# Rectangular observations for finite matrix gradients: observation classes and critical endpoints

Lluis Eriksson

10 November 2026 — manuscript 232, offline integrated source candidate

**Abstract**

We study which observation classes restore coercivity of actual finite forward gradient forms. Arbitrary finite symmetric onsite matrices admit coefficient-only common rectangular means with output dimension $O_D(h(1+F_c(A)))$, where $F_c(A)$ is a capped local eigenchannel budget and the finite penalty may depend on negative depth. On the actual flat coexact domain we derive the sharp universal common-partition rank bound $g(D-1)(R-1)$, including its annihilator formula and an attaining fibre family. Explicit coexact low fields make product order necessary for common full-vector partitions. Component-adapted rectangles instead give an additive $O_D(h+F_c(A))$ bound for component-block-diagonal coefficients and escape that obstruction on the same family. For constant internally coupled coefficients, fixed-basis channel-separable observations can again require super-additive actual coexact rank of sharp order $\Theta_D(n^{2D-2})$; a cheap mixed observation demonstrates why that obstruction is basis dependent. The scalar weak critical endpoint has matching worst-case logarithmic upper and lower order. Complete scalar BV, weighted-gap, tree, matrix, coexact and endpoint proofs are included. The mechanisms are classical; priority, sharp profilewise ranks, efficient algorithms and physical or continuum transfer are not established.

## 1 Objects, principal results and scope

A small number of spatial averages can restore coercivity even when local coefficients are very small or have negative channels. Two issues must be kept separate: how many averages are needed, and how large a penalty must be placed on them. We control the first by a capped coefficient budget while allowing the second to depend on the entire finite profile. At the scalar weak endpoint a logarithmic loss in the coefficient scale is unavoidable in the worst case.

### 1.1 The actual finite form and observation rows

Except in the standalone coexact rank theorem, $D \geq 3$, $n \geq 1$ and $h \geq 1$ are integers, $c > 0$,

$$X = (\mathbb{Z}/n\mathbb{Z})^D, \qquad V = n^D, \qquad P = D/2, \qquad r = Dh.$$

We represent $X$ by the fixed chart $\{0,\dots,n-1\}^D$. For $u : X \to \mathbb{R}^h$ the forward stack is

$$du(x) = \big(u(x+e_i) - u(x)\big)_{i=1}^{D} \in \mathbb{R}^r.$$

All additions in this expression are periodic. Given a finite real symmetric $r \times r$ matrix $A(x)$ at each original start $x$, define

$$Q_A(u) = \sum_{x \in X} du(x)^{\mathsf{T}} A(x) du(x), \qquad \|u\|_V^2 = V^{-1} \sum_x \|u(x)\|^2. \tag{1}$$

No common eigenbasis or sign assumption is imposed on $A(x)$. For a scalar profile $a > 0$ we write

$$Q_a(u) = \sum_{x,i} a(x)\|u(x+e_i) - u(x)\|^2.$$

In particular, the coefficient is always evaluated at the forward start; it is never moved to the endpoint or averaged over an edge.

An integer rectangle is a nonempty product of consecutive intervals in the fixed chart. It need not be a periodic rectangle crossing the chart boundary. For such a rectangle $B$ define the vector row

$$q_B(u) = V^{-1/2}|B|^{-1/2} \sum_{x\in B} u(x) = \sqrt{|B|/V}\ \mathrm{mean}_B\, u. \tag{2}$$

For a partition $\mathcal{R}$ put $C_{\mathcal{R}}u = (q_B(u))_{B\in\mathcal{R}}$, with the Euclidean product norm. Its output dimension is $h\#\mathcal{R}$; on the full field space its rows are independent, so its rank equals that dimension. A coefficient-only partition is fixed by the coefficient profile and deterministic geometric conventions before the field $u$ is chosen.

### 1.2 An arbitrary symmetric matrix theorem

The constants used below are intentionally conservative:

$$\begin{aligned} b &= 2^D, \quad J = 2D, \quad C_{\mathrm{ext}} = 1 + J\{2J(1+12J)+1\}, \\ G_D &= 2^{2D+10}D^2C_{\mathrm{ext}}^2, \qquad B_D = (2b+3)bG_D^P, \\ \alpha_D &= \frac{D}{2}\left(\frac{D}{D-2}\right)^{D-1}\left(\frac{3}{4}\right)^D. \end{aligned} \tag{3}$$

Here $G_D$ is the scalar weighted-gap constant and $B_D$ is the resulting tree-count constant. Let $\lambda_j(x)$, $1 \le j \le r$, be the eigenvalues of $A(x)$, counted with multiplicity. Let $\nu_0$ count all its nonpositive local eigenvalues over all $x$. Define

$$F_c(A) = \nu_0 + \sum_{x,j:\lambda_j(x)>0} \min\{1, \alpha_D c^P \lambda_j(x)^{-P}\}. \tag{4}$$

These are local coefficient eigenvalues, not the eigenvalues of the global operator associated with $Q_A$.

**Theorem 1.1** (Capped-budget rectangular observations)**.** *For every profile $A$ in* (1) *there is a coefficient-only partition $\mathcal{R}$ into integer rectangles such that*

$$h\#\mathcal{R} \le 2Dh\left[1 + B_D + \max\left\{\frac{B_D}{\alpha_D}, D+1\right\} F_c(A)\right]. \tag{5}$$

*There is a finite coefficient-dependent $K$ such that, simultaneously for all $u : X \to \mathbb{R}^h$,*

$$V^{-1}Q_A(u) + K\|C_{\mathcal{R}}u\|^2 \ge c\|u\|_V^2. \tag{6}$$

*One may take $K = c + 4DL$, with $L$ explicitly defined in* (28)*. Neither the partition nor $K$ depends on $u$.*

The budget caps the cost of each sufficiently small positive channel at one and assigns the same unit cost to each nonpositive channel. Its count thus does not grow with negative depth. The penalty $K$ can grow with that depth; (6) makes no assertion that $K = c$ for signed forms. The elementary singleton partition has $hV$ outputs and also admits a finite penalty for every finite symmetric form. One may choose whichever of that partition and Theorem 1.1 has fewer outputs. No optimal dimension, sharp coefficient, or efficient implementation is asserted.

### 1.3 The cost of a common full-vector partition

The dimension in Theorem 1.1 concerns a common partition with all $h$ component means on each spatial member. The following refinement shows why this measurement class can require a product of $h$ and the budget.

**Theorem 1.2** (Necessary product order for common partitions)**.** *Fix $D \geq 3$ and $c > 0$. For every pair of integers $m, h \geq 1$ there is an actual positive diagonal matrix profile on $X_n$, $n = 3m + 3$, with exactly $2Dm$ weak direction-labelled channels and*

$$2Dm \leq F_c(A) \leq 2Dm + 1.$$

*If a common partition $\mathcal{R}$ outputs all $h$ means* (2) *and satisfies* (6) *for every general field $u : X_n \to \mathbb{R}^h$ with any finite quadratic penalty coefficient, then*

$$\operatorname{rank} C_{\mathcal{R}} = h\#\mathcal{R} \geq hm \geq \frac{hF_c(A)}{2D + 1}. \tag{7}$$

*Thus no bound $C_D(h + F_c(A))$ holds uniformly in $h$ in this class: setting $h = m$ makes the ratio of necessary rank to $h + F_c(A)$ diverge. Together with Theorem 1.1, this gives matching worst-case $hF_c(A)$ order, up to dimensional constants, in joint $h, F_c(A)$ growth. The lower family consists of general vector fields and is not claimed to be coexact or divergence-free. It does not address component-dependent partitions, channel-adapted means, or postprocessed measurement classes.*

### 1.4 The scalar endpoint and the precise lower-bound domain

For a positive scalar profile let

$$\begin{aligned} \mu = V^{-1} \sum_x a(x), \qquad & Y(x) = \mu / a(x), \\ W(a, \mu) = \sup_{\tau > 0} \tau^P \mathbb{P}_X(Y > \tau), \quad & T(n, \mu) = \max\{1, n^D/(1 + \mu)^P\}, \end{aligned} \tag{8}$$

where $\mathbb{P}_X$ is normalized counting probability. The strict superlevel convention in this definition is used throughout.

**Theorem 1.3** (Weak endpoint upper bound)**.** *Suppose $1 \leq a \leq 2\mu$ and $W(a, \mu) \leq M_w$, with $\mu$ the actual arithmetic mean. For every $h \geq 1$ there is a coefficient-only rectangular partition with*

$$h\#\mathcal{R} \leq 2Dh \left\{1 + 2^P c^P B_D[1 + PM_w \log \mu]\right\} T(n, \mu) \tag{9}$$

*and*

$$V^{-1} Q_a(u) + c\|C_{\mathcal{R}} u\|^2 \geq c\|u\|_V^2 \quad (u : X \to \mathbb{R}^h). \tag{10}$$

To state the lower bound without an unspecified Hodge domain, fix $g \geq 1$. A flat one-cochain is $U = (U_1, \dots, U_D)$ with $U_j : X \to \mathbb{R}^g$; we regard it as an $h = Dg$ component field in $Q_a$. For antisymmetric two-cochains $B_{ij} = -B_{ji} : X \to \mathbb{R}^g$, define

$$D_i^- f(x) = f(x) - f(x - e_i), \qquad (\delta B)_j = -\sum_i D_i^- B_{ij}, \qquad \mathcal{H}_{\mathrm{coex}}(X) = \{\delta B : B \text{ antisymmetric}\}. \tag{11}$$

Backward divergence $\sum_j D_j^- U_j$ vanishes on this space by commutation and antisymmetry. The energy on it is precisely the scalar component-gradient $Q_a(U)$, with all its original forward starts; no weighted curl identity is used to define or replace that energy.

**Theorem 1.4** (Weak endpoint obstruction and worst-case order)**.** *Fix $D \geq 3$, $c > 0$ and $g \geq 1$. For every integer $m \geq 1$ there are $n_m$, a profile $a_m$ with actual mean $\mu_m$, and an $m$-dimensional real subspace $E_m \subset \mathcal{H}_{\mathrm{coex}}(X_{n_m})$ such that*

$$\begin{aligned} 1 \leq a_m \leq 2\mu_m, \qquad & W(a_m, \mu_m) \leq 1, \\ & Q_{a_m}(U) \leq (c/4) \sum_x \|U(x)\|^2 \quad (U \in E_m), \\ & T(n_m, \mu_m) \leq C(D, c), \qquad \mathbb{E}_X(\mu_m / a_m)^P \geq \frac{m}{2^P 16^D}. \end{aligned} \tag{12}$$

*Suppose an observation $C$ from a neighborhood of zero in $\mathcal{H}_{\mathrm{coex}}(X_{n_m})$ to a finite-dimensional real normed space $F$ is differentiable at zero and $C(0) = 0$. Suppose also that $\Phi(y) \leq K\|y\|^2$ near zero for some finite $K$, and that locally*

$$c \sum_x \|U(x)\|^2 \leq Q_{a_m}(U) + \Phi(C(U)). \tag{13}$$

*Then $DC(0)|_{E_m}$ is injective and $\dim F \geq m$. For this fixed-parameter sequence $m$ is comparable to $1 + \log \mu_m$ for large $m$. Thus the logarithmic loss in Theorem 1.3 has matching worst-case order. This is not a claim of optimal rank for every individual profile.*

The centered, differentiable and local quadratic assumptions are part of the lower theorem. Its penalty need not have a uniform bound as $m$ varies. A common division of energy and squared field norm by $V$ gives the same rank conclusion. The full critical moment grows along the lower family, so this family does not contradict the strong-moment scalar theorem.

### 1.5 Observation classes on the actual coexact domain

A coexact one-cochain has spatial components $j = 1, \dots, D$ and internal channels $\alpha = 1, \dots, g$. Thus $h = Dg$; we identify its real components with pairs $(j, \alpha)$. For one internal channel write $\mathcal{U}_0 = \mathrm{im}(d_1^*)$, where

$$(d_1 u)_{ij} = D_i^+ u_j - D_j^+ u_i \quad (i < j), \qquad D_i^+ f(x) = f(x + e_i) - f(x).$$

The adjoint uses counting inner products on independent $i < j$ two-cochain coordinates. With antisymmetric extension of those coordinates, $(d_1^* B)_j = \sum_i D_i^- B_{ji}$, exactly the domain (11). With $g$ channels the domain is $\mathcal{U} = \mathcal{U}_0^g$. A common volume normalization changes none of these images or ranks.

Three observation classes must be distinguished. A common full-vector partition retains every component mean on each member. A component-adapted map chooses its own rectangle

partition for each real component. A channel-separable linear map may mix all sites and spatial components within each fixed internal channel, but its rows do not mix distinct internal channels. Component-adapted rows belong to this last class. A global internal rotation changes what "channel-separable" means.

**Theorem 1.5** (Sharp common-partition coexact rank)**.** *For this theorem alone let $D \geq 2$, $n, g \geq 1$. For any partition into $R$ nonempty subsets, its full-vector mean map satisfies*

$$\operatorname{rank}(C_{\mathcal{R}}|_{\mathcal{U}}) \geq g(D-1)(R-1). \tag{14}$$

*The bound is sharp as a dimension-only bound: with $n = R$, coordinate-fibre partitions attain equality for every $R \geq 1$. No geometric hypothesis on the partition is required. For $D \geq 3$ there is a positive diagonal coefficient family, specified in* (36)*, for which every restoring common full-vector partition, with any penalty vanishing at zero, has*

$$\operatorname{rank}(C_{\mathcal{R}}|_{\mathcal{U}}) \geq \frac{h(D-1)}{4D(\kappa_D+1)} F_c(A), \qquad \kappa_D = 2D\,17^D. \tag{15}$$

*For that family $m \geq 4$, $g = m$, $h = Dm$ and $\kappa_D m \leq F_c(A) \leq \kappa_D m + 1$. Hence product order is necessary for actual restricted rank on the coexact domain, not just for the unrestricted output dimension $hR$.*

Define the dimensional constants

$$a_D = 2D(1 + B_D), \qquad b_D = 2D \max\{B_D/\alpha_D, D+1\}. \tag{16}$$

They must not be confused with the coefficient profile $a(x)$ or the balanced-tree branching bound $b = 2^D$.

**Theorem 1.6** (Additive component-adapted rectangles)**.** *Suppose $D \geq 3$ and $A(x)$ is block diagonal by the actual real field component: $A_{(i,k),(j,\ell)}(x) = 0$ for $k \neq \ell$. Its $D \times D$ blocks may be signed and may couple directions. There are coefficient-only rectangle partitions $\mathcal{R}_k$, with rows*

$$q_{k,B}(u) = V^{-1/2}|B|^{-1/2} \sum_{x \in B} u_k(x), \qquad B \in \mathcal{R}_k,$$

*and a finite coefficient-dependent $K \geq c$ such that*

$$N_{\mathrm{out}} = \sum_k \#\mathcal{R}_k \leq a_D h + b_D F_c(A), \qquad V^{-1} Q_A(u) + K\|C_{\mathrm{ad}} u\|^2 \geq c\|u\|_V^2. \tag{17}$$

*This is an all-field inequality. On the exact family of* (36)*, a choice of such maps has*

$$\operatorname{rank}(C_{\mathrm{ad}}|_{\mathcal{U}}) \leq 2a_D + b_D F_c(A) = O_D(m), \tag{18}$$

*whereas every restoring common full-vector partition has restricted rank at least $(D-1)m^2/4$. The component-block-diagonal coefficient hypothesis is retained; it is not asserted for arbitrary coupled profiles.*

The next obstruction concerns a larger observation class but coupled coefficients. Its internal dimension grows with the torus size.

**Theorem 1.7** (A fixed-basis channel-separable obstruction)**.** *Let $D \geq 3$, $n \geq 2$, $V = n^D$, $g = V$, $h = Dg$ and $d_0 = (D-1)(V-1)$. There is an explicit constant positive coupled profile* (61) *with*

$$h + F_c(A) \leq L_D V, \qquad L_D = D + D^2(1 + \alpha_D),$$

*for which every channel-separable linear observation $C$ restoring $c$-coercivity on $\mathcal{U}$ with any penalty $\Phi(0) = 0$ has actual rank*

$$\operatorname{rank}(C|_{\mathcal{U}}) \geq \frac{g(D-1)(V-1)}{2(1+8Dn^2)}. \tag{19}$$

*Consequently no uniform $C_D(h + F_c(A))$ bound holds in this fixed-basis channel-separable class, even for positive constant coefficients. For the very same profile, cross-channel mixing admits a restoring map of actual rank $d_0 = O_D(h + F_c(A))$ and penalty coefficient $c$. Thus this is neither an obstruction for all linear observations nor a basis-invariant impossibility. It gives no fixed-$g$ impossibility.*

For this exact family, Appendix C constructs allowed channel-separable real-mode rows and proves the matching minimum-rank order $r_{\min} = \Theta_D(n^{2D-2})$, with explicit dimensional constants in Theorem C.1. Its finite penalty is $K = c + (4DH)^2/c$, with $H = cn^2$. These rows may be nonlocal and are not claimed rectangular.

### 1.6 Classical predecessors, contribution and proof dependencies

This is an integrated successor source for the existing AIRR record `ARR-2026-7EPM3V3DP796BBRR`, following its exact version 14 (manuscript 226C). It consolidates the same finite-gradient question rather than presenting separate papers for its rank refinements. A new source is not the same reviewed artifact as its predecessor.

The scalar construction uses substantive classical predecessors: Cohen–DeVore–Petrushev–Xu [3], Section 6 and Theorem 6.1, explicitly organize dyadic rings through final, branching and chaining nodes; Brudnyi [2], Appendix II, gives weighted-tree/basic-path approximation mechanisms, with ring covers in Appendix I. Those constructions are pertinent prior art, not merely background mentions. The finite implementation here treats atomic coefficient mass, odd integer subdivision, original-start edge weights and retained intra-ring interfaces. The matrix application uses local spectral replacement and exact endpoint refinement; the coexact applications use explicit low fields and restricted image rank. These are descriptions of the proved adaptations, not proofs that no classical theorem already implies them. Historical priority, exact equivalence/non-equivalence and exhaustive predecessor coverage remain unestablished. The elementary forward maximum principle is explicitly attributed where it is used in Section 4.

Frank [1], Theorems 2.1 and 3.1, provides the classical Cwikel/CLR and operator-valued context. A weak multiplier condition there must be distinguished from the weak spatial coefficient distribution here. Our matrix proof invokes no global CLR estimate. The positive $\alpha_D$ normalizes the cap; another fixed positive normalization changes the corresponding constants. Spectral replacement, direct-sum spectra, BV, tree compression, layer cake and finite projection/rank principles are classical. No new general principle is claimed.

The central dependency map is as follows. Every named input is proved within this manuscript, including the shared bump used by two different coefficient families.

| Result | Proof dependencies and domain |
|---|---|
| Scalar foundation | Appendix A: integer collars → BV extension → original-start weighted ring gap → compressed heavy chains → rectangular means, retaining intra-ring interfaces. |
| Matrix upper | Scalar foundation → Section 2: positive replacement, endpoint refinement and full signed correction; all fields. |
| Common lower bounds | Section 3: general-field deltas. Section 4: closed-field decomposition, maximum principle and actual coexact low fields; these ranks are distinct. |
| Adapted upper/escape | The $h = 1$ matrix theorem, exact block spectra and the forward torus gap, in Section 5. |
| Weak endpoint | Scalar moment corollary plus layer cake for the upper; Lemma 6.1, geometric scales and derivative-kernel scaling for the actual coexact lower, in Section 6. |
| Coupled obstruction | Actual coexact dimension and row-space projections inside that domain and real-mode upper rows with a finite Green trace; Appendices B–C. Internal dimension grows and the observation basis is fixed. |

The full-vector common partition, component-adapted partition and fixed-basis channel-separable classes cannot be interchanged in a rank comparison. Restricted rank always means the actual image dimension of the stated map on the stated domain. In particular the coexact argument never substitutes the full-space count $hR$ for restricted rank. All results concern finite unit-spacing component-gradient forms. No interacting/covariant coefficient producer, variable-weight curl identity, continuum construction or Yang–Mills solution follows. External programme bookkeeping is omitted from the proof because it supplies no mathematical hypothesis or inference here.

## 2 Positive replacement and exact rectangular refinement

We prove Theorem 1.1. Every step uses the actual matrix profile on the forward stack in (1). In particular, positivity is used to discard wrap terms only after replacing the signed profile.

### 2.1 The spectral split and scalar inverse moment

Set

$$\sigma = \alpha_D^{1/P} c, \qquad S_x = 1_{(-\infty,\sigma]}(A(x))\mathbb{R}^r, \qquad N = \sum_x \dim S_x, \qquad I = \{x : S_x \neq \{0\}\}.$$

Here $1_{(-\infty,\sigma]}(A(x))$ is the spectral projector. A degenerate eigenspace at the cutoff is selected in full, so no choice of basis within it enters the construction. Every nonpositive channel is selected. A positive selected channel contributes exactly one to the capped budget, and an unselected channel contributes its uncapped inverse power. Consequently

$$F_c(A) = N + \alpha_D c^P \sum_{x,j:\lambda_j(x)>\sigma} \lambda_j(x)^{-P}, \qquad 0 \le N \le F_c(A), \quad |I| \le N. \tag{20}$$

Equality $\lambda = \sigma$ is included in $N$.

Put $T = \sigma + c(1 + N) > \sigma$. Replace every selected eigenvalue by $T$, leaving all unselected eigenvalues and every local spectral subspace unchanged, and denote the resulting positive

definite matrix by $A_T(x)$. For $N = 0$ this changes nothing. With $a_T(x) = \lambda_{\min}(A_T(x)) > 0$, pointwise matrix positivity gives

$$Q_{A_T}(u) \geq Q_{a_T}(u). \tag{21}$$

This controls the complete stacked vector, including all directional and internal mixing. It requires no commutation between distinct matrices. Moreover, the largest eigenvalue of $A_T(x)^{-P}$ is at most its trace, hence

$$\begin{aligned} c^P \sum_x a_T(x)^{-P} &\leq c^P \sum_x \operatorname{tr}(A_T(x)^{-P}) \\ &= c^P \sum_{\lambda_j(x)>\sigma} \lambda_j(x)^{-P} + N(c/T)^P \\ &\leq \frac{F_c(A) - N}{\alpha_D} + 1. \end{aligned} \tag{22}$$

The last inequality uses $T \geq c(1 + N)$ and $N/(1 + N)^P \leq 1$ for $P \geq 1$. When $N = 0$ the extra one is harmless.

### 2.2 The scalar partition is derived, not assumed

Theorem A.1 in Appendix A applies to every finite positive $a_T$, without an arithmetic-mean, lower-envelope or upper-envelope condition. Its partition $\mathcal{R}_0$, combined with (21), satisfies

$$c\|u\|_V^2 \leq V^{-1} Q_{A_T}(u) + c\|C_{\mathcal{R}_0} u\|^2, \tag{23}$$

by summing the same scalar inequality over all $h$ real components. If $R_0 = \#\mathcal{R}_0$, its exact count and (22) yield

$$\begin{aligned} R_0 &\leq 2D \left[1 + (2b + 3)b(cG_D)^P \sum_x a_T(x)^{-P}\right] \\ &\leq 2D \left[1 + B_D \left(1 + \frac{F_c(A) - N}{\alpha_D}\right)\right]. \end{aligned} \tag{24}$$

There is no auxiliary local-gap or extension hypothesis in this application: the scalar construction and its gap are proved in the appendix.

### 2.3 Isolating every affected endpoint

Let

$$Z = I \cup \bigcup_{i=1}^{D} (I + e_i), \qquad |Z| \leq (D + 1)|I| \leq (D + 1)N, \tag{25}$$

where the translations are periodic. We refine $\mathcal{R}_0$ so that each $z \in Z$ is a singleton member. Order $Z$ lexicographically in the fixed chart. If $z$ is not already a singleton, let $R$ be its current rectangle. Partition $R \setminus \{z\}$ according to the first coordinate in which a point differs from $z$, and according to whether it is smaller or larger in that coordinate. Earlier coordinates are fixed to those of $z$, later coordinates range over their original intervals, and the differing coordinate ranges below or above $z$. Each nonempty piece is an integer rectangle. There are at most $2D$ such pieces. Adding $\{z\}$ replaces one rectangle by at most $2D + 1$, increasing the count by at

most $2D$. Earlier isolated points lie in disjoint singleton members and remain intact. Thus the final partition $\mathcal{R}$ satisfies

$$R := \#\mathcal{R} \le R_0 + 2D|Z| \le R_0 + 2D(D+1)N. \tag{26}$$

A wrapping forward endpoint is simply its chart representative in $Z$; this causes no alteration of the rectangle argument.

For each original parent $B$ and its final pieces $B_k$,

$$q_B(u) = \sum_k \sqrt{|B_k|/|B|}\, q_{B_k}(u), \qquad \sum_k |B_k|/|B| = 1.$$

Hilbert-space Cauchy with these exact weights proves

$$\|C_{\mathcal{R}_0} u\|^2 \le \|C_{\mathcal{R}} u\|^2, \qquad V^{-1} \sum_{z \in Z} \|u(z)\|^2 \le \|C_{\mathcal{R}} u\|^2. \tag{27}$$

The second assertion holds because each $z \in Z$ contributes the singleton row $V^{-1/2} u(z)$. The remaining rows are still only component means on rectangles; no gradient or nonrectangular observation has been introduced.

### 2.4 Charging the complete signed correction

Define the finite number

$$L = \begin{cases} 0, & N = 0, \\ \max\limits_{x,j:\lambda_j(x) \le \sigma} (T - \lambda_j(x)), & N > 0. \end{cases} \tag{28}$$

The spectral construction gives $0 \le A_T(x) - A(x) \le L\mathrm{I}$, with zero difference for $x \notin I$. This is an operator inequality on all $r$ coordinates; in particular it bounds all rotated off-diagonal contributions together. It does not assign signs to individual cross terms. Therefore

$$\begin{aligned} 0 \le Q_{A_T}(u) - Q_A(u) &\le L \sum_{x \in I} \|du(x)\|^2 \\ &\le 2L \sum_{x \in I} \left[ D\|u(x)\|^2 + \sum_i \|u(x+e_i)\|^2 \right] \\ &\le 4DL \sum_{z \in Z} \|u(z)\|^2. \end{aligned} \tag{29}$$

Indeed each site has its own bad-start multiplicity at most $D$ and at most $D$ incoming direction-labelled starts. All their endpoints belong to $Z$. This counts parallel edges at period two separately, as required by the actual form. At $n = 1$ every difference vanishes and the estimate remains valid. No periodic term has been dropped from the signed correction.

Insert (29) into (23), divide the point sum by $V$, and use (27). This proves (6) with $K = c + 4DL$. The count is

$$\begin{aligned} hR &\le 2Dh \left[ 1 + B_D + \frac{B_D}{\alpha_D}(F_c(A) - N) + (D+1)N \right] \\ &\le 2Dh \left[ 1 + B_D + \max\{B_D/\alpha_D, D+1\} F_c(A) \right], \end{aligned}$$

using the nonnegativity of both $N$ and $F_c(A) - N$. This completes the proof of Theorem 1.1.

**Remark 2.1** (What the cap does and does not control)**.** *The only use of the budget after the local spectral split is the exact identity* (20)*. In particular its control of observation number is compatible with arbitrarily negative selected eigenvalues, since the penalty absorbs the actual finite correction. When $N = 0$, the correction vanishes and $K = c$. For a positive scalar profile the sharper direct scalar construction also has $K = c$ without selecting or replacing any channels. For completeness, choose $M = \max\{0, \max_x[-\lambda_{\min}(A(x))]\}$, so $M \geq 0$ and $A(x) \geq -M\mathrm{I}$ for all $x$. Then $Q_A(u) \geq -4DM \sum_x \|u(x)\|^2$, because $\sum_x \|du(x)\|^2 \leq 4D \sum_x \|u(x)\|^2$. For the singleton partition $\|C_{\mathcal{R}}u\|^2 = \|u\|_V^2$, so $K = c + 4DM$ proves its asserted finite-penalty alternative. The matrix construction establishes an existence bound with explicit constants, rather than a sharp spectral rank formula.*

# 3 Necessary product order for full-vector common partitions

We prove Theorem 1.2 by an explicit positive diagonal family. The proof concerns all vector fields, as in Theorem 1.1, rather than the smaller coexact domain of the endpoint lower theorem. It uses only local fields and finite-dimensional rank-nullity.

## 3.1 The coefficient family and every local channel

Let $n = 3m + 3$ and choose centers

$$x_j = (3j + 1, 1, \ldots, 1), \qquad 0 \leq j < m.$$

They lie away from the periodic wrap boundary. With $\sigma = \alpha_D^{1/P} c$, define

$$a_0 = \min\{c/(8D), \sigma/2\} > 0, \qquad M = rn^D, \qquad H = \sigma M > \sigma.$$

In the canonical direction/component basis $(i, k)$ of $\mathbb{R}^{Dh}$ let $A(x)$ be diagonal, with entries

$$\lambda_{i,k}(x) = \begin{cases} a_0, & k = 1 \text{ and } x \in \{x_j, x_j - e_i\} \text{ for some } j, \\ H, & \text{otherwise.} \end{cases} \tag{30}$$

At a center all $D$ direction channels for component one are weak; at an incoming start only its specified direction/component channel is weak. All eigenvalues are positive. For each fixed direction the starts $x_j$ and $x_j - e_i$ are distinct. Across centers their first coordinates lie in disjoint sets $\{3j, 3j + 1\}$. Channels with different direction labels remain distinct even when they share a start. There are therefore exactly $2Dm$ weak channels.

Since $a_0 \leq \sigma/2$, each weak channel has capped cost one in (4). Every other channel has capped cost $\alpha_D c^P H^{-P} = M^{-P}$. Thus, including all strong channels,

$$F_c(A) = 2Dm + (M - 2Dm)M^{-P}, \qquad 2Dm \leq F_c(A) \leq 2Dm + 1. \tag{31}$$

Indeed $0 \leq M - 2Dm \leq M$ and $P > 1$, so the tail is at most $M^{1-P} \leq 1$. Increasing $h$ increases $H$ and keeps this strong tail bounded; it introduces no uncontrolled $h$-sized budget contribution. The choice of $H$ is permitted by the arbitrary finite profile hypothesis.

### 3.2 The actual low space and the restriction rank

Let $U_j(x) = 1_{\{x_j\}}(x)e_1$, where $e_1 \in \mathbb{R}^h$ is the first canonical unit vector. These fields are orthonormal in counting units. Their only nonzero direction-labelled forward differences are $-e_1$ at $x_j$ and $+e_1$ at $x_j - e_i$, for each direction $i$. Both original starts are assigned coefficient $a_0$ by (30); there are no further wrap contributions. Consequently

$$Q_A(U_j) = 2Da_0 \le c/4.$$

No nearest-neighbor edge joins distinct centers, and their gradient-start supports are disjoint. Hence the quadratic cross terms vanish and

$$L_m = \operatorname{span}_{\mathbb{R}}\{U_0, \dots, U_{m-1}\}, \quad \dim L_m = m, \quad Q_A(U) \le (c/4)\sum_x \|U(x)\|^2 \quad (U \in L_m). \tag{32}$$

These are general vector functions. No divergence or coexact property is asserted for them.

For any common partition $\mathcal{R}$ into $R$ nonempty rectangles, its actual rows on this space are

$$q_B\left(\sum_j t_j U_j\right) = V^{-1/2}|B|^{-1/2}\left(\sum_{j:x_j\in B} t_j\right)e_1.$$

Each rectangle supplies at most one scalar functional on $L_m$, so $\operatorname{rank}(C_{\mathcal{R}}|_{L_m}) \le R$. If $R < m$, rank-nullity gives a nonzero $U \in L_m$ with $C_{\mathcal{R}}U = 0$. Coercivity with any finite quadratic penalty, after multiplication by $V$, would imply

$$c\sum_x \|U(x)\|^2 \le Q_A(U) \le (c/4)\sum_x \|U(x)\|^2,$$

a contradiction. Thus every successful common partition has $R \ge m$. The same argument works for any penalty on these linear observations with $\Phi(0) = 0$, without a differentiability assumption on that penalty.

The full field-space rank is exactly $hR$, not merely a padded output dimension. To see this, for each $B$ choose a site $z \in B$. A field supported at $z$ with value $\sqrt{V|B|}\,e_k$ has observation $e_k$ in row $B$ and zero in every other row. These fields show surjectivity onto all $hR$ row coordinates. Therefore necessary genuine rank is at least $hm$. Since $F_c(A) \le 2Dm + 1 \le (2D+1)m$, this proves (7).

### 3.3 Joint growth and the measurement-class restriction

Set $h = m$. The necessary rank is at least $m^2$, while $h + F_c(A) \le (2D+1)m + 1$. Their ratio tends to infinity for fixed $D, c$. Thus a dimension-only $C_D(h + F_c(A))$ bound fails for these common full-vector partitions, even with arbitrary finite penalty. The product lower bound and Theorem 1.1 have matching $hF_c(A)$ order along this family, since $F_c(A) \ge 2Dm$ makes $1 + F_c(A)$ comparable to $F_c(A)$. Sharp coefficients, additive constants and per-profile optimality are not determined.

The extra factor $h$ measures the cost of retaining every component mean on every common partition member; it does not count $h$ independent tests of the one-component low space. The argument does not apply to distinct componentwise partitions, projected or channel-adapted rows, or compressed postprocessed outputs. It requires no external theorem about the rank of those broader observation classes. These elementary rank and local-field arguments carry no claim of historical priority, a new general principle, formal verification, or physical transfer.

# 4 Sharp rank on the actual coexact domain

The general-field obstruction in Section 3 counts full row rank $hR$. On a coexact domain those rows have relations, so that count cannot simply be reused. We calculate their actual annihilator instead.

## 4.1 Closed fields and the annihilator formula

For one internal channel, $\mathcal{U}_0^\perp = \ker d_1$ by finite adjoint duality. We first prove the real decomposition

$$\ker d_1 = \{\text{constant vector fields}\} \oplus \operatorname{im} G, \qquad (G\phi)_j = D_j^+ \phi. \tag{33}$$

Complexify and use the finite characters $\exp(2\pi \mathrm{i} k \cdot x/n)$. At frequency $k$ put $q_j(k) = \exp(2\pi \mathrm{i} k_j/n) - 1$. The closed equations are $q_i \widehat{u}_j - q_j \widehat{u}_i = 0$. If $k \neq 0$, some $q_\ell \neq 0$, and the equations force $\widehat{u}_j = q_j \widehat{u}_\ell / q_\ell$ for all $j$, including those with $q_j = 0$. A coordinate-free choice of potential is

$$\widehat{\phi}(k) = \frac{\sum_j \overline{q_j(k)} \widehat{u}_j(k)}{\sum_j |q_j(k)|^2} \quad (k \neq 0), \qquad \widehat{\phi}(0) = 0.$$

For real $u$ this formula has conjugate symmetry, hence gives a real mean-zero $\phi$. Its gradient supplies every nonzero-frequency mode; the zero mode is a constant vector. Conversely constants and gradients are closed. A periodic gradient has zero component means, so its intersection with constant vectors is zero. The kernel of $G$ consists of constants because forward coordinate steps connect the torus. This proves (33), including $n = 1$. It also yields the useful real dimension identity

$$\dim \mathcal{U}_0 = DV - [D + (V - 1)] = (D - 1)(V - 1). \tag{34}$$

Real and complex dimensions here agree by complexification of a real map; there is no doubling from conjugate frequencies. Forward Nyquist symbols are $-2$, rather than the spurious zero of a central difference.

For a partition into $R$ nonempty sets, let $S$ be the $R$-dimensional scalar block-constant space and $W = S^D$. The adjoint range of its full mean map is $W$: all row scaling factors in (2) are nonzero. Define the real potential space

$$\mathcal{F} = \{\phi : \sum_x \phi(x) = 0,\ G\phi \in W\}.$$

Constants lie in $W$, and all gradient component means vanish. Therefore (33) gives the direct sum

$$W \cap \mathcal{U}_0^\perp = \{\text{constant vectors}\} \oplus G\mathcal{F}.$$

Since $G$ is injective on mean-zero functions, this intersection has dimension $D + \dim \mathcal{F}$. The adjoint of the restricted map is $P_{\mathcal{U}_0} C^*$. Finite rank-nullity consequently gives the exact formula

$$\operatorname{rank}(C|_{\mathcal{U}_0}) = DR - D - \dim \mathcal{F}. \tag{35}$$

This formula is for the actual restricted image, not its listed output size.

### 4.2 Forward maximum principle and sharpness

Let $\mathcal{L} = \sum_j D_j^+$. If a real $\phi$ satisfies $\mathcal{L}\phi = 0$, choose a site of global maximum. Every forward increment there is nonpositive and their sum is zero, so each is zero. Every forward neighbor is another maximum. Iterating these steps reaches every torus site, including inverse movements by periodic wrap. Thus $\phi$ is constant. At $n = 1$ this is immediate. This is the classical finite irreducible-chain maximum principle; see Levin–Peres, Lemma 1.16 [4]. Here $\mathcal{L} = D(P_{\text{walk}} - I)$ for the uniform forward-coordinate walk. Neither reversibility nor aperiodicity is needed. We claim no new maximum principle.

For $\phi \in \mathcal{F}$, every component of $G\phi$ lies in $S$. Thus $\mathcal{L}\phi \in S$, with global sum zero by telescoping. The maximum principle proves that

$$\mathcal{L} : \mathcal{F} \longrightarrow S \cap 1^{\perp}$$

is injective. The target has dimension $R - 1$, since the global-sum functional on $S$ is the nonzero weighted sum of block values with weights $|B|$. Hence $\dim \mathcal{F} \le R - 1$. Substitution into (35) proves $(D - 1)(R - 1)$ for one channel. The input and output maps for $g$ independent internal channels are literal direct sums, so ranks add and prove (14).

For sharpness set $n = R$ and $B_a = \{x : x_1 = a\}$. Every $\phi(x) = f(x_1)$ with $\sum_a f(a) = 0$ belongs to $\mathcal{F}$: its first forward derivative is constant on each fibre, and all others vanish. These potentials form an $(R - 1)$-dimensional space. The preceding upper bound gives equality in (35), and then equality in (14). At $R = 1$ the common global means vanish on the coexact domain and both sides are zero. Sharpness concerns the allowed family of tori with $n = R$, not every fixed pair $(n, R)$.

### 4.3 A complete fixed-size coexact coefficient family

For the coefficient consequence we retain $D \ge 3$ and $m \ge 4$. Let $L = 16$, $g = m$, $h = Dg$, $n = m(L + 3) + 3$ and $V = n^D$. Use the compact product field in (42), with internal unit vector $e_1 \in \mathbb{R}^g$, translated by

$$q_t = (1 + t(L + 3), 1, \dots, 1), \quad 0 \le t < m, \qquad B_t = q_t + \{0, \dots, L\}^D.$$

Lemma 6.1 and its preceding support proof apply to this exact field: the norm is nonzero, the unit-gradient Rayleigh bound is $K_D/L^2$, and every nonzero original forward start lies in $B_t$. The last first-coordinate endpoint is $m(L + 3) - 2 < n - 1$, and the other endpoints are $L + 1 < n - 1$. Successive boxes have two unused sites between them. Thus there are no wrap collisions and complete field and gradient-start supports are disjoint. Their span $E_m \subset \mathcal{U}$ is truly $m$-dimensional. Only spatial components $1, 2$ of internal channel 1 are active on $E_m$.

Put $r = Dh$, $M = rV$, $\sigma = \alpha_D^{1/P} c$ and define

$$\begin{aligned} a_0 &= \min\{\sigma/2, cL^2/(4K_D)\} > 0, \qquad H = \sigma M, \\ A_{(i,j,\alpha),(i',j',\alpha')}(x) &= \delta_{ii'}\delta_{jj'}\delta_{\alpha\alpha'}\lambda_{i,j,\alpha}(x), \\ \lambda_{i,j,\alpha}(x) &= \begin{cases} a_0, & x \in \bigcup_t B_t,\ j \in \{1, 2\},\ \alpha = 1, \\ H, & \text{otherwise.} \end{cases} \end{aligned} \tag{36}$$

The triple denotes gradient direction, spatial field component and internal channel. This defines an actual positive diagonal matrix on all $r = D^2 g$ gradient coordinates at every start. The

constant $K_D$ is the bump Rayleigh constant of Lemma 6.1, not the observation penalty. Every nonzero derivative of each bump has coefficient $a_0$, so on $E_m$

$$Q_A(U) \le a_0(K_D/L^2)\sum_x \|U(x)\|^2 \le (c/4)\sum_x \|U(x)\|^2. \tag{37}$$

All norm and energy cross terms vanish by disjoint supports.

The number of weak local direction/component/site channels is exactly $W_* = 2Dm(L+1)^D = \kappa_D m$, with $\kappa_D = 2D17^D$. This count includes zero-derivative entries assigned $a_0$; none are dropped. The boxes are disjoint and $r \ge 2D$, so $W_* \le M$. Each weak cap cost is one; every strong cost is $M^{-P}$. Hence the full budget is

$$F_c(A) = W_* + (M - W_*)M^{-P} \in [\kappa_D m, \kappa_D m + 1]. \tag{38}$$

The strong tail is bounded by $M^{1-P} \le 1$, using $P > 1$. No arithmetic-mean or envelope condition is asserted for this matrix family.

If a common partition with $R$ members restores coercivity on $\mathcal{U}$ with $\Phi(0) = 0$, its restriction to $E_m$ has rank at most $2R$. Indeed every row has only the two active real component coordinates there. If $2R < m$, a nonzero low-space kernel vector contradicts coercivity and (37), after multiplying by $V$. Thus $R \ge \lceil m/2 \rceil$. The sharp rank bound independently yields

$$\operatorname{rank}(C_{\mathcal{R}}|_{\mathcal{U}}) \ge g(D-1)(\lceil m/2 \rceil - 1) \ge g(D-1)m/4 = h(D-1)m/(4D).$$

Here $\lceil m/2 \rceil - 1 \ge m/4$ for $m \ge 4$. Using $F_c(A) \le (\kappa_D + 1)m$ proves (15) and completes Theorem 1.5. With $g = m$, necessary rank is at least $(D-1)m^2/4$ whereas $h + F_c(A) \le (D + \kappa_D + 1)m$. The all-field matrix upper bound restricts to $\mathcal{U}$, giving the matching worst-case product order up to dimensional constants. The universal rank theorem allows $D = 2$; this coefficient application retains $D \ge 3$ throughout.

# 5 Component-adapted rectangles and a structured escape

We prove Theorem 1.6 and compare maps on the same coefficient family and actual coexact domain, without changing either energy.

## 5.1 Block spectra and a common finite penalty

Under the component-block-diagonal hypothesis, a fixed permutation of gradient coordinates gives actual $D \times D$ blocks $A^k(x)$ and

$$Q_A(u) = \sum_{k=1}^{h} Q_{A^k}(u_k), \qquad F_c(A) = \sum_{k=1}^{h} F_c(A^k).$$

The spectral identity holds because the full multiset is the disjoint union of block spectra, including all nonpositive multiplicities and cap threshold equalities. No spatially varying rotation is performed. For each actual block apply Theorem 1.1 with $h = 1$. Here $h = 1$ means a single real scalar field with a $D$-coordinate gradient, not a single spatial direction. It supplies a partition $\mathcal{R}_k$ with

$$\#\mathcal{R}_k \le a_D + b_D F_c(A^k), \qquad V^{-1}Q_{A^k}(u_k) + K_k \sum_{B \in \mathcal{R}_k} |q_{k,B}(u)|^2 \ge cV^{-1}\sum_x |u_k(x)|^2,$$

where $K_k \geq c$ is finite. All these partitions are fixed from their blocks before $u$ is chosen. Taking $K = \max_k K_k$ and summing proves (17). There is no rescaling of $c$ or $V$ by $h$. The constants $a_D, b_D$ depend only on $D$. On the full field space each component's disjoint nonempty block rows are independent, and different components have independent coordinates. Thus actual full rank equals $N_{\text{out}}$; on a constrained domain it may decrease. The coefficient-block hypothesis is essential to this proof.

### 5.2 Root observations for the strong components

For $n \geq 2$ the unweighted forward torus gradient has first nonzero Laplacian eigenvalue

$$\lambda_1 = 4\sin^2(\pi/n) \geq 16/n^2. \tag{39}$$

Indeed Fourier symbols give eigenvalues $\sum_i 4\sin^2(\pi k_i/n)$; the least nonzero one has exactly one nonzero coordinate of smallest distance from zero. The chord bound $\sin t \geq 2t/\pi$ on $[0, \pi/2]$ gives the inequality. Conjugate pairing gives the same real Poincare bound. The constant component is handled separately, so no statement about a nonzero gap at $n = 1$ is needed. For the family (36),

$$\sigma/c = \alpha_D^{2/D} = (3/4)^2[(D/2)(D/(D-2))^{D-1}]^{2/D} \geq 9/16.$$

Consequently $H\lambda_1 \geq 9crn^{D-2} \geq c$. Writing $f = \text{mean}_X f + f_0$ gives

$$V^{-1}H\sum_{x,i} |D_i^+ f(x)|^2 + c|\,\text{mean}_X f|^2 \geq cV^{-1}\sum_x |f(x)|^2. \tag{40}$$

The observation here is exactly $q_X(f) = \text{mean}_X f$.

Only the two real components $(j, \alpha) = (1, 1), (2, 1)$ are weak. Apply the $h = 1$ matrix theorem to each, obtaining at most $2a_D + b_D(F_1 + F_2) \leq 2a_D + b_D F_c(A)$ rows in total. Choose only the root mean for each of the remaining $h - 2$ globally strong components. One maximum of $c$ and the two finite weak-component penalties yields an all-field inequality with at most $h - 2 + 2a_D + b_D F_c(A)$ outputs. Every real component of a coexact field has zero global sum, since it is a sum of periodic backward differences. Hence all these $h - 2$ strong root rows vanish identically on $\mathcal{U}$. Restricting the already proved inequality to that domain gives

$$\text{rank}(C_{\text{ad}}|_{\mathcal{U}}) \leq 2a_D + b_D F_c(A).$$

This does not split a coexact spatial vector into independently variable spatial components, nor assert independence of the remaining weak rows. It proves (18) and Theorem 1.6.

For this same family $hF_c(A) \geq D\kappa_D m^2$ and the adapted rank is at most $2a_D + b_D(\kappa_D m + 1)$. Their ratio tends to zero. In contrast, Theorem 1.5 forces common full-vector maps to have restricted rank at least $(D-1)m^2/4$. Thus no positive uniform $\gamma_D hF_c(A)$ lower bound extends to all component-adapted maps. Refining their different partitions to a common partition and adding every vector component changes the map and can add the very rank charged by the common-class lower bound. No optimal adapted rank, efficient algorithm or arbitrary coupled-matrix additive theorem is inferred.

## 6 The weak critical endpoint

The positive and negative results use the same scalar original-start energy. The positive estimate holds on every component field. The negative estimate is already forced on the precisely defined coexact subspace (11).

### 6.1 Finite layer cake and the upper estimate

Since $a \geq 1$, its actual mean obeys $\mu \geq 1$ and $0 < Y \leq \mu$. For each value $y$ in this finite distribution,

$$y^P = P \int_0^\mu \tau^{P-1} 1_{\{\tau<y\}}\, d\tau.$$

Summing over $X$ proves the exact identity

$$\mathbb{E}_X Y^P = P \int_0^\mu \tau^{P-1} \mathbb{P}_X(Y > \tau)\, d\tau \leq 1 + P M_w \log \mu. \tag{41}$$

The contribution on $(0,1)$ is at most one, and on $[1,\mu]$ we use $\mathbb{P}_X(Y > \tau) \leq M_w \tau^{-P}$. At $\mu = 1$ the second integral is zero; in fact the actual mean and $a \geq 1$ then force $a \equiv 1$. Atoms at strict cutoffs change no ordinary integral. The envelope $a \leq 2\mu$ is retained in the theorem but is not needed for this step.

Corollary A.2 in the appendix, with the actual moment $M = \mathbb{E}_X Y^P$, gives

$$h\#\mathcal{R} \leq 2Dh[1 + 2^P c^P B_D M] T(n,\mu),$$

and the normalized inequality (10) with penalty $c$. Substitution of (41) proves Theorem 1.3. Its geometry and gap are supplied by the complete scalar construction, rather than by a separate assumed certificate.

### 6.2 A compact product bump on the actual coexact domain

We now prove Theorem 1.4. Fix a unit vector $e \in \mathbb{R}^g$. Define a globally $C^1$ function

$$\beta(s) = \begin{cases} s^2(1-s)^2, & 0 \leq s \leq 1, \\ 0, & s \notin [0,1]. \end{cases}$$

We use $\beta$ for the bump to distinguish it from the critical exponent $P$. Both $\beta$ and $\beta'$ vanish at the endpoints. On $[0,1]$,

$$\beta'(s) = 2s(1-s)(1-2s), \qquad \beta''(s) = 2 - 12s + 12s^2.$$

Thus $0 \leq \beta \leq 1$, $|\beta'| \leq 2$, and $\beta'$ is globally 2-Lipschitz: its interior second derivative has absolute value at most two, and the zero extension has no derivative jump.

For an integer $L \geq 16$ divisible by 16, put

$$\psi_L(x) = \prod_{i=1}^D \beta(x_i/L), \qquad x \in \mathbb{Z}^D,$$

and define the one-cochain

$$U_1 = D_2^- \psi_L\, e, \qquad U_2 = -D_1^- \psi_L\, e, \qquad U_j = 0 \quad (j \geq 3). \tag{42}$$

Taking $B_{12} = \psi_L e$, $B_{21} = -\psi_L e$ and all other $B_{ij} = 0$ in (11) gives exactly (42), including its signs. Therefore $U = \delta B$ and

$$\sum_j D_j^- U_j = D_1^- D_2^- \psi_L e - D_2^- D_1^- \psi_L e = 0.$$

Every component has zero mean by telescoping. The same identities hold exactly on any torus large enough to contain this support away from wraps.

The one-dimensional lattice sequence $f(t) = \beta(t/L)$ is supported at $t = 1, \dots, L-1$, and $D^- f$ is supported at $t = 1, \dots, L$. It follows from the product formula that $U$ is supported in $\{0, \dots, L\}^D$ and vanishes on every lower coordinate face $x_i = 0$. If $D_i^+ U_j(x) \neq 0$, its original start belongs to this same box. Indeed the only potentially external incoming start is $x_i = -1$; its next endpoint is on the zero lower face, so both field values vanish. Upper outgoing starts $x_i = L$ can contribute and are included. In all other coordinates an external coordinate makes both terms zero. This verifies the boundary-start condition needed below with no interior-only energy convention.

**Lemma 6.1** (Finite bump constants)**.** *Let $\mathcal{E}_1(U) = \sum_{x,i,j} \|D_i^+ U_j(x)\|^2$ on $\mathbb{Z}^D$. Then*

$$\begin{aligned} \mathcal{E}_1(U) &\le U_D L^{D-4}, & U_D &= 32D\, 2^D, \\ \sum_x \|U(x)\|^2 &\ge \eta_D L^{D-2}, & \eta_D &= \frac{9}{2^{D+12} 256^{2(D-1)}}. \end{aligned} \tag{43}$$

*In particular, with $K_D = U_D/\eta_D$,*

$$\mathcal{E}_1(U) \le (K_D/L^2) \sum_x \|U(x)\|^2. \tag{44}$$

*Proof.* The Lipschitz bound on $\beta'$ bounds each same-coordinate second finite difference by $2/L^2$. For different coordinates the product of two first differences is bounded by $4/L^2$. All unused factors have magnitude at most one. At most $2D$ component/direction pairs are nonzero; each has at most $(L+1)^D \le 2^D L^D$ original starts by the support analysis. The common squared bound $16/L^4$ gives the first line of (43), including the upper outgoing starts.

For the norm, take $x_2 = L/8 + 1, \dots, L/4$, exactly $L/8$ sites, and $x_i \in [L/4, 3L/4]$ for every $i \neq 2$, at least $L/2$ sites each. On each backward step in coordinate two, the mean-value point belongs to $[1/8, 1/4]$, where $\beta' \ge 3/32$. On the other coordinates $\beta \ge 1/256$ on $[1/4, 3/4]$. Hence

$$\|U_1(x)\| \ge \frac{3}{32L} 256^{-(D-1)}$$

on all these sites. Squaring and multiplying by $(L/8)(L/2)^{D-1}$ gives the stated $\eta_D$. It is strictly positive, so division yields (44). □

### 6.3 Geometric scales, exact mean and the weak distribution

Choose an integer $k \ge 0$ such that $\varepsilon = 4^{-k}$ satisfies $\varepsilon K_D \le c/4$. This choice depends on $D, c$ and is fixed while $m$ varies. Put

$$L_0 = 2^{\max\{4,k\}}, \quad L_j = 2^j L_0 \ (0 \le j < m), \quad L_* = L_{m-1}, \quad n = 16 L_*, \quad H = \varepsilon L_*^2. \tag{45}$$

Every $L_j$ is divisible by 16, and $\varepsilon L_0^2 \ge 1$. Place boxes

$$B_j = q_j + \{0, \dots, L_j\}^D, \quad q_{0,1} = 1, \quad q_{j+1,1} = q_{j,1} + L_j + 3, \quad q_{j,i} = 1 \ (i \ge 2).$$

Their first-coordinate intervals have two unused sites between successive boxes. The rightmost endpoint is $\sum_j L_j + 3m - 2$. Since $\sum_j L_j < 2L_*$ and $m \le 2^{m-1} \le L_*/16$, this is less than

$3L_* + 1 < n$. Other coordinates also lie strictly inside the chart. Thus no support or gradient-start support collides periodically. Translate (42) into each $B_j$, obtaining $U^{(j)}$. Their field supports and their complete original gradient-start supports are disjoint.

Define the actual scalar coefficient by

$$a(x) = \begin{cases} \varepsilon L_j^2, & x \in B_j, \\ H, & x \notin \bigcup_j B_j. \end{cases} \tag{46}$$

The boxes are disjoint and the top box has the same coefficient as the background. In particular $1 \le a \le H$. Define its mean by its actual sum:

$$\mu = V^{-1} \sum_x a(x) = H - V^{-1} \sum_j |B_j|(H - \varepsilon L_j^2). \tag{47}$$

This is not a freely adjusted normalization. With

$$C_{\mathrm{vol}} = \frac{2^D}{1 - 2^{-D}},$$

the geometric series and $(L_j + 1)^D \le 2^D L_j^D$ imply

$$\sum_j |B_j| \le C_{\mathrm{vol}} L_*^D, \qquad V^{-1} \sum_j |B_j| \le C_{\mathrm{vol}}/16^D < 1/2.$$

It follows that

$$H/2 \le \mu \le H, \qquad \mu \ge 1, \qquad 1 \le a \le H \le 2\mu. \tag{48}$$

The separate inequality $\mu \ge 1$ follows directly from $a \ge 1$, including when $H/2 < 1$.

We check $W(a, \mu) \le 1$ with strict cutoffs. If $0 < \tau \le \mu/H$, then $\tau \le 1$ and $\tau^P \mathbb{P}_X(Y > \tau) \le 1$. If $\tau > \mu/H$, the background and top box are absent. The boxes contributing to $Y > \tau$ satisfy $\varepsilon L_j^2 < \mu/\tau$. If there is any such box, let $j$ be the largest qualifying index. They form an initial segment and have total volume at most $C_{\mathrm{vol}} L_j^D$. Consequently

$$\begin{aligned} \tau^P \mathbb{P}_X(Y > \tau) &\le C_{\mathrm{vol}} (\mu/\varepsilon)^P / V \\ &\le C_{\mathrm{vol}} L_*^D / (16 L_*)^D < 1. \end{aligned} \tag{49}$$

If no box qualifies the probability is zero. Equality at a coefficient atom excludes that box under the strict criterion and cannot increase the bound. This proves the weak distribution assertion for every $m$, including $m = 1$.

### 6.4 The low space and every finite local quadratic penalty

Every nonzero gradient of $U^{(j)}$ has its original start in $B_j$, where (46) has coefficient $\varepsilon L_j^2$. All wrap terms vanish for these nonwrapping supports. Thus the complete energy, not merely its interior contribution, is

$$Q_a(U^{(j)}) = \varepsilon L_j^2 \mathcal{E}_1(U^{(j)}) \le \varepsilon K_D \sum_x \|U^{(j)}(x)\|^2 \le (c/4) \sum_x \|U^{(j)}(x)\|^2. \tag{50}$$

The norm bound in Lemma 6.1 makes every $U^{(j)}$ nonzero. Disjoint field supports give norm orthogonality; disjoint gradient-start supports give vanishing quadratic-energy cross terms. The same inequality therefore holds on

$$E_m = \operatorname{span}_{\mathbb{R}}\{U^{(0)}, \dots, U^{(m-1)}\} \subset \mathcal{H}_{\mathrm{coex}}(X), \qquad \dim E_m = m.$$

This proves an actual low subspace, with no local eigenfunction or gap assumption. From (45) and (48),

$$T(n, \mu) \le \left(16\sqrt{2/\varepsilon}\right)^D. \tag{51}$$

The right side is independent of $m$.

For clarity the finite-dimensional obstruction is proved directly. If $DC(0)|_{E_m}$ has a nonzero kernel vector $U$, differentiability and $C(0) = 0$ imply $C(tU) = o(t)$. For sufficiently small $t$, the assumed upper bound on the penalty gives

$$\Phi(C(tU)) \le K\|C(tU)\|^2 = o(t^2).$$

Apply (13) to $tU$, use the low-space inequality, and divide by $t^2$. Letting $t \to 0$ would imply

$$c\sum_x \|U(x)\|^2 \le (c/4)\sum_x \|U(x)\|^2,$$

a contradiction. Therefore the derivative is injective on $E_m$ and $\dim F \ge \operatorname{rank}(DC(0)|_{E_m}) = m$. No uniformity of $K$ between profiles was used. Positivity or differentiability of $\Phi$ is unnecessary; only its stated finite quadratic upper bound is used. Common normalization of energy and field norm changes none of this finite-dimensional argument.

Finally, disjoint boxes, $|B_j| \ge L_j^D$ and $\mu \ge H/2$ give

$$\mathbb{E}_X(\mu/a)^P \ge V^{-1}\sum_{j=0}^{m-1} L_j^D \left(\frac{\mu}{\varepsilon L_j^2}\right)^P = \frac{m(\mu/\varepsilon)^P}{V} \ge \frac{m}{2^P 16^D}. \tag{52}$$

This establishes every part of (12) and explains why the strong-moment theorem is not contradicted.

### 6.5 The matching order and its quantifiers

For fixed $D, c$, $\varepsilon$ and $L_0$ remain fixed. Equations (45) and (48) imply

$$\log \mu = 2(m-1)\log 2 + \log(\varepsilon L_0^2) + \theta_m, \qquad -\log 2 \le \theta_m \le 0. \tag{53}$$

Thus $m/(1 + \log \mu) \to 1/(2 \log 2)$. The lower bound just proved forces output dimension at least a positive fixed-parameter multiple of $1 + \log \mu$ for large $m$. Theorem 1.3, with $h = Dg$, $M_w = 1$ and (51), gives an upper bound of the same order on these profiles. In particular, no bound proportional to $T(n, \mu)\, o(\log \mu)$ can hold uniformly over the weak-bound class for the centered differentiable/local-quadratic observations under discussion. This proves the final assertion of Theorem 1.4.

The conclusion is a worst-case comparison across profiles. It does not state that each profile needs a logarithm or identify its minimal rank. The lower construction also does not cover a nonzero observation baseline with an unrelated positive penalty at that baseline, nondifferentiable encodings, or penalties outside the stated local quadratic class. Those assumptions matter to the argument rather than merely to its presentation.

# 7 Model assessments and review context

Lluis Eriksson is the named author. Codex, exposed as an OpenAI system based on GPT-6, assisted drafting and consolidation. The exact served backend and reasoning-effort setting are unknown. The authoring role has read the retained predecessor review to address its editorial findings; this is not a fresh numerical review or a request to improve its score.

At the source-preparation freeze of 10 November 2026, no PDF of this integrated source had been produced or assessed. The following gate status records that dated freeze, not a statement about later document production. The separate source arguments incorporated here have qualified analytical source-audit reports, but at that freeze the combined exact source still required its own whole-source audit. Their AI roles share account/system, parent framing and retained context; independence is not established. No blind, human or independent-model review is claimed. In particular the changed coupled source 231B received its own exact-source audit after the earlier source 231A; the later acceptance is not inferred from the earlier one. Its initially pending integration gate was closed only after the actual complete 231B report arrived and was read. The matching upper construction 233 likewise has a separate exact-source report. These are source readings, not numerical assessments or formal certificates.

### The predecessor assessment belongs to its exact artifact

The predecessor is manuscript 226C, source commit

`2d3237501569f6afac8a44d524a5d6db59f63c39`.

Its exact 20-page, 419,361-byte PDF has raw-file SHA-256

`28cb2a7502e05cd8588aa6294f07a7eca`
`5e0e452cce1bf5e4f09e9ec3c89e1a6`.

That artifact's first numerical AI assessment assigned scientific importance 3.25/10, rendered as 3 native stars, with correctness/rigor/novelty/ significance/reproducibility criteria 4/4/2/3/3 on their five-star scales, and recommended minor revision. These numbers describe that old PDF only. They are not inherited by the new source or any future PDF. Its original complete scientific report has raw-byte SHA-256

`bd00e31013a1bdbf1f2766b25665af17ab`
`7792aae8fcdff3c50670393a1d7272`.

The original report, response JSON, context, report-hash record and dated author response are retained byte-for-byte in the accompanying input packet. The original response omitted a native version UUID; a later documented projection added only that binding and changed no scientific judgment. No new UUID, receipt, publication or assessment is created by this preparation.

The predecessor review requested a closer comparison with CDPX Section 6 and Brudnyi Appendix II, a clearer dependency guide, removal or definition of external programme shorthand, and improved reference-page economy. This source identifies those constructive predecessors explicitly, adds the proof dependency table, defines each observation class and coefficient family, and keeps programme bookkeeping outside mathematical inferences. The bibliography is consolidated with a compact standard layout. Its actual page economy and equation layout remain unobserved until rendering; no visual fix or new score is presumed. The additions are substantive sharp coexact rank, structured escape and coupled matching-order results, rather than an attempt to regrade an unchanged paper. Conservative constants, unestablished priority, absent practical algorithms and profile-dependent penalties remain limitations.

### Complete evidence remains distinct from the proofs

`SOURCE-INPUTS.json` lists exact captured mathematical sources, the complete listed source-audit reports and their context disclosures, the clarification to the dimension range of the coexact coefficient application, and the predecessor's original assessment/author response. Its hashes identify raw bytes with no normalization. The scalar, matrix, general-field and weak-endpoint proof files are retained from exact 226C captures. Prior long execution histories remain external evidence, not premises of any new theorem. No selected old diagnostic, normal/optimized agreement, or source verdict is represented as an all-size executable proof, formal certification, two-fresh-clone replay or execution of the new arguments.

The complete earlier 226B source rejection for an auxiliary missing nonnegative-depth condition and the later corrected 226C source report are preserved unchanged among the listed companions. The corrected optional singleton derivation here retains an explicitly nonnegative spectral depth. Their original verdicts keep their original referents. No separate original pre-224 scalar-audit report is claimed to have been identified. The new Green-trace constant $\mathfrak{G}_D$ is deliberately distinct from the inherited weighted-gap constant $G_D$; the source 233 audit requested that notation repair on consolidation.

This preparation performed editing, binary source capture, hashes and static text checks only. It produced no scientific execution, build, Lean result, rendered PDF, native submission or numerical grade. Full combined-source audit, document production and visual checks remain separate gates.

# A Coefficient-selected rectangular means for arbitrary scalar profiles

This appendix gives the complete scalar construction used in the main proofs. It uses the original positive-start scalar coefficient and retains the edges crossing rectangular observation pieces inside a ring. It does not identify a variable-weight curl/divergence form with that gradient form. The balanced-ring, BV and tree mechanisms are classical; historical originality of this finite specialization is unestablished.

## A.1 Statement and actual observations

Let integers $D \geq 3$, $n \geq 1$, and $c > 0$ be fixed, and put $X = \{0, \dots, n-1\}^D$. For any finite scalar profile $a(x) > 0$ define

$$P = D/2, \quad b = 2^D, \quad J = 2D, \qquad C_{\mathrm{ext}} = 1 + J\{2J(1+12J)+1\}, \qquad G_D = 2^{2D+10} D^2 C_{\mathrm{ext}}^2.$$

Write $S(A) = \sum_{x \in A} a(x)^{-P}$ and $t = (cG_D)^{-P}/b$. Every internal edge below keeps its original positive start, and all sums and norms initially use counting measure. Here $\mathcal{E}_a(f; \Omega)$ denotes the sum over every chart-internal edge with both endpoints in $\Omega$, weighted by its original positive-start value $a(x)$; $\mathcal{E}_1$ has unit weights. Thus the counting row $\widehat{q}_B$ used at the end of this appendix equals $V^{1/2} q_B$ in the normalized notation of (2).

**Theorem A.1.** *There is a partition $\mathcal{R}$ of $X$ into nonempty axis-aligned integer rectangles, determined by $a$ and fixed geometric conventions alone, such that*

$$\#\mathcal{R} \leq 2D\{1 + (2b+3)S(X)/t\},$$

*and, simultaneously for every real field $f$,*

$$\sum_{i,x:\ x,x+e_i\in X} a(x)|f(x+e_i)-f(x)|^2 + c\sum_{B\in\mathcal{R}}\frac{1}{|B|}\Big(\sum_{x\in B} f(x)\Big)^2 \ge c\sum_{x\in X}|f(x)|^2. \tag{54}$$

*The original periodic energy also satisfies this inequality, by adding its nonnegative wrap terms. For every integer $h \ge 1$ real components the same partition has output dimension $h\#\mathcal{R}$ and penalty $\Phi(y) = c\|y\|^2$.*

**Corollary A.2.** *For any reference value $\mu \ge 1$, if $n^{-D}\sum_x(\mu/a(x))^{D/2} \le M$, the actual normalized rectangular means have output dimension at most $C_{\mathrm{rect}}T(n,\mu)$, with $K = c$ and*

$$C_{\mathrm{rect}} = 2Dh\{1 + (2b+3)b(cG_D)^P 2^P M\}.$$

*The exact arithmetic mean and envelope $1 \le a \le 2\mu$ are unnecessary for this sufficiency. The inverse-moment hypothesis is still assumed; it is not deduced from an arithmetic mean.*

*Proof.* $S(X) \le Mn^D\mu^{-P} \le 2^P M\max\{1, n^D/(1+\mu)^P\}$, since $\mu \ge 1$. Insert $1/t = b(cG_D)^P$ into the theorem. Apply its scalar inequality to each of the $h$ real components. Dividing energy and squared norms by $V = n^D$ gives observation rows $V^{-1/2}|B|^{-1/2}\sum_B f = \sqrt{|B|/V}\ \mathrm{mean}_B f$. The all-field inequality restricts to divergence-zero fields, and linear means with the displayed quadratic penalty supply the original local data. □

## A.2 Balanced integer nodes and actual collars

Split every coordinate interval with at least two sites into its floor/ceiling halves; a singleton remains unchanged. Use coordinate order $1,\dots,D$, lower before upper, and lexicographic order on children. Each node has at most $b$ children, and all side lengths differ by at most one. At global depth $k$ the nonempty intervals have lengths $h_k$ or $h_k+1$, until they are singletons, where $h_k = \max\{1, \lfloor n/2^k\rfloor\}$.

Consider a proper descendant $I$ of an ancestor $Q$, at global depth $k$, and set $h_0 = h_k$. Every side of $I$ belongs to $[h_0, 2h_0]$. Every positive clearance from $I$ to the boundary of $Q$ is a union of depth-$k$ intervals, so is at least $h_0$. This uses the global $h_0$, not the smallest side of this particular $I$: at odd sizes that substitution can be false. If every side of $Q$ is at least two, a proper descendant is shorter in every direction, so each axis has at least one positive clearance. If the smallest side of $Q$ is one, every side is one or two; each proper descendant is a corner singleton. That small case is treated separately below.

For the large case put $\Omega = Q\setminus I$. Expand $I$ by exactly $h_0$ on every face with positive clearance, leaving a boundary-touching face unchanged, and call the resulting rectangle $B$. Its side lengths $\ell_i$ lie in $[2h_0, 4h_0]$. The collar $A = B\setminus I$ is covered by the actual positive-clearance face slabs: normal width $h_0$ and all other sides those of $B$. There are between $D$ and $J = 2D$ slabs, each contained in $\Omega$. For slabs $S_i, S_j$ with distinct normals, direct product counting gives

$$\frac{|S_i|}{|S_i\cap S_j|} = \frac{\ell_j}{h_0} \le 4, \qquad \frac{|S_i|}{|S_j|} = \frac{\ell_j}{\ell_i} \le 2. \tag{55}$$

Choose the first slab as root; attach every distinct-normal slab directly to it, and its opposite slab, if present, through the first slab with a different normal. This tree has depth at most two.

### A.3 An unweighted BV extension, with every hole-crossing edge

For a finite product rectangle $S$ of longest side at most $4h_0$, write $V_1(S) = \sum_S |f - \mathrm{mean}_S f|$ and let $\mathrm{TV}_S(f)$ be the sum of $|df|$ over its actual internal edges. Then

$$V_1(S) \le 2h_0 \mathrm{TV}_S(f). \tag{56}$$

Indeed on an interval with $m$ sites, averaging pair paths bounds the coefficient of its $r$th edge by $2r(m-r)/m \le m/2$. Commuting coordinate-average projections contract $L^1$; telescoping their product and applying this interval inequality fibrewise proves (56), without an additional factor $D$.

Let $m_i$ be the mean on $S_i$ and let $c_0$ be the root mean. Averaging on the actual intersection and using (55) gives $|S_i||m_i - m_j| \le 4\{V_1(S_i) + V_1(S_j)\}$. For a depth-two path $i, p, r$, the second edge costs the volume ratio $|S_i|/|S_p| \le 2$, so

$$|S_i||m_i - m_r| \le 4V_1(S_i) + 12V_1(S_p) + 8V_1(S_r) \le 12 \sum_j V_1(S_j).$$

Direct paths satisfy the same safe upper bound. Covering $A$ and counting each actual edge at most $J$ times yields, with $P_D = 2J(1 + 12J)$,

$$\sum_A |f - c_0| \le (1 + 12J) \sum_j V_1(S_j) \le P_D h_0 \mathrm{TV}_A(f). \tag{57}$$

At each positive-clearance hole face take the adjacent outside layer of width $h_0$, with tangential sides those of $I$. It is contained in $A$. On a normal fibre its nearest value satisfies

$$|f(\mathrm{first}) - c_0| \le h_0^{-1} \sum_{\mathrm{fibre}} |f - c_0| + \sum_{\mathrm{fibre}} |df|.$$

Sum over all actual faces and use (57). The sum of hole-crossing jumps to the constant $c_0$ is at most $J(P_D + 1)\mathrm{TV}_A(f)$. Fill $I$ by $c_0$ and leave $f$ unchanged on $\Omega$, obtaining $\widetilde{f}$ on $Q$. Internal hole differences vanish; the added variation consists precisely of these crossing faces. Therefore

$$\mathrm{TV}_Q(\widetilde{f}) \le C_{\mathrm{ext}} \mathrm{TV}_\Omega(f). \tag{58}$$

Boundary-touching faces create no exterior edge; at $h_0 = 1$ the fibre gradient sum is zero. No coefficient is extended, reflected, averaged or reassigned. In particular (58) makes no weighted quadratic trace-extension claim.

### A.4 Actual critical Sobolev and weighted gap on a ring

Here are the finite details needed to pass from BV to the original energy. For $T \subset \{0, \dots, m-1\}^D$ with $|T| \le m^D/2$, the occupied-fibre argument and the discrete projection inequality give

$$|\partial T| \ge \sum_i |\pi_i T| - D|T|/m \ge D|T|^{(D-1)/D} - D|T|/m \ge c_{\mathrm{iso}} |T|^{(D-1)/D}, \quad c_{\mathrm{iso}} = D(1 - 2^{-1/D}) \ge \tfrac{1}{2}.$$

Each occupied nonfull fibre supplies an edge, while full fibres number at most $|T|/m$. The projection inequality $|T|^{D-1} \le \prod_i |\pi_i T|$ is the finite Loomis–Whitney inequality; it follows from

iterated Hölder. A singleton cube has only zero variance. Median parts, finite coarea and Minkowski consequently imply

$$\inf_z \|u - z\|_{D/(D-1)} \le (2/c_{\mathrm{iso}})\mathrm{TV}(u)$$

on the cube. On a rectangle of aspect at most two, fold a cube of its longest side onto it: in each coordinate first visit the interval in order, repeat its final endpoint and reflect until the prescribed length. Every vertex and edge has at most two preimages per coordinate; the initial product block maps identically onto the rectangle. Restriction and the variation multiplicity at most $2^D$ give the safe constant $C_{\mathrm{full}} = 2^{D+1}/c_{\mathrm{iso}}$ for that rectangle, still unweighted.

Apply this bound and (58) to an indicator on $\Omega$. For every real $z$, $\|1_T - z\|_{s,\Omega} \ge \frac{1}{2}\min\{|T|, |\Omega\setminus T|\}^{1/s}$, where $s = D/(D-1)$. Thus $\Omega$ has relative edge isoperimetric constant

$$c_\Omega = c_{\mathrm{iso}}/(2^{D+2}C_{\mathrm{ext}}).$$

Coarea gives $\|g\|_s \le c_\Omega^{-1}\mathrm{TV}_\Omega(g)$ for nonnegative $g$ supported on at most half the vertices. With a median $m_f$ of $f$, take $w = (f - m_f)_+$ or $(m_f - f)_+$ and put

$$\alpha = 2(D-1)/D, \qquad q = 2D/(D+2), \qquad q' = 2D/(D-2).$$

Then $\alpha s = 2$ and $(\alpha - 1)q' = 2$. The power-difference inequality, Hölder and degree at most $2D$ yield

$$\mathrm{TV}(w^\alpha) \le 2\alpha D^{1/q'}\|dw\|_q \|w\|_2^{\alpha-1}.$$

In fact the $q'$th power of the endpoint factor sums to at most $2^{q'}D\sum_x w(x)^2$. Divide only when $w \ne 0$. The positive and negative truncated $q$-edge energies sum to at most the original one; since $2/q \ge 1$, their squared bounds combine to give

$$\sum_\Omega |f - \mathrm{mean}_\Omega f|^2 \le K_\Omega^2\Big(\sum_{\Omega\ \mathrm{edges}} |df|^q\Big)^{2/q}, \qquad K_\Omega = 2\alpha D^{1/q'}/c_\Omega. \tag{59}$$

Only now use weights, on the original retained positive-start edges:

$$\Big(\sum |df|^q\Big)^{2/q} \le \Big(\sum_{\Omega\ \mathrm{edges}} a(\mathrm{start})^{-D/2}\Big)^{2/D}\mathcal{E}_a(f;\Omega) \le \{DS(\Omega)\}^{2/D}\mathcal{E}_a(f;\Omega).$$

Here $q/(2-q) = D/2$, and at most $D$ retained edges start at one vertex. Since $\alpha \le 2$, $c_{\mathrm{iso}} \ge 1/2$ and $2/q' + 2/D = 1$, $K_\Omega^2 D^{2/D} \le 2^{2D+10}DC_{\mathrm{ext}}^2 \le G_D$. Full balanced rectangles obey the same safe bound without removing a hole.

For the small nodes with a singleton direction, $Q$, or $Q$ minus a corner, is a connected graph with at most $b$ vertices whenever nonempty. Pair paths of length at most $b-1$ give variance at most $b^2\mathcal{E}_1$. Each internal edge satisfies $a(\mathrm{start})^{-D/2} \le S(\Omega)$, so $\mathcal{E}_1 \le S(\Omega)^{2/D}\mathcal{E}_a$. As $G_D \ge b^2$, this proves in every case

$$\sum_\Omega |f - \mathrm{mean}_\Omega f|^2 \le G_D S(\Omega)^{2/D}\mathcal{E}_a(f;\Omega). \tag{60}$$

A singleton has zero variance. Empty groups are never output.

### A.5 Coefficient-only heavy-chain compression and complete count

Give each vertex mass $a(x)^{-P}$ and call a tree node heavy if its mass is strictly greater than $t$. If the root is light, output it. Otherwise let $L$ be the number of terminal heavy nodes and $B_h$ the number with at least two heavy children. Terminal heavy nodes are disjoint, so $L \le S(X)/t$. Finite tree counting gives $B_h \le L-1$. After contracting unary paths, there are at most $B_h+L-1$ skeleton edges and possibly a root-leading chain; hence at most $B_h + L \le 2L-1$ maximal unary chains. The bound is on chains, not their unrestricted total length.

At a unary node $Q \to Q'$ the increment $Q \setminus Q'$ is the union of at most $b-1$ light children, of mass at most $(b-1)t$. Traverse a maximal chain from ancestor toward descendant, closing a group when the cumulative mass first reaches $t$. A closed group is a balanced ancestor ring $Q_i \setminus Q_j$ of mass in $[t, bt)$; a residual is output only if nonempty, and has mass less than $t$. Across all chains the disjoint closed groups number at most $S(X)/t$, while residuals number at most $2L-1$. At branching or terminal heavy internal nodes output each light child; at a heavy singleton terminal output the singleton. Such cubes/singletons number at most $b(B_h + L) \le b(2L-1)$.

These groups partition $X$: a unary ring excludes the endpoint heavy subtree, which is then handled at the endpoint; branching processing recurses into every heavy child and outputs all remaining children. Thus their total number is at most

$$S(X)/t + (2L-1) + b(2L-1) \le 1 + (2b+3)S(X)/t.$$

Each nonsingleton group has mass at most $bt = (cG_D)^{-D/2}$, so (60) gives residual gap at least $c$. Heavy singleton mass may exceed this threshold, but its variance is zero and its mean captures its entire norm.

### A.6 Rectangular rows while retaining intra-ring interfaces

Partition $Q \setminus I$ by the first coordinate outside $I$: preceding coordinates belong to $I$, the current coordinate is below or above $I$, and later coordinates fill $Q$. This gives at most $2D$ nonempty integer rectangles. Cubes and singletons require one rectangle. With $\widehat{q}_A(f) = |A|^{-1/2} \sum_A f$, Cauchy gives $\widehat{q}_\Omega(f)^2 \le \sum_{B \subset \Omega} \widehat{q}_B(f)^2$. The actual regional gap therefore implies

$$\mathcal{E}_a(f;\Omega) + c \sum_{B \subset \Omega} \widehat{q}_B(f)^2 \ge c\|f\|_\Omega^2.$$

Crucially $\mathcal{E}_a(f;\Omega)$ retains all edges between rectangular pieces of this same ring; no separate gap is asserted for each piece. Distinct regions are disjoint, so summing discards only nonnegative edges between different regions. The full energy dominates their internal sum. The count of rectangles is at most $2D$ times the region count, proving Theorem A.1. The retained intra-ring interfaces are part of the regional energy throughout.

## B Coupled coefficients and fixed-basis channel separation

This appendix proves Theorem 1.7. The direct-sum observation hypothesis is imposed in a specified internal basis; the coefficient profile deliberately couples those channels in that basis. The conclusion is not invariant under a change of separability basis.

### B.1 The actual constant profile and its full budget

Fix $D \geq 3$, $n \geq 2$, $V = n^D$ and $g = V$. The domain is $\mathcal{U} = \mathcal{U}_0^g$ and its one-channel dimension is $d_0 = (D-1)(V-1) > 0$ by (34). There are $h = Dg$ real field components and $r = D^2 g$ gradient coordinates. Let $v = (1, \dots, 1)/\sqrt{g} \in \mathbb{R}^g$ and put

$$\begin{gathered} a = c/(16D), \quad H = cn^2, \quad B = HI_g - (H-a)vv^{\mathsf{T}}, \\ A_{(i,j,\alpha),(i',j',\beta)} = \delta_{ii'}\delta_{jj'}B_{\alpha\beta}. \end{gathered} \tag{61}$$

Here $i$ is the gradient direction, $j$ the spatial one-cochain component, and $\alpha$ the internal channel. This is a constant positive matrix with $0 < A \leq HI$; its actual energy is

$$Q_A(z) = \sum_{x,i,j,\alpha,\beta} B_{\alpha\beta} D_i^+ z_j^\alpha(x) D_i^+ z_j^\beta(x).$$

The weak polarization $v$ has eigenvalue $a$, and $v^\perp$ has eigenvalue $H$. There are $D^2$ copies of $B$ at each start, so the full local multiplicities are $D^2$ weak and $D^2(g-1)$ strong. None can be removed when the tested field domain is restricted.

Since $\sigma/c = \alpha_D^{2/D} \geq 9/16$ and $a/c = 1/(16D)$, every weak channel has cap cost one. A strong channel has cost $\min\{1, \alpha_D n^{-D}\}$. Thus the exact full budget is

$$F_c(A) = D^2 V[1 + (g-1)\min\{1, \alpha_D n^{-D}\}] \leq D^2(1+\alpha_D)V. \tag{62}$$

The minimum is retained in the identity; the inequality uses $g = V = n^D$. The strong tail is allowed to have order $V$, and is not claimed bounded independently of volume. Adding $h = DV$ gives $h + F_c(A) \leq L_D V$.

### B.2 A projection inside the actual coexact space

Write the channel-separable observation as $C = \bigoplus_{\alpha=1}^g C_\alpha$. Each $C_\alpha$ is an arbitrary finite real linear map on one channel, restricted to $\mathcal{U}_0$; within that channel it may mix all sites and spatial components. Let $R_\alpha$ be the counting orthogonal projection in $\mathcal{U}_0$ onto the representer space of the restricted rows. This is a projection within the actual domain, not an unrestricted spatial row-space projection. Therefore

$$R = \bigoplus_\alpha R_\alpha : \mathcal{U} \to \mathcal{U}, \quad \operatorname{rank} R = r_C := \operatorname{rank}(C|_{\mathcal{U}}), \quad \ker R = \ker(C|_{\mathcal{U}}).$$

The internal direct sum is legitimate. No decomposition of $\mathcal{U}_0$ by its individual spatial components is assumed.

Define the isometry $T : \mathcal{U}_0 \to \mathcal{U}$ by $(Tf)_\alpha = v_\alpha f$. Channel separation gives

$$T^*RT = g^{-1}\sum_\alpha R_\alpha, \qquad \operatorname{tr}(T^*RT) = r_C/g. \tag{63}$$

This real symmetric positive contraction has a real unit vector $f \in \mathcal{U}_0$ with Rayleigh quotient at most its trace divided by dimension. As $R$ is an orthogonal projection, there is consequently $f$ with

$$t := \|RTf\|^2 \leq r_C/(gd_0), \qquad 0 \leq t \leq 1.$$

Set $u = (I-R)Tf$. Then $u \in \mathcal{U}$, $Cu = 0$ and $\|u\|^2 = 1 - t$. The witness may depend on the proposed observation map. Projections need not preserve locality, gradients, spatial components or weak polarization; their preservation of $\mathcal{U}$ is enough.

### B.3 The kernel witness and the necessary rank

For any periodic field, shifts are isometries, whence

$$\sum_i \|D_i^+ z\|^2 \le 4D\|z\|^2.$$

This counts every original forward start and wrap, including period two, with no further internal-dimension factor. The constant weak polarization and the bound $A \le HI$ give, respectively,

$$Q_A(Tf) \le 4Da, \qquad Q_A(RTf) \le 4DHt.$$

Positivity implies $Q_A(x-y) \le 2Q_A(x) + 2Q_A(y)$, hence $Q_A(u) \le 8Da + 8DHt$. If

$$V^{-1}Q_A(z) + \Phi(Cz) \ge c\|z\|_V^2 \quad (z \in \mathcal{U}), \qquad \Phi(0) = 0,$$

apply this at the actual kernel vector $u$ and multiply by $V$:

$$c(1-t) \le 8Da + 8DHt.$$

Since $8Da = c/2$ and $H = cn^2$, this yields

$$t \ge \frac{c - 8Da}{c + 8DH} = \frac{1}{2(1 + 8Dn^2)}.$$

Combining with (63) proves (19). No division by $1-t$ occurred; $t = 1$ is harmless. The penalty may be nonlinear, and no differentiability, quadratic growth or uniform penalty bound was used. Observations remain linear.

For fixed $D \ge 3$ the lower-bound expression has order $n^{2D-2}$, whereas $h + F_c(A) \le L_D n^D$. Explicitly,

$$\frac{r_C}{h + F_c(A)} \ge \frac{(D-1)(V-1)}{2L_D(1 + 8Dn^2)} \longrightarrow \infty.$$

The divergence has order $n^{D-2}$, including $D = 3, 4$. This concerns growing $g = V$, a specified internal basis and finite unit-spacing tori. It supplies no fixed-$g$ or $D = 2$ impossibility and no physical mass-gap inference.

### B.4 The same profile admits a cheap mixed observation

Let $Wz = \sum_\alpha v_\alpha z_\alpha$. Since $WT = I$, this map surjects $\mathcal{U}$ onto $\mathcal{U}_0$. Observe its values at every site and spatial component with rows $V^{-1/2}(Wz)_j(x)$. The listed output has $DV$ coordinates but actual restricted rank is exactly $d_0$. The orthogonal splitting

$$z = TWz + z_{\text{strong}}$$

stays in $\mathcal{U}$, since constant internal combinations commute with the codifferential. They also commute with forward differences. The split is orthogonal in norm and in coefficient energy. Every real component of $z_{\text{strong}}$ has zero mean, so the actual torus gap (39) gives

$$Q_A(z_{\text{strong}}) \ge H(16/n^2)\|z_{\text{strong}}\|^2 = 16c\|z_{\text{strong}}\|^2.$$

The weak energy is nonnegative, and $\|C_{\text{weak}} z\|^2 = V^{-1}\|Wz\|^2 = V^{-1}\|TWz\|^2$. Therefore

$$V^{-1}Q_A(z) + c\|C_{\text{weak}} z\|^2 \ge c\|z\|_V^2.$$

Its actual rank $d_0 = O_D(V) = O_D(h + F_c(A))$ is additive in order on this family. A fixed global internal rotation exposing $v$ makes the cheap channel visible. Thus no all-linear or basis-free obstruction follows. The profile in its original basis is not component-block-diagonal, so Theorem 1.6 is not contradicted. This completes Theorem 1.7. Appendix C proves a matching order for this precise family, rather than an all-profile upper bound.

# C Matching channel-separable order for the coupled family

The obstruction in Appendix B is sharp in order for its precise family and row class. We construct a nonlocal, channel-separable map from real coexact Laplacian modes. This supplies no general-profile upper bound and no rectangularity claim for these new rows.

Let $r_{\min}$ be the minimum actual restricted rank among finite channel-separable linear maps on $\mathcal{U}$ for which some penalty vanishing at zero restores $c$-coercivity for the fixed profile (61). This class is nonempty: observe a complete basis of $\mathcal{U}$, whose squared observations with coefficient $c$ control its norm since the energy is positive. Possible ranks are integers between zero and $\dim \mathcal{U}$, so the minimum rank exists. Define a separate finite Green-trace constant

$$\mathfrak{G}_D = D\, 3^{D-1}/8. \tag{64}$$

It is distinct from the scalar weighted-gap constant $G_D$ in (3).

**Theorem C.1** (Matching order in the fixed-basis class)**.** *For $D \geq 3$, $n \geq 2$, $g = V = n^D$ and the exact profile* (61)*,*

$$\frac{D-1}{4(1+8D)} n^{2D-2} \leq r_{\min} \leq (D-1)(1+4\mathfrak{G}_D) n^{2D-2}. \tag{65}$$

*The upper map constructed below has a finite quadratic penalty with $K = c + (4DH)^2/c$. Thus $r_{\min} = \Theta_D(n^{2D-2})$ for this family and observation class. This is not an exact leading constant, efficient algorithm, all-profile optimum or uniform-penalty result.*

## C.1 Real modes and the exact energy decomposition

The scalar forward torus Laplacian $\mathcal{L} = \sum_i (D_i^+)^* D_i^+$ acts componentwise on spatial vectors. It is real symmetric and commutes with curl and its adjoint, hence preserves $\mathcal{U}_0$. Its eigenvalue at nonzero frequency $k$ is

$$\lambda(k) = 4 \sum_i \sin^2(\pi k_i/n) > 0.$$

There are $D-1$ coexact polarizations at every nonzero frequency and none at zero, by (33). Conjugate pairs give $2(D-1)$ real modes for those two frequencies, and self-conjugate frequencies give $D-1$ real modes. Equivalently the real symmetric restriction has a real counting-orthonormal eigenbasis $f_s$, $1 \leq s \leq d_0 = (D-1)(V-1)$, with this same full eigenvalue multiset. There is no further real/complex factor of two. The torus estimate (39) and the symbols give

$$16/n^2 \leq \lambda_s \leq 4D.$$

Expand $z_\alpha = \sum_s z_{\alpha,s} f_s$ and put $z_s = (z_{\alpha,s})_{\alpha=1}^g$. Orthonormality gives

$$\|z\|^2 = \sum_s \|z_s\|^2, \qquad Q_A(z) = \sum_s \lambda_s z_s^{\mathsf{T}} B z_s. \tag{66}$$

For the second identity, $\langle df_s, df_t\rangle = \langle f_s, \mathcal{L} f_t\rangle = \lambda_t \delta_{st}$; the constant internal matrix $B$ does not alter that calculation. It also holds within degenerate eigenspaces. No independence of the individual spatial components as coexact domains is asserted.

### C.2 Selected real coordinates and a strict kernel margin

For each real mode let

$$k_s = \min\left\{g, \left\lceil \max\left\{0, \frac{g(2c/\lambda_s - a)}{H - a}\right\}\right\rceil\right\}. \tag{67}$$

Observe $V^{-1/2} z_{\alpha,s}$ for $\alpha \le k_s$. Each row uses one internal channel and may mix sites and spatial components. The row representers $e_\alpha f_s$ are an orthonormal subset on $\mathcal{U}$, so the actual restricted rank is exactly

$$r_C = \sum_s k_s. \tag{68}$$

The rows are selected from the fixed profile/mode data before the field. They extend to all fields by the same inner products, but the inequality claimed here is on $\mathcal{U}$.

In a kernel field, the first $k_s$ internal coordinates of mode $s$ vanish. If $k_s = g$ nothing remains. Otherwise the remaining principal block of $B$ is

$$H I_{g-k_s} - \frac{H - a}{g}\, 11^{\mathsf{T}},$$

whose least eigenvalue is $a + (H - a)k_s/g$. The ceiling in (67) implies $\lambda_s[a + (H - a)k_s/g] \ge 2c$ whenever its positive argument applies. If the maximum argument is zero, $a\lambda_s \ge 2c$ already. If the outer truncation applies there is no unobserved part of that mode. Thus (66) proves the strict kernel bound

$$Q_A(x) \ge 2c\|x\|^2 \quad (x \in \ker(C|_{\mathcal{U}})). \tag{69}$$

These are deliberately chosen whole real Laplacian modes; no commutation of an arbitrary observation projection with the energy was assumed.

### C.3 A finite penalty on all of the actual domain

Let $R$ be the counting orthogonal projection onto the selected row space inside $\mathcal{U}$, and split $z = x + y$, $x = (I - R)z$, $y = Rz$. Then $x$ is in the kernel, $x \perp y$, and $\|Cz\|^2 = V^{-1}\|y\|^2$ exactly. The positive energy operator on $\mathcal{U}$ has norm at most $M_* = 4DH$, from $A \le HI$ and the full gradient norm bound $4D$. Dropping $Q_A(y) \ge 0$, applying (69) and bounding the cross term gives

$$\begin{aligned} Q_A(x + y) &\ge 2c\|x\|^2 - 2M_*\|x\|\|y\| \\ &\ge c\|x\|^2 - (M_*^2/c)\|y\|^2. \end{aligned}$$

The last step is Young's inequality. Hence the finite choice $K = c + M_*^2/c$ yields

$$V^{-1} Q_A(z) + K\|Cz\|^2 \ge c\|z\|_V^2.$$

The margin is $2c$, not merely $c$: one unit of $c$ was spent controlling the cross term. No uniform or optimal penalty is inferred.

### C.4 Ceiling count and the finite Green trace

Since $a/H = 1/(16Dn^2) < 1/2$, $H - a \geq H/2$. Whether the maximum argument is positive, zero or truncated,

$$k_s \leq 1 + \frac{4g}{n^2\lambda_s}, \qquad r_C \leq d_0 + \frac{4g}{n^2}\sum_s \lambda_s^{-1}.$$

The ceiling cost is $d_0$, rather than $gd_0$. Choose centered integer representatives of each frequency with $|k_i| \leq n/2$, taking only one Nyquist representative in each even coordinate. The sine chord estimate gives $\lambda(k) \geq 16|k|_2^2/n^2$. For $N = \lfloor n/2 \rfloor$, overcount nonzero frequencies by every integer point on max-norm shells $\ell = 1, \ldots, N$. Each shell has at most

$$(2\ell+1)^D - (2\ell-1)^D \leq 2D(2\ell+1)^{D-1} \leq 2D\,3^{D-1}\ell^{D-1}$$

points. The first inequality follows by integrating $Dt^{D-1}$ over the interval of length two. On that shell $|k|_2 \geq \ell$. Since $D \geq 3$, $\sum_{\ell=1}^{N} \ell^{D-3} \leq N^{D-2} \leq n^{D-2}$, including $D = 3$. Therefore the complete finite sum satisfies

$$\sum_{k\neq 0} \lambda(k)^{-1} \leq \frac{n^2}{16}\,2D3^{D-1}\sum_{\ell=1}^{N}\ell^{D-3} \leq \mathfrak{G}_D V. \tag{70}$$

Odd sizes and Nyquist overcounting only enlarge this upper estimate. Multiplying by the real coexact multiplicity gives $\sum_s \lambda_s^{-1} \leq (D-1)\mathfrak{G}_D V$. Thus

$$r_C \leq (D-1)(V-1) + 4(D-1)\mathfrak{G}_D gV/n^2 \leq (D-1)(1+4\mathfrak{G}_D)n^{2D-2}.$$

Here $g = V = n^D$ and $V \leq n^{2D-2}$. The estimate need not beat the trivial full dimension at small sizes; the actual truncation still ensures $r_C \leq gd_0$.

Finally Appendix B applies to every allowed map and any penalty vanishing at zero. With $V-1 \geq V/2$ and $1+8Dn^2 \leq (1+8D)n^2$, its necessary bound is at least $[(D-1)/(4(1+8D))]n^{2D-2}$. Combining proves (65). The mixed map of the preceding appendix still has actual rank $d_0 = O_D(V)$ on this same profile. Matching the separable order therefore strengthens the fixed-basis comparison without producing an all-linear, basis-free, fixed-$g$, all-profile or physical obstruction. The Fourier, principal-block, Young and finite-shell mechanisms used here are classical; no historical-priority or formal-execution claim is made.

## References


[1] R. L. Frank, *Cwikel's theorem and the CLR inequality*, arXiv:1206.3325v1 (2012), Theorems 2.1 and 3.1. `https://arxiv.org/abs/1206.3325v1`.

[2] Yu. Brudnyi, *Nonlinear Piecewise Polynomial Approximation and Multivariate BV spaces of a Wiener–L. Young Type. I*, arXiv:1511.03971v1 (2015). `https://arxiv.org/abs/1511.03971v1`.

[3] A. Cohen, R. DeVore, P. Petrushev and H. Xu, *Nonlinear approximation and the space $BV(\mathbb{R}^2)$*, American Journal of Mathematics **121** (1999), 587–628. Author PDF: `https://people.math.sc.edu/pencho/Publications/BV9707.pdf`.

[4] D. A. Levin and Y. Peres, with contributions by E. L. Wilmer, *Markov Chains and Mixing Times*, second edition, American Mathematical Society, 2017, Lemma 1.16. Author PDF: `https://pages.uoregon.edu/dlevin/MARKOV/markovmixing.pdf`.