A\rangle \leq 0$ for every cochain. For $n \geq 2$ the positive edge is separate from all negative contributions, so $w^+ = A_0\mathbf{1}_{p=(i,(o,j))}$ and $b_{ij}^+ = A_0/n^d$. Choosing $A_0 = n^d n^2$ gives $b_{ij}^+ = n^2$.

If the original local coercivity hypotheses hold for this actual repair, they imply the same hypotheses with zero repair, since $Q_w(A) \leq \langle A, K_{0,n} A\rangle$. Applying Theorem 9.1 with zero weights again forces $\dim F_n \to \infty$. The selected normalized positive-part mean is $1 + n^{-2}$ and does not tend to zero. This last limit interpretation is an elementary analytic consequence of the formal exact mean identity, not a separate negated-limit oracle endpoint. This single sparse cycle does not make every directional criterion fail: other eligible means are zero. Its formal endpoints establish the selected mean and derived order, not an all-directions counterexample.

There is, however, a finite sum that covers every eligible pair. This extension is analytic here. For every j and every $i \neq j$, use the preceding cycle with $k = j$, the same site o and amplitude A_0, and add all its weight functions to obtain $\overline{w}$. Linearity of the exact weighted-energy sum gives $\langle A, R_{\overline{w}} A\rangle \leq 0$. For every off-diagonal pair $i \neq j$, direct inspection of the four indicators gives

$$\overline{w}(i,(x,j)) = A_0 \mathbf{1}_{x=o} - 3A_0 \mathbf{1}_{x=o+e_j}.$$

Every diagonal component has only nonpositive contributions. For $n \geq 2$ the two displayed sites are distinct, so each eligible pair has $\overline{b}^{+}_{ij} = A_0/n^d$. With $A_0 = n^d n^2$, *all* off-diagonal positive-part means equal n^2, while the derived domination by zero still gives the original rank obstruction. Thus no eligible mean condition in Theorem 10.1 applies to this sum. The conclusion follows from its actual energy order, not from an assumed comparison field. The finite sum and its all-pairs mean calculation are not silently included in the single-cycle formal oracle reports.

The cancellation mechanism is classical, rather than a new general spectral principle. Restricted to the four active vertices, the negative of this scalar repair has three positive path conductances $3A_0$ and a direct negative conductance $-A_0$. For $A_0 > 0$ the path resistance is $1/A_0$, so this is the critical case of Zelazo and Bürger's Theorem III.3 [12]. Isolated vertices contribute zero blocks; the internal vector coordinates sum scalar copies. This comparison attributes the mechanism; it does not establish originality of the whole consolidated dossier.

The conditional obstruction is nonvacuous. An analytic check takes $F_n = E_n$, $C_n = I$ and $\Phi_n(y) = \|y\|^2 - \langle y, R_w y\rangle \geq 0$ (or replace w by $\overline{w}$). The operator norm gives the upper quadratic bound $(1 + \|R_w\|)\|y\|^2$, while $Q_w(A) + \Phi_n(A) = \langle A, K_{0,n} A\rangle + \|A\|^2 \geq \|A\|^2$. This example has growing output, as the theorem requires; it is not an additional formal endpoint or a compressed observation construction. None of these signed-weight statements proves a volume-density bound, identifies an interacting Hessian, or supplies a physical mass gap.

11 Full volume rank from signed total coefficient means

The directional strip in Section 9 detects divergence of the observation dimension. A full dimensional coexact packet yields a positive volume density under a stronger, explicitly stated mean hypothesis. Original weights may have either sign. The relevant quantity is the spatial mean of $\max\{1 + w, 0\}$, not an ordinary signed mean or a positive part taken after averaging. The extra 1 is the actual flat Hodge coefficient, already absorbed into the envelope; it must not be added a second time.

In this section write $D \geq 2$ for the number of lattice directions, $V_n = n^D$, and E_n for the actual real cochain space on $(\mathbb{Z}/n\mathbb{Z})^D$, with fixed $N_c \geq 2$. Use field components $0, 1$ and the internal real plane representing a complex scalar from Section 2. Put $D_i A_l(x) = A_l(x+e_i) - A_l(x)$ and retain the original form

$$Q_{w_n}(A) = \langle A, K_{0,n} A\rangle + \sum_{x,i,l} w_n(i,(x,l)) \|D_i A_l(x)\|^2.$$

The actual all-cochain flat Hodge identity gives $Q_{w_n}(A) = \sum_{x,i,l}(1 + w_n(i,(x,l)))\|D_i A_l(x)\|^2$.

Theorem 11.1 (Finite full volume mean budget)**.** *Let $1 \leq R$ and $2R < n$. Assume $m \geq 0$ and*

$$\sum_x \max\{1 + w_n(i,(x,l)), 0\} \leq V_n m \qquad (0 \leq i < D,\ l = 0, 1). \tag{12}$$

No mean bound on the other field components, original weight lower bound or prescribed uniform pointwise cap is assumed. Let $C_n : E_n \to F_n$ be Fréchet differentiable at zero, $C_n(0) = 0$, with F_n finite dimensional. Suppose $\Phi_n(y) \le K_n\|y\|^2$ for $\|y\| < \delta_n$, $\delta_n > 0$, and the original small divergence-free coercivity inequality holds with $c \ge 0$ and some radius $r_n > 0$. Set

$$H = 2R(2R+1)^{D-1}, \qquad \Lambda(k) = 4\sum_{i=0}^{D-1} \sin^2(\pi k_i/n), \qquad \mathcal{K} = \{1,\dots,R\} \times \{-R,\dots,R\}^{D-1}.$$

Then

$$c(H - \dim_{\mathbb{R}} F_n) \le m \sum_{(k,q)\in\mathcal{K}\times\{0,1\}} \Lambda(k) \le mH\,\frac{4\pi^2 DR^2}{n^2}. \tag{13}$$

Proof. For $k \in \mathcal{K}$ use the actual scalar character χ_k and the coexact field with components given by its backward differences in directions $1, 0$, with opposite signs. Commuting differences give actual divergence zero. Since $k_0 \neq 0 \pmod n$, the normalization denominator is positive. Counting Fourier orthogonality and both internal real quadratures produce an orthonormal family of H actual real fields, spanning S_n. The condition $2R < n$ prevents aliasing of their scalar frequencies; no scalar real-valued conjugacy constraint is imposed on the internal plane. Every vector of S_n is supported on field components $0, 1$.

Define actual coefficients $v_n(i,(x,l)) = \max\{1+w_n(i,(x,l)), 0\}$ for $l = 0, 1$, and zero otherwise, and the positive form $P_n = G^* V_n^{\mathrm{wt}} G$ with those coefficients. The notation V_n^{wt} here denotes the weight multiplication operator, not lattice volume. Then $P_n \succeq 0$ on all cochains and $Q_{w_n}(A) \le \langle A, P_n A\rangle$ on S_n. This comparison is needed only on the actual packet image, and is not asserted on every cochain or every divergence-free field.

Write $L_n = DC_n(0)$. The same shrinking-ray proof as in Section 9, now followed by the proved envelope comparison, gives $c\|v\|^2 \le \langle v, P_n v\rangle$ on $S_n \cap \ker L_n$. Its dimension is at least $H - \dim F_n$. Complete an orthonormal basis of that kernel within S_n. Positivity of P_n controls the complementary trace terms. Thus $c(H - \dim F_n) \le \mathrm{tr}(P_n|_{S_n})$, where the restriction means the genuine orthogonal compression.

For an individual normalized packet $U_{k,q}$, its squared component amplitudes are constant in x and are $V_n^{-1}a_{k,l}$ for $l = 0, 1$, where $a_{k,l} \ge 0$ and $a_{k,0} + a_{k,1} = 1$. Consequently

$$\langle U_{k,q}, P_n U_{k,q}\rangle = V_n^{-1} \sum_{i=0}^{D-1}\sum_{l=0}^{1} |\chi_k(e_i) - 1|^2 a_{k,l} \sum_x v_n(i,(x,l)) \le m \sum_i |\chi_k(e_i) - 1|^2 = m\Lambda(k).$$

Sum over all H basis vectors. Finally every signed integer frequency satisfies $|k_i| \le R$, and $|\sin t| \le |t|$ gives $\Lambda(k) \le 4\pi^2 DR^2/n^2$. All D derivative directions remain in this sum even though only two field components are occupied. □

Theorem 11.2 (One positive density on the full integer tail)**.** *Fix D, N_c, $c > 0$ and $m \ge 0$. For each integer n under consideration, allow its own $F_n, C_n, L_n, \Phi_n, K_n, \delta_n, r_n, w_n$ satisfying the original local hypotheses and* (12)*, with these same c, m. There are $k \ge 3$ and $\varepsilon > 0$, depending only on D, m, c, such that*

$$\dim_{\mathbb{R}} F_n \ge \varepsilon n^D \qquad (n \ge k)$$

whenever those finite-side hypotheses hold. If the original family hypotheses hold only for $n \ge n_0$, the conclusion applies for $n \ge \max\{n_0, k\}$. One may take $\varepsilon = 1/(2k^D)$ for any integer $k \ge 3$ satisfying $2m\,4\pi^2 D/k^2 \le c$. No common quadratic penalty constant, differentiability neighborhood or coercivity radius is required.

Proof. Choose such a k by the Archimedean property. For each $n \geq k$ put $R = \lfloor n/k \rfloor$. Integer division yields

$$R \geq 1, \qquad kR \leq n < k(R+1) \leq 2kR, \qquad 2R < kR \leq n.$$

Hence the actual packet is admissible and $R/n \leq 1/k$. The finite theorem gives $c(H - \dim F_n) \leq (c/2)H$, so $H \leq 2 \dim F_n$. Also $H \geq (2R)^D \geq (n/k)^D$. Combining these gives the displayed density. Each shrinking-ray argument was applied at a fixed n before using these integer inequalities. No limit of derivatives or local neighborhoods is taken. □

Proposition 11.3 (Exact trace average and a sharper analytic constant)**.** *For $R \geq 1$, put $A_R(x) = \sum_{r=1}^{R} \sin^2(rx)$. Any quotient below involving $\sin x$ is used where it is nonzero; at removable singularities use its continuous extension, equivalently the defining finite sum. The actual multiplier average in* (13) *is*

$$\begin{aligned} B_{D,R,n} &= \frac{1}{H} \sum_{k,q} \Lambda(k) = 4A_R(\pi/n) \frac{2DR+1}{R(2R+1)}, \\ A_R(x) &= \frac{R}{2} - \frac{\sin(Rx)\cos((R+1)x)}{2\sin x}, \\ B_{D,R,n} &\leq \frac{2\pi^2}{3n^2}(R+1)(2DR+1). \end{aligned}$$

For fixed $0 < \alpha < 1/2$ and $R = \lfloor \alpha n \rfloor$, the original family hypotheses therefore imply the analytic bound

$$\liminf_{n\to\infty} \frac{\dim F_n}{n^D} \geq (2\alpha)^D \left[1 - \frac{m}{c}\, 2D \left(1 - \frac{\sin(2\pi\alpha)}{2\pi\alpha}\right)\right].$$

Its elementary quadratic lower estimate replaces the bracket by $1 - 4\pi^2 Dm\alpha^2/(3c)$. This improves the low-frequency quadratic coefficient of the maximum-coordinate estimate by a factor three.

Proof. The positive first coordinate contributes $4A_R/R$ to the average; each of the $D-1$ centered coordinates contributes $8A_R/(2R+1)$. Both real quadratures occur in numerator and denominator. The cosine sum follows by summing $2\sin x \cos(2rx) = \sin((2r+1)x) - \sin((2r-1)x)$ and telescoping; $\sin^2(rx) = (1 - \cos(2rx))/2$ gives the formula for A_R. Bounding $\sin^2(rx) \leq r^2x^2$ and summing r^2 gives the finite upper bound, including its lower-order terms. Floor inequalities give $H/n^D \to (2\alpha)^D$, and $\sin x/x \to 1$ gives $B_{D,R,n} \to 2D[1 - \sin(2\pi\alpha)/(2\pi\alpha)]$. Divide the finite rank bound by n^D and take the lower limit. Negative right sides are merely weak bounds. For $m = 0$, taking $\alpha \uparrow 1/2$ gives lower density one. □

These are classical Fourier and finite trace arguments, with the standard discrete-torus spectrum already present in the literature [13]. The exact average and its improved constant are analytic calculations here, not extra formal endpoints or sharp optimal observation-rank claims. They do not identify an interacting background or a continuum Yang–Mills spectral gap.

12 Why dimension is not a sufficiency criterion

The following analytic statement is included to prevent the reverse inference from a necessary rank budget. It is not a new Lean-verified endpoint here. The proof below is the classical singular-Schur range condition [6, 7], expressed as boundedness below on each measurement fiber; we claim no new general matrix criterion. Let E be a finite-dimensional real inner-product space,

$H = H^*$ and $B = H - cI$. Fix linear $C : E \to F$, and decompose $E = N \oplus W$ with $N = \ker C$, $W = N^\perp$. In this decomposition write

$$B = \begin{pmatrix} A & D \\ D^* & G \end{pmatrix}.$$

Proposition 12.1. *A finite real-valued function Φ with $\Phi(0) = 0$ can make $\langle x, Bx\rangle + \Phi(Cx) \geq 0$ for every x if and only if $A \succeq 0$ and $\ker A \subseteq \ker D^*$. In that event a quadratic penalty suffices.*

Proof. Taking $x \in N$ proves $A \succeq 0$. For $z \in \ker A$ and fixed $w \in W$, the value at $tz + w$ contains $2t\langle z, Dw\rangle$, while $C(tz + w) = Cw$ is fixed. Boundedness below for all real t forces this coefficient to vanish.

Conversely, $\ker A \subseteq \ker D^*$ implies $D(W) \subseteq \operatorname{ran} A$. Let $A^\dagger$ be the inverse on $\operatorname{ran} A$ and zero on $\ker A$. Completing the square gives

$$\langle n + w, B(n + w)\rangle = \langle n + A^\dagger Dw, A(n + A^\dagger Dw)\rangle + \langle w, (G - D^*A^\dagger D)w\rangle.$$

The restriction $C|_W$ is injective; extend its inverse on $C(W)$ to a linear map $T : F \to W$. The quadratic function $\Phi(y) = \langle Ty, (D^*A^\dagger D - G)Ty\rangle$ cancels the second term and vanishes at zero. If F has a chosen Euclidean norm, a sufficiently large nonnegative multiple of $\|y\|^2$ also suffices, by finite-dimensional norm bounds. □

For a concrete counterexample let $H = \operatorname{diag}(0, 2)$, $c = 1$ and $C(x, y) = x + y$. On $\ker C$, $\langle x, (H - I)x\rangle = 0$, and C detects the negative eigendirection of $H - I$. Nevertheless, on $(t, 1 - t)$ the repaired form is $1 - 2t + \Phi(1)$, unbounded below for every finite $\Phi(1)$. Kernel nonnegativity and low-mode detection alone do not suffice; cross-coupling must be checked.

13 Operator scaling and the actual consumer

In the repository, the flat fixed-volume covariance endpoint uses $K_0 + aQ^*Q$ with the physical counting ℓ^2 norm and the same unscaled flat Hodge operator. More precisely, the declared precision shell includes $-\Sigma$. Specializing every perturbation and its summable majorant to zero and taking $a > 0$ gives the available certified coercivity lower bound $\min(1, a)/C_P$, with $C_P > 0$ from the full-periodic block-Poincare estimate. This is not a claim that this value is optimal. At fixed d, N_c, N' the block codomain is fixed and finite. Inspection therefore identifies the unscaled flat base route with the setting of Theorem 6.1, with $\Phi_L(y) = a_L\|y\|^2$ and no fine repair. Even allowing arbitrary scalar a_L or rescaling Q_L does not yield a positive uniform full-space constant. This is a source-level specialization in this text; no new Lean adapter to that consumer is claimed. At the new source anchor below, the module `PhysicalGaugeFixedPrecision` defines `flatGaugeFixedPrecisionCLM`; its explicit-Poincare coercivity theorem is recorded by its full name and specialization in the linked declaration map.

The inspected endpoint is explicitly fixed-volume. Its positive rate may depend on L, N' and C_P, so there is no contradiction. Nor does the obstruction exclude all possible uniform Combes–Thomas arguments. Its locality radius and matrix-entry bounds also vary with L; coercivity alone would not settle the rate. An interacting Hessian, a growing measurement budget, a different sector norm or a physically rescaled operator must each be audited separately.

In particular, writing $\ell = L + 1$ so that $n_L = \ell N'$, replacing K_0 by $\ell^2 K_0$ changes the band threshold exactly to $4\pi^2 dR^2/(N')^2$. The proof above no longer drives it below every fixed $c > 0$.

This observation is an escape from this proof, not an established coercivity theorem for the programme's physical operator. No ultraviolet activity estimate, continuum limit, Wilson-Hessian identification or Yang–Mills mass gap is obtained.

14 Formal evidence and publication context

The source-tree anchor for the new physical-band and uniform-rank results is

`e31b428917aa767b4a1aad097db5cf92fb9cd23e`.

The original factorized-kernel and repair development in Section 3 instead has the distinct tested source anchor

`7b25bfb39357406cf61880e3ac0428d7d1ba3449`.

Its exact modules are not present in the new main tree merely because a reproduction supplement is present there. The two full preserved focal runs at that older SHA agree in compiled object, oracle log and source manifest; their six named reports use only the three standard axioms. The reproduced factorized endpoint is dimension four. The general-dimensional witness calculation, trigonometric closed form, spectral multiplicity, liminf and witness-line statements are analytic here; they are not all silently included in that formalization claim. The external declaration map keeps both anchors. The public source repository is THE-ERIKSSON-PROGRAMME on GitHub. The preserved historical v5 map is `declaration-map.json` and the preserved evidence is `reproducibility/`. Exact source SHAs, not moving branch heads, identify the formal claims. That parent map does not certify the new full volume extension. The map packaged with this manuscript records each added anchor and the separately reproduced full-tail extension. The accompanying exact-artifact report binds the final source and PDF hashes.

The geometric and Fourier modules are ordinary proof dependencies, not assumptions about a newly supplied band. Each previously verified stage identified below has complete same-source fresh-clone Colab evidence; the new extension boundary is stated separately. For the uniform endpoints the tested source is

`17ca395d4e9de992b1c575454a112da9e6f97e86`.

Its two runs agree in compiled object, raw oracle log and source-manifest hashes. The standalone budget lemma uses `propext, Quot.sound`; the other four uniform-module reports also use `Classical.choice`. No extra axiom is reported. Actual-log acceptance and three real-field rejection tests pass under normal and optimized Python. These tests concern evidence handling, not mathematical independence.

The additional whole-space gradient/Hodge identity and constructed weighted repair bounds use tested source

`dfd49d0b0730a4d36f58a4cbb7a17d6a41fb60bc`.

The full-space tangent endpoints in Section 7 use tested source

`81cf07b2e0d30543682c4073a4a2d82950d638b7`.

Each has two complete same-source fresh Colab Linux reconstructions with matching compiled object, raw oracle and LF-source-manifest hashes. For the tangent source, all three named endpoints use only `propext`, `Classical.choice` and `Quot.sound`. Its oracles and pinned Colab

reproduction drivers are separate from the manuscript render. Failed source candidates and their full logs are retained in the programme evidence archive. Later, separately reproduced sources formalize the actual implicit-function curve construction at

`ef81f028ce6e7e8aaca7e9229eb00aa683752a41`.

The local quadratic-energy regular-sector obstruction uses

`9bb90cfd77ec5cea5c4447667fe8baf8c21429a1`.

The displayed sector proofs remain readable analytic arguments; their uniform consequence is not silently included in those four-endpoint reports. No theorem about a nonlinear sector is inferred from a full-space no-go alone.

Short labels and an external declaration map replace long wrapped Lean names in this manuscript. The companion source map records exact names, source SHAs, and evidence directories. Reproduction commands remain in the existing copyable notebooks rather than being split into misleading line fragments. The fixed-map Schur criterion, its counterexample calculation and consumer specialization above are analytic or source-inspection arguments; the whole manuscript is not described as machine checked.

The actual mean-weight finite/asymptotic theorem of Section 9 uses tested source

`ac60c0df885d8866329a839e2b4e73c90595d806`.

Its two own fresh Colab Linux clones 076/076b pass six exact standard-axiom reports and all nine stages. Compiled module SHA-256 is

`28fcc004795a1d8ab74ca87b5e67cdf90f557551bc02b1bb912da38324c26e50`.

Raw oracle and LF-source-manifest hashes also match; full records and the strict mutation-tested validator are preserved in the mean-weight reproduction package. The proof composes the actual character strip, its exact trace, positivity, local derivative-kernel bound and multiplier estimate; these are derived lemmas, not desired conclusions supplied as hypotheses. Nonnegative weights, fixed positive c and the quadratic penalty upper bound remain. The constructed quadratic-weight mean observation instead uses tested source

`7f37d6437a8f69a1d3fddf3dc0bcfad74c6bf9b6`.

It has two same-source fresh-clone reproductions 067/067b with six exact standard-axiom reports and matching compiled, raw-oracle and LF-source hashes. The formal output is an actual fixed finite-dimensional type; the displayed numeric dimension and nonoptimal path explanation are not additional oracle endpoints. The penalty examples and regularity tradeoff are analytic only. Classical quadratic-form, Fourier and trace mechanisms, as discussed in [10], establish no new general principle or priority here.

The signed-weight results in Section 10 have distinct exact formal anchors. The positive-part construction, actual energy order and finite/asymptotic sufficient condition use

`810785b17b2f1bb14dd6caa5435865eb51a86ef3`.

Its fresh own Colab Linux clones 077/077b pass six exact standard-axiom reports. The actual zero-signed-mean fixed-output family instead uses

`497e3b4ef4046cee25b1069b8aa0fd8c0cfaa178`.

Its clones 079/079b pass ten exact reports, including actual Hodge coercivity and the constructed family on one fixed finite output type. The individual square-cycle order, selected positive-part mean and rank consequence use

`4c5728d147721523f5a521011351ac4bf2d3cf86`.

Its clones 081/081b pass ten exact reports. In every pair the compiled module, raw oracle and LF-source-manifest hashes match, with real mutation controls under normal and optimized Python. The declaration map records exact names and hashes, and links the complete reproduction packages. Failed probes 078 and 080, their sources, full diagnostics and corrections remain preserved; no passing result is inherited across source changes.

The individual cycle has zero positive-part means in other pairs. Its formal selected-pair identity is not an all-directions counterexample. The finite sum covering all off-diagonal pairs, its displayed mean calculation, the negative hyperplane diagnostic, numerical output dimension, normalized critical-mean limit interpretation and full-state identity-observation non-vacuity check are analytic here. They are not additional oracle endpoints. No density conclusion from ordinary signed spatial means alone has been established. The code indexes positive side length by $k + 1$, equivalent to the positive integer indexing used in the text.

14.1 Evidence boundary for the full volume extension

The actual normalized image and genuine compression trace are separately reproduced at source `9a8946b7`, the whole-image support and constructed positive total envelope at `88445ac6`, and the original local nonlinear finite rank composition at `3a2d5305`. The actual numerical bound and constructed uniform dilation density have two matching own Colab reproductions at `031ca274`, with five exact standard-axiom reports and full manifests. These are distinct source artifacts, not grades inherited by this manuscript.

The full integer-tail source `067fed75` now has two fresh own Colab reproductions, 094/094b, with all nine stages exiting zero and five exact standard-axiom reports. Their compiled module, raw oracle and LF-source manifest hashes match. Real mutation controls pass under both normal and optimized Python. The earlier browser loss and absent saved output are preserved separately; no outcome is reconstructed for that earlier event. The live session manager reported no active sessions before and after these fresh runs. This is execution reproducibility, not independent scientific review. The exact averaged multiplier and sinc limit remain analytic calculations. Preparation and rendering alone do not constitute a completed exact-artifact review.

Historical model assessments and review context

This candidate was prepared by the Codex agent in the author's continuing programme chat, with access to earlier objections, source audits and supplied reproduction records. The served backend model identifier and reasoning setting are not exposed, and no technical isolation or scientific independence is claimed for this preparation.

The public v2 of this record has canonical PDF SHA-256

`b37f2fe86b7a8a526ad02f8b28491c300a9ff72d5653536bc5532e8c9468ed8f`.

A subagent explicitly authorized by the author and configured as `gpt-6-astra` assessed that exact eight-page artifact after first preserving the earlier candidate's objections. Actual backend

identity and reasoning setting were not exposed. The workflow used the same account, provider and programme context; it was not blind or organizationally independent. The complete earlier and corrected-artifact reports remain in the v2 package. The corrected-artifact report regarded O1–O5 as resolved and retained O6: no independent complete Lean rebuild or dependency-closure audit. No numerical score or human peer review was supplied.

A historical public artifact is AIRR v5, a 33-page Working paper released on 4 November 2026 under CC BY 4.0, with CC0 metadata. Its canonical PDF SHA-256 is

`69d802bfea8f8f23bfe6aba03cdc4835feac777e024c939040c7f099a11462c8`.

It is public and citable, but has not completed AIRR screening and is not editorially accepted or human peer reviewed. Earlier versions and their complete reports remain preserved. The author is also the AIRR founder/operator; no independent editorial decision is inferred from publication.

The historical revision136 corrected verified weight-symbol and cross-reference defects in that artifact; those changes alone were editorial. This consolidated successor additionally supplies the classical local and geometric proofs, corrects a false one-cluster-per-leaf inference, and incorporates the separately scoped actual physical embedding and exact finite penalty calculation. A newly rendered PDF will be a different artifact: its hash, visual review, assessment and any necessary deposit permission must be recorded separately. No public version number, private submission number, license grant, historical review score or acceptance is assigned by this source.

15 From necessary rank budgets to sufficient observations

The preceding sections establish necessary information budgets and their formal and analytic boundaries. The continuation below asks when actual coefficient distributions suffice to construct observations of the same order. It is integrated into this revision, not a second paper or deposit. For these sections set $D = d \geq 3$ and $g = N_c^2 - 1$, $h = Dg$. The symbol $\mu \geq 1$ denotes the scalar coefficient mean budget here; it does not replace the earlier symbol $m = N_c^2 - 1$ for internal dimension. The field is denoted U instead of A. All norms and original forward-edge occurrences use the same counting, unit-spacing flat convention. Restrictions on scalar coefficients and the distinction between spectral coordinates and rectangular means are stated explicitly below.

Historical main-text and continuation artifacts were assessed on different PDF bytes. The historical integrated v5 has a public exact-artifact AI report, described in Section 26. That report is historical evidence for v5, not an assessment of this changed candidate. A first exact-PDF review must follow rendering and inspection of this candidate.

16 Actual objects and the original observation problem

Fix $D \geq 3$, $N_c \geq 2$, $c > 0$, and the unit-spacing torus $X_n = (\mathbb{Z}/n\mathbb{Z})^D$. The actual real adjoint coefficient space has dimension $g = N_c^2 - 1$; there are D field components, so write $h = Dg$. For the actual forward difference $d_i u(x) = u(x + e_i) - u(x)$, take one scalar total coefficient $a(x) = 1 + w_{i,x,j} > 0$ shared by all derivative directions and field components. Use counting norms and the actual componentwise flat energy

$$Q_w(U) = \mathcal{E}_a(U) = \sum_{x,i,j} a(x) \|d_i U_j(x)\|^2. \tag{14}$$

The existing flat Hodge identity identifies these forms; its formal source is listed in Section 25. All continuation adapters remain analytic. There is no substitution of a covariant or interacting derivative in (14).

The original local data consist of a differentiable map C at zero, with $C(0) = 0$, into an actual finite-dimensional real normed space F, and a penalty satisfying $\Phi(y) \le K\|y\|^2$ near zero, such that

$$c\|U\|^2 \le Q_w(U) + \Phi(C(U)) \tag{15}$$

for sufficiently small actual divergence-zero fields. "Local" here means a configuration-space neighbourhood; spatial support of C is an additional condition. The constant c is fixed, while observation maps and neighbourhoods may depend on n.

A useful necessary condition follows directly. If an actual J-dimensional divergence-zero subspace has $Q_w(U) \le \lambda\|U\|^2$ with $\lambda < c$, then $\dim F \ge J$. Otherwise a nonzero U in that subspace lies in $\ker DC(0)$. Along tU, $C(tU) = o(t)$ and the penalty is at most $o(t^2)$. Dividing (15) by t^2 and taking $t \to 0$ contradicts $\lambda < c$. No rank or band certificate is supplied as an extra hypothesis in this argument.

17 The critical global moment theorem

Put $p = D/2 > 1$. For $\mu \ge 1$ suppose that

$$n^{-D} \sum_{x \in X_n} (\mu/a(x))^p \le M < \infty. \tag{16}$$

Theorem 17.1. *Under* (16)*, there is an actual real linear spectral observation C such that*

$$c\|U\|^2 \le Q_w(U) + c\|C(U)\|^2 \quad \textit{for every actual cochain } U,$$

and

$$\dim F \le C_* \max\{1, (n/\sqrt{1+\mu})^D\}, \tag{17}$$

where, with $I_p = (p-1)^2/[p(p+1)]$ and $\Gamma_p = \frac{p}{p-1} I_p^{-(p-1)}$, one safe choice is

$$C_* = h\left[1 + \Gamma_p (3/4)^D D^{p+1} c^p M 2^p\right].$$

Thus $\Phi(y) = c\|y\|^2$, $K = c$ and any positive neighbourhood radii supply the original local data. In particular the assertion holds on the actual divergence-zero subspace.

The original subclass $1 \le a \le 2\mu$ with exact arithmetic mean μ is included, but those extra normalizations are not needed for sufficiency. The distribution condition (16) remains additional; it is not derived from an arithmetic mean.

17.1 A finite normalized Fourier lemma

For a finite abelian group of size V, let $\mathcal{F}$ be its counting-unitary Fourier matrix, whose entries have modulus $V^{-1/2}$. For a frequency multiplier g put

$$L = V^{-1} \sup_{t>0} t^p \#\{k : |g(k)|^2 > t\}.$$

For a nonnegative matrix B, write $\mathsf{n}(s, B)$ for its number of eigenvalues strictly greater than s, counted with ordinary multiplicity.

Lemma 17.2. *For $K = M_f \mathcal{F}^* M_g \mathcal{F}$ and $s > 0$,*

$$\mathsf{n}(s, K^*K) \le \Gamma_p L s^{-p} \sum_x |f(x)|^{2p}.$$

The trace and eigenvalue counts are not volume-normalized.

Proof. Remove phases by unitaries and set $A = \mathcal{F}^* M_{|g|^2} \mathcal{F}$. If $L = 0$ the assertion is immediate. Otherwise let $B_0 = pL/(p-1)$ and let A^+ be the inverse on $\operatorname{ran} A$, zero on its kernel. For $0 \le \gamma \le A$ write $\rho(x) = \langle \delta_x, \gamma \delta_x \rangle$. Positivity makes γ annihilate $\ker A$. On the range let $P_E = \mathbf{1}_{\{A^+ \le E\}}$ and $R_E = \mathbf{1}_{\{A^+ > E\}}$. The weak-frequency tail gives

$$\langle \delta_x, P_E A P_E \delta_x \rangle = V^{-1} \sum_{|g(k)|^2 \ge 1/E} |g(k)|^2 \le B_0 E^{p-1}.$$

Indeed the tail above t is t times its count plus the integral of counts above t; the resulting bound is $VLp\,t^{1-p}/(p-1)$. Non-strict thresholds follow by limits. The triangle inequality for $\gamma^{1/2}\delta_x$ split by $P_E + R_E$ yields

$$\langle \delta_x, R_E \gamma R_E \delta_x \rangle \ge \big(\sqrt{\rho(x)} - \sqrt{B_0} E^{(p-1)/2}\big)_+^2.$$

Exact spectral layer integration consequently gives

$$\operatorname{Tr}(A^+ \gamma) \ge I_p B_0^{-1/(p-1)} \sum_x \rho(x)^{p/(p-1)}. \tag{18}$$

The scalar integral reduces to $\int_0^1 (1 - t^{(p-1)/2})^2 dt = I_p$. Now take P to be the spectral projection of $A^{1/2} M_{|f|^2} A^{1/2}$ above s, with rank $\nu > 0$, and set $\gamma = A^{1/2} P A^{1/2}$. Then $0 \le \gamma \le A$, $\operatorname{Tr}(A^+\gamma) \le \nu$, and

$$s\nu < \sum_x |f(x)|^2 \rho(x) \le \Big(\sum_x |f(x)|^{2p}\Big)^{1/p} \Big(\sum_x \rho(x)^{p/(p-1)}\Big)^{(p-1)/p}.$$

Insert (18) and raise to p. The constant becomes $B_0 I_p^{-(p-1)} = \Gamma_p L$. This also explicitly handles a zero-frequency kernel. □

The density method is classical and follows the Rumin–Frank approach [16, Section 2]. The proof above supplies its finite normalization and kernel adapter; it is not a claim of a new Cwikel mechanism or independent discovery.

17.2 Actual lattice frequencies and the weighted inverse

For $n \ge 2$ let $L_0 = d^*d$ and work first on the actual counting-mean-zero scalar space H_0. Its Fourier symbols are

$$z_i(k) = e^{2\pi \mathrm{i} k_i/n} - 1, \qquad \lambda(k) = 4 \sum_i \sin^2(\pi k_i/n).$$

Choose centred integer representatives r_i with $|r_i| \le n/2$, using only one representative at an even endpoint. Concavity of sine implies

$$\lambda(k) \ge 16|r|^2/n^2, \qquad |g_i(k)| := |z_i(k)/\lambda(k)| \le n/(4|r|), \quad k \ne 0.$$

Set $g_i(0) = 0$. If $|g_i| > \tau$, then $0 < |r| < R = n/(4\tau)$. For $R \le 1$ there are none, and for $R > 1$ there are at most $(3R)^D$ representatives. Thus

$$n^{-D} \sup_{t>0} t^{D/2} \#\{|g_i|^2 > t\} \le (3/4)^D. \tag{19}$$

This is uniform in n, including $n = 2$, not an imported infinite-volume bound.

Let $L_a = d^* M_a d$. On H_0 it is positive definite. For $f \in H_0$, the flow $j_* = M_a d L_a^{-1} f$ minimizes $\langle j, M_{a^{-1}} j\rangle$ among $d^* j = f$: write $j = j_* + k$, $d^* k = 0$, and the cross term is $\langle dL_a^{-1} f, k\rangle = 0$. Inserting the explicit feasible flow $j = dL_0^{-1} f$ gives the actual quadratic-form comparison

$$L_a^{-1} \le \sum_{i=1}^{D} T_i^* T_i, \qquad T_i = M_{a^{-1/2}} d_i L_0^{-1}. \tag{20}$$

Extend L_0^{-1} by zero on constants. Lemma 17.2 and (19) bound each $\mathsf{n}(s, T_i^* T_i)$ by

$$\Gamma_p (3/4)^D s^{-p} \sum_x a(x)^{-p}.$$

For nonnegative matrices, the count of their sum above t is no larger than the sum of their counts above t/D: intersect the complements of their high spectral subspaces and apply min–max. No commutation is required. With $t = 1/c$ in (20), inversion on H_0 and the single scalar constant mode yield

$$\#\{\lambda(L_a) < c\} \le 1 + \Gamma_p (3/4)^D D^{p+1} c^p M n^D / \mu^p. \tag{21}$$

No spectral-band dimension is assumed in this derivation.

Completion of Theorem 17.1. Choose an actual real counting-orthonormal eigenbasis of the real symmetric matrix L_a. Retain the coordinates with eigenvalue $< c$, separately in each of the h real field coordinates. Complexification used in the Fourier proof preserves real eigenvalue multiplicities; it does not double them. Energy controls coefficients at eigenvalues $\ge c$, while $c\|C(U)\|^2$ controls the rest. Equation (21), $p = D/2$, and $(1 + \mu)/\mu \le 2$ give (17). For $n = 1$ take the identity on all h coordinates. This constructs the actual global spectral observation. □

18 Every lower positive inverse exponent is insufficient

Theorem 18.1. *For every fixed $0 < r < D/2$, the analogue of Theorem 17.1 is false uniformly under the conditions*

$$\mathrm{mean}((\mu/a)^r) \le 2, \qquad 1 \le a \le 2\mu, \qquad \mathrm{mean}(a) = \mu.$$

Proof. Choose a fixed integer $\ell > 3D/c$, and put $P = 2\ell + 2$, $B = (2\ell + 1)^D$. The triangular window $t(s) = (1 - |s|/\ell)_+$ has finite sums

$$T = \sum t^2 = (2\ell^2 + 1)/(3\ell), \quad V = \sum (d^- t)^2 = 2/\ell, \quad W = \sum (d^+ d^- t)^2 = 6/\ell^2.$$

For a product window $\psi_y(x) = \prod_i t(x_i - y_i)$ and nonzero actual adjoint H, the actual curl $U_1 = d_2^- \psi_y H$, $U_2 = -d_1^- \psi_y H$, other components zero, has backward divergence zero. Its field

support and all nonzero forward-gradient start sites lie in the box $y + [-\ell, \ell]^D$. Direct separable sums give

$$\|U\|^2 = 2VT^{D-1}\|H\|^2, \qquad \frac{\mathcal{E}_0(U)}{\|U\|^2} = \frac{W}{V} + (D-1)\frac{V}{T} = \frac{3}{\ell} + \frac{6(D-1)}{2\ell^2+1} \le \frac{3D}{\ell} < c.$$

Take sufficiently large multiples n of P, set $\mu = n^2$ and $J = \lfloor n^{D-2r}/B \rfloor$. These J boxes fit on the separated P-grid eventually, since $P^D \le Bn^{2r}$. They satisfy $J \ge n^{D-2r}/(2B) \to \infty$. Put $s = BJ/n^D \le \mu^{-r}$ and set $a = 1$ on their union, $a = A = (\mu - s)/(1 - s)$ elsewhere. Eventually $s \le 1/2$, so

$$\operatorname{mean}(a) = \mu, \quad 1 \le a \le 2\mu, \quad A \ge \mu, \qquad \operatorname{mean}((\mu/a)^r) \le s\mu^r + 1 \le 2.$$

Each curl's gradients start only in its weak box. The actual J-dimensional span has that same Rayleigh quotient, because both field and gradient supports are disjoint. Section 16 forces $\dim F \ge J$, whereas the target maximum at $\mu = n^2$ is 1. This proves the refutation using original nonlinear data. □

Together the two theorems give the sharp inverse-moment exponent for this specified family of global scalar conditions. They do not identify all sufficient coefficient criteria, prove a physical coefficient hypothesis, or impose spatial locality on C.

19 Spatial means and the critical dyadic obstruction

For a partition $\mathcal{P}$ into integer rectangles, its actual normalized means are

$$C_{\mathcal{P}}(U)_{R,j} = \sqrt{|R|}\,\operatorname{mean}_R U_j, \qquad \dim F = h\#\mathcal{P}.$$

A full dyadic cube tree splits each selected cube into all 2^D equal children and retains unsplit siblings. The origin of the torus coordinate cube may be translated. This class excludes arbitrary anisotropic rectangle cuts and compressed observations.

Proposition 19.1. *There are fixed critical-moment profiles with target maximum* 1 *for which every coercive full dyadic cube-tree mean partition has* $\Omega(\log n)$ *leaves, even after choosing a translated origin. The same profiles admit an explicit constant-size rectangular mean partition satisfying the global inequality for every actual cochain.*

Proof. Keep $\ell > 3D/c$ fixed. Choose a dyadic $L = 2^q \ge 4\ell + 2$, put $\gamma = \max(1, c/4)$, $n = 2^J$ with $J \ge q + 1$, $\mu = \gamma n^2$, and $S = \{0, \dots, L-1\}^D$. Set $a = 1$ on S and $a = A = (\mu - s)/(1 - s)$ elsewhere, $s = L^D/n^D$. Then $1 \le a \le 2\mu$, its mean is exactly μ, and

$$\operatorname{mean}((\mu/a)^{D/2}) \le 1 + L^D\gamma^{D/2} =: M_0.$$

All these parameters except n are fixed, and the target maximum is 1.

In the hypothetical complete level $j_* = J - q$ of any translated dyadic grid, cells have side L. Each coordinate interval of S meets at most two such cyclic grid intervals, so one intersection has integer length at least $L/2 \ge 2\ell + 1$. Taking these intersections independently places a whole gradient-start box for a compact curl inside S and inside one level-j_* cell. If the proposed partition had maximum depth at most j_*, that cell would lie in one actual leaf. The curl then

has zero means on every leaf, since each component has total sum zero, yet its weighted Rayleigh quotient is $< c$. This contradicts coercivity. A single zero-extended window needs no packing divisibility assumption on n; $n \geq 2L$ suffices here. The field may depend on the partition, since coercivity is asserted for every field. Thus maximum depth is at least $J - q + 1$, and the full-arity tree identity gives

$$\#\mathcal{P} \geq 1 + (2^D - 1)(J - q + 1). \tag{22}$$

For a successful unrestricted partition, isolate all L^D vertices of S as singletons, and use the D disjoint complement rectangles

$$R_i = [0, L-1]^{i-1} \times [L, n-1] \times [0, n-1]^{D-i}.$$

Coverage follows from the unique first coordinate outside $[0, L-1]$. Every internal edge starts at a strong vertex with coefficient A. On a length-l interval, counting variance and path Cauchy–Schwarz give

$$\sum \|u - \bar{u}\|^2 \leq \frac{l(l-1)}{4A} \sum A \|du\|^2.$$

Indeed the path resistance is at most $(l-1)/A$, and a cut at t is crossed by $t(l-t) \leq l^2/4$ pairs in the variance identity. Product-variance tensorization on original fibres yields a rectangle constant $n^2/(4A)$, with no extra factor D. Singletons have zero residual, internal edge sets are disjoint, and omitted edges have positive energy. Since $cn^2/(4A) \leq c/(4\gamma) \leq 1$, exact orthogonal decomposition proves

$$c\|U\|^2 \leq Q_w(U) + c\|C_{\mathcal{P}}(U)\|^2, \qquad \#\mathcal{P} = L^D + D.$$

These rectangles have potentially unbounded aspect ratio; they leave the cube-tree class. □

The general critical moment question for unrestricted rectangular means remains open. Changing only the stopping rule within full cube trees cannot resolve it. Classical weighted Sobolev inequalities [17] also yield block observations for inverse moments of finite order strictly greater than $D/2$; the dossier contains the explicit balanced-box proof. Proposition 19.1 explains why its critical tree-counting loss cannot be removed merely by choosing a better cube-tree stopping criterion.

20 What a spatial observation proof must still establish

The preceding positive spectral result does not construct a rectangular mean map. Three further checks prevent substitutions for that missing construction. They concern local geometry, the penalty coefficient, and support of observations, respectively. They are collected here as limitations of the same dossier, not as additional papers.

20.1 Critical inverse mass does not remove rectangle geometry

For a rectangle R of side lengths l_i, set $V = \prod_i l_i$, $L = \max_i l_i$, $\kappa = L^2/V^{2/D}$ and $\eta_R = \sum_R a^{-D/2}$. The endpoint Sobolev estimate on balanced cubes, with $q = 2D/(D+2)$, transfers to arbitrary rectangles by repeated coordinate reflection and gives

$$\sum_R |u - \bar{u}_R|^2 \leq A_D D^{2/D} \kappa \eta_R^{2/D} \mathcal{E}_{a,R}(u). \tag{23}$$

Here the energy uses original internal edges and their actual starting-vertex weights. The full reflection multiplicities and graph adapters to the classical estimate [17, Section 2.2] are proved here in Section 22; this is an analytic adapter, not a new Lean theorem.

The factor cannot be discarded even for uniform coefficients. With $u(x) = x_i - (L-1)/2$ along a longest side, direct sums give

$$\sum_R u^2 = V(L^2-1)/12, \qquad \mathcal{E}_{a,R}(u) = aV(L-1)/L.$$

Their ratio is $L(L+1)/(12a) \geq \kappa \eta_R^{2/D}/12$. On the thin rectangle of lengths $(L, 1, \ldots, 1)$, choosing $a = L^{2/D}$ makes $\eta_R = 1$ while this ratio diverges. Thus compressing a cube-tree chain into thin rectangles cannot simply reuse a shape-independent inverse-mass certificate. This prescribed-block test does not defeat every global partition: omitted crossing edges and a different partition may help. The successful partition in Proposition 19.1 is not contradicted.

20.2 Fixed and free penalties are different obligations

Let C be normalized means on a fixed partition of a finite real Euclidean space. Then $CC^* = I$; $P = C^*C$ is the orthogonal block-mean projection and $Q = I - P$. For the actual positive stiffness L_a, write the blocks on $\operatorname{ran} P \oplus \operatorname{ran} Q$ as

$$A = PL_aP, \qquad B = PL_aQ, \qquad D_0 = QL_aQ.$$

Positivity implies $\ker A \subseteq \ker B^*$: otherwise the energy of $tp + v$ would have a nonzero linear coefficient with zero quadratic coefficient in t. Completing the square with the finite spectral pseudoinverse $A^\dagger$ shows

$$L_a + cP \geq cI \quad \Longleftrightarrow \quad D_0 - B^*A^\dagger B \geq cI \text{ on } \operatorname{ran} Q. \tag{24}$$

Indeed the first inequality is exactly $L_a - cQ \geq 0$, and its quadratic form is $\langle p + A^\dagger Bv, A(p + A^\dagger Bv)\rangle + \langle v, (D_0 - cI - B^*A^\dagger B)v\rangle$. This is classical shorted-operator algebra [15]; it supplies no spatial partition or target count. In particular $D_0 \geq cI$ alone is insufficient.

Proposition 20.1. *On the four-cycle with edge weights $(1, 6, 4, 1)$ and blocks $\{0, 1\}, \{2, 3\}$, coercivity holds on zero block sums with $c = 19/5$, but $L_a + cP \geq cI$ fails.*

Proof. For $v = (x, -x, y, -y)$, with the counting Euclidean norm,

$$\mathcal{E}_a(v) - \frac{19}{5}\|v\|^2 = \frac{17}{5}\left(x + \frac{35}{17}y\right)^2 + \frac{84}{85}y^2 \geq 0.$$

The actual normalized observation is

$$Cu = \left((u_0 + u_1)/\sqrt{2},\ (u_2 + u_3)/\sqrt{2}\right).$$

For $u = (19, -9, -5, -5)/14$, direct exact evaluation gives

$$\mathcal{E}_a(u) = 52/7, \quad \|u\|^2 = 123/49, \quad \|Cu\|^2 = 25/49,$$

and hence $\mathcal{E}_a(u) + (19/5)\|Cu\|^2 - (19/5)\|u\|^2 = -6/35$. The weights have mean 3 and lie between 1 and 6. Ten focal Lean theorems preserve the constructed weights, normalized observations, projection identities, kernel square identity and exact negative witness; Section 25 identifies the successful immutable source. Those ten endpoints concern the scalar four-cycle. Section 24 separately derives its actual physical embedding and the all-field kernel premise; only its stated fixed-penalty endpoints have paired formal evidence. □

For the original class allowing an arbitrary finite K, put $J = D_0 - cI$. The necessary and sufficient condition is instead

$$\exists K \geq 0 : \; L_a + KP \geq cI \quad \Longleftrightarrow \quad J \geq 0 \text{ and } B \ker J = 0. \tag{25}$$

Necessity follows by testing $v \in \ker C$, then $v + tp$ with $v \in \ker J$. For sufficiency, $B \ker J = 0$ implies $B^* p \in \operatorname{ran} J$. Completing the square in $v + J^\dagger B^* p$ leaves $A + (K - c)I - BJ^\dagger B^*$, so $K = c + \|BJ^\dagger B^*\|$ suffices. In particular a proved all-cochain margin $D_0 \geq (c + \epsilon)I$ gives $K = c + \|L_a\|^2/\epsilon$. For actual $a \leq 2\mu$, the elementary edge estimate $|r - s|^2 \leq 2|r|^2 + 2|s|^2$ gives $\|L_a\| \leq 8D\mu$, including the parallel edge occurrences at $n = 2$. A margin $2c$ therefore permits $K = c + 64D^2\mu^2/c$, without increasing the number of observations. No margin or partition is proved by this conditional implication. Ordinary block averaging need not preserve divergence, so an estimate only on the physical kernel cannot silently replace the all-cochain hypothesis in this decomposition.

Even free K needs care at equality. On the same cycle let $s = \sqrt{85}$, $c_* = (17 - s)/2$ and $v = (7, -7, 6 - s, s - 6)$. Its kernel defect at c_* is zero, since for general $(x, -x, y, -y)$ it equals $(s - 6)[x + 7y/(s - 6)]^2$. For $p = (\delta/2, \delta/2, -\delta/2, -\delta/2)$ the full defect is

$$\mathcal{E}_a(v + p) + K\|C(v + p)\|^2 - c_*\|v + p\|^2 = (7 + K - c_*)\delta^2 - 10(13 - s)\delta.$$

For each finite $K \geq 0$ a sufficiently small positive δ makes it negative. This additional equality example is exact analytic algebra and finite diagnostic evidence, not one of the ten Lean theorems of Proposition 20.1.

20.3 Sparse point samples cannot replace stable averages

For uniform $a = \mu$, if an observation depends only on vertex values on a set S, $|S| = s$, and $C(0) = 0$, then the scalar field $u = 1$ off S and 0 on S has zero observation. At most $2Ds$ forward-edge occurrences cross S, so $\mathcal{E}_\mu(u) \leq 2D\mu s$. Consequently all-cochain coercivity requires

$$s \geq \frac{cn^D}{2D\mu + c}.$$

For an actual divergence-zero test, fix a component direction j and let the field be constant along every j-line, equal to zero on each line meeting S and one on the other lines, in one fixed internal coordinate. There are at most s hit lines. Its backward divergence is zero, its norm squared is at least $n(n^{D-1} - s)$, and transverse crossing energy is at most $2(D - 1)\mu n s$. Thus necessarily

$$s \geq \frac{cn^{D-1}}{2(D - 1)\mu + c}.$$

If every transverse line is hit, the second bound is immediate without a nonzero test. These bounds give the required obstruction: at $a = \mu = n^2$ and critical moment $M = 1$, bounded samples fail for all cochains when $D \geq 3$ and for these physical tests when $D \geq 4$. For the same uniform coefficients one global mean succeeds for $0 < c \leq 16$, by the nonzero Fourier gap $n^2\lambda(k) \geq 16$. This comparison rules out a point-sampling shortcut, not rectangular averages or singletons supplemented by complementary block means. A nonlinear local observation is tested on sufficiently small multiples of these fields.

21 Constructive rectangular bounds and their limits

Retain $\mu \geq 1$, $1 \leq a \leq 2\mu$, exact mean μ, (16), and $p = D/2$. All assertions in this section remain analytic.

21.1 A general threshold construction

For an interval of length l, the identity $\mathrm{Var}(u) = l^{-1} \sum_{x<y} |u(x) - u(y)|^2$ and path Cauchy–Schwarz give coefficient at most $l(l-1)/4$ per edge: each crossing pair has distance at most $l-1$, and at most $l^2/4$ pairs cross any edge. Product variance tensorization therefore yields, for a rectangle of maximal side L,

$$\sum_R |u - \mathrm{mean}_R\, u|^2 \leq \frac{L(L-1)}{4} \sum_{\text{internal forward edges of } R} |du|^2. \tag{26}$$

There is no additional factor D in this tensorization.

Fix $\tau \geq 1$. If $\tau < c/4$ use singletons. Otherwise set $\ell = \min\{n, \lfloor 2\sqrt{\tau/c} \rfloor\}$ and start with a grid whose sides have length at most ℓ. Isolate each vertex with $a < \tau$ as a singleton. Its current rectangle minus that vertex splits into at most $2D$ rectangles by the first coordinate differing from that vertex. Delete empty pieces and preserve earlier singletons. Remaining rectangles have $a \geq \tau$ and side at most ℓ. Equation (26), nonnegative omitted edges and normalized-mean orthogonality give

$$c\|U\|^2 \leq \mathcal{E}_a(U) + c\|C_{\mathcal{P}} U\|^2, \quad (C_{\mathcal{P}} U)_{R,j} = \sqrt{|R|}\, \mathrm{mean}_R\, U_j \tag{27}$$

for ALL cochains. Singleton residuals vanish, including $n = 1$; parallel edge occurrences for $n = 2$ are counted separately, and wrap edges can be omitted positively. With $B_c = (1 + \sqrt{c})^D$,

$$|\mathcal{P}| \leq B_c \max\{1, n^D/\tau^p\} + 2D\#\{a < \tau\}, \quad \#\{a < \tau\} \leq M n^D \tau^p/\mu^p.$$

The singleton branch satisfies this count since $B_c \tau^{-p} > 1$ there. Taking $\tau = \sqrt{\mu}$ yields the weaker general bound

$$\dim F \leq h(B_c + 2DM) \max\{1, n^D/\mu^{D/4}\}.$$

21.2 Every global threshold can lose the target order

Let $\gamma = \max\{1, c/4\}$, $n = t^2$, $t \geq 2$, $\mu = \gamma n^2$, $b = \gamma n$, $S = \{0, \ldots, t-1\}^D$, and $f = n^{-p}$. Set $a = b$ on S, $a = A = (\mu - fb)/(1-f)$ outside. Then $1 \leq a \leq 2\mu$, the mean is μ, and the critical moment is less than $2 - f$. The target maximum is one. For any global $\tau \geq 1$, consider the preceding algorithm with any translated initial grid or excision order, retaining grid boundaries and every excised singleton. If $\tau \leq b$, no vertex is weak and the grid volume bound gives

$$|\mathcal{P}| \geq n^D/(2\sqrt{b/c})^D = (c/(4\gamma))^p n^{D/2},$$

or the singleton branch is larger. For large n the grid side is less than n. If $\tau > b$, all $t^D = n^{D/2}$ weak vertices remain distinct singletons. Every such threshold therefore loses the target order.

For the SAME profile retain the whole weak cube and split its complement by the first coordinate outside it, giving $D + 1$ rectangles. On S, $ct^2/(4b) = c/(4\gamma) \leq 1$; outside, $A \geq \mu$ and $cn^2/(4A) \leq 1$. Equation (26) proves (27) for this constant-size partition. This is a mechanism obstruction, not a counterexample to every rectangular partition.

21.3 Supplied scaled cubes attain the target order

Suppose additionally that disjoint integer cubes S_k of side L_k, nonwrapping in one common root, satisfy for fixed $\beta \ge 1$

$$\begin{aligned} cL_k(L_k-1)/4 \le a(x) \le \beta L_k^2 \quad (x \in S_k), \\ a(x) \ge cn(n-1)/4 \quad (x \notin \cup_k S_k). \end{aligned} \tag{28}$$

These are extra coefficient and geometric assumptions. Each cluster charges at least β^{-p} to $\sum_x a(x)^{-p}$; thus $J \le \beta^p M n^D/\mu^p$.

Section 23 supplies a self-contained construction for precisely these disjoint integer cubes. Use quarter-offset solid cubes, then a compressed corner-potential subdivision. Its at most $2^{D+1}(2^D-1)J$ terminal rectangles have no cube corner in their interiors. This does not imply that a terminal rectangle meets only one cube: the proper-intersection and covering bounds give at most $T_D = 2^{2D-1}$ cubes per terminal rectangle. Refining by all their facet coordinates gives at most $(2T_D+1)^D$ subrectangles per terminal rectangle. Consistent boundary assignment followed by taking integer points yields an actual partition; every surviving rectangle is inside one supplied cube or outside all supplied cubes. Fragmentation is permitted. Consequently, with the explicit dimension-only constant

$$A_D = 2^{D+1}(2^D-1)(2^{2D}+1)^D, \qquad |\mathcal{P}| \le A_D \max\{1, J\},$$

$$\dim F \le hA_D(1+2^p\beta^p M)\max\{1, n^D/(1+\mu)^p\}.$$

The earlier one-cluster-per-leaf inference is not used. The multiple-object stress family and the boundary adapter are included in the cited section. Inside a cluster fragment maximal side is at most L_k, outside at most n. Equations (28) and (26) prove (27). Projection separation and a high-budget restriction are unnecessary. This uses classical geometry; it does not derive the supplied cubes from the original moment.

21.4 The exact cube decomposition is not universal

For fixed c, β take arbitrarily large integer $n \ge 3$, $\mu = n$, $W = \{e_1, -e_1\}$, $f = 2/n^D$, with $a = 1$ on W and $a = A = (n-f)/(1-f)$ elsewhere. Then $n < A \le 2n$, mean $a = n$, and the critical moment is less than $2n^{-p} + 1 - f \le 2$. Choose also $n > \max\{\beta, 8\beta/c, 1 + 8/c\}$ and $n \ge c$. Every vertex has $a < cn(n-1)/4$, so any family satisfying (28) must cover zero. A side-one cube containing zero would require $A \le \beta < n$. A cube of side $L \ge 2$ contains at least one of $e_1, -e_1$, in every translated root. Its lower bound forces $L^2 \le 8/c$, while its upper bound at zero requires $A \le \beta L^2 \le 8\beta/c < n < A$. Even allowing overlap cannot repair this exact decomposition.

The SAME profile nevertheless admits a strong grid with $\ell = \min\{n, \lfloor 2\sqrt{A/c} \rfloor\}$ and two singleton excisions. Equation (27) holds and

$$\dim F \le h\big[2^p(1+\sqrt{c})^D + 4D\big]\max\{1, n^D/(1+\mu)^p\}.$$

There are at most $\lceil n/\ell \rceil^D + 4D$ blocks. If $\ell < n$, $\ell \ge \sqrt{A/c}$; if $\ell = n$, there is one grid box. Remaining rectangles have coefficient A and weak residuals are singletons. The universal supplied-cube route is refuted. The original unrestricted rectangular question is neither proved nor refuted.

Self-contained proof continuations. The next three sections complete the local and supplied-geometry arguments and exhibit one actual physical observation map. The constants denoted A_D are local to their sections: the Sobolev and partition constants differ. Counting measure and unit lattice spacing are used throughout. These continuations do not derive supplied geometry from the inverse moment.

22 A finite proof of the critical rectangle estimate

This supplement derives the balanced-cube endpoint used by the rectangle adapter, rather than importing a local Sobolev estimate as an external input. The method is classical: finite entropy proves the projection inequality, which yields relative edge isoperimetry, layer sets, and a power argument. We claim neither a new general Sobolev principle nor a solution of the open critical rectangular-partition problem. This is part of the existing dossier.

Fix $D \geq 3$, and put

$$p = D/2, \quad q = \frac{2D}{D+2}, \quad q' = \frac{2D}{D-2},$$
$$\alpha = \frac{2(D-1)}{D}, \quad c_D = D(1-2^{-1/D}), \quad K_D = \frac{2\alpha D^{1/q'}}{c_D}.$$

An internal axis edge is an unordered adjacent pair, counted once. For its positive-axis orientation $e = (x, x+e_i)$, write $du(e) = u(x+e_i) - u(x)$.

22.1 Finite projection inequality, with proof

For a nonempty finite $A \subset \prod_{j=1}^{D} X_j$, let π_i delete coordinate i. Then

$$|A|^{D-1} \leq \prod_{i=1}^{D} |\pi_i A|. \tag{29}$$

Here is a finite proof, including the entropy facts needed for this use. For a finite probability vector r, define $H(r) = -\sum r \log r$ with $0 \log 0 = 0$. Concavity follows from $(-t \log t)'' = -1/t < 0$ and continuity at zero. Consequently mixing conditional distributions cannot decrease their average entropy: $H(Y \mid Z) \geq H(Y \mid Z, W)$. Factoring finite joint probabilities proves the chain rule $H(X) = \sum_j H(X_j \mid X_1, \ldots, X_{j-1})$. Also $H(Y) \leq \log |\operatorname{supp} Y|$: Jensen's inequality for log gives $\sum_y r_y \log(1/r_y) \leq \log \sum_y r_y / r_y$ on the positive support. For X uniform on A, $H(X) = \log |A|$, while

$$\begin{aligned}\sum_i H(X_{-i}) &= \sum_i \sum_{j \neq i} H(X_j \mid X_{<j} \setminus X_i) \\ &\geq \sum_i \sum_{j \neq i} H(X_j \mid X_{<j}) = (D-1) H(X).\end{aligned}$$

The notation in the conditioning set means omit X_i only when $i < j$. Bounding each $H(X_{-i})$ by $\log |\pi_i A|$ and exponentiating proves (29). No projection inequality is an unproved hypothesis in the sequel.

22.2 Relative isoperimetry on a finite cube

Let $Q_s = \{0, \dots, s-1\}^D$, $s \geq 2$, and let $\partial_E A$ consist of internal cube edges crossing A and its complement. If $0 < |A| \leq |Q_s|/2$, then

$$|\partial_E A| \geq c_D |A|^{(D-1)/D}. \tag{30}$$

Indeed, put $a = |A|$, $P_i = |\pi_i A|$. At most a/s occupied fibers in direction i are full. Every occupied nonfull fiber contains an internal crossing edge, because the fiber is a connected interval. Edges in distinct fibers/directions are different. Thus projection inequality and AM–GM give

$$|\partial_E A| \geq \sum_i (P_i - a/s) \geq D a^{(D-1)/D} - Da/s \geq D(1 - 2^{-1/D}) a^{(D-1)/D}.$$

Empty A satisfies the same inequality trivially. The boundary is the EDGE boundary, not the vertex boundary in a separate graph theorem.

22.3 Endpoint counting-norm Sobolev–Poincare inequality

For every real u on Q_s, including $s = 1$,

$$\sum_{Q_s} |u - \bar{u}_{Q_s}|^2 \leq K_D^2 \left(\sum_{e \subset Q_s} |du(e)|^q \right)^{2/q}. \tag{31}$$

For $s = 1$ both sides vanish. For a nonnegative v supported on at most half the cube, layer-set decomposition and the triangle inequality in $\ell^{D/(D-1)}$, followed by (30), imply

$$\|v\|_{D/(D-1)} \leq \int_0^\infty |\{v > t\}|^{(D-1)/D} \, dt \leq c_D^{-1} \int_0^\infty |\partial_E \{v > t\}| \, dt = c_D^{-1} \sum_e |dv(e)|.$$

The integrals are finite sums of step lengths between values of v; the last equality follows edgewise from the length of the interval between its two endpoint values. Apply this inequality to v^α. For $a, b \geq 0$, the fundamental theorem of calculus gives

$$|a^\alpha - b^\alpha| \leq \alpha |a - b| (a^{\alpha-1} + b^{\alpha-1}).$$

Since $\alpha D/(D-1) = 2$ and $(\alpha - 1)q' = 2$, Holder and the degree bound $\deg_{Q_s}(x) \leq 2D$ give

$$\begin{aligned} \|v\|_2^\alpha &\leq \frac{\alpha}{c_D} \Big(\sum_e |dv(e)|^q \Big)^{1/q} \Big(\sum_{e=\{x,y\}} (v(x)^{\alpha-1} + v(y)^{\alpha-1})^{q'} \Big)^{1/q'} \\ &\leq \frac{2\alpha D^{1/q'}}{c_D} \Big(\sum_e |dv(e)|^q \Big)^{1/q} \|v\|_2^{\alpha-1}. \end{aligned}$$

The second inequality uses $(a+b)^{q'} \leq 2^{q'-1}(a^{q'} + b^{q'})$ and $\sum_e (v(x)^2 + v(y)^2) \leq 2D\|v\|_2^2$. If $v = 0$ the conclusion is immediate; otherwise division proves $\|v\|_2 \leq K_D \|dv\|_q$.

Choose a counting median m of u and set $v_+ = (u - m)_+$, $v_- = (m - u)_+$. Both supports have size at most half the cube. For each edge,

$$|dv_+(e)|^q + |dv_-(e)|^q \leq |du(e)|^q.$$

For endpoint values on the same side of m this is equality in the nonzero part; on opposite sides it is $a^q + b^q \le (a+b)^q$. Since $2/q \ge 1$, addition of the two squared estimates gives

$$\|u - \bar{u}_{Q_s}\|_2^2 \le \|u - m\|_2^2 \le K_D^2 \big(\|dv_+\|_q^q + \|dv_-\|_q^q\big)^{2/q} \le K_D^2 \|du\|_q^2.$$

This proves (31) without an external graph, ball-adapter, coefficient, or boundary-condition hypothesis.

22.4 Reflection with all aspect-ratio multiplicities

Let $R = \prod_i \{0, \dots, l_i - 1\}$, $V = \prod_i l_i$, $L = \max_i l_i$, and $\kappa = L^2/V^{2/D}$. If $V = 1$ the desired variance bound is zero. Otherwise $L \ge 2$. Put $s = 2L + 1$, $r = s^D/V$ and fold Q_s coordinatewise onto R by the endpoint-repeating pattern of period $2l_i$,

$$0, 1, \dots, l_i - 1, l_i - 1, \dots, 1, 0, 0, 1, \dots.$$

For $l_i = 1$ the fold is constant. Since $s/(2l_i) > 1$, each original coordinate value has at least $2\lfloor s/(2l_i) \rfloor \ge s/(2l_i)$ preimages and at most $s/l_i + 2 \le 3s/l_i$ preimages. An original internal edge has at most the same upper count in its own direction; each full period traverses it twice. Adjacent folded values either coincide or are an original internal edge, possibly with reversed orientation. Thus for the unweighted extension $\widetilde{u}$,

$$\mathrm{Var}_{Q_s}(\widetilde{u}) \ge 2^{-D} r \, \mathrm{Var}_R(u), \qquad \sum_{e \subset Q_s} |d\widetilde{u}(e)|^q \le 3^D r \sum_{e \subset R} |du(e)|^q.$$

The variance inequality holds by applying the preimage lower bound at the mean of $\widetilde{u}$, and then minimizing over constants on R. Apply (31). Since $2/q - 1 = 2/D$ and $s \le 3L$,

$$\mathrm{Var}_R(u) \le A_D \kappa \left(\sum_{e \subset R} |du(e)|^q \right)^{2/q}, \qquad A_D = K_D^2 2^D 3^{2D/q+2}. \tag{32}$$

The factor r is not bounded independently of aspect ratio; it was kept until it became $s^2/V^{2/D}$. Singleton coordinate sides cause no problem.

22.5 Positive original edge weights and field coordinates

For arbitrary positive ORIGINAL edge coefficients b_e, Holder on R with exponents $2/q$ and $2/(2-q)$ gives

$$\left(\sum_e |du|^q \right)^{2/q} \le \Big(\sum_e b_e |du|^2 \Big) \Big(\sum_e b_e^{-q/(2-q)} \Big)^{(2-q)/q}.$$

Here $q/(2-q) = D/2 = p$. Consequently

$$\mathrm{Var}_R(u) \le A_D \kappa \Big(\sum_{e \subset R} b_e^{-D/2} \Big)^{2/D} \sum_{e \subset R} b_e |du(e)|^2. \tag{33}$$

No reflected weights were introduced: reflection was completed before this Holder step. For $b_{(x,x+e_i)} = a(x) > 0$, each vertex supplies at most D internal starting edges. Thus the inverse edge mass is at most $D \sum_R a^{-D/2}$, proving exactly the scalar-start estimate

$$\mathrm{Var}_R(u) \le A_D D^{2/D} \kappa \Big(\sum_R a(x)^{-D/2} \Big)^{2/D} \mathcal{E}_{a,R}(u).$$

For any fixed finite real internal/field dimension, apply this scalar inequality to each coordinate and sum; its coefficient factor is common. There is no additional internal-dimension factor and no replacement of original edges by enlarged-region energy.

22.6 Necessity of geometry and status

For uniform $a > 0$ and a longest coordinate of length $L \geq 2$, put $u(x) = x_i - (L-1)/2$. Direct sums give

$$\mathrm{Var}_R(u) = V(L^2-1)/12, \qquad \mathcal{E}_{a,R}(u) = aV(L-1)/L,$$

so their ratio is $L(L+1)/(12a) \geq L^2/(12a)$. Since $\kappa(\sum_R a^{-D/2})^{2/D} = L^2/a$, the geometric dependence cannot be uniformly dropped even for uniform weights. This is not a global counterexample to a partition algorithm or to the physical divergence-zero observation problem.

For attribution and comparison see S. Andres, J.-D. Deuschel and M. Slowik, *Harnack inequalities on weighted graphs and some applications to the random conductance model*, arXiv:1312.5473v6, Section 2.2, (2.5)–(2.8). Their local normalized inequalities are classical prior context; this proof does not import their theorem, graph assumptions or Harnack conclusions. The finite projection/entropy, coarea and power steps above are classical, not a new general method. The gain is a complete finite derivation of the specific input previously left external. No Lean execution, independent review, original-priority claim, new physical hypothesis removal or new AIRR paper follows from this supplement. The earlier separately frozen source PDF and its own reports/permissions remain historical artifacts. This fragment is integrated in the present changed artifact, which does not inherit their assessments.

23 A finite rectangular partition of disjoint lattice cubes

The linear partition for supplied disjoint cubes can be derived directly, without importing a general fat-object or low-density theorem. We use the classical compressed cube-splitting strategy of de Berg, UU-CS-1998-12, Section 2.1 and Lemma 2.1, and prove the required special-case geometry and all discrete/energy adapters below. This is not an original general BSP method. No time-complexity, measured implementation or Lean claim is made.

Fix $D \geq 3$ and an integer $n \geq 1$. In the common root $X = \{0, \dots, n-1\}^D$, take pairwise disjoint nonempty integer cubes

$$S_k = \prod_i \{b_{ki}, \dots, b_{ki} + L_k - 1\}, \qquad 1 \leq k \leq J.$$

Set

$$Q_k = \prod_i [b_{ki} - 1/4, b_{ki} + L_k - 3/4].$$

These are positive-side closed solid cubes, including $L_k = 1$. They are strictly disjoint: disjoint integer intervals in a separating coordinate leave a real gap at least $1/2$. Their integer points are exactly S_k, and each such point is interior to Q_k. All corners of the Q_k lie strictly inside $[-1, n]^D$.

23.1 Proper-intersection corner bound

Let C be any axis-aligned cube of side $s > 0$ with no corner of any Q_k in its interior. Then at most 2^D of the Q_k meet $\operatorname{int} C$. Indeed, if Q_k of side $a < s$ properly meets C, each of its

coordinate intervals has an endpoint in the interior of the corresponding interval of C. Choosing these endpoints gives a corner of Q_k in $\operatorname{int} C$, a contradiction. Thus every properly meeting Q_k has $a \geq s$. In each coordinate its interval then contains at least one endpoint of the interval of C; otherwise it would lie strictly inside that shorter interval. Hence Q_k contains a corner of C. Strict disjointness makes these corner choices injective, proving the bound.

The qualifier "properly" is essential. Closed boundary-only contact must not be counted here. For $C = [0,1]^D$ and any integer $K \geq 1$, put $\epsilon = 1/(4K+4)$ and

$$Q_j = [-\epsilon, 0] \times [(2j+1)\epsilon, (2j+2)\epsilon] \times [\epsilon, 2\epsilon]^{D-2}, \qquad 0 \leq j < K.$$

These K strictly disjoint small cubes touch C but none meets its interior or has a corner in its interior. Thus the corresponding closed-contact bound would be false. This boundary guard does not assert that de Berg's theorem is false; our construction below explicitly uses open-cell intersections and interior lattice points.

23.2 A multiple-object stress family

The corner-free property does not imply that a query cube meets only one input cube. This can fail even for the quarter-offset solids of equal-side integer cubes, not merely for arbitrary real objects. For $D \geq 2$, put $n = 4$, $C = [1,2]^D$, and, for $\sigma \in \{0,1\}^{D-1}$, set

$$b_\sigma = (2\sigma_1, \ldots, 2\sigma_{D-1}, 1), \quad L_\sigma = 2, \qquad S_\sigma = \prod_{i=1}^{D} \{b_{\sigma i}, b_{\sigma i} + 1\}.$$

The 2^{D-1} integer cubes are pairwise disjoint: distinct binary words have a coordinate with the disjoint integer intervals $\{0,1\}$ and $\{2,3\}$. Their solids have, in that coordinate, intervals $[-1/4, 5/4]$ and $[7/4, 13/4]$, with a gap $1/2$. In the last coordinate every solid has interval $[3/4, 9/4]$. Consequently every solid properly meets C, while no solid corner is interior to C: every such corner has last coordinate $3/4$ or $9/4$. Thus a universal proper-occupancy bound for corner-free query cubes in this exact lattice class must be at least 2^{D-1}. This neither proves sharpness of 2^D nor asserts that this particular C occurs as a leaf of our root-based recursion. It rules out replacing the occupancy argument by a one-object assumption. The finite unit139 baseline leaves all had occupancy at most one, so those runs did not test this phenomenon.

23.3 Constructing linearly many vertex-free rectangles

Let V be the set of all corners of the Q_k, so $N = |V| = 2^D J$. For $J = 0$ use the single root rectangle and skip this construction. For $J > 0$ begin with the cube $[-1, n]^D$. At a current cube C keep exactly $V_C = V \cap \operatorname{int} C$. If V_C is empty, stop at that cube. Otherwise consider its 2^D open half-side children.

If the points are not all in one open child, perform all midpoint cuts and recurse on the nonempty child point sets. Points on the cuts are removed from subsequent interior sets. If all points lie in one open child, reflect and translate coordinates so the common corner is zero and $C = [0, s]^D$. Put

$$a = \max\{x_i : x \in V_C,\ 1 \leq i \leq D\}, \qquad C' = [0, a]^D.$$

Then $0 < a < s/2$. Successively cut at the D upper facet planes of C', producing C' and the D boxes

$$E_i = [0, a]^{i-1} \times [a, s] \times [0, s]^{D-i}, \qquad 1 \leq i \leq D.$$

Use disjoint boundary ownership for these boxes. Only C' can contain points of V in its interior, and at least one point of V_C lies on its boundary. Stop at the other boxes and recurse inside C' if needed. Every intermediate or terminal cell is full-dimensional; a boundary ownership convention never creates a separate hyperplane cell.

Here is a direct finite termination and size proof. Among active cells, let $n_C = |V_C| > 0$ and use the nonnegative integer potential

$$\Psi = \sum_{C \text{ active}} (2n_C - 1).$$

Initially $\Psi = 2N - 1$. If a processed cell loses r points to new cut boundaries and has t nonempty child point sets, the decrease is $2r + t - 1$. In the shrinking step $r \geq 1$. In a midpoint step either $r \geq 1$ or $t \geq 2$. Thus every event decreases Ψ by at least one, so there are at most $2N - 1$ events. Each event creates at most $2^D - 1$ additional leaves, since a shrinking step creates only D extra leaves. The number M of terminal rectangles satisfies

$$M \leq 1 + (2N - 1)(2^D - 1) \leq 2^{D+1}(2^D - 1)J. \tag{34}$$

Tracking exactly $V \cap \operatorname{int} C$ ensures that every terminal rectangle has no corner of an input cube in its interior. This argument is a construction and a decreasing integer potential, not a supplied partition certificate. It does not prove a fast implementation bound.

23.4 A uniform occupancy bound for the terminal rectangles

Terminal cubes meet at most 2^D input cubes properly, by the corner bound. A terminal noncube has the form E_i above after a coordinate reflection. It is contained in the slab $[0, s]^{i-1} \times [a, s] \times [0, s]^{D-i}$. Cover this slab by 2^{D-1} closed cubes of side $s - a$: in coordinate i use $[a, s]$, and in every other coordinate choose either $[0, s - a]$ or $[a, s]$. These intervals cover $[0, s]$ since $a < s/2$. Every covering cube lies inside the parent cube and outside the interior of C', and its interior contains no point of V. The corner bound therefore applies separately to each covering cube.

If an input cube properly meets E_i, the intersection contains an open set. A point of that set avoiding the finitely many covering-facet planes belongs to the interior of a covering cube. Thus all proper intersections are counted by the cover, even if some particular point lies on a cover boundary. Consequently each terminal rectangle properly meets at most

$$T_D = 2^{D-1}2^D = 2^{2D-1} \tag{35}$$

input cubes. The covering argument, including its boundary qualification, is proved here; no arbitrary-box unclutteredness is assumed.

23.5 From open geometry to an actual integer partition

For each terminal rectangle list the at most T_D cubes meeting its interior. Refine by all their facet coordinates. There are at most $2T_D$ cut coordinates per axis, so at most $(2T_D + 1)^D$ refined rectangles per terminal cell. All cuts are axis-parallel. Now assign every splitting hyperplane to one child consistently, intersect with X, and delete empty integer sets. The result is a disjoint cover of X by integer rectangles. Taking integer points and deleting empty sets does not increase the count. More explicitly, use the uniform convention $[l_i, h_i)$ in every axis of every cell. A

coordinate split at t then assigns the plane to the upper child. The two intervals $[l,t)$ and $[t,h)$ partition $[l,h)$ exactly, even when t is an integer. By induction through every cut, these half-open cells partition $[-1,n)^D$; intersecting with X gives

$$R=\prod_i(\{\lceil l_i\rceil,\dots,\lceil h_i\rceil-1\}\cap\{0,\dots,n-1\}).$$

This formula is also valid for negative or integer endpoints. An empty axis makes the entire integer cell empty; no lattice point is assigned twice or lost when those cells are deleted. The outer upper face at n contains no point of X.

To check membership rigorously, suppose an assigned integer point x lies in S_k. It is interior to Q_k, while its assigned terminal rectangle has x in its closure. A small ball about x lies in Q_k and meets the interior of that full-dimensional terminal rectangle. Therefore Q_k was in its list. None of the new facet coordinates is an integer. Hence each refined integer rectangle is entirely inside a listed S_k or disjoint from every listed S_k. An unlisted cluster cannot contain one of its integer points, by the same closure argument. Strict disjointness then implies that each integer rectangle is inside one cluster or outside every cluster.

In particular, with the explicit dimension-only constant

$$A_D=2^{D+1}(2^D-1)(2^{2D}+1)^D,\qquad |\mathcal{P}|\le A_D\max\{1,J\},\tag{36}$$

we have derived the required partition. The constant is deliberately loose and no sharpness or algorithm running-time claim is made. For $n=1$ and for singleton input cubes all constructions remain valid.

23.6 The actual coefficient and fixed-penalty consequence

Suppose, in addition, that for fixed $c>0$ and $\beta\ge1$ the actual positive scalar starting coefficients obey

$$\begin{aligned}a(x)&\ge cL_k(L_k-1)/4\quad(x\in S_k),\\ a(x)&\ge cn(n-1)/4\quad(x\notin\cup_kS_k),\\ a(x)&\le\beta L_k^2\quad(x\in S_k).\end{aligned}$$

These coefficient/cube assumptions are retained, not inferred from the critical moment. On an interval of l integer vertices, pairwise variance, path Cauchy–Schwarz and at most $l^2/4$ pairs through one cut give

$$\sum|u-\bar u|^2\le\frac{l(l-1)}{4}\sum_{\text{edges}}|du|^2.$$

For completeness, product variance tensorization follows from the commuting coordinate averages $P_iu=l_i^{-1}\sum_{t=0}^{l_i-1}u(x_1,\dots,t,\dots,x_D)$: the telescoping orthogonal decomposition of $I-\prod_iP_i$ gives a sum bounded by $\sum_i\|u-P_iu\|_2^2$. Applying the interval bound on its fibers proves

$$\operatorname{Var}_R(u)\le\frac{L(R)(L(R)-1)}{4}\sum_{e\subset R}|du(e)|^2.$$

Inside a cluster $L(R)\le L_k$, and outside all clusters $L(R)\le n$. The displayed coefficient lower bounds therefore give $c\operatorname{Var}_R(u)\le\mathcal{E}_{a,R}(u)$. Singleton residuals are zero, so no zero lower bound is divided by. Sum over all field/internal coordinates and all partition cells. Original internal

positive-axis edge occurrences are disjoint; omitted inter-cell and periodic-wrap occurrences have nonnegative energy. Thus the ACTUAL normalized means

$$(C_{\mathcal{P}}U)_{R,j} = \sqrt{|R|}\,\operatorname{mean}_R U_j$$

satisfy

$$c\|U\|^2 \le \mathcal{E}_a(U) + c\|C_{\mathcal{P}}U\|^2$$

for every cochain, with output dimension $h|\mathcal{P}|$, $h = D(N_c^2 - 1)$. This uses an all-cochain variance estimate, not the invalid inference from a kernel-only margin to fixed penalty c. Means need not preserve the divergence constraint. The inequality consequently also holds on that constraint subspace without changing $C_{\mathcal{P}}$.

If $p = D/2$, $\mu \ge 1$, and the original critical mass bound is $\sum_X a^{-p} \le Mn^D/\mu^p$, then each supplied cube contributes at least $L_k^D/(\beta L_k^2)^p = \beta^{-p}$ to this mass. Hence $J \le \beta^p Mn^D/\mu^p$, and (36) gives

$$\dim F \le hA_D(1 + 2^p\beta^p M)\max\{1, n^D/(1+\mu)^p\}.$$

This proves the restricted supplied-cube consequence with an explicit constant and without a fat-object BSP theorem as a black box. It does not construct the supplied cubes from the critical moment, remove either scale bound or the strong complement bound, resolve the unrestricted critical rectangular question, or establish a covariant/continuum gap. No new originality or independent-discovery claim is made: the compressed splitting strategy is attributed classical geometry, here derived and adapted completely for this special class. The earlier separately frozen reviewed source PDF and its full reports and pending permission remain unchanged.

24 Physical cycle obstructions for an actual observation map

This part closes a limitation of the scalar four-cycle example: the obstruction is now realized inside the actual divergence-zero flat physical cochain domain. The kernel estimate below is proved for all actual cochains, not just the special fields used as negative witnesses. The conclusion concerns this observation and this fixed penalty. It does not obstruct all partitions or the optimal-rank question from the preceding parts.

24.1 The actual form and rectangular observations

Fix $D \ge 3$, $N_c \ge 2$ and $X = (\mathbb{Z}/4\mathbb{Z})^D$. Write $V = \mathbb{R}^{N_c^2-1}$ with its standard Euclidean norm and $A = (A_\ell(x))_{x\in X, 1\le\ell\le D}$ for an actual one-cochain. All norms use counting measure. Choose the actual total coefficients

$$\alpha = (1, 6, 4, 1), \qquad a(x) = \alpha_{x_1}, \qquad w(i, x, \ell) = \alpha_{x_1} - 1$$

for every derivative direction i and physical component ℓ. Define actual forward and backward differences by

$$D_i^+ f(x) = f(x + e_i) - f(x), \qquad D_i^- f(x) = f(x) - f(x - e_i).$$

The flat curl and backward divergence, with one curl for each unordered pair, are

$$(\operatorname{curl} A)_{ij} = D_i^+ A_j - D_j^+ A_i \quad (i < j), \qquad \operatorname{div}^- A = \sum_i D_i^- A_i.$$

Their counting quadratic form and the original componentwise repair are

$$H_0(A) = \sum_{i<j,x} \|D_i^+ A_j(x) - D_j^+ A_i(x)\|^2 + \sum_x \Big\| \sum_i D_i^- A_i(x) \Big\|^2,$$
$$R_w(A) = \sum_{i,x,\ell} w(i,x,\ell) \|D_i^+ A_\ell(x)\|^2.$$

In operator notation the repair is $\langle A, G^* W G A \rangle$ for the actual stacked gradient G and diagonal multiplication W. These formulas fix the energy on paper; no comparison certificate is supplied.

Deriving the flat identity. Periodic reindexing gives

$$\sum_x \langle D_i^+ f, h \rangle = -\sum_x \langle f, D_i^- h \rangle.$$

Coordinate shifts commute. Applying this formula twice and real symmetry therefore yields

$$\sum_x \langle D_i^+ f, D_j^+ g \rangle = \sum_x \langle D_j^- f, D_i^- g \rangle.$$

Also $\sum_x \|D_i^- f\|^2 = \sum_x \|D_i^+ f\|^2$ by translation. Expand each curl square with $f = A_j$, $g = A_i$: it supplies the two off-diagonal gradient norms and the negative mixed terms. Expansion of the divergence square supplies the diagonal norms and the opposite mixed terms. Cancellation for each $i < j$ proves, for every cochain,

$$H_0(A) = \sum_{i,x,\ell} \|D_i^+ A_\ell(x)\|^2.$$

No divergence condition or lift ansatz was used. Define $H = H_0 + R_w$. Adding the displayed repair, with $1 + w(i,x,\ell) = \alpha_{x_1}$, gives

$$H(A) = \sum_{x \in X} \sum_{i=1}^{D} \sum_{\ell=1}^{D} \alpha_{x_1} \|A_\ell(x + e_i) - A_\ell(x)\|^2. \tag{37}$$

The cancellation was unweighted; the varying repair is added afterwards. We do not move α through summation by parts or assert a weighted curl–divergence identity. This is the paper counterpart of `finite_flat_curl_div_energy` and `componentwiseWeightedRepair_energy` in the pinned formal source. Thus the coefficient is the original shared scalar total coefficient, not a substituted energy or an interacting Hessian.

Put $I_0 = \{0,1\}$ and $I_1 = \{2,3\}$. For every transverse $y \in (\mathbb{Z}/4\mathbb{Z})^{D-1}$ the blocks $I_b \times \{y\}$ are nonempty axis-aligned rectangles. They are disjoint and cover X. For each physical component ℓ and each standard internal coordinate k, define the actual normalized means

$$(CA)_{b,y,\ell,k} = \frac{1}{\sqrt{2}} \sum_{t \in I_b} A_\ell(t,y)_k. \tag{38}$$

The output includes *all* these coordinates. If $h = D(N_c^2 - 1)$, its dimension is $2h4^{D-1}$; no optimality is asserted.

The spatial mean of a is $\mu = 3$, and $1 \le a \le 2\mu$. For $p = D/2$ its original normalized inverse moment is

$$\frac{1}{|X|} \sum_{x \in X} \left(\frac{\mu}{a(x)} \right)^p = \frac{2 \cdot 3^p + (1/2)^p + (3/4)^p}{4} =: M_D < \infty. \tag{39}$$

This is an actual member of the stated bounded-moment coefficient class when the allowed bound is at least M_D. Its fixed finite size and the large output dimension do not disprove an asymptotic optimal-rank order.

24.2 A fixed-penalty counterexample with a proved kernel premise

For a real cyclic vector $u = (u_0, u_1, u_2, u_3)$ set

$$E(u) = (u_1 - u_0)^2 + 6(u_2 - u_1)^2 + 4(u_3 - u_2)^2 + (u_0 - u_3)^2,$$
$$|u|^2 = \sum_{t=0}^{3} u_t^2, \qquad O(u) = \frac{(u_0 + u_1)^2 + (u_2 + u_3)^2}{2}.$$

Choose a unit internal coordinate vector T. The actual lift is

$$(Ju)_2(x) = u_{x_1} T, \qquad (Ju)_\ell(x) = 0 \quad (\ell \neq 2).$$

Its backward divergence is zero: its only nonzero component is constant in coordinate 2. It is injective and direct summation gives

$$\|Ju\|^2 = 4^{D-1}|u|^2, \quad H(Ju) = 4^{D-1}E(u), \quad \|CJu\|^2 = 4^{D-1}O(u). \tag{40}$$

The last equality uses the entire output in (38), not only a selected nonzero coordinate.

Kernel premise on all actual fields. For $v = (x, -x, y, -y)$, direct algebra yields

$$E(v) = 11x^2 + 14xy + 23y^2, \qquad |v|^2 = 2x^2 + 2y^2,$$
$$E(v) - \frac{19}{5}|v|^2 = \frac{17}{5}\left(x + \frac{35}{17}y\right)^2 + \frac{84}{85}y^2 \geq 0. \tag{41}$$

If $CA = 0$, both pair sums vanish for every actual scalar profile $t \mapsto A_\ell(t, y)_k$. Apply (41) to each such profile and sum over all y, ℓ, k. These profiles account for the full norm of A. Their energies are exactly the $i = 1$ term of (37); every other derivative term is nonnegative because $\alpha_t \geq 1$. Therefore

$$CA = 0 \quad \Longrightarrow \quad H(A) \geq \frac{19}{5}\|A\|^2 \qquad \text{for every actual cochain } A. \tag{42}$$

No divergence condition, scalar-lift ansatz or supplied coercivity certificate was used in this proof.

Failure on an actual divergence-zero field. For $u = (19, -9, -5, -5)/14$ one has

$$E(u) = \frac{52}{7}, \qquad |u|^2 = \frac{123}{49}, \qquad O(u) = \frac{25}{49}.$$

Consequently the actual divergence-zero field Ju satisfies

$$H(Ju) + \frac{19}{5}\|CJu\|^2 - \frac{19}{5}\|Ju\|^2 = -\frac{6}{35}4^{D-1} < 0. \tag{43}$$

Equations (42) and (43) prove the complete fixed-map statement: kernel coercivity does not imply the same coefficient as a fixed coarse penalty, even on the actual flat physical divergence-zero domain. They do not rule out larger penalties or different observations.

24.3 The exact finite penalty at the subsharp target

The failure just proved concerns $K = 19/5$. Here we determine exactly which larger penalties work for this same observation and target. This subsection is a paper proof with an exact rational coefficient check; it is not part of the paired Lean theorem below.

Every $u \in \mathbb{R}^4$ has the unique decomposition

$$u = m(1,1,1,1) + (x,-x,y,-y) + (\delta/2,\delta/2,-\delta/2,-\delta/2),$$

where $m = (\sum_t u_t)/4$, $x = (u_0 - u_1)/2$, $y = (u_2 - u_3)/2$ and $\delta = (u_0 + u_1 - u_2 - u_3)/2$. Direct expansion gives

$$\begin{aligned} |u|^2 &= 4m^2 + 2x^2 + 2y^2 + \delta^2, \qquad O(u) = 4m^2 + \delta^2, \\ E(u) &= 11x^2 + 14xy + 23y^2 + 7\delta^2 - 10\delta(x+y). \end{aligned}$$

Set $X = x - 25\delta/2$ and $Y = y + 75\delta/14$. For every real K,

$$\begin{aligned} E(u) + KO(u) - \frac{19}{5}|u|^2 = 4\left(K - \frac{19}{5}\right)m^2 + \frac{17}{5}\left(X + \frac{35}{17}Y\right)^2 \\ + \frac{84}{85}Y^2 + \left(K - \frac{1138}{35}\right)\delta^2. \end{aligned} \tag{44}$$

Since $1138/35 > 19/5$, this proves sufficiency of $K \geq 1138/35$ on every scalar profile. Necessity follows from

$$u_* = (13,-12,-41/7,34/7), \qquad O(u_*) = 1, \quad |u_*|^2 = \frac{18174}{49}, \quad E(u_*) = \frac{67468}{49}.$$

It has $m = 0$, $\delta = 1$ and $X = Y = 0$, hence defect $K - 1138/35$.

Apply the scalar inequality to every profile $t \mapsto A_\ell(t,y)_k$ and sum over all y, ℓ, k. Their squared norms and observations exhaust the actual full norms, and their energies exhaust the first-direction contribution in (37). The remaining directions are nonnegative. Conversely Ju_* is actual divergence zero and multiplies the scalar defect by 4^{D-1}. Therefore the following are equivalent:

1. $K \geq 1138/35$;
2. $H(A) + K\|CA\|^2 \geq (19/5)\|A\|^2$ for every actual cochain;
3. the same inequality for every actual divergence-zero cochain.

This is a necessity-and-sufficiency statement for the fixed map, not an obstruction to different partitions. In particular, a larger finite penalty repairs the subsharp target even though $K = 19/5$ fails. The sharp kernel target in the next subsection behaves differently.

A second calculation records the classical block elimination behind the constant. The kernel residual matrix and mixed vector are

$$M = \begin{pmatrix} 17/5 & 7 \\ 7 & 77/5 \end{pmatrix}, \quad \det M = 84/25, \quad b = (5,5), \quad M^{-1}b = (25/2, -75/14).$$

Thus $b^T M^{-1} b = 250/7$ and the residual coefficient is $K + 16/5 - 250/7 = K - 1138/35$, agreeing with (44). No new Schur-complement principle is claimed.

The unit151 exact verifier constructs E, O and $|u|^2$ from the actual four-coordinate linear forms and compares every coefficient of the polynomial in (m, x, y, δ, K), rather than sampling values. Normal and optimized Python in one own Colab Linux runtime each accepted five checks and rejected three real coefficient mutations; both exits were zero and their JSON outputs were identical. The 8794-byte archive has ten members and nine checked inner hashes. It and the full hand derivation are included in the reproduction package. The driver took about 0.56 seconds; the runtime was released after recovery and exact billed consumption is unknown. This check is neither a new Lean theorem nor independently authored verification. Its source is

`564b84230478cb8d40505cb67985470e9bae699d.`

24.4 Sharp free-penalty extension and its distinct status

The following is an analytic continuation, outside the paired Lean result. Put $s = \sqrt{85}$ and $c_* = (17 - s)/2$. Since $9 < s < 10$, the exact identity

$$E(v) - c_*|v|^2 = (s - 6)\left(x + \frac{7y}{s - 6}\right)^2$$

holds on the same block kernel. The nonzero vector $v_* = (7, -7, 6 - s, s - 6)$ attains equality. The all-profile argument above also establishes the sharp kernel constant for the actual form in paper proof, and the lift establishes sharpness.

For $p_\delta = (\delta/2, \delta/2, -\delta/2, -\delta/2)$, the two vectors are orthogonal, $Cv_* = 0$, and $O(p_\delta) = |p_\delta|^2 = \delta^2$. For any finite $K \geq 0$, direct expansion gives

$$\begin{aligned} E(v_* + p_\delta) + KO(v_* + p_\delta) - c_*|v_* + p_\delta|^2 \\ = (7 + K - c_*)\delta^2 - 10(13 - s)\delta. \end{aligned}$$

The positive number $q_K = 7 + K - c_*$ permits the choice $\delta = 5(13 - s)/q_K$, giving the strictly negative value $-25(13 - s)^2/q_K$. Lifting multiplies this complete defect by 4^{D-1} and preserves actual divergence zero. Thus no finite quadratic penalty repairs this same observation at the sharp kernel endpoint. A penalty $\Phi(z) \leq K\|z\|^2$ only near zero is also defeated: scale the negative witness by a sufficiently small nonzero real number. Both radius conditions can be met, and its defect remains a negative multiple of the square of that scale. This is not a strict-margin claim.

These exact analytic identities and arguments are supplied here, but their complete physical sharp-endpoint/free-penalty statement has *not* been verified in Lean in the paired execution described next. Finite checks in dimensions three and four, including $K = 0, 1, 100$, are diagnostics only, not the universal theorem or independent review.

24.5 Formal scope and the paired reproduction record

The fixed-penalty result, including the actual full-kernel premise and all transport/comparison identities used to connect it, was compiled in two separately allocated Colab Linux CPU sessions with high RAM. Each started with a fresh clone at the same immutable source

`0c706827dcbb11859e56bb6c7dc981eb37a02e69.`

The final theorem is named `physicalCycleFour_kernel_fixed_penalty_counterexample` in `PhysicalCycleFullForm.lean`. It asserts the proved all-cochain kernel bound together with failure of the fixed-penalty inequality on the actual divergence-zero domain. Four preceding printed endpoints identify longitudinal energy, the full actual energy sum, its comparison, and the kernel estimate. They do not assume transport certificates or posit a structure containing the conclusion.

All five actual dependency readouts in each clone contain only `propext`, `Classical.choice` and `Quot.sound`. The normal and optimized validator in each clone accepted all five readouts and rejected three actual oracle mutations. All actual decimal exit sentinels were zero. Lean was version 4.29.0-rc6 with Mathlib pinned in the included manifest. The recovered theorem source, oracle source, compiled binary, oracle output, validator source, normal and optimized outputs and driver have identical hashes across the two fresh executions. The raw per-run reports still say first-clone pass and paired false: the identical driver executes one clone per session. They remain unmodified; the separate aggregate records the completed pair.

The first successful driver took about 374.31 seconds and the second 400.28 seconds. These are not billed connection times. Both environments were disconnected after complete evidence recovery. No Lean compilation, oracle or numerical campaign ran on Windows. Separate fresh executions are not independently authored proofs or independent scientific review. The previous failed attempt, its exact source and full compiler log are preserved: mixed finite-index inference and nested sum type inference were corrected without changing the theorem statements.

The package includes both public evidence archives. Their comparison is recorded in the aggregate file:

```
paired-full-form148.json
```

Full hashes and source bindings are in `inputs.json`. Each archive has 45 checked inner manifest members and excludes third-party installers and downloaded binary assets. Their complete originals, including those assets, are retained in the private evidence record. The compiled binary hash begins `b0f8ff8654d79d02`, and the full value is in the paired comparison file.

Attribution and scope of novelty. The scalar square identities and negative vectors predate this physical embedding in the dossier. The present consolidation adds the explicit actual-domain embedding and the formally derived all-field comparison; it claims no new Schur-complement principle, priority or exhaustive novelty finding. Carlson, Haynsworth and Markham's classical paper, *A Generalization of the Schur Complement by Means of the Moore–Penrose Inverse*, SIAM Journal on Applied Mathematics 26 (1974), 169–175, doi:10.1137/0126013, supplies algebraic context. The publisher abstract and bibliography were inspected in the source investigation; the full paywalled article was not read. The direct finite proof here carries the mathematical claim.

25 Evidence provenance and verification limits

The analytic sources, their exact revisions and the executable finite diagnostics are collected in a hash-allowlisted companion package. Its map gives full SHAs and SHA256 values; short forms below are identifiers, not substitute hashes.

Item	Source revision	Verification status
Actual flat Hodge identity	`dfd49d0b0730`	earlier formal source
Critical spectral theorem	`527fcdcd119a`	analytic adapters
Dyadic spatial separation	`fb65a67af360`	analytic adapters
Higher moment block proof	`6dcb893155e5`	analytic adapters
Exact finite fixture code	`4bb460e1537e`	own Colab109
Full diagnostic report	`602ae3f0b59a`	finite rational batch

Colab109 checked four actual profiles on $(\mathbb{Z}/2\mathbb{Z})^4$ using rational arithmetic. Each normal and optimized Python mode accepted four profiles and rejected 16 real matrix, inverse or flow mutations, with matching outputs and exit code zero. It checked the constant mode, forward symbols, actual pseudoinverse identities, feasible/minimizing flow energy, comparison inertia, strict low-spectrum multiplicities and noncommuting counts. There are 12 distinct density examples, evaluated again for each profile: 48 evaluations per mode, not 48 independent examples.

All eight files in the diagnostic archive's inner manifest were rehashed. Runtime execution was recorded at 17:13:06.300923–17:13:16.201862 UTC on 2 November2026, with automatic runtime release and zero sessions verified afterwards. Rounded resource units do not prove zero cost. No local Windows mathematical execution occurred. This single embedded-SHA finite batch is not two fresh matching clones, a Lean proof, a theorem about arbitrary dimensions or odd lattices, or independent mathematical review. This manuscript assembly adds no new Lean execution. Operational repository checks on unchanged Lean files do not verify these new analytic proofs.

There is also subsequent, strictly focal formal evidence. Source `41f8492c6190` supplies the ten four-cycle theorems and their oracle. Own Colab118 compiled that same source in two fresh clones; both actual oracle logs and compiled module hashes agree. Each of four normal/optimized validator runs accepted ten readouts and rejected three real corruptions. All readouts use only `propext`, `Classical.choice`, and `Quot.sound`. The full archive has 69 manifest files and is included with the failed Colab116/117 compiler attempts. Source hashes, complete logs and exact reports are allowlisted. This verifies Proposition 20.1's finite identities, not the geometric, shorted-operator, free-penalty or general rectangular adapters. Own Colab115 separately checked 18 finite diagnostics and eight real mutations in each mode of two fresh clones. Finite examples do not establish their general claims. Section 21 consolidates analytic sources at `0fa0d91000f6`, `312f476ea6e5`, `f52ae344f61f`, `1850be2aa6d5` and `b348f4488b3a`. The original spatial proofs and the self-contained Sobolev and cube arguments remain analytic. The physical fixed-penalty endpoints have the separately pinned paired formal record in Section 24. The exact subsharp threshold has an exact rational coefficient check, not a Lean certificate; the sharp extension is analytic. Rendering does not upgrade any verification status. Earlier exact PDFs and assessments remain preserved.

26 Model assessments and review context

This is a changed exact-artifact successor of the public v11 working paper. Its exact predecessor PDF SHA256 is

`68aae0aa0c93a2718abcc3f489f7ccdcb5cdc3477b2c93472cfa308e0c540a6b`.

Its full numerical AI assessment is retained with that immutable artifact at the v11 assessment record. The v11 report gave 4.00/10 overall, with five criterion stars 4/4/3/3/3 and a minor-revision recommendation. Those figures describe only that 88-page predecessor; they are not a

score for this candidate. The separately assessed eight-page supplement from unit 186 is also a historical source artifact. Neither its score nor the parent's score or resolved-objection status transfers to the resulting changed PDF. All earlier reports, objections and operator dispositions remain preserved with their own exact versions. No new public version identity or receipt is assigned by this manuscript.

Preparation, analytical checking and the new finite adapters were performed in the continuing Codex programme context. Exact served backend identity and technical history isolation are not attested. This is not independent review. The new model report, when completed, must be retained in full and bound to the resulting exact PDF hash, with factual model identity, provider, supplied materials, prior exposure and review limitations. Publication requires a genuine numerical overall/criterion assessment and visible stars for this exact artifact; an ungraded report or an inherited parent score is insufficient. No grade is assigned by this paragraph.

The complete version11 report and all its objections remain attached to that exact public artifact. Objections concerning unfinished originality comparison, un-replayed historical execution and the open global/physical questions remain limitations here. The mean-zero minimum penalty in Section 29 is now explicitly labelled, and its full-space counterpart includes the constant-mode constraint. The authoritative chronology is version11 followed by this unassigned candidate; references to earlier frozen artifacts below describe those historical objects only. The grouped shell and geometric sections have source-level reasoning reviews and finite checks, which are not the numerical exact-PDF assessment of this changed manuscript.

The principal limitations remain substantive. Classical Fourier, Dirichlet-form comparison, Schur, variance and refinement mechanisms limit originality claims; a complete priority comparison is unfinished. Some endpoints have earlier separately scoped formal proofs, while these spatial continuations are analytical with selected exact finite reproduction. Their coefficient and geometry hypotheses are not physical coefficient producers. A global spectral observation is not a local rectangular construction; the unrestricted critical rectangular question remains open. Finite selected cases do not certify every lattice size or replace an independent audit. AI evaluations, working-paper deposits, editorial acceptance and human peer review are distinct. No interacting or continuum Yang–Mills gap or reduced programme obligations are claimed.

27 The single remaining bounded question

For fixed $D \geq 3, c > 0, N_c \geq 2, M$, do all actual scalar profiles $1 \leq a \leq 2\mu$, exact mean μ, satisfying (16), admit (17) using normalized means on an unrestricted axis-aligned integer rectangular partition? The original penalty may have any finite locally quadratic upper coefficient; a construction with $K = c$ is stronger. A negative answer must defeat every such partition and finite penalty, not just a prescribed algorithm or sufficient certificate. The additions in Sections 29, 30 and 35 solve fixed-gap and partially invariant subclasses and exclude several incomplete approaches, while giving successful observations on the same counterexample families. They do not settle this question or create another one. The physical continuum problem remains separate.

28 Matrix gradients and critical rectangular penalties

This grouped continuation retains its actual finite conventions. Its mathematical dossier is pinned at `66580c59`; the earlier sections and their verification referents are preserved. Its scalar rectangular optimum is not asserted to be the full restricted coexact optimum.

28.1 Actual finite objects and observation scope

Let $D \geq 3$, $n \geq 2$, $h \geq 1$, $V = n^D$, $r = Dh$, and $p = D/2$. On $X = (\mathbb{Z}/n\mathbb{Z})^D$ use the counting inner product. For $u : X \to \mathbb{R}^h$, define

$$du(x) = (u(x + e_i) - u(x))_{i=1}^D, \qquad Q_A(u) = \sum_x du(x)^T A(x) du(x), \quad L_A = d^* A d,$$

where $A(x)$ is a real symmetric $r \times r$ matrix. A channel is one local coefficient eigenvector, counted with multiplicity. Distinct matrices need not commute. All derivatives have unit spacing. A physical derivative *nd* requires an explicit rescaling of every energy and threshold.

The number $N_{<c}(L_A)$ counts all eigenvalues strictly below $c > 0$, including negative and zero ones, with ordinary real multiplicity. Spectral coordinates on this low space give a real observation C and penalty

$$\Phi(y) = \sum_{\lambda<c} (c - \lambda) y_\lambda^2, \qquad Q_A(u) + \Phi(Cu) \geq c\|u\|^2. \tag{45}$$

This is an existence result for an actual, generally nonlocal observation. It does not construct the unresolved local rectangular observation. For signed forms, $c\|y\|^2$ alone need not be sufficient.

28.2 The positive matrix estimate

Put

$$\alpha_D = \frac{D}{2} \left(\frac{D}{D-2} \right)^{D-1} \left(\frac{3}{4} \right)^D .$$

Theorem 28.1. *If every $A(x)$ is positive definite, then*

$$N_{<c}(L_A) \leq h + \alpha_D c^p \sum_x \operatorname{tr} A(x)^{-p}. \tag{46}$$

Proof. On the mean-zero space H_0, write $L_0 = d^* d$ and $J = d L_0^{-1}$. For any $b \in H_0$, the flow $j_* = A d L_A^{-1} b$ has $d^* j_* = b$. Every other feasible flow $j = j_* + k$ has $d^* k = 0$, and expansion gives

$$\langle j, A^{-1} j \rangle = \langle b, L_A^{-1} b \rangle + \langle k, A^{-1} k \rangle.$$

Taking the actual feasible flow Jb proves $L_A^{-1} \leq J^* A^{-1} J$. No coefficients commute in this argument.

The unitary counting Fourier transform on r components has $L_1(\ell_2^r) \to L_\infty(\ell_2^r)$ norm $V^{-1/2}$. The symbol of J is $z(k)/\lambda(k)$ with $z_i(k) = e^{2\pi i k_i/n} - 1$ and $\lambda(k) = 4 \sum_i \sin^2(\pi k_i/n)$; set $J(0) = 0$. It is rectangular $r \times h$. Embed $\mathbb{R}^h$ isometrically in $\mathbb{R}^r$ by I and use the square symbol $G(k) = J(k) I^*$. Thus $\|G(k)\| = \lambda(k)^{-1/2}$ for $k \neq 0$. Writing $\mathcal{F}_r$ for the counting transform, the actual factorization is $J = \mathcal{F}_r^* M_G \mathcal{F}_r I$. Centered frequency representatives b satisfy $\lambda(k) \geq 16|b|^2/n^2$. Counting the cube containing $0 < |b| < n/(4t)$ gives

$$\sup_{t>0} t^D \#\{k : \|G(k)\| > t\} \leq (3/4)^D V.$$

When the radius is at most one the set is empty; otherwise $(2\lfloor R\rfloor+1)^D \le (3R)^D$ suffices.

Apply Frank's operator-valued Cwikel theorem [18, Theorem 3.1] with $q=D$, $f=A^{-1/2}$, this square multiplier and counting measures. Its constant is $\frac{D}{2}(\frac{D}{D-2})^{D-1}V^{-1}$, so the displayed frequency V cancels exactly once. It bounds the number of singular values of $A^{-1/2}J$ strictly above s by $\alpha_D s^{-D}\sum_x \operatorname{tr} A(x)^{-p}$. The right factor $\mathcal{F}_r I$ is an isometry and cannot increase singular-value counts. Apply min-max to the inverse comparison at threshold $1/c$ on H_0, then add the h constant modes. Complexification of a real symmetric matrix preserves multiplicities. □

Input and interpretation. The operator-valued estimate is classical; this proof fixes its finite normalization, padding and actual form dictionary. It is not a new general Cwikel theorem. A matrix arithmetic mean alone does not imply its inverse moment hypothesis. For the formal flat cochain application one may set $h=D(N_c^2-1)$, but an interacting Hessian or weighted curl/divergence form must first be proved equal or comparable to this Q_A.

28.3 Spectral repair and the capped budget

Theorem 28.2. *Let $A(x)$ now be arbitrary real symmetric matrices. In local orthonormal eigenbases choose a set S of channels containing all nonpositive ones. Let P_S be their orthogonal projection and $R(S)=\operatorname{rank}(P_S d)$. Then*

$$N_{<c}(L_A) \le \min\left\{ hV,\ h+R(S)+\alpha_D c^p \sum_{j\notin S} \lambda_j^{-p} \right\}. \tag{47}$$

Writing ν_0 for the total multiplicity of nonpositive local eigenvalues, one consequently has

$$N_{<c}(L_A) \le \min\{hV,\ h+F_c(A)\}, \quad F_c(A)=\nu_0+\sum_{\lambda_j>0}\min\{1,\alpha_D c^p \lambda_j^{-p}\}. \tag{48}$$

Proof. Replace selected eigenvalues by T larger than zero and all of them. The resulting field A_T is positive definite and

$$L_{A_T}-L_A = d^* P_S(T\mathrm{I}-A)P_S d \ge 0, \quad \operatorname{rank}(L_{A_T}-L_A)=R(S).$$

The rank equality follows because $T\mathrm{I}-A$ is strictly positive on $\operatorname{im} P_S$. Finite-rank min-max gives $N_{<c}(L_A)\le N_{<c}(L_{A_T})+R(S)$: intersect the low space of L_A with the kernel of the update and the orthogonal complement of the low space of L_{A_T}. A dimension larger than $R(S)$ would give a nonzero vector whose identical quadratic forms are both strictly below and at least $c\|u\|^2$. Theorem 28.1 and $T\to\infty$ prove (47). For (48), use $R(S)\le |S|$ and select all positive channels with $\alpha_D c^p\lambda_j^{-p}\ge 1$ as well. □

A strict extension of the full moment criterion. Let $m\ge 1$ and $T_m=\max\{1, V/(1+m)^p\}$. If $F_c(A)\le KT_m$, then $N_{<c}(L_A)\le (h+K)T_m$. Every former SPD budget

$$\frac{1}{rV}\sum_x \operatorname{tr}((mA(x)^{-1})^p) \le M$$

implies this capped budget, with $K=2^p\alpha_D c^p r M$. Conversely choose $m=n^2$, one actual diagonal edge channel $\epsilon=n^{-4}$, and all others m. Now $T_m=1$ and $F_c(A)\le 1+\alpha_D c^p r$, whereas the normalized full inverse moment is at least n^{2D}/r and diverges. These bounds are derived for this explicit family; no new moment bound is assumed.

Sparse contamination and its limitation. Selecting every channel with $\lambda < m$ gives

$$N_{<c}(L_A) \le h + R_m + \alpha_D c^p r V m^{-p}, \qquad R_m = \operatorname{rank}(\mathbf{1}_{(-\infty,m)}(A)d).$$

An actual selected-channel count $k_m \le KT_m$ suffices since $R_m \le k_m$. If $A(x) \ge m\mathrm{I}$ outside a specified set B, then $k_m \le r|B|$. Neither condition follows from an arithmetic mean. This particular prescribed-threshold count criterion does not subsume every full moment family: $A = (m/2)\mathrm{I}$ has normalized inverse moment 2^p but every channel lies below m. The two particular sufficient criteria are incomparable; the general capped criterion (48) does include the old one.

Actual rank and unavoidable costs. For diagonal direction/component channels the selected rows are graph incidences. If κ_a counts all components, including isolated vertices, of the selected-edge graph for component a, then $R(S) = \sum_a (V - \kappa_a)$: its kernel consists precisely of fields constant on those components. Parallel rows at $n = 2$ do not change that kernel. This formula is not applied to rotated local eigenchannels. Isolating q nonadjacent vertices by zeroing their incident edges gives at least $q + 1$ kernel modes. Replacing these zeros by $0 < \epsilon < c/(2D)$ yields a q-dimensional strict low subspace of delta functions. The defect-count cost cannot simply be deleted.

In (45), a generic bound is $K \le c + 4Db$ with $b = \max_x \|A_-(x)\|$, since $\|du\|^2 \le 4D\|u\|^2$. One negative channel $-B$ with all other channels m, at $n \ge 3, h = 1$, has the trial vector $v = \delta_y - \delta_x$ on that edge and

$$\frac{Q_A(v)}{\|v\|^2} = -2B + (2D - 1)m.$$

Any normalized linear observation with penalty bounded by $K\|y\|^2$ therefore needs $K \ge c + 2B - (2D - 1)m$. Sparsity alone does not bound penalty.

28.4 Stable signed families from actual resistance

This section is independent of the Cwikel input. Its mechanism is classical effective resistance and finite-rank congruence [19, 20]. Let e be one actual direction/component row and $v = d^*e$. With all coefficients m except this one set to $-B$, $B \ge 0$,

$$H = mL_0 - (m + B)vv^*, \qquad L_0 = d^*d.$$

On H_0 let $G = L_0^{-1}$. Translations and coordinate permutations make every scalar row leverage equal. Their sum is $\operatorname{tr}(dGd^*) = V - 1$, so

$$\rho = v^*Gv = \frac{V - 1}{DV}. \tag{49}$$

Components decouple; the same value holds for every h. The DV rows are counted individually even at $n = 2$.

Proposition 28.3. *The exact threshold is $B_* = m(DV/(V - 1) - 1)$. On H_0, H is positive definite for $B < B_*$, has one zero direction at equality, and has exactly one negative direction above it. If $m = n^2$, $B = \theta m$ and $0 \le \theta < D - 1$, its mean-zero gap is at least*

$$g = 16(D - 1 - \theta)/D > 0.$$

For every $0 < c \le g$, the normalized mean observation $C(u) = V^{-1/2} \sum_x u(x) \in \mathbb{R}^h$ and $\Phi(y) = c\|y\|^2$ give coercivity with optimal rank h, uniformly in n.

Proof. Congruence by $G^{1/2}$ gives $m\mathrm{I} - (m+B)ww^*$ with $\|w\|^2 = \rho$, proving the phase claim. The equality witness is Gv. Cauchy–Schwarz in the L_0 energy gives $|v^*u|^2 \le \rho\langle u, L_0 u\rangle$, hence $H \ge [m-(m+B)\rho]L_0$ on H_0; equality holds on Gv. For the stated scaling the coefficient is at least $m(D-1-\theta)/D$. The actual first nonconstant Fourier eigenvalue of L_0 is $4\sin^2(\pi/n) \ge 16/n^2$. This proves the gap. Decompose into constants and H_0 to prove the displayed observation inequality. The constant kernel forces rank at least h. □

For $\theta > 0$, negative coefficient depth grows as θn^2, yet the penalty is uniformly c. A bounded-depth hypothesis is sufficient for the generic estimate, not necessary for every stable signed family. The actual form is positive despite an indefinite coefficient field; this does not contradict the fixed-m, $B \to \infty$ penalty obstruction above.

28.5 Coupled channels and a nonlinear observation obstruction

Choose q actual rows of d as Z, and $W = \mathrm{diag}(w_j) > 0$. The modified form is exactly $H = mL_0 - Z^*WZ$. Define

$$\Gamma = ZGZ^*, \qquad \mathcal{C} = mW^{-1} - \Gamma, \quad N_0 = h(V-1).$$

Schur elimination of either positive diagonal block in $\left[\begin{smallmatrix} mL_0 & Z^* \\ Z & W^{-1} \end{smallmatrix}\right]$ on $H_0 \oplus \mathbb{R}^q$ proves

$$n_-(H) = n_-(\mathcal{C}), \quad n_0(H) = h + n_0(\mathcal{C}), \quad n_+(H) = N_0 - q + n_+(\mathcal{C}). \tag{50}$$

Indeed the two congruences give the identity

$$\mathrm{inertia}(H|_{H_0}) + (q,0,0) = (N_0,0,0) + \mathrm{inertia}(\mathcal{C}/m).$$

No disjoint-cycle or independent-row assumption occurs. In particular, H is positive semidefinite with exactly h constant zeros if and only if $\mathcal{C}$ is positive definite. General signed matrix criteria preexist in [21]; our Γ uses the positive reference inverse.

Example 28.4. *Take $D = 3, n = 2, h = 1, m = 4$, and change the first two direction rows at one vertex separately or together to coefficient $-84/11$. Then*

$$\Gamma = \begin{pmatrix} 7/24 & 5/48 \\ 5/48 & 7/24 \end{pmatrix}, \quad w = 128/11, \quad w/m = 32/11.$$

The diagonal follows from (49); the off-diagonal Walsh sum is $(1/8)(1/2+1/3) = 5/48$, with both forward orientations counted. Each single channel is below its critical magnitude $68/7$ and has full inertia $(7,0,1)$. Together the reduced matrix has diagonal $5/96$ and off-diagonal $-5/48$, so its eigenvalues are $-5/96, 5/32$. Equation (50) gives joint inertia $(6,1,1)$.

Each single form has mean-zero gap at least $16(1-(32/11)(7/24)) = 80/33 > 1$. The same scalar mean observation with $\Phi(y) = |y|^2$ repairs either single form at $c = 1$. The joint form cannot be repaired by any scalar observation $F : \mathbb{R}^8 \to \mathbb{R}$ differentiable at zero with $F(0) = 0$ and $0 \le \Phi(z) \le K|z|^2$ near zero. Its negative eigenvector and constant vector span a two-dimensional nonpositive space. The derivative of F has a nonzero kernel vector v there. Then $F(tv) = o(t)$ and $\Phi(F(tv)) = o(t^2)$, whereas $Q_A(tv) \le 0$, contradicting coercivity at $c = 1$ for small nonzero t.

This refutes the shortcut from individual to joint control, not the full matrix resistance criterion or any theorem retaining a disjoint-cycle hypothesis. The q-row perturbation also gives $N_{<c}(H) \le h + q$ for $0 < c \le 16m/n^2$. Extra zero or negative modes provide the necessary rank lower bound in (50); they do not alone bound the penalty coefficient.

28.6 Critical constants for rectangular mean observations

For a disjoint partition $\mathcal{P}$ of a finite counting space, its normalized means are $(Cf)_R = |R|^{-1/2}\sum_{x\in R} f(x)$. They satisfy $CC^* = \mathrm{I}$ and $C^*C = P$, the orthogonal projection onto functions constant on each cell. This concrete normalization matters for the penalty.

Theorem 28.5. *Let $L = L^*$ on a finite real Hilbert space, $c > 0$, and C a coisometry onto $\mathbb{R}^q$, $q \geq 1$. On $\mathcal{K} = \ker C$ and $\mathcal{O} = \operatorname{im} C^*$, write*

$$L - c\mathrm{I} = \begin{pmatrix} A_0 & X \\ X^* & B_0 \end{pmatrix}.$$

A finite $k \geq 0$ satisfying $L + kC^*C \geq c\mathrm{I}$ exists if and only if

$$A_0 \geq 0, \qquad X^*(\ker A_0) = 0. \tag{51}$$

When these hold, its smallest value is

$$k_{\min} = \max\{0, \lambda_{\max}(X^*A_0^\dagger X - B_0)\}.$$

Here $A_0^\dagger$ is the Moore–Penrose inverse on $\mathcal{K}$.

Proof. The restriction to $\mathcal{K}$ forces $A_0 \geq 0$. For $z \in \ker A_0$ the augmented form at $tz + y$ has linear term $2t\langle z, Xy\rangle$ and no quadratic term in t. Nonnegativity for every real t forces $X^*z = 0$. This is equivalent to $\operatorname{im} X \subseteq \operatorname{im} A_0$. Under that condition the full form at $z + y$ is exactly

$$\langle z + A_0^\dagger Xy, A_0(z + A_0^\dagger Xy)\rangle + \langle y, (B_0 + k\mathrm{I} - X^*A_0^\dagger X)y\rangle.$$

Finite spectral decomposition proves the criterion and the minimum. □

This is the classical generalized Schur criterion [22, Theorem 1], applied to actual means; it is not a new general matrix theorem. Strict positivity of A_0 suffices, whereas a weak bound at the target requires the extra compatibility in (51).

Example 28.6. *On $(\mathbb{Z}/2\mathbb{Z})^3$ take the same scalar coefficient in every gradient direction,*

$$a(x) = 6 + (-1)^{x_1} + 4(-1)^{x_2}, \qquad L = d^*ad.$$

Its values are $1, 3, 9, 11$, each twice, and its mean is $m = 6$; hence $1 \leq a \leq 2m$. Choose the four actual rectangles $R_{yz} = \{0,1\} \times \{y\} \times \{z\}$ and $(Cf)_{yz} = (f(0,y,z) + f(1,y,z))/\sqrt{2}$. The best coercivity constant of L on $\ker C$ is 16, but no finite quadratic penalty for these means attains the same constant on the full space. Every $0 < c < 16$ is attainable. Thus the supremum 16 is not attained.

Proof. For $S \subseteq \{1,2,3\}$ let $\chi_S(x) = (-1)^{\sum_{i\in S} x_i}/\sqrt{8}$. These are an orthonormal real Walsh basis. Retain all 24 oriented forward rows, including opposite rows at period two. The actual adjoint calculation gives

$$\begin{aligned} L\chi_S = 24|S|\chi_S &+ 4\big(|S| - \mathbf{1}_{1\in S}\big)\chi_{S\triangle\{1\}} \\ &+ 16\big(|S| - \mathbf{1}_{2\in S}\big)\chi_{S\triangle\{2\}}. \end{aligned}$$

Indeed $d_i\chi_S = -2\chi_S$ when $i \in S$, and the corresponding adjoint term is $2(a + T_i a)\chi_S$; absent directions contribute zero. In the kernel basis $(\chi_1, \chi_{12}, \chi_{13}, \chi_{123})$ the compression has the two blocks

$$\begin{pmatrix} 24 & 16 \\ 16 & 48 \end{pmatrix}, \qquad \begin{pmatrix} 48 & 32 \\ 32 & 72 \end{pmatrix}.$$

After subtracting 16I, the first is positive semidefinite with kernel vector $(-2, 1)$, and the second is positive definite, with first pivot 32 and determinant 768. This proves the complete kernel bound, not just a bound on selected test vectors.

Set $v = -2\chi_1 + \chi_{12}$ and $u = \chi_2$. Direct evaluation gives

$$Cv = 0, \quad \|v\|^2 = 5, \quad \|Cu\|^2 = \|u\|^2 = 1, \quad (L - 16\mathrm{I})v = 4u, \quad \langle u, (L - 16\mathrm{I})u\rangle = 8.$$

The augmented form on $u + tv$ is therefore $8 + 8t + k$, negative for $t < -(8 + k)/8$. No finite k works. For $c < 16$ the compressed block is strictly positive and Theorem 28.5 supplies a finite penalty; for $c > 16$, v violates coercivity already in the kernel. □

A quantitative penalty obstruction. Write $\delta = 16 - c > 0$. On the orthonormal pair $(v/\sqrt{5}, u)$ the augmented form has matrix

$$\begin{pmatrix} \delta & 4/\sqrt{5} \\ 4/\sqrt{5} & 8 + \delta + k \end{pmatrix}.$$

Its determinant gives the necessary bound

$$k \geq \max\{0, 16/(5\delta) - 8 - \delta\}. \tag{52}$$

This lower bound need not equal the full optimum. It proves divergence as the target approaches 16 from below without numerical extrapolation. If $F(0) = 0$, $DF(0) = C$ and $0 \leq \Phi(y) \leq k\|y\|^2$ near zero, local coercivity would imply $L + kC^*C \geq c\mathrm{I}$ by substituting sf, dividing by s^2 and taking $s \to 0$. Higher order terms with this same derivative do not repair the failure at $c = 16$.

The exact penalty rather than only a lower bound. For this same scalar example, the full optimum at $c = 16 - \delta$, $0 < \delta < 16$, is

$$k_{\min}(c) = \max\left\{c, \frac{16(8 + \delta)}{\delta(40 + \delta)} - 8 - \delta\right\}. \tag{53}$$

Here is a complete eight-dimensional check. Order the kernel basis as $(\chi_1, \chi_{12}) \oplus (\chi_{13}, \chi_{123})$ and the observed basis as $(\chi_\varnothing, \chi_2) \oplus (\chi_3, \chi_{23})$. At target 16 the kernel, observed and cross blocks of $L - 16\mathrm{I}$ split into

$$A_1 = \begin{pmatrix} 8 & 16 \\ 16 & 32 \end{pmatrix}, \quad B_1 = \mathrm{diag}(-16, 8), \quad X_1 = \begin{pmatrix} 0 & 0 \\ 0 & 4 \end{pmatrix},$$
$$A_2 = \begin{pmatrix} 32 & 32 \\ 32 & 56 \end{pmatrix}, \quad B_2 = \begin{pmatrix} 8 & 16 \\ 16 & 32 \end{pmatrix}, \quad X_2 = \mathrm{diag}(4, 8).$$

At $c = 16 - \delta$, add $\delta\mathrm{I}$ to each A_i, B_i, leaving X_i unchanged. The first observed Schur block is diagonal, with entries

$$16 - \delta, \qquad \frac{16(8 + \delta)}{\delta(40 + \delta)} - 8 - \delta.$$

The second observed Schur block cannot exceed $(16-\delta)\mathrm{I}$: $A_2 - 4\mathrm{I}$ has first pivot 28 and determinant 432, while $B_2 \geq 0$; hence $X_2^*(A_2+\delta\mathrm{I})^{-1}X_2 \leq 16\mathrm{I}$. The exact Schur minimum in Theorem 28.5 therefore proves both necessity and sufficiency of (53), including all eight dimensions. The branch transition is

$$\delta_0 = \frac{\sqrt{3529}-59}{3}: \quad k_{\min}(c) = c \text{ for } \delta \geq \delta_0,$$

and the other branch is larger for $0 < \delta < \delta_0$. Comparing the entries reduces to $3\delta^2+118\delta-16 \geq 0$. Partial fractions give the sharper critical expansion

$$k_{\min}(c) = \frac{16}{5(16-c)} - \frac{192}{25} + O(16-c) \qquad (c \uparrow 16). \tag{54}$$

The previous lower bound (52) alone is not the full optimum.

The general finite mechanism. For the blocks of Theorem 28.5 at the least kernel eigenvalue $c_* > 0$, let P project onto $\ker A_0$. At $c = c_* - \delta$,

$$k_{\min}(c) = \max\{0, \lambda_{\max}(X^*(A_0+\delta\mathrm{I})^{-1}X - B_0 - \delta\mathrm{I})\}.$$

Finite spectral splitting writes the inverse as $\delta^{-1}P$ plus a bounded remainder converging to $A_0^\dagger$. The variational formula for the largest eigenvalue consequently gives

$$\lim_{\delta\downarrow 0} \delta k_{\min}(c_* - \delta) = \|PX\|^2.$$

If $PX \neq 0$, the penalty diverges as $\|PX\|^2/\delta + O(1)$ and the critical target fails. If $PX = 0$, it converges to the finite critical Schur minimum. This uses classical finite Schur and resolvent principles, not a claim of a new general matrix theorem.

28.7 Transfer to actual flat coexact cochains

The preceding obstruction also holds on the actual finite coexact domain, with a proved form dictionary. Fix $N_c \geq 2$ and an actual coordinate unit $e \in \mathbb{R}^{N_c^2-1}$. At the trivial background, a physical one-cochain is $A_j(x) \in \mathbb{R}^{N_c^2-1}$, $j = 1,2,3$, with counting inner product. Use one unordered plaquette copy for each $i < j$:

$$(d_1 A)_{ij} = d_i A_j - d_j A_i, \qquad \operatorname{div} A = \sum_j (A_j - T_j^{-1} A_j).$$

Expansion of curl squares and cyclic summation by parts cancel the cross terms against divergence, proving the all-cochain identity

$$\|d_1 A\|^2 + \|\operatorname{div} A\|^2 = \sum_{i,j,x} \|d_i A_j(x)\|^2.$$

These are the actual repository conventions, not a supplied comparison. Take the actual componentwise repair coefficient $w(i,(x,j)) = a(x) - 1 \geq 0$. The total Hodge-plus-repair form is consequently

$$\begin{aligned} Q_w(A) &= \|d_1 A\|^2 + \|\operatorname{div} A\|^2 + \sum_{i,j,x} w(x)\|d_i A_j(x)\|^2 \\ &= \sum_{i,j,x} a(x)\|d_i A_j(x)\|^2. \end{aligned} \tag{55}$$

It is a direct sum of $3(N_c^2-1)$ scalar copies. Componentwise normalized means on the same four bars have output dimension $12(N_c^2-1)$, and the whole ambient kernel has coercivity constant 16 by the two blocks in Example 28.6. Restriction to coexact fields preserves the inequality; coisometry on the restricted domain is not assumed.

Proposition 28.7. *At target* 16*, no finite quadratic penalty for these actual componentwise rectangle means supplies coercivity even on* $\operatorname{im} d_1^*$*. The same is true for local nonlinear observations with derivative equal to those restricted means and a locally quadratic penalty. On that domain every* $0 < c < 16$ *is attainable, and the necessary bound* (52) *still applies.*

Proof. Construct actual fields $U_3 = ue$, $V_3 = ve$ from Example 28.6, with all other components zero. They are isometric copies of u, v and independent of x_3, so their actual backward divergence vanishes. More explicitly, use unordered two-cochains with only

$$B_{23}^U = -\tfrac{1}{2}\chi_2 e, \qquad B_{13}^V = (\chi_1 - \tfrac{1}{2}\chi_{12})e$$

nonzero. At side two $d_i^* = T_i^{-1} - \mathrm{I} = T_i - \mathrm{I}$. The adjoint of the displayed curl puts $d_i^* B_{ij}$ in component j and $-d_j^* B_{ij}$ in component i. Since $d_3^* B^U = d_3^* B^V = 0$ and each relevant character has multiplier -2, these actual potentials satisfy $d_1^* B^U = U$ and $d_1^* B^V = V$. Thus the entire line $U + tV$ is coexact.

Equation (55) and the actual means give the same augmented value $8 + 8t + k$. Scaling a violating coexact field towards zero also excludes locally bounded quadratic penalties and the stated nonlinear observations. Strict positivity on the ambient kernel supplies a finite penalty for every $c < 16$, which can then be restricted. The coexact field V excludes $c > 16$. The same two-dimensional determinant proves the bound. □

Scope and source of the transfer. The pinned repository gives the actual form identities in these two sources:

`componentwiseFlatGradient_eq_hodge`
`YangMills/RG/ComponentwiseFlatHodge.lean`

`actualTotalStiffness_energy`
`YangMills/RG/ActualPacketSupportEnvelope.lean`

Their proofs and the actual curl/divergence definitions were read; no new Lean compilation is claimed. Unit 166 checks the newly constructed physical stencils rather than relabelling a scalar diagnostic. For derivatives nd_i at $n = 2$, energies and targets scale by four, so 16 becomes 64 with rescaled potentials. This is finite rescaling, not a physical mass-gap claim.

The example has finite inverse moments at this one size. It refutes weak kernel sufficiency for these prescribed rectangles at the same target; it does not rule out another rectangular partition with the desired rank, or establish an unbounded-size violation of the critical rank-order conjecture. Only the trivial-background form has been identified here.

The sharp coexact leading coefficient. Let $k_{\mathrm{coex}}(c)$ be the minimum penalty for the actual means restricted to $\operatorname{im} d_1^*$. The witness pair just proved and the ambient direct-sum identity give, for $\delta = 16 - c \in (0, 16)$,

$$\max\{0, 16/(5\delta) - 8 - \delta\} \le k_{\mathrm{coex}}(c) \le k_{\min}(c).$$

Multiplying by δ and using (53) proves

$$\lim_{c\uparrow 16}(16-c)k_{\mathrm{coex}}(c)=\frac{16}{5}.$$

No coisometry is assumed for the restricted mean map. This proves the same leading coefficient on the actual coexact domain, not equality of the full ambient and restricted optima or of their next constant terms.

28.8 Reproduction and limitations

The source repository is THE ERIKSSON PROGRAMME PR186, stacked on PR185. All paths below are under `reproducibility/matrix-moment-sweep-20261106/`. Full source for this consolidation is pinned at

`66580c59ea9ef5f3c0374b339a165c9f89173db8`.

Units 162 and 163 each used two distinct own CPU allocations, four actual fresh sparse Git clones at one pinned source, and eight normal/optimized Python processes. Outputs matched within each unit. Both runtimes were deleted and disconnected after evidence export. Full stdout, stderr, exit sentinels, Git logs, exact source bytes and manifests are in the corresponding `evidence162-*.zip` and `evidence163-*.zip` archives.

Unit	Executed source prefix	Fixture cases	Mutations per process
162	`f644c093`	four	seven types
163	`378e879b`	four	seven types across four cases

The 163 fixtures are $(D,n,h)=(3,2,1),(3,3,1),(3,3,2),(4,2,2)$. They check actual Green and leverage identities, below/equal/above phase inertia, equality witnesses, gap and mean-observation coercivity, and coupled pair inertia. Unit 162 covers zero and near-zero stars, one large negative channel and rationally rotated local eigendirections. Every acceptance check raises explicitly on failure in both ordinary and `-O` execution. There are 28 fixture-mutation pairs per 163 process, not 28 distinct types. The two unit 163 recipe times were 53.83 and 29.98 seconds; requested-to-release times were 179.71 and 102.74 seconds. Billed compute units were not exposed.

Units 165 and 166 each used two distinct own CPU allocations, four actual fresh pinned Git clones and eight ordinary/optimized processes. Unit 165 checks the entire scalar torus matrix, the rectangular projection and three distinct mutations per process. Unit 166 constructs the actual 72×72 physical matrix at $D=3,n=2,N_c=2$, with 72 unordered curl rows, 24 backward divergence rows and 216 gradient rows. It checks the full Hodge-plus-repair identity, both actual adjoint potentials, divergence zero and the rectangular witnesses. Four distinct physical mutations are rejected per process: omit divergence, double the unordered curl, reverse a potential sign and put the field in the wrong component. Both 166 archives have 55 verified manifest subjects. Their recipe times are 5.28 and 8.19 seconds; requested-to-release times are 46.23 and 49.10 seconds. Both runtimes were disconnected after export. Billed compute units remain unknown. The analytic all-N_c coordinate embedding and nonlinear limit are not extra executed cases.

Finite replay checks these finite matrices, not the universal analytical statements. Arbitrary q, unequal weights, the all-size capped-family example and the nonlinear derivative consequence are

analytical arguments, not extra executed fixtures. Earlier 159 diagnostics used source downloads rather than fresh Git clones; their original limitation remains in their own records. No Lean compilation certifies this appendix. A full independently prepared audit of the imported Fourier adapter and the application form dictionary remains necessary before broader claims.

Unit 168 derives the exact penalty and critical residue analytically. Unit 169 then builds the actual scalar spatial matrix from all 24 forward rows and the actual pair-average projection using rational arithmetic. Generic full 8×8 positive-semidefinite elimination checks feasibility; separate rational row reduction finds observed null vectors, which prove that every strictly smaller penalty fails. Thus excess feasible penalties are rejected rather than labelled optimal. Nine distinct rational δ values cover both branches and bracket the transition without rounding. Six mutation types per process reject the old lower bound used as the optimum, an omitted constant-mode floor, doubled residue, excess penalty, erased actual coupling and collapsed opposite rows. Two own CPU allocations, four fresh pinned Git clones and eight ordinary or optimized processes give identical payloads. Each complete archive has 53 verified manifest subjects and 16 real zero-exit sentinels. Recipe times were 3.99 and 3.36 seconds; requested-to-release times were 36.32 and 41.92 seconds, including provisioning and export. Both runtimes were disconnected after export; billed units are unknown. Complete records are in `evidence169-*.zip` and `reproduction169-summary.json`. These finite checks do not certify every real δ, a new physical 72×72 replay, the universal resolvent lemma or the coexact limiting argument. Those remain the displayed analytic proofs. No new Lean build certifies these extensions.

28.9 Model assessments and review context

This appendix was prepared and checked analytically in the continuing Codex session that prepared its sources and diagnostics. The application does not expose the exact backend model identity for that preparation. There is no claim of independence, blinding, human peer review or editorial acceptance. Any assessment of this exact PDF is retained separately with its artifact hash and factual review context. An assessment of an earlier PDF does not transfer to this artifact.

The retained objections are substantive: the Cwikel mechanism and Schur principles are classical, so independent originality is unestablished; the observation from spectral coordinates is generally nonlocal; sparse-count and full-moment criteria must not be conflated; a diagonal graph-rank formula does not cover rotated eigenchannels; negative depth is neither universally dispensable nor universally necessary; the full off-diagonal interaction must be retained; critical kernel compatibility cannot be omitted at the target; finite diagnostics do not prove all-size statements; and no interacting physical form equality has been supplied. These limits are part of the result, not editorial acceptance conditions represented as satisfied.

The original positive critical rectangular construction remains open. The common-refinement route remains stopped after three unsuccessful units. The recorded PR29 baseline stays 20/41 with TermSource 0. This appendix adds supporting finite mathematical results and useful counterexamples without changing those records. This historical source continuation is integrated here; it does not itself assign a deposit, acceptance or new verified public release.

29 Local geometry of coarse observations

The preceding critical inverse-moment theorem constructs a global spectral observation. It does not construct a spatial partition. We now consolidate four complementary statements: an abstract limitation of spectral data, a local coexact point-support obstruction, a positive rectangular construction for a restricted coefficient class, and a limitation of minimum-coefficient certificates even with sharp unweighted constants. The original unrestricted critical rectangular question remains open.

Throughout the actual lattice statements, $D \geq 3$, $p = D/2$, $X = (\mathbb{Z}/n\mathbb{Z})^D$, $N = n^D$ and $h = D(N_c^2 - 1)$. The energy is the componentwise start-weighted form

$$\mathcal{E}_a(U) = \sum_{x,i,j} a(x)|D_i U_j(x)|^2, \qquad D_i U_j(x) = U_j(x + e_i) - U_j(x).$$

We do not silently substitute a variable-coefficient weighted curl form. For a rectangular partition $\mathcal{P}$, the actual normalized mean map is

$$C_{\mathcal{P}}(U)_{R,j} = \sqrt{|R|}\,\mathrm{mean}_R\, U_j, \qquad \dim F = h\#\mathcal{P}.$$

A global inequality with penalty $c\|C_{\mathcal{P}}(U)\|^2$ gives the original differentiable near-zero data with $K = c$. A kernel obstruction below applies to every finite locally upper-quadratic penalty, rather than only to this particular coefficient.

29.1 An isospectral obstruction to forgetting coordinate geometry

This first construction is deliberately a nonlocal scalar graph, not an admissible nearest-neighbour Hodge profile. Let n be even and set $e = N^{-1/2}\mathbf{1}$, $v = N^{-1/2}(-1)^{x_1+\cdots+x_D}$. For $0 < \alpha < c < \beta$ put

$$L = \beta(I - ee^*) - (\beta - \alpha)vv^*.$$

Its eigenvalues are 0 on e, α on v, and β on their orthogonal complement. On $e^\perp$ there is exactly one mode below c, and the actual scalar observation $\langle v, f\rangle$ with penalty coefficient $c - \alpha$ repairs the target. For $x \neq y$ the edge weight is

$$w_{xy} = \frac{\beta + (\beta - \alpha)\sigma(x)\sigma(y)}{N} > 0, \qquad \sigma(x) = (-1)^{x_1+\cdots+x_D}.$$

The row sums vanish, so this is a connected positive graph Laplacian. For $\alpha = c/2$, $\beta = 2c$, the weight ratio is only seven.

Let Π project onto the span of normalized cell indicators and let $\tau = \|\Pi v\|^2$. Constants belong to this span. If $\tau < 1$, the actual unseen mean-zero vector $u = (I - \Pi)v$ has

$$\|u\|^2 = 1 - \tau, \qquad \frac{\langle Lu, u\rangle}{\|u\|^2} = \alpha + (\beta - \alpha)\tau.$$

For a rectangle, the alternating sum in every cyclic coordinate interval has magnitude at most one. Consequently, for q cells,

$$\tau = \frac{1}{N}\sum_R \frac{(\sum_R \sigma)^2}{|R|} \leq \frac{q}{N}.$$

The exact finite quadratic-penalty criterion is stricter than a nonnegative kernel margin. Put $d = \beta - c$ and $r = \beta - \alpha$. On $e^{\perp}$, $L - cI + K\Pi = dI + K\Pi - rvv^{*}$. Since $dI + K\Pi > 0$, Cauchy–Schwarz after its square root shows that positivity is equivalent to

$$r\left[\frac{1-\tau}{d} + \frac{\tau}{d+K}\right] \le 1.$$

Solving gives

$$\text{some finite } K \text{ exists} \quad \Longleftrightarrow \quad r\tau > c - \alpha, \qquad K_{\min}^{e^{\perp}} = \frac{(\beta - c)(c - \alpha)}{(\beta - \alpha)\tau - (c - \alpha)}.$$

Here $K_{\min}^{e^{\perp}}$ is the optimum on the mean-zero subspace. On the full space the constant mode additionally requires $K \ge c$, so the optimum is $\max\{c, K_{\min}^{e^{\perp}}\}$. In particular, at $\tau = 1$ the restricted value $c - \alpha$ is below the full-space minimum; the constant field is an exact witness to the difference. At equality the second summand stays positive for every finite K. A locally upper-quadratic penalty cannot outperform this quadratic condition. Thus the displayed positive-weight example requires $q > N/3$ for every successful rectangular mean partition. Its restricted mean-zero observation rank is $q - 1$.

Replace the checkerboard signs by signs constant on two equal half-slabs. A vertex permutation conjugates the two graphs: their complete spectra and unlabelled weighted graphs coincide. The two half-slab means observe the sole low direction and admit a finite penalty. Therefore even the unlabelled graph, not just the eigenvalue count, omits the coordinate geometry needed for rectangular observations. This is an analytical counterexample to a locality-free conversion; it is not a counterexample to the original local coefficient question or a new Schur principle.

29.2 Growing point support on actual coexact fields in dimension three

The earlier line-constant point test can be strengthened using a local coexact correction, with an essential change of coefficient scale. For even $n \ge 4$, define the rational triangular profile and its periodic primitive by

$$f(t) = 1 - 4\,\mathrm{dist}_{\mathbb{Z}/n\mathbb{Z}}(t, 0)/n, \qquad F(0) = 0, \quad F(t) = -\sum_{r=1}^{t} f(r).$$

Then $\sum f = 0$, $|Df| = 4/n$, $D^{*}F = f$ and $\sum f^2 = n/3 + 8/(3n)$. The positive-pair curl and backward adjoint satisfy

$$(d_1 U)_{ij} = D_i U_j - D_j U_i, \qquad (d_1^{*} G)_1 = -D_2^{*} G_{12}, \quad (d_1^{*} G)_2 = D_1^{*} G_{12}.$$

Put $B_{12}^{0}(x) = F(x_1)e_0$ in one fixed internal unit coordinate, with other entries zero. Then $U^0 = d_1^{*} B^0$ has only component two, equal to $f(x_1)e_0$. It is explicitly coexact and divergence zero, and

$$\|U^0\|^2 \ge N/3, \qquad \mathcal{E}_{\mu}(U^0) = 16\mu N/n^2.$$

For an arbitrary sampled set S of s vertices, join two samples when their periodic maximum-coordinate distance is at most four. Let S_α be the resulting components, of sizes s_α, with representatives p_α. Write $N_r(S_\alpha)$ for the radius-r neighbourhood and set

$$\chi_\alpha = \mathbf{1}_{N_1(S_\alpha)}, \qquad (G_\alpha)_{12}(x) = \chi_\alpha(x)[F(x_1) - F((p_\alpha)_1)]e_0, \qquad U = U^0 - \sum_{\alpha} d_1^{*} G_\alpha.$$

This is an actual potential, not an assumed right inverse of divergence. At every sample, χ_α equals one at the site and its two backward neighbours, so $d_1^* G_\alpha$ agrees with U^0 in component two and is zero in component one. Other components contribute zero there. Thus $U_i(x) = 0$ for every $x \in S$ and every i.

Distinct $N_2(S_\alpha)$ are disjoint and their total size is at most $5^D s$. Outside them $U = U^0$, hence, with $C_D = 5^D$,

$$\|U\|^2 \geq N/3 - C_D s.$$

A component graph path has at most $s_\alpha - 1$ steps of cyclic length at most four. The unit increment bound for F gives $|F(x_1) - F((p_\alpha)_1)| \leq 4s_\alpha$ on $N_1(S_\alpha)$. Each correction has at most two components of magnitude $8s_\alpha$, so its squared norm, summed over disjoint supports, is at most $128 C_D \sum_\alpha s_\alpha^3$. Expanding finite differences gives the unit flat Hodge identity, and $\mathcal{E}_\mu(A) \leq 4D\mu\|A\|^2$. The quadratic inequality for U^0 minus the correction therefore gives

$$\mathcal{E}_\mu(U) \leq 32\mu N/n^2 + A_D \mu \sum_\alpha s_\alpha^3, \qquad A_D = 1024 D 5^D. \tag{G1}$$

All support and energy constants have been derived, without a separation or correction-operator hypothesis on the samples.

Proposition 29.1. *Fix $c > 0$ and take $a = \mu_n = cn^2/384$ for sufficiently large even n. Any successful point-supported original observation on the actual coexact domain must have sampled support*

$$s \geq \min\left\{ \frac{n^D}{12 C_D}, (64/A_D)^{1/3} n^{(D-2)/3} \right\}.$$

The same conclusion holds when only the derivative at zero factors through the point values. In particular support diverges in $D = 3$.

Proof. The profile satisfies $1 \leq a \leq 2\mu_n$, its mean is μ_n, and its normalized critical inverse moment is exactly one. If $s \leq N/(12C_D)$, then $\|U\|^2 \geq N/4 > 0$. Point-supported maps with $C(0) = 0$ annihilate tU; derivative-supported maps satisfy $C(tU) = o(t)$. Divide the local coercivity inequality by t^2 and use the quadratic upper bound on its penalty. Since $32\mu_n N/n^2 = cN/12$, (G1) implies $\sum s_\alpha^3 \geq (64/A_D) n^{D-2}$. Finally $\sum s_\alpha^3 \leq s^3$. If the initial restriction on s fails, the first alternative already holds. □

The family scale is $\mu_n = cn^2/384$, not the older $\mu = n^2, c \leq 16$ coexact regime, where a uniform Fourier gap can already suffice without observations. No fine-space repair is covered here. On this same new family, observing all real Fourier modes with centred $|k|^2 < 24$ has dimension at most $h9^D$ and gives target c with penalty coefficient c: the omitted modes obey $\mu_n \lambda_n(k) \geq c|k|^2/24$. Thus this separates bounded point support from bounded global spectral rank, not from rectangular averages or point samples plus complementary means. The arbitrary-cluster exponent is a safe bound, not claimed sharp.

29.3 A fixed-gap rectangular construction with arbitrary weak geometry

Lemma 29.2. *Suppose $a \geq A > 0$ outside any set S of k vertices, with $a > 0$ everywhere. Put $L = \min\{n, \max\{1, \lfloor \sqrt{A/c} \rfloor\}\}$. There is an actual rectangular partition satisfying*

$$c\|U\|^2 \leq \mathcal{E}_a(U) + c\|C_{\mathcal{P}}(U)\|^2 \quad \text{for every cochain,} \qquad \#\mathcal{P} \leq \lceil n/L \rceil^D + 2Dk. \tag{G2}$$

No shape, connectedness or separation of S is assumed.

Proof. Start with coordinate intervals of length at most L. To isolate a sample in its current rectangle, cut just before and after its first coordinate, retain the side slabs, and continue only in the central width-one slice. Repeat in each direction. At most $2D$ actual binary cuts isolate that vertex; each increases the leaf count by one. Previously isolated vertices stay singleton leaves. Induction gives disjoint coverage, the asserted count, and every nonsingleton leaf avoids S. This is not a product of all coordinate cut sets or a full equal-child cube tree.

On a length-l path, the counting variance identity, path Cauchy–Schwarz and the pair count $r(l-r) \le l^2/4$ give

$$\sum |u-\bar{u}|^2 = \frac{1}{l}\sum_{r<s} |u_r - u_s|^2 \le \frac{l(l-1)}{4}\sum |Du|^2 \le \frac{l^2}{4}\sum |Du|^2.$$

Commuting coordinate-averaging projections P_i satisfy $I - \prod_i P_i \le \sum_i (I - P_i)$ as quadratic forms; their simultaneous zero-one eigenvalues prove it. Tensorization bounds each leaf's variance by $L^2/4$ times its unweighted internal-edge energy. Every retained positive edge starts outside S, hence has coefficient at least A. If $L \ge 2$ this bounds c times residual by one quarter of internal weighted energy; if $L = 1$ every residual is zero. Sum the disjoint internal edge sets and discard nonnegative cross-leaf/wrap edges. The exact orthogonal mean decomposition proves (G2). Arbitrary aspect ratios cause no extra hypothesis. □

A safe mesh count is

$$\lceil n/L \rceil^D \le 4^D \max\{1, c^p\} \max\{1, N/A^p\}.$$

For $A/c < 4$, $L = 1$ and $N \le (4c)^p N/A^p$; otherwise, unless $L = n$, $L \ge \sqrt{A/c}/2$ and $\lceil n/L \rceil \le 1 + 2n\sqrt{c/A}$. These estimates also include small sizes and all-singleton cases.

Theorem 29.3. *Fix $B \ge 1$ and $0 < \theta \le 1$ independently of n, μ. In addition to the original critical-moment profile hypotheses, assume*

$$\text{every } x \text{ satisfies } a(x) \le B \text{ or } a(x) \ge \theta\mu. \tag{G3}$$

Then the original rectangular inequality holds globally with $K = c$ and

$$\dim F \le h 2^p \left[4^D \max\{1, c^p\}\theta^{-p} + 2DMB^p\right] \max\{1, N/(1+\mu)^p\}.$$

Proof. For $S = \{a \le B\}$ the moment hypothesis derives $|S| \le MB^p N/\mu^p$. Off S, (G3) gives $a \ge \theta\mu$. Insert these facts in (G2), use $\theta \le 1$ and $\max\{1, N/\mu^p\} \le 2^p \max\{1, N/(1+\mu)^p\}$. □

In particular every original profile taking values $\{1, H\}$ is included: its actual mean gives $H \ge \mu$. Arbitrary weak-set geometry replaces the earlier single weak cube. This does not cover every two-valued family whose lower value grows with n; (G3) is additional and is not derived from the critical moment alone.

The limitation is explicit. At $c = 1$, set $n = t^2$, $\mu = t^4$, $b = t^2$, let the weak cube have side t and fraction $s = t^{-D}$, and set $a = b$ there and $H = (\mu - sb)/(1-s)$ elsewhere. The exact mean is μ, $1 \le a \le 2\mu$, and the critical moment is less than two. For the specifically prescribed mesh of side $\max(1, \lfloor\sqrt{\tau/c}\rfloor)$, followed by isolation of exactly $a < \tau$, the count is at least t^D: if $\tau \le b$, its mesh side is at most t; if $\tau > b$, all t^D weak sites are isolated. This is not a numerical lower bound for every sufficient uniform mesh. For even t and $\tau = b$, a mesh of side $2t$ has only $(t/2)^D$

cells and passes the actual bound $L(L-1)/(4b) = 1 - 1/(2t) < 1$. This corrects the broader sentence in the preserved revision187 and retains its full adverse review. Nevertheless the weak cube itself and its D first-coordinate-outside complement slabs form $D+1$ rectangles. Their variance coefficients are at most $t^2/(4b) = 1/4$ or $n^2/(4H) \le 1/4$, respectively. They therefore give the actual target with $K = 1$. This defeats the named threshold method, not unrestricted rectangular means.

29.4 Actual detours and the failure of minimum-only certificates

We next leave coefficient minima behind, using the classical path comparison mechanism[23]. Let a weak set W have periodic maximum-coordinate separation at least three, let $a \ge H > 0$ off W, and allow arbitrary nonnegative coefficients on W. Denote by $\mathcal{E}_{\mathrm{str}}$ and $\mathcal{E}_{\mathrm{wk}}$ the unweighted edge-square sums with positive starts outside and inside W.

Lemma 29.4. *For $D \ge 2, n \ge 3$ under these hypotheses,*

$$\mathcal{E}_{\mathrm{wk}} \le 3\mathcal{E}_{\mathrm{str}}, \qquad \mathcal{E}_0 \le 4\mathcal{E}_{\mathrm{str}}, \qquad \mathcal{E}_a \ge (H/4)\mathcal{E}_0. \tag{G4}$$

Proof. Let $\sigma(i) = i + 1$ cyclically. Bypass every weak-start edge $x \to x + e_i$ by the actual three-edge path

$$x \longrightarrow x - e_{\sigma(i)} \longrightarrow x + e_i - e_{\sigma(i)} \longrightarrow x + e_i.$$

Its positive starts are $x - e_{\sigma(i)}$ twice in different directions, and $x + e_i - e_{\sigma(i)}$. They are strong by separation. The first edge is traversed backwards, so its coefficient is not evaluated at x.

For a fixed positive edge (y, k), the three possible weak sources are

$$y + e_k, \qquad y + e_{\sigma(k)}, \qquad y - e_{\sigma^{-1}(k)} + e_k.$$

For each type its direction i is uniquely fixed. These candidate sources are distinct and pairwise at periodic maximum-coordinate distance at most two. At most one is weak. Thus every strong positive edge is used at most once, counting uses rather than just distinct paths. Cauchy–Schwarz along each three-step path and summation prove $\mathcal{E}_{\mathrm{wk}} \le 3\mathcal{E}_{\mathrm{str}}$. The remaining assertions follow directly, componentwise. Periodic wraps are already included in the path identities. □

The Fourier torus gap gives

$$\mathcal{E}_a(U) \ge \frac{H}{4}\frac{16}{n^2}\|U - \mathrm{mean}_X\, U\|^2. \tag{G5}$$

If $4H/n^2 \ge c$, one global normalized mean gives the original inequality with $K = c$, on all cochains. No supplied flow, comparison or kernel margin remains as an unproved assumption.

For a nonsingleton rectangle R, let $\lambda_1^0(R)$ be the sharp unweighted internal Neumann gap after cutting each coordinate interval into a path. A standard sufficient local test is

$$\min_R(a)\lambda_1^0(R) \ge c. \tag{G6}$$

Singletons have zero residual. On a longest side of length l, the actual mean-zero ramp $u(r) = r - (l-1)/2$ has energy $l - 1$ and counting squared norm $l(l^2-1)/12$. Tensor copies cancel, giving

$$\lambda_1^0(R) \le 12/[l(l+1)] < 12/l^2. \tag{G7}$$

Thus the following obstruction applies even to the best unweighted constant, not just a loose path or trace estimate.

Proposition 29.5. *For every fixed $D \geq 3$ and $M > 1$ there are original admissible profiles at $c = 1, \mu = n^2$ for which a single rectangular mean gives original data with $K = 1$, but every partition passing (G6) has an unbounded number of cells. No fixed multiplicative loosening of (G6) rescues its uniform rank order. This does not refute the original rectangular question.*

Proof. Choose an integer $m \geq 4$ with $1 + m^{-D} \leq M$, and take large multiples t of m. Put $n = t^2$, $\mu = t^4$, $b = t^2$ and

$$W = \{0, mt, 2mt, \dots, n - mt\}^D, \quad k = (t/m)^D, \quad s = \frac{1}{m^D t^D},$$
$$a = b \text{ on } W, \qquad H = \frac{\mu - sb}{1 - s} \text{ elsewhere.}$$

Then $H \geq \mu$, $1 \leq a \leq 2\mu$, and the arithmetic mean is exactly μ. The normalized critical inverse moment is

$$s(\mu/b)^p + (1 - s)(\mu/H)^p = m^{-D} + (1 - s)(\mu/H)^p < 1 + m^{-D} \leq M.$$

Weak sites have periodic separation mt. By (G5), $\mathcal{E}_a$ is at least four times global mean-zero squared norm, so one mean suffices. The target rank scale is exactly one.

Every weak-containing nonsingleton leaf passing (G6) has, by (G7), $l^2 < 12b = 12t^2$, hence $l < \sqrt{12}t < mt$. It contains at most one weak site; a singleton does too. Therefore every passing partition has $\#\mathcal{P} \geq k = (t/m)^D \to \infty$, with arbitrary cut locations and translated cyclic intervals allowed. For the test $\eta \min_R(a)\lambda_1^0(R) \geq 1$ with any fixed $\eta > 0$, choose also $m > \sqrt{12\eta}$. The same proof applies. □

If a local test retains periodic edges on a full-period side, the ramp bound (G7) must not be used on that cycle. Its first Fourier mode instead gives a gap upper bound $4\pi^2/l^2 < 40/l^2$. Choosing $m > \sqrt{40\eta}$ gives the same conclusion uniformly. Alternatively, for every fixed m, a weak-containing full-period side fails that test for sufficiently large t, since its upper bound is $4\eta\pi^2/t^2 < 1$. No periodic edge was discarded in the actual global comparison (G4)–(G5).

A separate example clarifies non-necessity of the scalar moment for an individual family. Let n run through multiples of three, take $W = \{0, 3, \dots, n - 3\}^D$, $s = 3^{-D}$, $a = 1$ on W, $\mu = n^2$ and $H = (\mu - s)/(1 - s)$ elsewhere. The original mean and coefficient bounds hold, and (G5) again gives one-cell data at $c = K = 1$, while

$$\text{mean}((\mu/a)^{D/2}) \geq 3^{-D} n^D \longrightarrow \infty.$$

This does not remove the moment assumption from its unrestricted sufficiency theorem or contradict its sharp exponent among universal scalar inverse-moment assumptions. Here separation is a verified feature of the explicit profiles, not a consequence of an arbitrary moment bound.

29.5 Reproduction and remaining limitations

The analytical sources are the pinned notes at `aaa9e278f2c7` (nonlocal graph), `bfc4ada98dd9` (coexact point support), `ee2b86c0e62b` (fixed-gap rectangles) and `29fc9366bf4e` (detours and minimum certificates). Classical comparison is attributed to Section2B and Example2.1 of[23], which already use short bypasses and edge congestion. No new comparison principle or independent originality is claimed. The present finite start-weighted and coordinate-geometry adapters are derived in the proofs above.

Replay	Executed source	Cases	Rejected mutations
178	`115fedf62ad2`	17	7
180	`1cf122830b45`	12	5 plus2 solver attacks
182	`4422b0e19acf`	11	7 plus2 solver attacks

Each replay used four fresh pinned clones across two own CPU Colab allocations and eight byte-identical normal/optimized payloads. Source hashes, full stdout/stderr, actual decimal child exit codes, archive manifests, resource records and disconnections are retained in the reproducible package. Windows performed editing, git and hash/readback work, not scientific execution. In180 the two host fingerprints matched; in178 and182 they differed. Neither observation establishes independent model review or physical-machine independence.

Replay178 checks the actual coexact potential, every pin component, backward divergence, norm and Hodge/gradient identities. Replay180 includes ten exact 64×64 positivity checks and two intermediate island cases; no large full matrix was claimed. Replay182 includes four exact small positivity checks, the full path/preimage/edge-use dictionary, sparse large-family coefficient checks and the path-versus-cycle ramp exception. The nonlocal graph proof has no corresponding executed replay in this consolidation. Finite selected cases do not prove all-size statements, formalize these additions, establish novelty or solve the unrestricted rectangular question.

A useful research direction is to derive actual weighted local energy or capacity estimates. The counterexamples above concern the specified eigenvalue-only, point-support, full-tree, trace, threshold-singleton and minimum-coefficient methods in isolation; they do not exclude all hybrids or classify every possible positive mechanism. These scoped negative results preserve the successful partitions on the same examples. They make no statement about interacting gauge forms, a continuum limit or a physical mass gap. Section 30 supplies such an actual-energy construction for arbitrary one-coordinate laminates and proves a separate barrier for strong Green-moment certificates. Extending it to unrestricted transverse variation is partly supplied by Section 35; the fully variable case remains open.

30 Rectangular means under critical inverse moments

This grouped continuation belongs to the existing coarse-information paper dossier [24]. It combines the complementary research units 184 and 185 rather than proposing a deposit for each corollary. Let $D \geq 3$, $n \geq 2$, $X = (\mathbb{Z}/n\mathbb{Z})^D$, $N = n^D$ and $p = D/2$. All differences have unit spacing and all norms use counting measure. Set

$$\mathcal{E}_a(u) = \sum_{x \in X} \sum_{j=1}^{D} a(x)|u(x + e_j) - u(x)|^2, \qquad a(x) > 0.$$

For a disjoint rectangular partition $\mathcal{P}$ define the actual normalized arithmetic means $C_{\mathcal{P}}u = (\sqrt{|R|}\,\bar{u}_R)_{R \in \mathcal{P}}$. Our positive construction gives, simultaneously for all fields,

$$c\|u\|^2 \leq \mathcal{E}_a(u) + c\|C_{\mathcal{P}}u\|^2, \qquad c > 0. \tag{56}$$

For a fixed h-component field the output dimension is $h\#\mathcal{P}$ over the selected scalar field; complex dimension must be doubled when counted over the reals. Scalar estimates sum over components, and hence restrict to any actual coexact subspace without assuming an abstract surrogate.

The original question asks for a target-order partition under $\mu = \mathrm{mean}(a) \geq 1$, $1 \leq a \leq 2\mu$ and $\mathrm{mean}((\mu/a)^p) \leq M$, with cell count bounded by a constant depending on D, c, M times $\max(1, N/(1+\mu)^p)$. We solve here only the subclass $a(x) = \alpha(x_1)$. Lamination is an explicit retained hypothesis. The Green-moment results below concern the cost of a sufficient certificate, not a lower bound for every successful observation or for coexact rank.

31 A constructive theorem for arbitrary laminates

Put $d = D - 1$ and

$$r = (2/3)^d, \qquad B_D = \frac{2\,6^d}{1 - r}.$$

Theorem 31.1. *For every positive sequence $\alpha(0), \ldots, \alpha(n-1)$, set $a(x) = \alpha(x_1)$. There is a disjoint axis-aligned rectangular partition independent of u that satisfies* (56) *and*

$$\#\mathcal{P} \leq 2^d + B_D(c/2)^p \sum_{x \in X} a(x)^{-p}. \tag{57}$$

If the original critical inverse-moment assumptions hold, its output rank is at most

$$h\big[2^d + B_D(c/2)^p 2^p M\big] \max\big(1, N/(1+\mu)^p\big).$$

The upper bound $a \leq 2\mu$ is not needed for (57)*. Target order here means the requested upper order, not a matching lower bound for every individual coefficient.*

31.1 Actual local energy estimate

For an interval I of length m, write $R_I = \sum_{x \in I} \alpha(x)^{-1}$. Include its last starting vertex in this sum: although its longitudinal edge may leave the cut interval, its coefficient is needed for transverse averaging. Internal rectangle energies keep only edges within cut paths; periodic closing and cross-cell edges are discarded as nonnegative contributions.

Lemma 31.2. *If $R = I \times Z$ and every transverse side has length at most m, then*

$$\operatorname{Var}_R(u) := \sum_R |u - \bar{u}_R|^2 \le \frac{mR_I}{2}\mathcal{E}_{a,R}(u). \tag{58}$$

The constant $mR_I/2$ is an inverse-gap bound.

Proof. Let $v(y) = m^{-1}\sum_{x\in I} u(x,y)$. Orthogonal decomposition gives

$$\operatorname{Var}_R(u) = \sum_{y\in Z} \operatorname{Var}_I(u(\cdot,y)) + m\operatorname{Var}_Z(v).$$

The pairwise variance identity and weighted Cauchy along a path give $\operatorname{Var}_I \le (m-1)R_I\mathcal{E}_1/2 \le mR_I\mathcal{E}_1/2$. For an unweighted path of length ℓ, the same identity bounds the coefficient of its kth edge by $k(\ell-k) \le \ell^2/4$: use the path Cauchy bound with its length and then bound that length divided by ℓ by one. Tensorizing these path inequalities gives $\operatorname{Var}_Z(v) \le m^2\sum_{y,j>1}|D_j v(y)|^2/4$; the tensorized constant is the largest coordinate constant, not their sum. Weighted Cauchy yields

$$|D_j v(y)|^2 \le \frac{R_I}{m^2}\sum_{x\in I}\alpha(x)|D_j u(x,y)|^2.$$

Thus $\operatorname{Var}_R \le (mR_I/2)\mathcal{E}_1 + (mR_I/4)\mathcal{E}_\perp$, proving the lemma. This is the step where a common coefficient sequence on all fibres is used. □

31.2 Partition algorithm and critical charge

Start from the interval $\{0,\dots,n-1\}$. Split a node of length $m \ge 2$ into consecutive lengths $\lfloor m/2\rfloor, \lceil m/2\rceil$ exactly when $mR_I > 2/c$. Accept equality. Never split a singleton. For each final interval of length m, partition every transverse coordinate into intervals of length m with a possible shorter last interval. Their products form an actual disjoint cover. A nonsingleton leaf satisfies $c\operatorname{Var}_R \le \mathcal{E}_{a,R}$ by Lemma 31.2; singleton longitudinal leaves produce only singleton cells, with zero variance. Summing and using $\|u\|^2 = \sum_R \operatorname{Var}_R(u) + \sum_R |R||\bar{u}_R|^2$ proves (56).

For the count set $b(I) = \sum_I \alpha^{-p}$ and $w(I) = m^{-d}$. Hölder gives $R_I \le m^{1-1/p}b(I)^{1/p}$. Every split therefore satisfies

$$w(I) < (c/2)^p b(I), \qquad d = 2p-1. \tag{59}$$

Let T consist of internal nodes with no internal children. These intervals are disjoint. Assign each internal node to a terminal internal descendant. Each child has size at most $2/3$ of its parent, also when sizes are odd. Nodes assigned to a fixed terminal node lie on its ancestor path at distinct depths, so their weights sum to at most $w(J)/(1-r)$. Consequently

$$\sum_{I \text{ internal}} w(I) \le \frac{(c/2)^p}{1-r}\sum_{x_1}\alpha(x_1)^{-p}.$$

Each child also has size at least $m/3$. Telescoping leaf weights over all splits gives

$$\sum_{I \text{ leaf}} w(I) \le n^{-d} + 2\,3^d \sum_{I \text{ internal}} w(I).$$

The exact transverse cost of a leaf is $\lceil n/m\rceil^d \le (2n/m)^d$. Multiply the last display by $(2n)^d$ and use $\sum_X a^{-p} = n^d\sum_{x_1}\alpha^{-p}$ to obtain (57). Finally $\sum_X a^{-p} \le MN/\mu^p$ and $\mu^{-p} \le 2^p(1+\mu)^{-p}$ prove the stated target bound. This completes the proof of Theorem 31.1 without a logarithmic charge at each tree level.

31.3 A thin weak layer beyond minimum certificates

Take integer $t \geq 2$, $n = t^D$, $\mu = n^2$, $b = n^{2-2/D}$, and

$$\alpha(0) = b, \qquad \alpha(i) = H = \frac{n\mu - b}{n-1} \quad (i > 0).$$

Then $1 \leq \alpha \leq 2\mu$, its exact mean is μ, and

$$\operatorname{mean}((\mu/a)^p) = 1 + (1 - 1/n)(\mu/H)^p < 2. \tag{60}$$

The uniform bound for this unbounded sequence is $M = 2$; the moments tend to two, so a single fixed $M < 2$ is not asserted. Moreover

$$nR_{\text{root}} = n^{-1+2/D} + \frac{(n-1)^2}{n^2 - b/n} < 2.$$

Thus at $c = 1$ the algorithm accepts the root and uses one mean. Yet $b \to \infty$ and $b/\mu \to 0$, excluding every fixed coefficient-gap condition $a \leq B$ or $a \geq \theta\mu$ for fixed $B, \theta > 0$ eventually.

Consider instead any partition whose cells pass the minimum-only sufficient test $(\min_R a)\lambda_1^0(R) \geq 1$, where λ_1^0 is the actual unweighted internal gap. For a singleton the mean-zero condition is vacuous. Every cell meeting the weak layer has minimum b. A linear path test gives $\lambda_1^0 \leq 12/\ell^2$ in a cut coordinate of length $\ell \geq 2$; a Fourier test gives $\lambda_1^0 < 40/\ell^2$ in a full cycle coordinate. Thus each transverse side of such a cell is at most $\sqrt{40b}$. Covering the n^d weak-layer vertices forces

$$\#\mathcal{P} \geq 40^{-d/2} n^{d/D}.$$

This is a limitation of that certificate, since the preceding one-cell construction works on the same coefficients.

32 The cell cost of strong Green moments

Now take homogeneous $a = \mu = n^2$. These original coefficient data have $M = 1$ and target order one. On a rectangle R use cut coordinate paths, or retain a periodic edge only when its coordinate spans the full period. Let L_R be its actual weighted Laplacian and $G_R = L_R^\dagger$, zero on the constant kernel. For $s > 0$ let

$$T_s(R) = \operatorname{tr}(G_R^s) = \sum_{\lambda_j(L_R) > 0} \lambda_j(L_R)^{-s}.$$

For a singleton set $T_s = 0$. The strong certificate $T_s(R) \leq 1$ indeed implies $\operatorname{Var}_R(u) \leq \mathcal{E}_{a,R}(u)$, because its first positive inverse eigenvalue raised to s is one summand. Define $q_s(n)$ as the minimum number of cells in any disjoint rectangular partition passing this test. No hierarchy, aspect ratio or placement restriction is imposed.

Theorem 32.1. *For fixed $D \geq 3$ and $s > 0$,*

$$q_s(n) = \begin{cases} \Theta_{D,s}(n^{D-2s}), & 0 < s < D/2, \\ \Theta_D(\log n), & s = D/2, \\ \Theta_{D,s}(1), & s > D/2. \end{cases}$$

The actual gap and a weak critical counting certificate both allow one root cell on this family. This phase diagram concerns sufficient certificate cost, not necessary observation rank.

32.1 Finite spectra and all-partition lower bounds

The cut path of length ℓ has eigenvectors $\cos(\pi k(x+1/2)/\ell)$ and eigenvalues $4\sin^2(\pi k/(2\ell))$, $0 \le k < \ell$. The endpoint identity follows from cosine addition and $\sin(\pi k) = 0$. Distinct eigenvalues give a full orthogonal basis. A full cycle has the usual Fourier eigenvalues $4\sin^2(\pi k/n)$. At $n = 2$ retain both positive-start periodic edges; its nonzero eigenvalue is four, while the cut path eigenvalue is two. Tensor products give the cut-rectangle spectrum

$$\lambda_k = 4n^2 \sum_{i=1}^{D} \sin^2 \frac{\pi k_i}{2\ell_i}, \qquad 0 \le k_i < \ell_i, \quad k \ne 0. \tag{61}$$

Adding full-period edges increases ordered eigenvalues by min-max and keeps the constant kernel. Comparing eigenvalues suffices for every $s > 0$; we do not assume an operator-power order valid for arbitrary powers.

Every such internal graph has $\lambda_{\max} \le 4Dn^2$. Hence, with $V = |R|$, $T_s(R) \ge (V-1)/(4Dn^2)^s$. A passing partition satisfies

$$q \ge \frac{n^D}{1 + (4Dn^2)^s}. \tag{62}$$

This proves the subcritical lower bound for arbitrary shapes.

For the endpoint set $p = D/2$, $m = \min_i \ell_i$, $J(m) = \max(0, \lfloor \log_2 m \rfloor - 2)$ and $A_D = 2^{-D}(256D)^{-\lceil D/2 \rceil}$. We claim

$$T_p(R) \ge A_D (V/n^D) J(m). \tag{63}$$

It is vacuous if $m < 8$. Otherwise for $j = 0, \dots, \lfloor \log_2 m \rfloor - 3$ put $t_j = 2^{j+1}/m$ and select the disjoint mode blocks $\lceil t_j \ell_i \rceil \le k_i < \lceil 2t_j \ell_i \rceil$ for every i. Their sizes are at least $V t_j^D / 2^D$: for $x = t_j \ell_i \ge 2$, $\lceil 2x \rceil - \lceil x \rceil \ge x/2$. All selected frequencies have $k_i/\ell_i < 2t_j \le 1/2$ and satisfy $\lambda_k < 256 D n^2 t_j^2$, under both path and retained-cycle spectra. Their contribution is at least $A_D V/n^D$ per block. This proves (63), including arbitrary aspect ratios.

In a partition with q cells, those of volume at least $N/(2q)$ cover at least half of X. Each of their sides is at least $n/(2q)$. Summing (63) over these cells and using $T_p \le 1$ gives

$$q \ge \frac{A_D}{2} \max\{0, \lfloor \log_2(n/(2q)) \rfloor - 2\}.$$

For $n \ge 2^{16}$ this implies $q \ge (A_D/8) \log_2 n$. Indeed if $q \ge \sqrt{n}$ this is immediate; otherwise the expression in braces is at least $(\log_2 n)/2 - 4 \ge (\log_2 n)/4$.

32.2 Matching actual meshes

Use a Cartesian mesh with all side lengths at most an integer $M \le n$, allowing shorter last intervals. From (61) and $\sin t \ge 2t/\pi$ on $[0, \pi/2]$, group modes by $j = \max_i k_i$ to get

$$T_s(R) \le D 2^{D-1} 4^{-s} (M/n)^{2s} \sum_{j=1}^{M-1} j^{D-1-2s}. \tag{64}$$

There are at most $D(j+1)^{D-1} \le D 2^{D-1} j^{D-1}$ modes in the group. When $2s < D$, the sum is at most a constant depending on D, s times M^{D-2s}. Choose $M = \max(1, \lfloor \varepsilon n^{2s/D} \rfloor)$ with fixed sufficiently small $\varepsilon > 0$. Then every trace passes, and the cell count is $O_{D,s}(n^{D-2s})$. Floor-zero cases use singletons. When $2s = D$, the sum is at most $1 + \log M$. Choosing

$M = \max(1, \lfloor \varepsilon n/(1+\log n)^{1/D} \rfloor)$ gives $O_D(1+\log n)$ cells. When $2s > D$, the sum is at most $1 + (2s-D)^{-1}$; choose $M = \max(1, \lfloor \varepsilon n \rfloor)$ to get a constant cell count. Retained full-period edges only decrease the traces. These constructions and the lower bounds prove Theorem 32.1.

32.3 Successful weak endpoint and exact finite distinctions

For the cut root, $\lambda_k \geq 4|k|^2$, so its true positive gap is at least four. Define

$$W_p = \sup_{\lambda > 0} \lambda^{-p} \#\{0 < \lambda_j(L_X) < \lambda\}.$$

There are no such eigenvalues for $\lambda < 4$; for $\lambda \geq 4$ their number is at most $(1 + \sqrt{\lambda}/2)^D \leq \lambda^p$. Thus $W_p \leq 1$. The same holds after adding periodic edges. Integrality of this counting function excludes any eigenvalue below one, so this weak criterion really certifies the successful one-mean inequality. The weak/strong ideal distinction is classical [28]; its use does not solve the original coefficient-to-rectangle construction.

For $D = 4, n = 4$, the periodic coordinate spectrum at $a = 16$ is $0, 32, 64, 32$. The root eigenvalues $32m$ have multiplicities $\binom{8}{m}$. Its gap is 32, but $T_1(X) \geq 56/96 + 70/128 = 217/192 > 1$; sixteen cut length-two boxes instead have $T_1 = 103/384$ each. At the critical endpoint an exact separate example takes $n = 2$, $\mu = 4$, $c = 16$ and test $16^p T_p(R) \leq 1$. In this exact example retain both positive-start periodic edges in every full coordinate. With all full coordinates cut instead, a two-vertex cell has eigenvalue eight and test value $2^p > 1$, so the exact minimum below would be 2^D; the asymptotic orders are unchanged. A rectangle with r full coordinates has spectrum $16m$ with multiplicity $\binom{r}{m}$, so this test passes exactly when $r \leq 1$. Therefore its minimum cell count is 2^{D-1}, whereas the actual one-mean inequality works with $K = c = 16$. For $D = 4$ the rejected root value is $865/144$. Changing a fixed c changes the constants in the phase diagram, not its orders. The ordinary identity $\sum_{x<y} R_{\mathrm{eff}}(x,y) = |R| \operatorname{tr}(G_R)$ follows by expanding $(e_x - e_y)^T G_R (e_x - e_y)$ and using $G_R \mathbf{1} = 0$; it only identifies the all-pairs resistance sum with the $s = 1$ test, not every resistance method.

33 A coexact falsifier for representative fibres

Lamination cannot be replaced by sampling one favourable fibre. In $D = 3, n = 4$ let $c = 15/2$ and set $a = 1$ on $\{0,1,2\}^3$, $a = 64$ elsewhere. Its exact mean is $\mu = 2395/64$, $1 \leq a \leq 2\mu$, and

$$\operatorname{mean}((\mu/a)^{3/2}) < \frac{27}{64}\frac{1178}{5} + \frac{37}{64} = \frac{31991}{320} < 100.$$

Here $\mu < 38$ and $\sqrt{\mu} < 31/5$. The fibre through $(3,3)$ in the transverse coordinates has all coefficients 64 and would pass $nR = 1/4 < 2/c = 4/15$, leading to a false one-cell conclusion.

Let $G = \delta_{(1,1,1)}$ and use backward differences to define $U_1 = D_2^- G$, $U_2 = -D_1^- G$, $U_3 = 0$. This is the codifferential of an actual antisymmetric plaquette potential, up to its global sign convention; backward divergence vanishes by commuting differences. Both nonzero components have zero mean and squared norm two, so $\|U\|^2 = 4$. An adjacent opposite delta pair has forward-gradient energy six in its own direction and four in each other direction. The total for U is 28. Every nonzero forward-gradient starting vertex lies in the weak cube, so the actual weighted energy is also 28. Consequently $\mathcal{E}_a(U) = 28 < 30 = c\|U\|^2$, while the global mean penalty vanishes. This refutes the representative-fibre mechanism even on the actual coexact domain. A different partition remains allowed by the original problem, so that problem has not been refuted.

34 Literature comparison and verification boundaries

Weighted path Cauchy estimates and weighted Hardy theory are classical; Muckenhoupt [25] is relevant context. Adaptive tree approximation has an extensive literature, including Binev–DeVore [26]. Our tree charge is proved explicitly rather than imported from a near-best approximation theorem. Publisher/author metadata for those two references were checked; unavailable full-text retrievals are not represented as complete comparisons.

Nochetto–Otárola–Salgado [27], Lemma 3.1 and Theorem 5.7, give constructive weighted Poincaré and anisotropic rectangular interpolation estimates in Muckenhoupt-weighted Sobolev spaces. Their constants depend on the weight characteristic and mesh geometry. Their inspected results are not the finite counting-norm observation theorem proved here, nor a proof of its uniform constants from only a global inverse moment. Conversely this distinction is not proof of originality: broader weighted approximation and adaptive-partition literature still requires comparison. In particular no claim is made to invent weighted Poincaré inequalities, spectral counting, adaptive trees, or weak Schatten ideals. The weak/strong distinction is separately attributed to the established spectral literature [28].

The universal mathematical arguments above are analytical, not Lean verified. Source units 184 and 185 have exact rational finite checks at pinned Git commits. Unit 184 has four matching normal/optimized payloads from two fresh Colab clones, six rejecting mutations, five critical matrix controls and six mode-block cases. Unit 185 has four matching payloads from two fresh clones, six global matrix controls, 72 balanced-tree cases, five weak-layer instances and six rejected intentional errors. It checks the actual coexact energy28 versus30. These are selected finite tests, not an all-size computation or independent replication. Full stdout, stderr, actual exit codes, hashes, failures and resource records are in the source evidence directories. No scientific computation ran on Windows.

34.1 Model assessments and review context

The exact eight-page source supplement from unit 186 was assessed once, with its complete report and all objections preserved. That historical assessment concerns only its own PDF bytes. This integrated successor is changed and must receive its own full exact-PDF numerical assessment with visible overall and criterion stars before publication. No grade or resolved-objection status is inherited from the supplement or public v8. The finite periodic-edge convention is now explicit above, and copyable source and evidence locators are supplied below. Earlier limitations on execution, priority and physical scope remain in force.

The source 185 proof also received a complete adversarial source audit. Both source-development context and historical reviewer participation are disclosed: OpenAI Codex, configured gpt-6-astra context, effective served backend and reasoning setting unexposed, shared historical account and provider context. Neither a new exact artifact nor a new conversation establishes blindness or organizational independence. Full reports, objections and author responses are retained separately with exact hashes. No source audit is represented as a numerical assessment of this PDF, human peer review, formal proof or editorial acceptance.

34.2 Reproducibility and remaining question

The source proof for185 has SHA256:

```
cee3556c1ff9dd99a69b3d62b20f65288189
26cca2734b2728b51ce9e4fdd407
```

The corresponding executed commit is `dfc2205dcfb8642bf5639c8433e4bdbff38075a3`; source184 was executed at `bedf55c593994aeba2cf47454c2f608870ecb871`. This integrated continuation's source and rendered-file hashes belong in its separate build manifest; a PDF cannot contain its own final hash. The research evidence is preserved at the pinned programme repository [29]. The exact source proofs and replay entrypoints are `papers/flat-coarse-rank/research/2026-11-07_green-moment-partitions184.txt`, `papers/flat-coarse-rank/research/2026-11-07_laminated-rectangles185.txt`, `scripts/green_moment_partition184.py` and `scripts/laminated_rectangles185.py`. Complete raw replay archives are also included in this integrated source package.

The sole unrestricted critical rectangular question remains open. A proposed extension must derive genuine transverse or all-fibre control from its coefficient hypotheses; one favourable fibre is insufficient. The strong Green test alone is abandoned at and below the critical exponent as a uniform target-order mechanism by Theorem 32.1. These findings do not reopen previously stopped routes, change PR29's 20/41 and TermSource0 baseline, or identify a finite sequence bound with a continuum or interacting Yang–Mills mass gap.

35 Rectangular observations with fewer than D variable coordinates

The one-coordinate dependence in the laminate theorem can be weakened to arbitrary dependence on any $k < D$ coordinates. The construction below uses the complete finite Sobolev proof already included in this manuscript. It retains at least one exactly invariant coordinate. No local spectral estimate or successful partition is added as an unproved input.

Put $p = D/2$, fix $1 \le k < D$, and set $d = D - k$. In this section define

$$\rho = (2/3)^d, \qquad C_D = 4A_D D^{1/p}, \qquad F_{D,k} = \frac{2^D 3^d}{1-\rho},$$

where A_D is the explicit constant in (32). These constants are not optimal.

Theorem 35.1. *Let $D \ge 3$, $n \ge 2$, $c > 0$, and let $a(x) = \alpha(x_1, \dots, x_k) > 0$ on $(\mathbb{Z}/n\mathbb{Z})^D$. There is a disjoint axis-aligned integer rectangular partition $\mathcal{P}$, independent of the field, such that for every scalar real or complex field*

$$c\|u\|^2 \le \mathcal{E}_a(u) + c\|C_{\mathcal{P}}u\|^2, \qquad C_{\mathcal{P}}u = (\sqrt{|R|}\,\overline{u}_R)_{R\in\mathcal{P}},$$

and

$$\#\mathcal{P} \le 2^d + F_{D,k}(cC_D)^p \sum_X a^{-p}. \tag{65}$$

If $\mu = \operatorname{mean}(a) \ge 1$ and $\operatorname{mean}((\mu/a)^p) \le M$, then for an h-component field the output dimension is at most

$$h\big[2^d + F_{D,k}(cC_D)^p 2^p M\big] \max\{1, n^D/(1+\mu)^p\}.$$

The energy inequality holds coordinatewise on all fields, and hence on an actual coexact subspace. Output dimensions are over the selected scalar field; double them for complex fields counted over the reals.

35.1 Local energy derived from the original coefficients

For a cut rectangle R with volume V, longest side L and $\kappa = L^2/V^{2/D}$, the established finite inequality (33) gives

$$\operatorname{Var}_R(u) \le A_D D^{1/p} \kappa \Big(\sum_R a^{-p}\Big)^{1/p} \mathcal{E}_{a,R}(u). \tag{66}$$

Indeed Hölder uses $q = 2D/(D+2)$ and $q/(2-q) = p$, and each original starting vertex contributes at most D internal edges. Reflection has already been completed in the unweighted proof; no reflected weights, enlarged-region energies or assumed Poincaré constant enter here. A singleton has zero variance. Internal cut edges are a subset of the actual periodic positive starts, including when $n = 2$; closing and cross-cell contributions may be discarded as nonnegative.

35.2 A balanced base tree and actual extrusion

The base domain is $\{0, \dots, n-1\}^k$. For a base rectangle Q, let $m(Q)$ be its longest integer side and $b(Q) = \sum_Q \alpha^{-p}$. Set $\theta = (cC_D)^{-p}$. Split a node precisely when $m \ge 2$ and $m^d b(Q) > \theta$. Each non-singleton coordinate is bisected into consecutive floor and ceiling halves; singleton coordinates stay unchanged. Accept equality and never split a base singleton.

All side lengths at a node differ by at most one, by induction from the root cube. Bisecting m and $m-1$ preserves this property, also for mixed lengths one and two. At least one parent side has length m, so every child has longest side m' satisfying

$$m/3 \le \lfloor m/2 \rfloor \le m' \le \lceil m/2 \rceil \le 2m/3 \qquad (m \ge 2). \tag{67}$$

The tree terminates, has at most 2^k children per internal node, and all base sides at a node with $m \ge 2$ are at least $m/2$.

For each leaf Q, partition each of the d invariant coordinates into $r = \lceil n/m \rceil$ consecutive balanced intervals of lengths $\lfloor n/r \rfloor$ or $\lceil n/r \rceil$. Every such length lies in $[m/2, m]$. The upper bound follows from $n \le rm$. For $r \ge 2$, $n \ge (r-1)m + 1$ implies $\lfloor n/r \rfloor \ge \lfloor (m+1)/2 \rfloor \ge m/2$; for $r = 1$, $m = n$. Thus every full rectangle $R = Q \times Z$ has $\kappa \le 4$. These rectangles partition the actual torus representatives, with no wrapping or overlap. If $m = 1$, every resulting full rectangle is a singleton.

Exact invariance, rather than a sampled favourable fibre, gives

$$\sum_R a^{-p} = |Z| b(Q) \le m^d b(Q).$$

For a non-singleton leaf, (66) therefore yields $c \operatorname{Var}_R(u) \le \mathcal{E}_{a,R}(u)$. Singleton leaves need no mass acceptance condition, since their residual variance is zero. Summing and using orthogonal mean decomposition proves the actual inequality with $K = c$ in Theorem 35.1 for all fields simultaneously.

35.3 Terminal charging and the critical rank count

Write $w(Q) = m(Q)^{-d}$. Every internal node satisfies

$$w(Q) < \theta^{-1} b(Q).$$

Let T be internal nodes with no internal children. These nodes are pairwise disjoint in the base lattice. Assign each internal node to any terminal internal descendant. Nodes assigned to a fixed terminal J lie on its ancestor path at distinct depths. By (67), ancestor weights decay by at least the factor $\rho = (2/3)^d < 1$, whence

$$\sum_{Q \text{ internal}} w(Q) \le \frac{1}{1-\rho} \sum_{J \in T} w(J) \le \frac{\theta^{-1}}{1-\rho} \sum_{\text{base}} \alpha^{-p}. \tag{68}$$

Only terminal internal nodes are charged as disjoint; all internal nodes are not incorrectly treated as disjoint.

Every child has weight at most 3^d times its parent's weight. Exact leaf-weight telescoping and the arity bound give

$$\sum_{Q \text{ leaf}} w(Q) \le n^{-d} + 2^k 3^d \sum_{Q \text{ internal}} w(Q).$$

A base leaf creates exactly $\lceil n/m \rceil^d \le (2n/m)^d$ full cells. Multiplying by $(2n)^d$, applying (68), and using $n^d \sum_{\text{base}} \alpha^{-p} = \sum_X a^{-p}$ proves (65). The accepted root case has one cell and no internal charges; singleton leaves remain in the telescoping identity. Finally the moment condition gives $\sum_X a^{-p} \le M n^D / \mu^p \le 2^p M n^D / (1+\mu)^p$, proving the stated output budget. This completes the analytical proof.

35.4 Verification, attribution and remaining boundary

The exact source188 proof is pinned at `0178b02bf26381633726edb82758e1c7ccb9aefd`; its complete involved-model source audit and all limitations are retained under `papers/flat-coarse-rank/research/evidence188/` in the source map above. The audit found no blocking mathematical gap. It did not perform a formal rebuild or scientific replay. Separate finite tree, geometry, inverse-mass and rank diagnostics were then executed at `b9d9dfdee043a62ec5d30926e215f8a581507b70`: two fresh own Linux Colab clones, each in normal and optimized modes, produced four identical payloads. Each checked 648 rational fixtures and rejected six actual mutated controls. All twelve child exit files contain zero. These checks are not local spectral computations, all-size proof or independent research replication; their coefficients are restricted to squares to keep inverse powers rational. The all-size claim rests on the analytical argument above.

Weighted Sobolev and adaptive-tree mechanisms are classical. The targeted primary comparison in source187–188 includes Cruz-Uribe, Dal and Rodney, *Degenerate Sobolev and Poincaré inequalities via extrapolation*, arXiv:2608.20498v1, Theorem 4.5 (`https://arxiv.org/html/2608.20498v1`), and Claros, Pérez and Zheng, *From generalized Poincaré to Poincaré–Sobolev inequalities via self-improving methods*, arXiv:2606.08556v1, Appendix C (`https://arxiv.org/html/2606.08556v1`). Their inspected continuum and summability statements are context, not imported premises or a priority certificate. Broader approximation literature still requires comparison.

Model assessments and review context. The source audit used the historically involved OpenAI Codex agent with a configured gpt-6-astra context. Effective served backend and reasoning setting were unexposed; shared history, workspace, account and provider context do not establish blindness or independence. Its full report, L1–L5 limitations and factual inspection scope remain preserved. Neither that source audit nor the full revision187 PDF assessment supplies a grade for this changed PDF. This complete artifact requires its own numerical overall and criterion stars, with all objections retained; no human peer review or editorial acceptance is inferred.

At $k = D$, the argument has $d = 0$ and $\rho = 1$, so the ancestor sum in (68) is no longer controlled. This is a limitation of the present method, not a counterexample to general rectangular means. Approximate invariance and arbitrary real-coefficient implementation or numerical stability are not established. Constants are nonoptimal upper-order bounds, not minimum-rank claims. The sole unrestricted critical rectangular question remains open, and physical interacting/continuum questions stay separate.

36 A weaker detour exclusion and an actual coexact boundary

This section refines the sufficient geometry in (G4), without claiming a new comparison principle or resolving the unrestricted critical-moment question. All edges retain their positive starts. Classical path-length and congestion comparison is attributed to [23].

36.1 Distance two suffices in dimension at least three

Let $X = (\mathbb{Z}/n\mathbb{Z})^D$, $n \geq 3$, and let $W \subset X$ be the weak set. Distinct points in W have periodic maximum-coordinate distance at least two. Write $\mathcal{E}_{\mathrm{wk}}$ and $\mathcal{E}_{\mathrm{str}}$ for the unweighted edge-square

sums with starts in W and $X \setminus W$. Suppose $a \geq H > 0$ outside W and $a \geq 0$ on W.

Theorem 36.1. *For $D \geq 3$ the actual energies satisfy*

$$\mathcal{E}_{\mathrm{wk}} \leq 3\mathcal{E}_{\mathrm{str}}, \qquad \mathcal{E}_0 \leq 4\mathcal{E}_{\mathrm{str}}, \qquad \mathcal{E}_a \geq (H/4)\mathcal{E}_0.$$

For $D = 2$ the same construction gives the sufficient, not asserted sharp, bounds $\mathcal{E}_{\mathrm{wk}} \leq 6\mathcal{E}_{\mathrm{str}}$ and $\mathcal{E}_a \geq (H/7)\mathcal{E}_0$.

Proof. Let σ be a single cycle through all D directions. For a weak positive edge $x \to x + e_i$, use

$$x \longrightarrow x - e_{\sigma(i)} \longrightarrow x + e_i - e_{\sigma(i)} \longrightarrow x + e_i.$$

Both positive-start vertices $x - e_{\sigma(i)}$ and $x + e_i - e_{\sigma(i)}$ are at periodic distance one from x, so are strong. The first step is reversed: its coefficient belongs to $x - e_{\sigma(i)}$, not x. Individual path edges are distinct for $n \geq 3$.

For a fixed positive edge (y, k) the complete list of possible weak sources, with the corresponding direction i, is

$$\begin{array}{c|c} y + e_k & \sigma^{-1}(k) \\ y + e_{\sigma(k)} & k \\ y - e_{\sigma^{-1}(k)} + e_k & \sigma^{-1}(k). \end{array}$$

Each row specifies at most one actual use. In $D \geq 3$, the directions $k, \sigma(k), \sigma^{-1}(k)$ are distinct. Every pair of these candidate sources differs by a nonzero vector with coordinates in $\{-1, 0, 1\}$, including their periodic interpretation at $n = 3$. Thus all pairwise distances are one, and at most one candidate is weak. Every strong edge is used at most once. Three-step Cauchy–Schwarz gives the factor three. For $D = 2$, the first candidate is incompatible with both other candidates; the latter two may coexist, so use multiplicity is at most two. This gives the factor six. Add $\mathcal{E}_{\mathrm{str}}$ and use $\mathcal{E}_a \geq H\mathcal{E}_{\mathrm{str}}$. □

The relaxation from distance three in (G4) is genuine: on $n = 4$, $W = \{0, 2e_1\}$ satisfies distance two but not distance three. At $n = 3$ both conditions force at most one weak vertex; the stated bounds still hold. The two-dimensional factor six is not proved optimal.

The torus gap $4\sin^2(\pi/n) \geq 16/n^2$ shows that, for $D \geq 3$ and $4H/n^2 \geq c$, one spatial-cell mean supplies the desired estimate with $K = c$. For a field with h real components this observation has h real outputs, not one scalar output. In $D = 2$ the sufficient condition is $16H/(7n^2) \geq c$. These comparisons apply componentwise and therefore on actual coexact restrictions of the component-gradient form.

36.2 Independence between neighbours is insufficient

Take even $n = 2m \geq 4$ and $D \geq 2$. Set $s(x) = \sum_j x_j \pmod n \in \{0, \ldots, n-1\}$ and $W = \{s \text{ even}\}$. No two weak vertices are nearest neighbours, including across periodic wraps. Put $a = \varepsilon > 0$ on W and $a = H > 0$ elsewhere. For real $z \in \mathbb{R}^m$ with $\sum_r z_r = 0$, define

$$f_z(x) = z_{j(s(x))}, \qquad j(s) = \lceil s/2 \rceil \pmod m.$$

Odd layer $2r - 1$ is paired with even layer $2r$, including $n - 1, 0$. Every layer has n^{D-1} vertices. Strong positive starts stay in a pair, whereas weak starts move from z_r to z_{r+1}. Hence

$$\mathcal{E}_{\mathrm{str}}(f_z) = 0, \qquad \mathcal{E}_{\mathrm{wk}}(f_z) = Dn^{D-1} \sum_r |z_{r+1} - z_r|^2,$$

$$\|f_z\|^2 = 2n^{D-1} \sum_r |z_r|^2, \qquad \mathrm{mean}_X\, f_z = 0.$$

The cyclic sums retain both positive starts when $m = 2$. Since the cyclic difference energy is at most $4\sum_r |z_r|^2$,

$$\frac{\mathcal{E}_a(f_z)}{\|f_z\|^2} = \frac{D\varepsilon}{2}\frac{\sum_r |z_{r+1} - z_r|^2}{\sum_r |z_r|^2} \le 2D\varepsilon. \tag{D190}$$

For nonzero mean-zero z, the weak energy is positive. Thus a finite comparison $\mathcal{E}_0 \le C\mathcal{E}_{\mathrm{str}}$ cannot follow merely from nearest-neighbour independence. Nor can $\mathcal{E}_a \ge \kappa H \mathcal{E}_0$ hold uniformly in H/ε for fixed $\kappa > 0$.

There is an actual coexact preimage. Regard f_z as a function of s. Its cyclic sum is zero. Put $\psi(0) = 0$ and $\psi(s) = \sum_{t=1}^{s} f_z(t)$ for $1 \le s < n$. Then $\psi(s) - \psi(s-1) = f_z(s)$, including the closing edge at $s = 0$. Set $B_{12}(x) = \psi(s(x))$, with all other antisymmetric components zero. The same codifferential convention used earlier gives

$$\begin{aligned} A_1(x) &= B_{12}(x) - B_{12}(x - e_2) = f_z(x), \\ A_2(x) &= -B_{12}(x) + B_{12}(x - e_1) = -f_z(x). \end{aligned}$$

and $A_i = 0$ for $i \ge 3$. This exhibits coexactness; backward divergence also cancels directly. Both the squared norm and the weighted component-gradient energy are twice their scalar counterparts. The map from real mean-zero z to A is injective, so its real dimension is $m - 1$. Its quotient obeys (D190). One fixed internal coordinate suffices. This does not identify an arbitrary variable-coefficient physical Hodge energy with the component-gradient form used here.

The real cosine $z_r = \cos(2\pi r/m)$ is nonzero and mean zero for every $m \ge 2$, and has quotient exactly $2D\varepsilon \sin^2(\pi/m) = 2D\varepsilon \sin^2(2\pi/n)$. More generally every vector of the exhibited real coexact space has quotient at most $2D\varepsilon$. Consequently if $c > 2D\varepsilon$, real-linear observation C and real-linear repair R with

$$\dim_{\mathbb{R}} \operatorname{range} C + \operatorname{rank}_{\mathbb{R}} R < m - 1$$

cannot supply the target inequality on that space, for any penalty with $\Phi(0) = 0$. Rank-nullity for (C, R) supplies a nonzero common-kernel vector; both the penalty and repair term vanish there. A complex-linear variant complexifies z and counts complex dimensions consistently. Real-linear observations on complex fields require real dimensions.

36.3 Moment boundary and evidence context

The two parity classes have equal size, so

$$\mu = (H + \varepsilon)/2, \qquad \operatorname{mean}\big((\mu/a)^{D/2}\big) = \tfrac{1}{2}(\mu/\varepsilon)^{D/2} + \tfrac{1}{2}(\mu/H)^{D/2}.$$

For fixed ε this is unbounded as $H \to \infty$. Even $\varepsilon = 1$, which meets $1 \le a \le 2\mu$, supplies no fixed moment constant M. The example therefore refutes the specified comparison route, not fixed-M rectangular sufficiency or Section 35.

The exact source190 and complete audit are retained in the programme research archive, at immutable commit `007cf5a7ca98ba99d9e187d075bbcd43557cb16f`. The involved source audit found no blocking gap. All five qualifications and the separate author response are preserved: component-valued means, consistent real/complex dimensions, nonoptimal $D = 2$ constants, the unbounded-moment boundary, and limits of verification. The contribution is a modest local refinement and route obstruction, not a new comparison principle or an established priority claim.

Model assessments and review context. The same historically involved Codex reviewer had a configured gpt-6-astra context; effective served backend and reasoning tier were unexposed. Shared history, account, workspace and provider do not establish blindness or independence. The source audit was analytical. Its prepared Linux finite diagnostics and paired-clone recipe remain unexecuted at consolidation; no finite replay or formal check is inferred. Neither that audit nor any predecessor's assessment supplies numerical stars for this changed full PDF. The exact final PDF requires its own full numerical overall and five-criterion assessment, with all objections retained. No human peer review, editorial acceptance, interacting result or continuum mass gap follows.

37 Limits of Neumann and harmonic boundary transfer

This section consolidates the analytic negative tests of source units 192–194. It concerns scalar finite positive-start energies and actual rectangles. The positive partial-invariance construction above is unchanged. None of the results below refutes the unrestricted rectangle-count question.

37.1 Exact boundary energy

Let a connected finite graph have positive edge conductances and energy $E(v) = \sum_{\{i,j\}} c_{ij}(v_i - v_j)^2 = v^T L v$, each edge counted once. For a nonempty proper retained set R write $B = V \setminus R$ and $L_{RR} = L_{\mathrm{int}} + D_{\mathrm{cut}}$, where D_{cut} is the diagonal of cut conductances at retained vertices.

Proposition 37.1. *The grounded block L_{BB} is positive definite, and*

$$\min_{v|_R = u} E(v) = E_{\mathrm{int}}(u) + u^T Q_R u,$$
$$Q_R = D_{\mathrm{cut}} - L_{RB} L_{BB}^{-1} L_{BR}.$$

The form Q_R is positive semidefinite and annihilates constants.

Proof. Extend a vector on B by zero on R. If its energy vanishes, connectedness makes it constant and hence zero. This proves positive definiteness of L_{BB}. For $h = -L_{BB}^{-1} L_{BR} u$, completion of the square gives

$$E(u, z) = u^T \Lambda_R u + (z - h)^T L_{BB} (z - h), \qquad \Lambda_R = L_{RR} - L_{RB} L_{BB}^{-1} L_{BR}.$$

Subtracting the fixed internal energy before minimization leaves the nonnegative sum of cut and B-internal edge energies. Thus $Q_R \succeq 0$. A constant trace has a constant zero-energy extension, so $Q_R 1_R = 0$. □

This is the classical Schur or Dirichlet-to-Neumann framework [31]. The identity derives the correction but supplies no quantitative moment bound: the inverse depends on the entire eliminated graph. Merely naming Q_R does not remove a hypothesis of a spatial theorem.

37.2 Internal energy does not control exact extension

On $X = (\mathbb{Z}/n\mathbb{Z})^D$, $D \geq 3$, $n \geq 3$, use

$$E_a(v) = \sum_{x,i} a(x) |v(x + e_i) - v(x)|^2, \quad N = n^D, \quad p = D/2, \quad \mu = N^{-1} \sum_x a(x).$$

Positive starts define unique unordered edges at these sizes. Internal Neumann energy retains only edges with both endpoints in the rectangle. For $H > 1$, set $a(0) = 1$ and $a(x) = H$ otherwise. Then

$$\mu = H - (H-1)/N, \qquad H/2 \le \mu < H, \qquad 1 \le a \le 2\mu, \quad M_a = \frac{\mu^p}{N} + (1 - N^{-1})(\mu/H)^p.$$

For $R = \{0, e_1\}$ the sole internal edge has coefficient one. A mean-zero trace $(1, -1)$ has norm squared two and internal energy four, hence sharp local Neumann gap two.

Proposition 37.2. *For every exact extension v of that trace, $E_a(v) \ge 4 + 4H/3$. Thus no size-independent internal-energy-only exact extension estimate with constant $C(D, M_0)$ holds for every admissible profile, for any fixed $M_0 > 1$.*

Proof. The path $0 \to -e_2 \to e_1 - e_2 \to e_1$ uses three distinct strong edges, with positive starts $-e_2$ (directions two and one) and $e_1 - e_2$ (direction two). Cauchy–Schwarz gives energy at least $4H/3$ on this path. The disjoint direct weak edge contributes four. For $H = n$,

$$1 < M_a < 1 + n^{-p} \longrightarrow 1.$$

Indeed $H^p/N = n^{-p}$ bounds its first term, while its second is below one; strict Jensen gives the lower bound. Every fixed $M_0 > 1$ admits an unbounded tail, but the extension-to-internal-energy ratio is at least $1 + n/3$. This argument includes nonlinear extension selection rules. □

The proposition does not include $M_0 = 1$, when strict convexity forces constant coefficients. Bounds allowed to depend on size, mean or a growing target scale are different assertions.

For completeness the two-terminal conductance also satisfies

$$\Lambda_R = g \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}, \qquad 1 + H/3 \le g \le D(H+1), \qquad Q_R = (g-1) \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}.$$

Symmetry and zero row sums give the displayed form of Λ_R. The direct edge plus bypass prove the lower bound for arbitrary endpoint values by the same telescoping calculation. For the upper bound, extend the value at e_1 constantly to all vertices other than zero. Only the D outgoing weight-one edges and D incoming weight-H edges at zero contribute. They are distinct for $n \ge 3$, proving the upper bound. Thus the boundary-correction to internal-energy ratio is at least $H/3$; for $H = n$ it too has no uniform bound for any fixed $M_0 > 1$.

A contrasting uniformly coercive family. For the different choice $H = n^2$, $H^p = N$, so $M_a < 2$ and $M_a \to 2$; the fixed admissibility bound is $M = 2$, not a fixed smaller bound. Here $\max\{1, N/(1+\mu)^p\} \le 2^p$, while the global mean-zero gap is at least four. To see the latter directly, split unit-coefficient energy as $E_0 = E_w + E_s$, the weak starts being the D edges from zero. Choose one cycle σ through all D coordinate directions and replace edge $(0, i)$ by

$$0 \to -e_{\sigma(i)} \to e_i - e_{\sigma(i)} \to e_i.$$

Every replacement start is strong. For a fixed positive edge (y, k), its possible weak sources in the first, middle and last segment are

$$y + e_k, \qquad y + e_{\sigma(k)}, \qquad y - e_{\sigma^{-1}(k)} + e_k.$$

They are pairwise distinct modulo n for $D \ge 3, n \ge 3$: the three coordinate indices are distinct. Only source zero is weak, so an edge is used at most once. Hence $E_w \le 3E_s$, $E_0 \le 4E_s$ and $E_a \ge HE_0/4$. Nonzero torus eigenvalues are at least $4\sin^2(\pi/n) \ge 16/n^2$, proving

$$E_a(v) \ge 4\|v - \mathrm{mean}_X\, v\|^2.$$

Thus for $c = 3$ one global normalized mean with penalty coefficient $K = c$ suffices, despite the displayed local gap being two. For vector fields that cell carries one mean per component, not one scalar output in total. The near-one family $H = n$ instead has only the comparison lower bound $4/n$, and its target scale grows like n^p. These limits must not be interchanged. No local coexact witness is asserted.

37.3 Disjoint traces do not justify a sum of harmonic costs

The next example refutes the aggregate estimate

$$\sum_j (f|_{R_j})^T Q_{R_j}(f|_{R_j}) \le C(D, M)E_a(f), \tag{69}$$

where each form separately eliminates the entire complement of its own rectangle. Joint elimination is not the same construction.

First let $T = (\mathbb{Z}/n\mathbb{Z})^d$, $d \ge 3$, with even n and unit energy E_T. For all $y, z \in T$,

$$|w(y) - w(z)|^2 \le C_d E_T(w), \qquad C_d = d3^{d-1}/2. \tag{70}$$

Here is a complete finite proof. Use the normalized Fourier coefficients

$$\widehat{w}(k) = n^{-d} \sum_y w(y) e^{-2\pi i k\cdot y/n}.$$

Parseval gives $E_T = n^d \sum_k \lambda_k |\widehat{w}(k)|^2$. Choose representatives $-n/2 \le k_l < n/2$. Then $\lambda_k = 4\sum_l \sin^2(\pi k_l/n) \ge 16|k|^2/n^2$. Fourier inversion and Cauchy–Schwarz bound the point-difference constant by

$$n^{-d} \sum_{k\ne 0} \frac{|e^{2\pi i k\cdot y/n} - e^{2\pi i k\cdot z/n}|^2}{\lambda_k} \le \frac{n^{2-d}}{4} \sum_{k\ne 0} |k|^{-2}.$$

The shell with $\max_l |k_l| = r$ has at most

$$(2r+1)^d - (2r-1)^d \le 2d3^{d-1} r^{d-1}$$

representatives. Its inverse-square contribution is at most $2d3^{d-1}r^{d-3}$. Summing $1 \le r \le n/2$ gives at most $2d3^{d-1}n^{d-2}$, proving (70). The restriction $d \ge 3$ is essential to this uniform estimate.

Theorem 37.3. *For every fixed $D \ge 4$, there are positive profiles with fixed critical moment bound $M = 3$, disjoint actual rectangles, and one global mean-zero field for which the ratio of the left side of* (69) *to $E_a(f)$ tends to infinity.*

Proof. Take even $n \ge 6$, $d = D - 1$, $p = D/2$, $m = n/2$, $H = n^{2/D}$ and $X = (\mathbb{Z}/n\mathbb{Z}) \times T$, writing vertices (t, y). Put $a = 1$ on the two planes $t = 0, m$, except $a(0, 0) = H$; put $a = H$ elsewhere. There are $W = 2n^d - 1$ weak vertices. Since $H^p = n$ and $W/n^D < 2/n \le 1/3$,

$$\tfrac{2}{3}H \le \mu < H, \quad 1 \le a \le 2\mu, \qquad 1 \le M_a = \mu^p[W/n^D + (1 - W/n^D)H^{-p}] < 3.$$

For each $y \neq 0$ let $R_y = \{(0,y),(1,y)\}$. These rectangles are disjoint and their internal coefficient is one. For arbitrary v, telescope the endpoint difference through

$$(0,y) \to (n-1,y) \to (n-1,0) \to (0,0) \to (1,0) \to (1,y).$$

The first, third and fourth changes use distinct longitudinal edges of weight H. The two other changes are within the entirely strong transverse layers $n-1$ and one. Their squared differences are bounded by C_d times the respective layer energy divided by H, by (70). These five energy sets are pairwise disjoint within this one comparison: $y \neq 0$ distinguishes the two incoming edges, and layers and directions distinguish all others. Weighted Cauchy–Schwarz with weights $1, C_d, 1, 1, C_d$ gives, for $K_d = 3 + 2C_d$,

$$|v(0,y) - v(1,y)|^2 \leq K_d E_a(v)/H.$$

The minimized two-point energy is g_y times the squared endpoint difference, so $g_y \geq H/K_d$ and its boundary correction is $g_y - 1$ times that squared difference.

Let $f = 1$ on layers $0, m+1, \ldots, n-1$ and $f = -1$ on layers $1, \ldots, m$. The two sets each contain m layers, so its mean is zero. Only longitudinal starts at zero and m contribute. One has weight H and the remaining $2n^d - 1$ have weight one, giving exactly

$$E_a(f) = 4(2n^d - 1 + H), \qquad \sum_{y \neq 0} (f|_{R_y})^T Q_{R_y} (f|_{R_y}) \geq 4(n^d - 1)(H/K_d - 1).$$

Eventually $H \geq 2K_d$. Since $H \leq n^d$ and $n^d - 1 \geq n^d/2$, the ratio is at least $H/(12K_d)$ and diverges. All constants depend only on the fixed dimension. The moments tend to three, not a fixed smaller bound. □

The lower-bound certificates reuse the hole and transverse layers across different trace problems. This does not force the actual minimizers to use that hole rather than parallel weak cuts. Disjoint retained vertices do not make independently optimized global energies disjoint. The original target count remains comparable to the number of pairs:

$$2^{-p} n^d \leq \frac{n^D}{(1+\mu)^p} \leq (3/2)^p n^d.$$

Thus the example is not a counting counterexample. There is no $D = 3$ conclusion from this proof and no obstruction to a joint reduction, partition-dependent selection, overlap compensation or a growing-factor bound. No coexact, rank or physical continuum result is obtained.

37.4 Attribution, assessment and unresolved mechanism

Weighted extension theorems have their own weight and domain hypotheses; the constants in Chua's Theorem 1.4 explicitly depend on the weights [30]. Our finite counterexamples do not contradict it. The retained-resistance invariance under Kron reduction in [31] concerns one retained set, not a sum over different eliminations. Fourier, path comparison and Schur reduction are classical. No exhaustive novelty or priority claim is made for these diagnostic families. Full source proofs and involved audits192–194 are retained at the immutable programme commits recorded in the source manifest. Each configured reviewer was the same historically involved Codex `gpt-6-astra` agent with shared history, provider/account and workspace; served backend and reasoning effort were not attested. These were complete source audits, not blind

or independent peer reviews, scientific replay, Lean verification, PDF assessments or numerical stars. All limitations and author responses remain in the evidence packet.

The main fully variable fixed-moment actual-rectangle question remains open. Cellwise internal-only extension, a near-one-moment rescue, and independent-cell harmonic-cost additivity have explicit falsifiers. A joint partition or direct local internal-energy argument must derive its count without assuming any of those failed transfers. No main-theorem hypothesis has been removed by these negative tests. No earlier PDF score or public Working-paper status transfers to this changed manuscript.

38 Joint trace discrepancies and constructive binary partitions

A failed comparison need not prevent a successful choice of observations. We first strengthen the preceding harmonic-cost diagnostic, then record a binary scalar corollary of the existing threshold construction. The corollary uses actual internal energy, not separate exterior minimizers. The unrestricted intermediate-coefficient question remains open.

38.1 A growing-dimensional discrepancy

Use the family of Theorem 37.3 on the subsequence $n = s^{2D}$, $H = s^4$, with even $s \geq 2$ and $D \geq 4$. Put $d = D - 1$, $m = n/2$ and $T = (\mathbb{Z}/n\mathbb{Z})^d$. The exact actual mean and moment bounds there still hold: $\frac{2}{3}H \leq \mu < H$, $1 \leq a \leq 2\mu$ and $1 \leq M_a < 3$. For $y \neq 0$ retain $R_y = \{(0, y), (1, y)\}$, and let U be their union. Define two distinct quadratic forms on real traces:

$$S_U(u) = \min_{f|_U = u} E_a(f), \qquad B_{\mathrm{sep}}(u) = \sum_{y \neq 0} Q_{R_y}(u|_{R_y}).$$

The grounded-block argument above makes S_U well defined with constant kernel; it is positive definite on the counting-mean-zero trace hyperplane. The separate summands minimize different global fields. Write $C_d = d3^{d-1}/2$, $K_d = 3 + 2C_d$ and $\beta_D = 4 + 2\pi^2 d$.

Proposition 38.1. *For each fixed $A > 0$ and all sufficiently large even s, any real-linear maps $C : \mathbb{R}^U \to F$ and $V : \mathbb{R}^U \to \mathbb{R}^U$, and any function Φ with $\Phi(0) = 0$, satisfying*

$$B_{\mathrm{sep}}(u) \leq AS_U(u) + \Phi(Cu) + \langle u, Vu \rangle$$

for all traces, or on a neighborhood of zero, must obey

$$\operatorname{rank} C + \operatorname{rank} V \geq L_s = (2s^{D-2} + 1)^{D-1} - 1 \geq n^{(D-2)(D-1)/(2D)}.$$

No symmetry or positivity of V, or continuity of Φ, is required.

Proof. The five-change comparison in the preceding theorem gives the individually minimized pair form $g_y(v_0 - v_1)^2$ with $g_y \geq H/K_d$; hence $Q_{R_y} = (g_y - 1)(v_0 - v_1)^2$. Set $R = s^{D-2}$, so $R^2 = n/H$ and $R < n/2$. The real Fourier space on T with frequencies $|k_i| \leq R$ has dimension $(2R+1)^d$: the frequency box is conjugation symmetric and contains no nonzero Nyquist character. Evaluation at zero is a nonzero functional since constants belong to the space. Its kernel G_R^0 therefore has dimension L_s. For $g \in G_R^0$ put $u_g(0, y) = g(y)$ and $u_g(1, y) = -g(y)$ for $y \neq 0$. These traces are mean zero; the map is injective because $g(0) = 0$, and $\|u_g\|^2 = 2\|g\|^2$. For $H \geq 2K_d$,

$$B_{\mathrm{sep}}(u_g) = 4 \sum_{y \neq 0} (g_y - 1)|g(y)|^2 \geq (H/K_d)\|u_g\|^2.$$

Here the coefficient g_y of the pair form is distinct from the field g.

Let $\sigma(t) = 1$ on $0, m+1, \dots, n-1$ and -1 on $1, \dots, m$. The extension $f_g(t,y) = \sigma(t)g(y)$ is globally mean zero, regardless of the mean of g. Its longitudinal energy is exactly $8\|g\|^2 + 4(H-1)|g(0)|^2 = 8\|g\|^2$. Its transverse energy is

$$\sum_y \left[(n-2)H + 2 + (H-1)\mathbf{1}_{y=0}\right] \sum_{l=1}^{d} |D_l g(y)|^2 \le nHE_T(g).$$

The derivatives $D_l g(0)$ have not been set to zero. Fourier support gives $E_T(g) \le 4\pi^2 dR^2\|g\|^2/n^2$. Therefore

$$S_U(u_g) \le E_a(f_g) \le (8 + 4\pi^2 d)\|g\|^2 = \beta_D \|u_g\|^2.$$

When $H > AK_d\beta_D$, every nonzero vector in this L_s-dimensional trace space violates $B_{\mathrm{sep}} \le AS_U$. If the combined rank of C, V were less than L_s, their joint kernel in that space would contain a nonzero vector. Both repair terms vanish there. Scaling that vector toward zero also handles a neighborhood-only inequality. Finally $L_s \ge s^{(D-2)(D-1)}$ gives the displayed power of n. □

The same subspace gives at least L_s generalized eigenvalues of B_{sep}/S_U at least $H/(K_d\beta_D)$ on the mean-zero trace hyperplane, by finite-dimensional min–max. This necessity is not shown sharp or sufficient. Its exponent is below the allowed n^{D-1} count for this family. It does not rule out growing-rank repairs or selectable partitions, and does not address $D = 3$ through this comparison.

38.2 Binary corollary of the existing threshold construction

The following is an explicit specialization of Section 21, not a new construction.

Proposition 38.2. *Let $D \ge 3, n \ge 3$ be integers, $c > 0, H \ge 1$, $p = D/2$, $N = n^D$, and let $a : (\mathbb{Z}/n\mathbb{Z})^D \to \{1, H\}$ have exact mean μ and $N^{-1}\sum_x (\mu/a(x))^p \le M$. There is a nonwrapping integer rectangle partition with normalized scalar means*

$$(Pf)_R = \frac{\sum_R f}{\sqrt{N|R|}}, \qquad \|f\|_N^2 = N^{-1}\sum_X |f|^2,$$

satisfying, for every real scalar f,

$$N^{-1}E_a(f) + c\|Pf\|_2^2 \ge c\|f\|_N^2, \tag{71}$$

$$\operatorname{rank} P \le 2^p\left[(1+\sqrt{c})^D + 2DM\right] \max\{1, (n/\sqrt{1+\mu})^D\}. \tag{72}$$

Here $E_a(f) = \sum_{x,i} a(x)|D_i f(x)|^2$ uses positive-start weights. No upper envelope $a \le 2\mu$ is required; H need not be attained.

Proof. Use the threshold construction of Section 21 with $\tau = H$. Its geometric and internal-energy argument uses only $a \ge 1$; the standing upper envelope in that section is unused in this argument. For $H > 1$ the set $\{a < H\}$ consists of coefficient-one vertices, whose actual moment contribution gives $\#\{a < H\} \le MN/\mu^p$. For $H = 1$ that set is empty. Thus the earlier rectangle count, with $B_c = (1+\sqrt{c})^D$ and $1 \le \mu \le H$, gives

$$J \le B_c \max(1, N/H^p) + 2DMN/\mu^p \le (B_c + 2DM)\max(1, N/\mu^p).$$

Dividing the scalar version of (27) by N gives (71), since $P = C_{\mathcal{P}}/\sqrt{N}$. The J rows have disjoint nonempty supports, so their rank is J. Finally $1 + \mu \le 2\mu$ proves (72). The singleton branch of the earlier construction also covers $H < c/4$. □

For constant $f = 1$, both squared normalized norms equal one and the energy is zero, so $K \geq c$ is necessary in this convention; the construction attains $K = c$ without claiming minimum rank. For r independent real components the means have rJ scalar coordinates, not J. The two-plane comparison family above is included: a failed prescribed comparison coexists with successful selectable observations.

38.3 A failed extension to intermediate values

For even n put $a = H/2$ on half the vertices and $a = H$ on the rest, with $H \geq 2$. Then $\mu = 3H/4$, $a \leq 2\mu$, and

$$M_a = \tfrac{1}{2}[(3/2)^p + (3/4)^p]$$

is independent of H. Nevertheless the attempted unit-cost bound $|\{a < H\}| \leq M_a N/\mu^p$ would require $H^p \leq 2^p + 1$, and fails for larger H. The actual moment does not pay unit cost per intermediate vertex. Rounding these coefficients down to one changes the hypotheses. This is a falsifier of that cardinal lemma, not of the target partition statement: the same field is uniformly strong at scale $H/2$. A derived multilevel construction remains necessary for the unrestricted question.

38.4 Model assessments and review context

The comparison continuation and earlier binary corollary are recorded in source units 197–199. Inventory correction201 identifies the binary result as a direct corollary already available from the earlier threshold section; source198 is a special case, and source201A reproduces earlier threshold results. None is counted here as a new mathematical advance. Revision202 removes the duplicate proof while retaining its historical sources and reports. Complete exact-source reports and preparing-agent responses are retained with the immutable manifest. Source198 supplies a special-case check rather than another paper. The separate Codex audit tasks were fresh `fork_turns=none` tasks with inherited configuration and shared filesystem/account. The exposed description was Codex based on GPT-6; precise served backend and reasoning effort were not exposed. No model, statistical or organizational independence, human peer review or formal verification is claimed. The full adverse findings remain attached, including scalar/component rank, unproved optimality or runtime complexity, and unestablished priority. No previous PDF grade transfers to this changed manuscript; an exact-PDF numerical assessment and visible stars are still required before a new publication-ready claim.

The Fourier, Schur, rank-nullity, Poincare and rectangle-cutting ingredients are classical. The comparison diagnostic and earlier threshold corollary separate a failed comparison from a successful binary-class construction; their originality is not established by this limited primary-source comparison. No physical hypothesis of Yang–Mills is removed. The continuum remains open; PR29 stays recorded at 20/41 with TermSource0, not newly certified by these source audits. The sole unrestricted critical rectangular question continues to require actual control of intermediate coefficient levels.

39 Local shell estimates and their sharp lower moment boundary

This section groups the local estimates from source units 204–205 and 216. They concern actual internal energy on concentric integer shells, rather than a supplied annular gap. They do

not construct or count the disjoint rectangular observations in Section 27. The lower moment estimate below pays a necessary thickness cost; it does not remove that cost from the earlier, thickness-uniform estimate.

Fix $D \geq 3$, integers $R \geq 1$ and $0 \leq r < R$, and set

$$A = \{-R, \dots, R\}^D \setminus \{-r, \dots, r\}^D, \quad t = R - r, \quad L = 2R + 1, \quad N = |A|, \quad J = 2D.$$

Only nonwrapping internal positive-axis edges are counted, once each. For $a : A \to (0, \infty)$ and $f : A \to \mathbb{R}$ put

$$V_A(f) = \sum_A |f - \bar{f}_A|^2, \qquad \mathcal{E}_{a,A}(f) = \sum_{x, x+e_i \in A} a(x)|f(x + e_i) - f(x)|^2.$$

The coefficient belongs to the original starting site x. No energy outside A, averaging of endpoint weights, or reflected weights is used.

Theorem 39.1 (Two local inverse moment estimates)**.** *There are explicit finite dimension-only constants U_D and C_D such that every f and positive a as above satisfy*

$$V_A(f) \leq U_D L^2 \left(\frac{1}{N} \sum_A a^{-D/2} \right)^{2/D} \mathcal{E}_{a,A}(f), \tag{73}$$

$$\begin{aligned} V_A(f) &\leq C_D \left(\sum_A a^{-(D-1)/2} \right)^{2/(D-1)} \mathcal{E}_{a,A}(f) \\ &\leq C_D J^{2/(D-1)} L^2 t^{2/(D-1)} \left(\frac{1}{N} \sum_A a^{-(D-1)/2} \right)^{2/(D-1)} \mathcal{E}_{a,A}(f). \end{aligned} \tag{74}$$

The statements include $t = 1$ and $r = 0$. Translation preserves them; off-center holes, several holes and wrapping identifications are not covered.

We give the proofs and constants. They use the finite counting cube theorem already proved in Section 22, not a new unproved graph or extension assumption. Coarea, variance decomposition and the final conductance Hölder step are classical; see also [17], Section 2.2. No priority finding for this particular shell formulation is made.

39.1 The actual cover and directional mean transport

Let $S_{i,+}$ restrict x_i to $[r + 1, R] \cap \mathbb{Z}$ and $S_{i,-}$ restrict it to $[-R, -r - 1] \cap \mathbb{Z}$, leaving every other coordinate in $[-R, R] \cap \mathbb{Z}$. These J slabs cover A; their common volume is $v = tL^{D-1}$. Slabs with distinct normal coordinates intersect in $t^2 L^{D-2}$ sites. Root a tree at $(1, +)$, attach all $J - 2$ slabs with other normals to it, and attach $(1, -)$ to $(2, +)$. Every root path has at most two edges, each joining distinct normals. Also

$$2tL^{D-1} \leq N = L^D - (L - 2t)^D \leq JtL^{D-1}. \tag{75}$$

The lower bound uses the two disjoint coordinate-1 slabs, and the upper bound the full cover.

For $q \geq 1$, an L-site interval I, and a nonempty subinterval T, Jensen and path Hölder give

$$|\bar{g}_I - \bar{g}_T|^q \leq \frac{1}{L|T|} \sum_{x \in I, y \in T} |g(x) - g(y)|^q \leq L^{q-1} \sum_{e \subset I} |dg(e)|^q.$$

Averaging over the same fibers of $B = I \times F$ therefore gives

$$|\bar{f}_B - \bar{f}_{T\times F}| \le L^{1-1/q}|F|^{-1/q}\left(\sum_{\substack{e\subset B\\ e\parallel I}} |df(e)|^q\right)^{1/q}. \tag{76}$$

For each side of a tree intersection, $|F| = v/L$. Its intersection mean is identical from either side. If $F_i = \sum_{e\subset S_i} |df(e)|^q$, then

$$|\bar{f}_{S_i} - \bar{f}_{S_j}| \le Lv^{-1/q}(F_i^{1/q} + F_j^{1/q}). \tag{77}$$

There is no inverse overlap-volume loss. This is a directional energy comparison, not the local best-constant error gluing statement refuted earlier.

39.2 The thickness uniform endpoint

Put $q = 2D/(D+2)$ and $H = \sum_i F_i$. The proved rectangle theorem (32) gives $V_{S_i}(f) \le A_D L^2 v^{-2/D} F_i^{2/q}$ with its explicit A_D. By the depth-two tree and (77), $|\bar{f}_{S_i} - \bar{f}_{S_{\mathrm{root}}}| \le 4Lv^{-1/q}H^{1/q}$. The cover and the exact variance decomposition about each slab mean yield

$$\begin{aligned} V_A(f) &\le \sum_i \left[V_{S_i}(f) + v|\bar{f}_{S_i} - \bar{f}_{S_{\mathrm{root}}}|^2\right] \\ &\le (A_D + 16J)L^2 v^{-2/D} H^{2/q}, \end{aligned}$$

using $1 - 2/q = -2/D$ and $\sum_i F_i^{2/q} \le H^{2/q}$. Writing $F_A = \sum_{e\subset A} |df(e)|^q$, we have $H \le JF_A$; and $N \le Jv$ implies $v^{-2/D} \le J^{2/D}N^{-2/D}$. Thus

$$V_A(f) \le (A_D + 16J)J^{2/q+2/D}L^2 N^{-2/D} F_A^{2/q}.$$

Hölder on precisely the retained original edges gives, for $p = q/(2-q)$,

$$F_A^{2/q} \le \mathcal{E}_{a,A}(f)\left(\sum_{e\subset A} a(\mathrm{start}\, e)^{-p}\right)^{1/p} \le D^{1/p}\mathcal{E}_{a,A}(f)\left(\sum_A a^{-p}\right)^{1/p}. \tag{78}$$

Here $p = D/2$, so (73) follows with $U_D = D^{2/D}(A_D + 16J)J^{2/q+2/D}$. This is the source205 bound, including its thickness-uniform conclusion, rather than a new result counted twice.

39.3 The lower dimensional slab endpoint

Put $d = D-1$, $p = d/2$, $q = 2d/(d+2)$. On a d-dimensional L-site cube Q, the earlier counting cube proof gives

$$V_Q(g) \le K_d^2\left(\sum_{e\subset Q} |dg(e)|^q\right)^{2/q}. \tag{79}$$

For $d \ge 3$, its explicit constant is

$$c_d = d(1 - 2^{-1/d}), \quad \alpha_d = 2(d-1)/d, \quad q' = 2d/(d-2), \qquad K_d = 2\alpha_d d^{1/q'}/c_d.$$

For $d = 2$, use $K_2 = 2$. To check this endpoint without a spurious infinite-conjugate power argument, the same finite projection proof gives $|\partial_E B| \geq (2 - \sqrt{2})|B|^{1/2} \geq \frac{1}{2}|B|^{1/2}$ for sets occupying at most half a cube. Finite coarea and the triangle inequality in ℓ^2 give $\|u\|_2 \leq 2 \sum_e |du|$ for a nonnegative u with such support. Split g at a counting median. Both positive and negative supports have size at most half, and their edge ℓ^1 energies add exactly to that of g. Squaring and adding proves (79) at $q = 1$. Singleton cubes have zero variance.

For a slab $B = Q_L^d \times \{1, \ldots, t\}$, let m_z be its layer means, and let T and U be its long- and short-direction q-energies. Exactly,

$$V_B(f) = \sum_z V_{Q_L^d}(f_z) + L^d \sum_z |m_z - \bar{m}|^2.$$

The first term is at most $K_d^2 T^{2/q}$. On a t-site interval, every pair difference is at most $t^{1-1/q}$ times the q-gradient norm. Jensen at its mean gives $\sum_z |m_z - \bar{m}|^2 \leq t^{3-2/q} (\sum |dm|^q)^{2/q}$. Jensen across each L^d-point layer gives $\sum |dm|^q \leq L^{-d} U$. Thus the second term is at most

$$L^{d-2d/q} t^{3-2/q} U^{2/q} = L^{-2} t^{2-2/d} U^{2/q} \leq U^{2/q}.$$

For $t = 1$ it is zero; otherwise $t \leq L$ proves the inequality. With $B_d = \max(K_d^2, 1)$ and $2/q \geq 1$, we have proved

$$V_B(f) \leq B_d (T + U)^{2/q}, \tag{80}$$

uniformly in thickness, using only actual internal edges.

39.4 Lower moment shell bound

Set $G = (\sum_{e \subset A} |df(e)|^q)^{1/q}$ and $h = L^{1-d/q} t^{-1/q} = L^{-d/2} t^{-1/q}$. By (76), each slab mean differs from a tree intersection mean by at most hG. Hence each root mean difference is at most $4hG$. Since $v = tL^d$, we have $vh^2 = t^{1-2/q} = t^{-2/d} \leq 1$. Using the same cover and exact decomposition, now with (80), proves

$$V_A(f) \leq J B_d G^2 + 16 J v h^2 G^2 \leq S_D G^2, \qquad S_D = J(B_d + 16).$$

Finally (78) has $p = d/2$ at this q. It proves (74) with $C_D = S_D D^{2/d}$. Equation (75) proves its normalized form. For $D = 3$, $S_3 = 120$, $C_3 = 360$, and the final normalized constant is 2160. These are explicit upper bounds, not optimal constants. Applying the scalar bounds to each real coordinate and summing adds no internal-dimension factor.

39.5 Sharp powers and an actual starting site falsifier

Proposition 39.2. *For arbitrary positive starting-site weights, the thickness power $2/(D-1)$ in the lower moment estimate cannot be replaced by a smaller power with a dimension-only constant. At thickness one, no exponent $p_0 < (D-1)/2$ gives a uniform $L^2(\mathrm{mean}_A\, a^{-p_0})^{1/p_0}$ variance bound. These are local scalar statements, not counterexamples to the fixed global $D/2$-moment rectangular question.*

Proof. For any $p_0 > 0$, let $x_0 = (R, R, 0, \ldots, 0)$, $W = \{x_0, x_0 - e_1, \ldots, x_0 - e_D\}$ and $h_0 = D + 1$. All these distinct sites belong to A, including when $t = 1$ and $R = 1$. The smallest shell

has $N \geq 3^D - 1 \geq 2h_0$ sites. Put $\varepsilon = (h_0/N)^{1/p_0}$, $a = \varepsilon$ on W and $a = 1$ elsewhere. Then $\mathrm{mean}_A\, a^{-p_0} = 2 - h_0/N < 2$. For $f = \mathbf{1}_{\{x_0\}}$,

$$V_A(f) = \frac{N-1}{N}, \qquad \mathcal{E}_{a,A}(f) = (2D-2)\varepsilon, \qquad \frac{V_A(f)}{\mathcal{E}_{a,A}(f)} \geq \frac{1}{4(D-1)} \left(\frac{N}{h_0}\right)^{1/p_0}.$$

There are D incoming and $D-2$ outgoing retained incident edges, and all their starting sites lie in W. Weakening x_0 alone would leave order-one incoming energy and would not prove this bound.

At $p_0 = (D-1)/2$, the volume brackets (75) show that the lower bound is of order $L^2 t^{2/(D-1)}$, while the normalized moment stays below two. Taking $r = 0$, $R \to \infty$ rules out any smaller thickness power. Taking $t = 1$, $R \to \infty$ and $p_0 < (D-1)/2$ instead gives growth of the ratio divided by L^2 at least as $L^{(D-1)/p_0 - 2}$, proving the second assertion.

One may even impose actual mean comparability: set $\widetilde{a} = a/\varepsilon$ and $\mu = \mathrm{mean}_A\, \widetilde{a}$. Since $\bar{a} \geq 1 - h_0/N \geq 1/2$, $1 \leq \widetilde{a} \leq 1/\varepsilon \leq 2\mu$, and $\mathrm{mean}_A(\mu/\widetilde{a})^{p_0} = \bar{a}^{p_0}(2 - h_0/N) < 2$. The indicator energy after rescaling is $2D-2$. However, at $p_0 = (D-1)/2$ its normalized $D/2$ moment is at least $2^{-D/2}(N/h_0)^{1/(D-1)}$, and diverges. Thus it does not obey the fixed global moment hypothesis of Section 27. No statement about every rectangular partition or every finite penalty follows. □

39.6 Evidence and the remaining conversion

The full analytical sources are units 205 and 216, with frozen source reviews accepting their stated local claims. The unit216 review read the entire proof and cube dependency; it used an existing shared-filesystem Codex task with earlier unit198 context, not a fresh isolated reviewer. The exposed description is Codex based on GPT-6; its exact hidden backend and effort are unknown. Independence and human peer review are not claimed. The exact source216 commit is

`8b0fd2f5a93606a3f94b3b1115a7d4cc1bb38337`.

The full report and diagnostic records are at `papers/flat-coarse-rank/research/intrinsic-shell216/`; the diagnostic is `scripts/shell_intrinsic_diagnostic216.py`. An owned CPU/high-memory Colab Linux run checked ten $D = 3$ geometries, 40 fields, 240 slab estimates, 960 fiber comparisons and ten negative controls for missing incoming starting sites. Normal and optimized modes had actual zero exits and identical output bytes. Its three-second child execution is finite diagnostic evidence, not an all-field or all-D machine proof. The runtime was stopped and separately read back inactive. No Lean theorem or fresh formal validation is asserted by this section.

The useful new local input is a weaker inverse moment on thin shells. To use it in the main question, favorable local moments, selection and count, and conversion to a disjoint admissible rectangle partition still must be derived from the original global coefficient assumptions. The slabs overlap and are not that partition. Historical stopped routes and their negative results remain in force. This section does not close the general question or change PR29's 20/41 and TermSource0 record.

40 Selected shell observations and geometric limits

This section groups one positive local construction with two sharp limits on extending it. All energies below use counting measure and original positive-axis starting-site coefficients. Write

$$V_Y(f)=\sum_Y |f-\overline{f}_Y|^2, \qquad E_{a,Y}(f)=\sum_{x,x+e_j\in Y} a(x)|f(x+e_j)-f(x)|^2.$$

No coefficient reflection or endpoint averaging is implicit. Normalized counting means are $q_Y(f)=|Y|^{-1/2}\sum_Y f$. Dividing every term by the ambient volume gives the probability-normalized convention used earlier. The local construction does not cover the whole domain. The first limit concerns a smaller internal-cell certificate; the second concerns an actual annular graph. Neither is a counterexample to the unrestricted full-domain question in Section 27.

40.1 Coefficient-selected unit shells with disjoint faces

Theorem 40.1. *Let $D\ge 3$, $d=D-1$, $P=D/2$, $p=d/2$, $q=2d/(d+2)$, $L\ge 1$ an integer, $z\in\mathbb{Z}^D$, and $\mu>0$. In the nonwrapping integer lattice, on $\Omega=\{x\in\mathbb{Z}^D : L\le \|x-z\|_\infty\le 2L-1\}$ let $a(x)>0$ and put $B=\sum_\Omega(\mu/a)^P$. There is a coefficient-only selected $k\in\{L,\dots,2L-1\}$ whose shell $A_k=\{x\in\mathbb{Z}^D : \|x-z\|_\infty=k\}$ has exactly $2D$ disjoint rectangular faces F, such that for every real f*

$$\sum_F V_F(f)\le \frac{C_D B^{2/D}}{\mu}\sum_F E_{a,F}(f)\le \frac{C_D B^{2/D}}{\mu}E_{a,A_k}(f). \tag{81}$$

Here $K_2=2$ and, for $d\ge 3$,

$$K_d=\frac{2\alpha_d d^{1/q'}}{d(1-2^{-1/d})}, \quad \alpha_d=\frac{2(d-1)}{d}, \quad q'=\frac{2d}{d-2}, \qquad C_D=4K_d^2 3^{2d/q}d^{2/d}.$$

In particular $C_3=2592$. If $cC_DB^{2/D}/\mu\le 1$, these actual means give

$$E_{a,A_k}(f)+c\sum_F |q_F(f)|^2\ge c\sum_{A_k}|f|^2.$$

The same assertion holds coordinatewise without an additional vector-dimension factor. It supplies local rank $2D$ per scalar coordinate only.

Proof. Set $b=\mu/a$, $W_j=\sum_{A_j} b^P$ and choose the least minimizer of W_j. Disjoint radial layers give $W_k\le B/L$. Relative to z, take

$$F_{i,\sigma}=[-k+1,k-1]^{i-1}\times\{\sigma k\}\times[-k,k]^{D-i}, \qquad \sigma\in\{-1,1\}.$$

The first coordinate attaining absolute value k assigns each shell site to a unique face. Each tangential side has $2k-1$ or $2k+1$ sites; the faces are nonempty, including $k=1$, and $|F|\le(4L)^d$.

We spell out the adapter to the complete counting cube Sobolev proof retained earlier in this manuscript. Fold the $(2k+1)$-site interval onto each $(2k\pm1)$-site tangential interval by the endpoint-repeating walk $0,1,\dots,\ell-1,\ell-1,\dots,0,0,\dots$. Every site has between one and three preimages. Every internal edge has at most three preimages: for $k\ge 2$ the walk takes at most two steps beyond its first traversal, while for $k=1$ the possible lengths are one and three.

The constant one-site map has no nonconstant edges. The product fold therefore multiplies an unweighted edge sum by at most 3^d. For its extension $\widetilde{f}$, minimizing over constants gives $V_F(f) \le V_{Q_{2k+1}}(\widetilde{f})$. Thus the proved cube inequality with exponent q gives

$$V_F(f) \le K_d^2 3^{2d/q} \left(\sum_{e \subset F} |df(e)|^q \right)^{2/q}.$$

This fold is entirely unweighted; reversing an edge transports no coefficient. Apply Hölder next on the original face edges:

$$\left(\sum_{e \subset F} |df|^q \right)^{2/q} \le \frac{d^{1/p}}{\mu} E_{a,F}(f) \left(\sum_F b^p \right)^{1/p}.$$

Every site starts at most d such edges. Since $p < P$,

$$\left(\sum_F b^p \right)^{1/p} \le |F|^{1/p - 1/P} W_k^{1/P} \le (4L)^{2/D} (B/L)^{2/D} \le 4B^{2/D}.$$

Here $d(1/p - 1/P) = 2/D$ is the radial cancellation. Summing over faces uses disjoint retained edges. Finally $\sum_{A_k} |f|^2 = \sum_F V_F(f) + \sum_F |q_F(f)|^2$ proves the penalty assertion. All starting sites, including incoming ones, occur in the actual edge sums; no local moment has been assumed. □

The scale cannot be improved to $o(L^2)/\mu$ on homogeneous profiles. Indeed $a = \mu$ gives $B = (4L-1)^D - (2L-1)^D \le (4L)^D$. For every candidate radius and $f(x) = x_2 - z_2$, the face $F_{1,+}$ alone has variance $(2k+1)^d((2k+1)^2 - 1)/12$, whereas $E_{a,A_k}(f) \le \mu |A_k| \le 2D\mu(2k+1)^d$. Consequently

$$\frac{\sum_F V_F(f)}{E_{a,A_k}(f)} \ge \frac{L^2}{6D\mu} \ge \frac{B^{2/D}}{96D\mu}.$$

This uses the whole shell energy, including its face interfaces. It is an all-candidate-radii lower scale bound. A unit shell has only order L^{D-1} sites: its $2D$ observations do not cover the unselected layers or interior. Repeating independently on order L layers costs order L rows. A global covering and charging theorem at the original target count remains missing.

40.2 Sharp rank cost of discarding canonical interfaces

Theorem 40.2. *Fix $D \ge 3$, integer $s \ge 1$, $R = s^D$, $n = 2R + 1$, and $X = \{-R, \dots, R\}^D$, $N = n^D$. Let $a = A = s^2$ on the $2D$ punctured coordinate half-axes and $a = H = n^2$ elsewhere. Its actual mean μ satisfies $1 \le a \le 2\mu$, $\mu \ge 1$, and*

$$1 \le \frac{1}{N} \sum_X (\mu/a)^{D/2} < 2D + 1, \qquad \max\{1, (n/\sqrt{1+\mu})^D\} \le 2^{D/2}.$$

For the actual full energy one rectangular mean suffices:

$$E_{a,X}(f) + c_D |q_X(f)|^2 \ge c_D \|f\|_2^2, \qquad c_D = [64D^2(2D+1)]^{-1}. \tag{82}$$

For any coordinate permutation π, partition instead by the first nonzero coordinate, with the origin as a singleton. Let P_0 be its $2D+1$ normalized means and $E_{\mathrm{int},\pi}$ its retained internal-cell

energy. Set $\ell = 32D(D+1)s$ *and* $m_s = \lfloor (R-1)/(4\ell) \rfloor$. *If real-linear maps* C, V *and any* Φ *with* $\Phi(0) = 0$ *obey, for all* f *or near zero,*

$$E_{\mathrm{int},\pi}(f) + K\|P_0 f\|^2 + \Phi(Cf) + \langle f, Vf \rangle \geq c_D \|f\|^2,$$

then $\operatorname{rank} C + \operatorname{rank} V \geq 2m_s$. *No symmetry or positivity of* V *is needed, and* K *may be any finite scalar. This growing repair cost has matching order* $\Theta_D(s^{D-1})$ *by an explicit refinement.*

Proof. There are $w = 2DR$ weak sites. With $P = D/2$,

$$\mu = H - \frac{w(H-A)}{N}, \qquad S := \sum_X a^{-P} = 2D + 1 - w/N.$$

Since $N \geq 2R3^{D-1}$, $w/N \leq D/3^{D-1} \leq 1/3$ and $\mu \geq 2H/3$. The stated moment equals $(\mu/H)^P S < 2D+1$; its lower bound is Jensen. The target factor follows from $\mu \geq H/2$. The earlier complete cube proof and original-start Hölder give

$$V_X(f) \leq K_D^2 (DS)^{2/D} E_{a,X}(f) \leq 64D^2(2D+1) E_{a,X}(f).$$

For clarity, its constant satisfies $K_D \leq 8\sqrt{D}$: $D(1 - 2^{-1/D}) \geq 1/2$ follows from $(1 - 1/(2D))^D \geq 1/2$, while $\alpha_D \leq 2$ and $D^{1/q'} \leq \sqrt{D}$. The variance identity proves (82). Periodic wrap edges, if added, only help it.

The canonical cells are

$$B_{i,+} = \{x : x_{\pi(j)} = 0 \ (j < i), \ 1 \leq x_{\pi(i)} \leq R\},$$

with all later coordinates free in $[-R, R]$; define $B_{i,-}$ using $[-R, -1]$ and include $\{0\}$. These disjoint integer rectangles cover X. Their final pair are weak half-axes for every π. For $R \geq 2$, the mean-zero linear field $f_t = t - (R+1)/2$ on one final half-axis, zero elsewhere, has

$$P_0 f = 0, \qquad \frac{E_{\mathrm{int},\pi}(f)}{\|f\|^2} = \frac{12s^2}{R(R+1)} \longrightarrow 0.$$

In the actual energy, each transverse incoming edge starts at a two-nonzero-coordinate site and carries H, not A. Therefore any field supported on the final two half-axes has $E_{a,X}(f) \geq (D-1)H\|f\|^2$. The omitted interfaces explain the certificate failure without contradicting the full-energy bound.

To obtain the rank obstruction, use the signed tent on $0 \leq t \leq 4\ell$:

$$g(t) = \begin{cases} t & 0 \leq t \leq \ell, \\ 2\ell - t & \ell \leq t \leq 3\ell, \\ t - 4\ell & 3\ell \leq t \leq 4\ell. \end{cases}$$

It has zero endpoints and sum, and

$$\sum g^2 = 2\ell(2\ell^2 + 1)/3, \qquad \sum |dg|^2 = 4\ell.$$

Place m_s copies on consecutive intervals $[1 + 4\ell j, 1 + 4\ell(j+1)]$ of each final half-axis, indexed by its absolute coordinate. Nonzero supports and edge differences do not overlap; shared endpoints are zero. Their span $G_s \subset \ker P_0$ has dimension $2m_s$, and for every $f \in G_s$

$$E_{\mathrm{int},\pi}(f) = \rho_s \|f\|^2, \qquad \rho_s = \frac{6s^2}{2\ell^2 + 1} \leq \frac{3}{[32D(D+1)]^2} \leq c_D/4.$$

The last step is $4(D+1)^2 \geq 3(2D+1)$. If combined repair rank is smaller than $2m_s$, rank-nullity supplies nonzero $f \in G_s \cap \ker C \cap \ker V$. All repair terms vanish; scaling toward zero also treats a neighborhood inequality. This contradicts the prescribed lower bound. The statement is vacuous when $m_s = 0$, and asymptotically m_s has order s^{D-1}.

For the upper repair order, split each canonical cell's strong complement by its first nonzero later coordinate. There are $2(D-i)$ strong rectangles per $B_{i,\pm}$, hence $2D(D-1)$ in total, each with coefficient $H = n^2$ and largest side at most n. Split each of the $2D$ weak axes into $R/s = s^{D-1}$ intervals of length s and retain the origin. The total is

$$J_s = 1 + 2D(D-1) + 2Ds^{D-1}.$$

The proved interval-product inequality $V_B \leq L_B(L_B-1)\sum_{e\subset B} |df|^2/4$ gives $E_{a,B} \geq 4V_B$ for every nonconstant cell. Singletons have zero variance. Summing and using the normalized mean identity gives coercivity with penalty c_D on these J_s means even for the refined internal energy. That energy is at most $E_{\mathrm{int},\pi}$, so these rows also repair the original certificate (take $\Phi(Cf) = c_D\|P_{\mathrm{ref}} f\|^2$ and a nonnegative base penalty). This proves the matching order without an optimal multiplicative constant. □

This theorem stops transferring a shell gap to these canonical full-volume cells while dropping their interfaces for free. It applies to every order selected from a, but not to every rectangular partition. Indeed its one-cell full-energy solution is explicit. The scalar tents have not been proved coexact or divergence constrained.

40.3 An actual off-center annular slowdown of sharp order

Theorem 40.3. *For $D \geq 3$ and integer $s \geq 2$, put $d = D-1$, $L = s^D$, $Q = 4L+1$,*

$$X = \{0,\dots,4L\}^D, \quad I = \{1,\dots,2L\}^D, \quad \Omega = X \setminus I, \quad N = Q^D - (2L)^D.$$

On the plate $F = \{x_1 = 0,\ 1 \leq x_j \leq 2L\ (j>1)\}$ set $a = A = s^{2d}$, and set $a = H = Q^2$ elsewhere in Ω. Its actual mean obeys $\mu \geq 1$, $1 \leq a \leq 2\mu$, and

$$1 \leq M_a := N^{-1}\sum_{\Omega} (\mu/a)^{D/2} < 2(1+2^d).$$

Nevertheless the best Poincaré constant for the actual nonwrapping annular energy satisfies

$$\frac{s^2}{12d} \leq \sup_{f\ \mathit{nonconstant}} \frac{V_\Omega(f)}{E_{a,\Omega}(f)} \leq B_D s^2, \qquad B_D = 25J(1+8J), \quad J = 2D.$$

Thus the thickness-uniform centered critical-moment formula cannot extend to arbitrary off-center holes without a quantitative geometric condition. The full cube obtained by filling the hole admits one rectangular mean.

Proof. With $w = (2L)^d$ and $P = D/2$, exact counting gives

$$\mu = H - w(H-A)/N, \qquad S := \sum_{\Omega} a^{-P} = 2^d + (N-w)/Q^D.$$

Indeed $A^P = L^d$. Since $w/N \leq [2(2^D-1)L]^{-1} \leq 1/14$, $\mu \geq 13H/14 \geq H/2$. Jensen gives $M_a \geq 1$, and

$$M_a = \mu^P S/N \leq \frac{1+2^d}{1-2^{-D}} < 2(1+2^d).$$

Let $h(t) = \min(t-1, 2L-t)$ on $1 \le t \le 2L$ and put $f(0, x_2, \dots, x_D) = \prod_{j>1} h(x_j)$ on F, zero elsewhere. Direct finite sums give

$$\sum h = L(L-1), \qquad T := \sum h^2 = \frac{(L-1)L(2L-1)}{3}, \qquad \sum |dh|^2 = 2L-2.$$

The normal incoming edge starts outside X, and the outgoing edge enters I; both are absent. At every tangential plate boundary $h = 0$, so all strong incoming or outgoing boundary edges have zero difference. No intercell edge in Ω is discarded. Consequently

$$\|f\|^2 = T^d, \quad \sum f = [L(L-1)]^d, \quad E_{a,\Omega}(f) = Ad(2L-2)T^{d-1}.$$

Cauchy on the w-site support gives $V_\Omega(f) \ge \|f\|^2/2$, whereas $E_{a,\Omega}(f)/\|f\|^2 = 6dA/[L(2L-1)] \le 6d/s^2$. This proves the lower bound. Subtracting its mean yields a mean-zero witness with the same energy and variance. Meanwhile the multiplier in the centered formula is bounded independently of s:

$$Q^2(S/N)^{2/D} \le \left(\frac{1+2^d}{1-2^{-D}}\right)^{2/D}.$$

Thus no dimension-only constant can supply that off-center formula.

For the matching upper bound, cover Ω by $J = 2D$ actual slabs

$$S_{i,-} = \{x : x_i = 0\}, \qquad S_{i,+} = \{x : 2L+1 \le x_i \le 4L\},$$

with all other coordinates free in $[0, 4L]$. Their volumes are $v_- = Q^d$ and $v_+ = 2LQ^d$. Root a tree at $S_{1,+}$, attach every slab with different normal to that root, and attach $S_{1,-}$ to $S_{2,+}$. Paths have length at most two, join distinct normals, and every parent is fat. Descendant volume is no larger than the volume of any slab on its root path.

Write E_i^0 for a slab's unweighted internal energy, m_i for its mean, and $Z = \sum_i E_i^0$. The interval-product proof gives $V_{S_i} \le Q^2 E_i^0/4$. To control means without an inverse overlap-width factor, for any nonempty subinterval T_0 of a Q-site interval B, Jensen and path Cauchy give

$$|\overline{g}_B - \overline{g}_{T_0}|^2 \le \frac{1}{Q|T_0|} \sum_{x \in B, y \in T_0} |g(x) - g(y)|^2 \le Q \sum_{e \subset B} |dg(e)|^2.$$

Averaging over identical fibers of S_i in another slab's normal direction gives

$$|m_i - \overline{f}_{S_i \cap S_j}|^2 \le Q^2 E_i^0 / v_i.$$

The restricted width may be one or $2L$. Applying this on both sides of an intersection yields $|m_i - m_j|^2 \le 2Q^2(E_i^0/v_i + E_j^0/v_j)$. Along a root path use Cauchy, its length at most two and volume monotonicity; each intermediate energy occurs at most twice. Hence $v_i|m_i - m_{\mathrm{root}}|^2 \le 8Q^2 Z$. The cover and variance identity now give, for every real g,

$$\begin{aligned} V_\Omega(g) &\le \sum_i \{V_{S_i}(g) + v_i|m_i - m_{\mathrm{root}}|^2\} \\ &\le (1/4 + 8J)Q^2 Z \le J(1/4+8J)Q^2 E_\Omega^0(g) \le B_D s^2 E_{a,\Omega}(g). \end{aligned}$$

Here every actual edge is counted at most J times, $a \ge A$, $Q \le 5L$, and $L^2/A = s^2$. For $D = 3$, $B_3 = 7350$.

Finally fill I with coefficient H and extend the same plate field by zero. The exact identity is

$$E_{a,X}(f) = E_{a,\Omega}(f) + A\|f\|^2.$$

The restored normal edge starts on the weak plate, so the coefficient is A, not H. All other added differences are zero. On the full cube $\sum_X a^{-P} < 1 + 2^d$, and the complete cube bound used above gives

$$V_X(g) \le 64D^2(1 + 2^d)E_{a,X}(g).$$

Thus $c_X = [64D^2(1 + 2^d)]^{-1}$ and the one mean q_X supply $E_{a,X} + c_X|q_X|^2 \ge c_X\|g\|^2$ for every g. □

Exact centering has not been proved the weakest sufficient condition; quantitative clearance might suffice. The annular quotient of order s^{-2} is a finite scalar Neumann quotient, not a physical mass gap. The full-domain unrestricted question and coexact restrictions remain untouched. An initial unreviewed candidate note incorrectly called the restored edge strong and incoming. Its original and pre-freeze correction are retained; the theorem and its regression fixture use the weak outgoing start. No failed scientific execution was hidden or rerun to change that criterion.

40.4 Comparison, evidence and remaining question

Radial selection, folds, finite Sobolev/Hölder, product tents and rank-nullity are classical. The weighted method should be compared with Andres–Deuschel–Slowik [1], Section 3.1, Proposition 3.5 of arXiv:1306.2521v5: its conductance and cutoff setup is not this finite-face construction. Birman–Solomyak's adaptive cube partitions [2] do not supply an unsupported zero-smoothness partition theorem here. Brudnyi [3], Sections 1–2 and Appendix I, treats function-dependent approximation and a dyadic inner cube strictly inside an outer cube. This does not supply one coefficient-selected observation map valid for all fields, an original-start weighted adapter, or an arbitrary shifted-hole bound. No contrary claim about that covering theorem is made. Exact priority of these quantitative specializations is unestablished; this limited primary comparison is not exhaustive.

The frozen analytical source commits are

selected shell218	`e040475e37bf6b7440625df41c723738dbe3dbd8`
internal certificate219	`5298926f130f3d680dfaa633bbf82e780bd0af84`
off-center annulus220	`0c79cd67a002c24a305b70cf707c420aa1bd9ae3`

Their complete source reports, objections, supplied context and original finite transcripts are retained under the corresponding research folders. All three received accept-with-limits source verdicts in a continuing, shared-context Codex task; this is role-separated but not independent or blinded review. Exact hidden backend identity is unattested. Source218 required only a citation locator correction after its freeze, preserved as a bounded text delta; sources219 and220 required no mathematical correction after review. These reports are distinct from this changed PDF's required numerical assessment.

Actual owned Colab normal and optimized diagnostics agreed in each unit. Source218 checked 160 fields, 960 face/fold cases, 192 moment cases, 20 scale cases and three rejected negative controls. Source219 checked 60 small-grid fields, 1680 refined-cell cases, six linear witnesses, ten profiles, four low-space cases and three rejected negative controls. Source220 checked 20 fields on four small geometries, 120 slab bounds, 200 directional transports, 120 root bounds, five fields on

one full weighted profile, five filled-cube cases, six closed-profile cases and three rejected negative controls. These are selected finite rational diagnostics, not all-size machine proofs, new Lean results or two-fresh-clone terminal reproduction. Runtime allocation, real exits, source hashes and release/readback records are preserved; billed compute units are unknown.

One positive local input has been derived from the coefficients, and two specific extension routes have been stopped by proofs. What is still missing is an actual coefficient-derived full-domain rectangular partition and target-order charging theorem, or a counterexample defeating every allowed partition and finite penalty. Earlier stopped harmonic-cost and common-refinement routes remain stopped. PR29 remains at 20/41 and TermSource = 0. No interacting or continuum conclusion follows.

References

[1] S. Andres, J.-D. Deuschel and M. Slowik, *Invariance principle for the random conductance model in a degenerate ergodic environment*, Annals of Probability 43(4) (2015), 1866–1891; arXiv:1306.2521v5 (2019), Section 3.1, Proposition 3.5. `https://arxiv.org/abs/1306.2521v5`.

[2] M. Sh. Birman and M. Z. Solomyak, *Piecewise-polynomial approximations of functions of the classes* W_p^α, Math. USSR-Sbornik 2(3) (1967), 295–317, Sections 1–2. `https://doi.org/10.1070/SM1967v002n03ABEH002343`.

[3] Yuri Brudnyi, *Nonlinear piecewise polynomial approximation and multivariate BV spaces of a Wiener–L. Young type. I*, arXiv:1511.03971v1 (2015), Sections 1–2 and Appendix I. `https://arxiv.org/abs/1511.03971v1`.

[4] L. Eriksson, *Exact-Kernel Fourier Families Obstruct Every Penalty Factored Through a Block Map in a Flat Lattice Gauge Form*, ai.viXra 2608.0015v1 (2026), AIRR ARR-2026-7EPM3V3DP796BBRR v1. `https://airr.science/papers/ARR-2026-7EPM3V3DP796BBRR/versions/v1/`. Historical import, not peer reviewed or assessed by AIRR; checked 30 October 2026.

[5] J. Dodziuk, *Finite-difference approach to the Hodge theory of harmonic forms*, American Journal of Mathematics **98** (1976), 79–104. `https://doi.org/10.2307/2373615`. Institutional source: `https://academicworks.cuny.edu/gc_pubs/180/`.

[6] A. Albert, *Conditions for Positive and Nonnegative Definiteness in Terms of Pseudoinverses*, SIAM Journal on Applied Mathematics **17** (1969), 434–440. `https://doi.org/10.1137/0117041`.

[7] D. Carlson, E. Haynsworth and T. Markham, *A Generalization of the Schur Complement by Means of the Moore–Penrose Inverse*, SIAM Journal on Applied Mathematics **26** (1974), 169–175. `https://doi.org/10.1137/0126013`.

[8] J. Dimock, *Covariant Axial Gauge*, arXiv:1410.2798v2 (2014). `https://arxiv.org/abs/1410.2798v2`.

[9] Y. Kudryashov and Mathlib contributors, *Implicit function theorem*, Mathlib source at commit `07642720480157414db592fa85b626dafb71355b`. Pinned primary formal source, particularly `HasStrictFDerivAt.to_implicitFunction` and its exact level-set identity; inspected 30 October 2026. This attributes the classical parametrization; the separately pinned sector formalizations are identified in the evidence section.

[10] D. A. Spielman, *Graphs, Vectors, and Matrices*, Bulletin of the American Mathematical Society **54** (2017), 45–61. Author-hosted primary exposition: https://cs-www.cs.yale.edu/homes/spielman/PAPERS/bull1557.pdf.

[11] D. A. Levin and Y. Peres, with contributions by E. L. Wilmer, *Markov Chains and Mixing Times*, second edition, American Mathematical Society (2017), Chapter 13. Author-hosted text: https://pages.uoregon.edu/dlevin/MARKOV/markovmixing.pdf.

[12] D. Zelazo and M. Bürger, *On the Definiteness of the Weighted Laplacian and its Connection to Effective Resistance*, arXiv:1408.2187v1 (2014), Theorem III.3. https://arxiv.org/pdf/1408.2187.

[13] Y. Lin, S. Wan and H. Zhang, *Connection Laplacian on discrete tori with converging property*, arXiv:2403.06105v1 (2024), Theorem 1.2. https://arxiv.org/html/2403.06105v1. The sin-squared discrete-torus spectrum is attributed prior work, not a new general spectral principle here.

[14] M. de Berg, *Linear Size Binary Space Partitions for Uncluttered Scenes*, UU-CS-1998-12 (1998), Section2.1, Lemmas2.1 and2.4, Theorem2.5 and Corollary2.9. https://ics-archive.science.uu.nl/research/techreps/repo/CS-1998/1998-12.pdf.

[15] W. N. Anderson, Jr. and G. E. Trapp, *Shorted Operators. II*, SIAM J. Appl. Math. 28(1) (1975), 60–71. https://doi.org/10.1137/0128007.

[16] R. L. Frank, *Cwikel's theorem and the CLR inequality*, arXiv:1206.3325v1 (2012), Section2, Theorem2.1, Corollary2.2, Lemmas2.4–2.5. https://arxiv.org/abs/1206.3325.

[17] S. Andres, J.-D. Deuschel and M. Slowik, *Harnack inequalities on weighted graphs and some applications to the random conductance model*, arXiv:1312.5473v6 (2019), Section2.2, equation(2.5) and Proposition2.1. https://arxiv.org/abs/1312.5473.

[18] R. L. Frank, *Cwikel's theorem and the CLR inequality*, arXiv:1206.3325v1 (2012), Theorem 3.1. https://arxiv.org/abs/1206.3325v1.

[19] D. Zelazo and M. Bürger, *On the definiteness of the weighted Laplacian and its connection to effective resistance*, IEEE CDC (2014), Theorem III.3. https://arxiv.org/abs/1408.2187.

[20] D. A. Spielman, *Spectral Graph Theory Lecture 14*, 15 October 2018, Claim 14.6.2. https://www.cs.yale.edu/homes/spielman/561/lect14-18.pdf.

[21] W. Chen, J. Liu, Y. Chen, S. Z. Khong, D. Wang, T. Başar, L. Qiu and K. H. Johansson, *Characterizing the positive semidefiniteness of signed Laplacians via effective resistances*, IEEE CDC (2016), pp. 985–990. https://people.kth.se/~kallej/papers/network_cdc16chen.pdf.

[22] N. J. Higham, *What Is the Schur Complement of a Matrix?*, 1 June 2023, Theorem 1, generalized Schur complement. https://nhigham.com/2023/06/01/what-is-the-schur-complement-of-a-matrix/.

[23] P. Diaconis and L. Saloff-Coste, *Comparison Theorems for Reversible Markov Chains*, Annals of Applied Probability 3(3) (1993), 696–730, Section2B, Theorem2.1 and Example2.1. https://pi.math.cornell.edu/~lsc/papers/comp-rev.pdf.

[24] L. Eriksson, *Coarse information and finite rank obstructions for flat lattice Hodge forms*, existing AIRR dossier ARR-2026-7EPM3V3DP796BBRR, historical public version8 (the immediate predecessor of this candidate is version11). https://airr.science/papers/ARR-2026-7EPM3V3DP796BBRR/versions/v8/. Working paper with AI assessment, not editorial acceptance.

[25] B. Muckenhoupt, *Hardy's inequality with weights*, Studia Mathematica 44 (1972), 31–38. https://doi.org/10.4064/sm-44-1-31-38.

[26] P. Binev and R. DeVore, *Fast computation in adaptive tree approximation*, Numerische Mathematik 97 (2004), 193–217. https://doi.org/10.1007/s00211-003-0493-6.

[27] R. H. Nochetto, E. Otárola and A. J. Salgado, *Piecewise polynomial interpolation in Muckenhoupt weighted Sobolev spaces and applications*, arXiv:1402.1916v2 (2014), Lemma 3.1, Section 5.2 and Theorem 5.7. https://arxiv.org/abs/1402.1916v2.

[28] R. L. Frank, *Cwikel's theorem and the CLR inequality*, arXiv:1206.3325v1 (2012), Section2. https://arxiv.org/abs/1206.3325v1.

[29] L. Eriksson programme research evidence, units 184–185, pinned at commit 506004b3e7251aa881b08dde1a1bb14aff33e76f. https://github.com/lluiseriksson/THE-ERIKSSON-PROGRAMME. At the pinned commit, see papers/flat-coarse-rank/research/evidence185 and papers/flat-coarse-rank/research/2026-11-07_green-moment-reproduction184.txt. The complete unit 184 archive accompanies the integrated source package.

[30] S.-K. Chua, *Extension Theorems on Weighted Sobolev Spaces and Some Applications*, Canadian Journal of Mathematics 58(3) (2006), 492–528. Theorem 1.4. https://doi.org/10.4153/CJM-2006-021-0.

[31] F. Dörfler and F. Bullo, *Kron Reduction of Graphs with Applications to Electrical Networks*, primary author manuscript, Theorem III.8 and equation (27); related arXiv version1102.2950. https://motion.me.ucsb.edu/pdf/2011d-db.pdf.