On Green's Problem 55: a Gowers-type norm on $\mathbb{F}_p^n \times \mathbb{F}_p^n$

Nevil Wrigley

22 November 2026

hexagon:2611.00905v1

Abstract

Let $G = \mathbb{F}_p^n$, where p is an odd prime. Answering Problem 55 of B. Green's list [2], we prove an inverse theorem for the cube average on G^2 obtained by taking one multiplicative difference in each of the horizontal, vertical, and antidiagonal directions. If a one-bounded function f has cube average at least δ, then f correlates, by at least $p^{-(C/\delta)^C}$, with a product of three one-variable functions (of x, of y, and of $x+y$) and a bilinear phase. The constant C is absolute; in particular, the bound is independent of the dimension n.

In outline, the proof runs as follows. First, a positive kernel on each diagonal selects a space of dimension $O(\delta^{-1})$ which retains correlation with f, and on which translation along the diagonal acts unitarily. Second, we reduce this matrix-valued structure to scalar phases by constructing a translation-compatible Hermitian matrix field whose eigenvalues are well separated on most diagonals. The quantitative inputs here are a polynomial family of matrix tests that detects every nonscalar two-dimensional compression, and an argument with convex sums of eigenvalues showing that small spectral gaps occur only with polynomially small probability. Third, a Fourier localization argument and an elementary stability theorem for scalar two-cocycles express each scalar phase as a short combination of the required test functions. Throughout, the logarithm of every partition size remains polynomial in δ^{-1}; this is what yields a single-exponential bound. The paper is self-contained: all auxiliary results, including the construction of coordinates and the near-equality theorem, are proved here.

Contents

1 Introduction **3**

2 Rectangles, translation, and step approximation **5**

3 Spectral spaces along the diagonals **7**

4 Separating equivariant matrix fields **10**

5 Translation in a controlled eigenbasis **18**

6 Scalar phases on products of cosets **23**

7 Proof of the inverse theorem **31**

A Stability when the cube average is near one **34**

B Choosing coordinates at a prescribed accuracy **40**

C Two elementary matrix inequalities **44**

References **45**

1 Introduction

Let p be an odd prime and let $G = \mathbb{F}_p^n$. For a complex-valued function on an abelian group, write

$$\Delta_v f(z) = f(z)\overline{f(z+v)}.$$

All averages below use normalized counting measure. For $f : G^2 \to \mathbb{C}$, define

$$\beta(f) = \mathbb{E}_{x,y,h,k,t \in G} \Delta_{(h,0)} \Delta_{(0,k)} \Delta_{(t,-t)} f(x,y). \tag{1}$$

This quantity is real and nonnegative, as will follow from the rectangle identity (3) in Section 2; when $|f| \le 1$, it lies in $[0,1]$. As is customary for Gowers-type norms, we also write $D(f) = \beta(f)^{1/8}$.

For $u \in \mathbb{F}_p$, put $e_p(u) = \exp(2\pi i u/p)$. A *test* in the original coordinates is a function

$$a_1(x)a_2(y)a_3(x+y)e_p(B(x,y)),$$

where the $a_i : G \to \mathbb{C}$ have modulus at most one and $B : G \times G \to \mathbb{F}_p$ is bilinear. If moreover $|a_1| = |a_2| = |a_3| = 1$, we call the test a *test phase*. Let $\tau(f)$ be the supremum of the absolute correlation $|\mathbb{E}_{x,y} f(x,y)\overline{T(x,y)}|$ of f with tests T. Since the class of tests is closed under conjugation, it is immaterial whether the conjugate is placed on f or on the test.

Theorem 1.1. *There is an absolute constant $C \ge 1$ such that, for every odd prime p, every $n \ge 0$, every $0 < \delta \le 1$, and every function $f : G^2 \to \mathbb{C}$ with $|f| \le 1$,*

$$\beta(f) \ge \delta \quad \Longrightarrow \quad \tau(f) \ge p^{-(C/\delta)^C}.$$

The one-variable factors in the correlating test may be chosen to have modulus one, and the bilinear form may be chosen alternating.

This answers Problem 55 posed by Ben Green [2].

The form of the conclusion reflects the three directions. A phase depending only on x, only on y, or only on $x+y$ is annihilated (that is, sent to the constant function 1) by the vertical, horizontal, or antidiagonal difference, respectively, and a bilinear phase is annihilated after all three differences have been taken. Consequently every test phase has cube average one, so each type of factor in a test is invisible to the cube average.

The bilinear factor, in particular, cannot in general be absorbed into the three one-variable factors. To see this, fix a phase s on G^2 and a base point (X,Y), and consider the *normalized rectangle*

$$N_s(h,k) = \frac{s(X+h,Y+k)\, s(X,Y)}{s(X+h,Y)\, s(X,Y+k)}.$$

If $s(x,y) = a_1(x)a_2(y)a_3(x+y)$ with unimodular factors, then

$$N_s(h,k) = \frac{a_3(X+Y+h+k)\, a_3(X+Y)}{a_3(X+Y+h)\, a_3(X+Y+k)},$$

which is symmetric in h and k. If instead $s(x,y) = e_p(B(x,y))$, then $N_s(h,k) = e_p(B(h,k))$. On $\mathbb{F}_p^2 \times \mathbb{F}_p^2$, the alternating form $B(x,y) = x_1y_2 - x_2y_1$ therefore gives a bilinear phase that is not a product of phases of x, y, and $x+y$.

On the other hand, only the alternating part of a bilinear form is needed. Since p is odd, every bilinear form is the sum of a symmetric and an alternating form. If S is symmetric and $Q(x) = S(x,x)/2$, then

$$S(x,y) = Q(x+y) - Q(x) - Q(y),$$

so the symmetric phase $e_p(S(x,y))$ can be absorbed into the three one-variable factors. The assertion about their moduli follows by optimizing the factors one at a time: once the other factors are fixed, the correlation is linear in the remaining one-variable function, and the best one-bounded choice can be taken to be a phase.

1.1 Context

Directional uniformity norms arise naturally when Cauchy–Schwarz is applied to configurations whose increments lie in specified subspaces. Milićević formulates an inverse problem in this generality and discusses the horizontal, vertical, and antidiagonal directions in connection with squares [3]. The corresponding problem over $\mathbb{F}_5^n$, with the diagonal direction in place of the antidiagonal direction, appears as Problem 4.5 in Peluse's survey [4]; reflection in the second coordinate interchanges the two formulations. Exact solutions and qualitative stability for related partial difference equations also fit the cohomological framework developed by Austin [1]. The arguments below are finite and quantitative, and none of these works is needed as an input to the proof.

1.2 Outline of the proof

It is convenient to reflect the second coordinate (Section 2). The antidiagonal direction then becomes simultaneous translation $T_h(a,b) = (a+h, b+h)$, which preserves each diagonal $\{a-b=d\}$. The proof has four stages. Since the aim is a single-exponential bound, we indicate at each stage where the size of the bound is determined.

Stage 1: spectral spaces along the diagonals (Section 3). On each diagonal, a positive kernel built from f selects a space of dimension $O(\delta^{-1})$. Translation along the diagonal carries these spaces to one another unitarily, so that, in suitable coordinates, it acts by a unitary cocycle σ_h, and f retains a definite correlation with a fixed linear functional of the resulting vector field. We choose the coordinates at one prescribed accuracy, with the logarithm of their rectangular partition complexity bounded by a polynomial (Appendix B).

Stage 2: a Hermitian field with separated eigenvalues (Section 4). We seek a Hermitian matrix field A_* with $A_*(T_h z) = \sigma_h(z) A_*(z) \sigma_h(z)^*$ whose eigenvalues are separated by a polynomially small gap on most diagonals. There are three quantitative ingredients. A tensor argument shows that an explicit class of fields detects a nonscalar compression on any family of two-dimensional planes of positive measure (Lemma 4.1). A bounded covariance potential selects polynomially many such fields which, together, detect every plane on most diagonals (Lemma 4.2). Finally, convex sums of eigenvalues show that a random linear combination of the selected fields has a small gap only with polynomially small probability (Lemma 4.3).

Stage 3: scalarization (Section 5). The separated eigenvalues allow us to choose eigenvectors with controlled step approximations. Translation acts on each eigenvector by a scalar phase, so the correlating function becomes a sum of $O(\delta^{-1})$ terms, each a controlled coefficient times a scalar phase.

Stage 4: scalar phases (Section 6). A Fourier argument localizes each scalar phase to a product of cosets, where a near-extremal cube average yields a test phase by the stability theorem of Appendix A. Transport along the diagonals, followed by iteration on the remaining diagonals, gives a global approximation by a combination of tests. Its cost is at most p raised to a power polynomial in δ^{-1}. All partition sizes in the argument have polynomial logarithms, so they contribute only to this polynomial power, and the final bound is single-exponential.

Section 7 chooses all parameters and closes the argument. The order matters: every accuracy required of the initial coordinates is determined from δ alone, before those coordinates are

constructed. The appendices prove the three auxiliary ingredients used in the body: near-equality stability (Appendix A), controlled coordinates for the spectral spaces (Appendix B), and a matrix factorization inequality (Appendix C).

1.3 Notation and conventions

Throughout, p is an odd prime, $G = \mathbb{F}_p^n$, and $e_p(u) = \exp(2\pi \mathrm{i} u/p)$ for $u \in \mathbb{F}_p$. We write $\mathbb{T}$ for the unit circle; a *phase* is a function with values in $\mathbb{T}$, and a function is *one-bounded* if its values have modulus at most one. All averages $\mathbb{E}$ use normalized counting measure, and L^2 norms of functions on finite sets are normalized accordingly; a subscript such as $\|\cdot\|_{2,R}$ or $\mathbb{E}_R$ indicates normalized measure on the set R. Inner products are linear in the second argument.

Vector norms are Euclidean, and finite matrices carry the unnormalized Frobenius norm $\|M\|_{\mathrm{F}} = (\operatorname{Tr} M^* M)^{1/2}$, the operator norm $\|M\|_{\mathrm{op}}$, or the trace norm $\|M\|_1$, unless explicitly stated otherwise. For a matrix field $X_h(z)$ depending on $h \in G$ and $z \in G^2$, we write $\|X\|_{L^2_{h,z}(\mathrm{F})} = (\mathbb{E}_{h,z} \|X_h(z)\|_{\mathrm{F}}^2)^{1/2}$, and similarly for other norms and variables.

The letter C denotes a positive absolute constant which may change from one occurrence to the next. A quantity is *polynomial* in specified parameters $x_i \geq 0$ if it is at most $(C + \sum_i x_i)^C$. We use this convention only for a fixed number of compositions; in particular, it never hides an exponent depending on a partition size or on δ.

1.4 AI declaration

This paper was generated during several sessions with ChatGPT. The author sent it to Ben Green, who said that since September he had been writing a human-readable paper on a similar result, also obtained using AI, with James Leng and Sarah Peluse. They additionally use the result to obtain bounds for subsets of $\mathbb{F}_p^n \times \mathbb{F}_p^n$ not containing any square $(x, y), (x, y + d), (x + d, y), (x + d, y + d)$.

Ben Green said he was happy for my paper to be posted on Hexagon.

2 Rectangles, translation, and step approximation

2.1 Reflected coordinates and the rectangle identity

For the proof we make the change of variables $a = x$, $b = -y$, and retain the letter f for the reflected function. Put

$$z = (a, b), \qquad d = a - b, \qquad T_h z = (a + h, b + h).$$

The third direction is now simultaneous translation. A test becomes

$$u(a)v(b)w(a - b)e_p(B(a, b)). \tag{2}$$

Under this change of coordinates the antidiagonal increment $(t, -t)$ becomes (t, t), the vertical increment changes sign (which is immaterial, since it is uniformly distributed), and $x + y$ becomes $a - b$. Thus both β and τ are preserved, with the bilinear form changed accordingly. For a rectangle $w = (a_0, a_1, b_0, b_1)$, write

$$\Box g(w) = g(a_0, b_0)\overline{g(a_0, b_1)g(a_1, b_0)}g(a_1, b_1), \qquad \|g\|_{\Box}^4 = \mathbb{E}_w \Box g(w).$$

The latter expression is nonnegative: it is $\mathbb{E}_{b_0,b_1} |\mathbb{E}_a g(a, b_0)\overline{g(a, b_1)}|^2$. The definition makes sense equally on a product of two finite probability spaces, and in particular on a product of cosets.

For a subgroup $H \le G$ and a function F of several variables $x_1, \dots, x_k$, put

$$\mathcal{I}_H F(x_1, \dots, x_k) = \mathbb{E}_{h \in H} F(x_1 + h, \dots, x_k + h).$$

Applied to a function of a rectangle $w = (a_0, a_1, b_0, b_1)$, this translates all four coordinates, that is, it applies T_h to all four vertices (a_i, b_j). The operator $\mathcal{I}_H$ is an orthogonal projection in L^2: it is self-adjoint, and applying it twice has the same effect as applying it once. Directly from the definitions,

$$\begin{aligned} D(g)^8 := \beta(g) &= \mathbb{E}_h \left\| g\,\overline{T_h g} \right\|_\square^4 \\ &= \mathbb{E}_w \square g(w)\, \overline{\mathcal{I}_G \square g(w)} = \|\mathcal{I}_G \square g\|_2^2 . \end{aligned} \tag{3}$$

Here $(T_h g)(z) = g(T_h z)$. For the first equality, recall that in the reflected coordinates the third difference is $\Delta_{(h,h)}$, so $\beta(g)$ is the average over h of the cube average of $g\,\overline{T_h g}$ in the horizontal and vertical directions; writing $a_0 = a$, $a_1 = a + h'$, $b_0 = b$, $b_1 = b + k$ for the base point and the remaining increments h', k, that cube average is $\left\| g\,\overline{T_h g} \right\|_\square^4$. For the second, note that $\square(g\,\overline{T_h g})(w) = \square g(w)\,\overline{\square g(T_h w)}$, where T_h acts on all four vertices, and average over h. The last equality is the orthogonal projection identity $\langle \mathcal{I}_G F, F \rangle = \|\mathcal{I}_G F\|_2^2$. These formulas also hold on a product of cosets of H, with all increments restricted to H.

Lemma 2.1 (Rectangular Cauchy–Schwarz)**.** *For functions g_{ij} on $A_i \times B_j$, where the four sets have uniform probability measure,*

$$\left| \mathbb{E} \prod_{i,j=0}^{1} \mathcal{C}^{i+j} g_{ij}(a_i, b_j) \right| \le \prod_{i,j=0}^{1} \|g_{ij}\|_{\square(A_i, B_j)} .$$

Here $\mathcal{C}$ is complex conjugation. Also, if $|u|, |v| \le 1$, then

$$|\mathbb{E}_{a,b} g(a,b) u(a) v(b)| \le \|g\|_\square .$$

Proof. Average the first expression in each b-coordinate, and apply Cauchy–Schwarz in (a_0, a_1). Its square is at most the product, over $j = 0, 1$, of

$$R_j = \mathbb{E}_{a_0, a_1} \left| \mathbb{E}_{b \in B_j} g_{0j}(a_0, b) \overline{g_{1j}(a_1, b)} \right|^2 .$$

Expanding the square and applying Cauchy–Schwarz in the two b-variables gives

$$R_j \le \|g_{0j}\|_\square^2 \, \|g_{1j}\|_\square^2 .$$

This proves the first assertion. For the second, first apply Cauchy–Schwarz in a, expand, and then apply it in (b, b'). The resulting fourth-power bound is

$$|\mathbb{E} guv|^4 \le \mathbb{E}_{b,b'} \left| \mathbb{E}_a g(a,b) \overline{g(a,b')} \right|^2 = \|g\|_\square^4 .$$

□

2.2 Finite averages and rectangular partitions

A function on G^2 is *L-step* if it is constant on the products of two partitions of G, each with at most L classes. For a vector- or matrix-valued function, the same pair of partitions must determine all coordinates. Refining finitely many partitions multiplies their class counts; equivalently, it adds their logarithms. All vector norms are Euclidean and all finite matrix Frobenius norms are unnormalized, except where explicitly stated otherwise. Inner products are linear in the second argument.

Lemma 2.2 (Sampling a Hilbert-space average)**.** *Let X be a random vector in a Hilbert space, with $\mathbb{E}\|X\|^2 \le 1$. For every positive integer q, some sample $X_1, \dots, X_q$ satisfies*

$$\left\| q^{-1} \sum_{i=1}^{q} X_i - \mathbb{E}X \right\|^2 \le q^{-1}.$$

The same assertion holds for Hilbert-valued functions, using their L^2 space as the Hilbert space.

Proof. Take independent samples. Expanding the squared norm, the mixed terms have mean zero after subtracting $\mathbb{E}X$. Thus

$$\mathbb{E}\left\| q^{-1} \sum_{i} (X_i - \mathbb{E}X) \right\|^2 = q^{-1}\mathbb{E}\|X - \mathbb{E}X\|^2 \le q^{-1}.$$

At least one sample attains the average bound. □

Lemma 2.3 (Approximating product averages)**.** *Suppose that*

$$h(a,b) = \mathbb{E}_\omega u_\omega(a) v_\omega(b), \qquad |u_\omega|, |v_\omega| \le 1.$$

For every $0 < e \le 1/2$, there is a one-bounded L-step function h_e such that

$$\|h - h_e\|_2 \le e, \qquad \log L \le Ce^{-2}\log(C/e).$$

Proof. Use Lemma 2.2 with $q = \lceil 4e^{-2} \rceil$ to approximate the average by q products, with error at most $e/2$. Round each one-variable factor to a point of an $e/8$-net in the closed unit disk. Such a net has at most $(C/e)^2$ points, by a square grid followed by nearest-point projection to the disk. Changing the two factors costs at most $e/4$ in each product. Refine the q rounding partitions in each variable. Their class count is at most $(C/e)^{2q}$, and the average of the rounded products remains one-bounded. These observations give the assertion. □

We will repeatedly use two further elementary facts.

(i) Conditional expectation onto a partition does not increase the L^2 distance to any function measurable with respect to that partition, and it maps functions with values in a closed convex set (such as the unit disk, or the unit ball of a matrix norm) to functions with values in the same set.

(ii) (Telescoping.) If $x_1, \dots, x_k$ and $y_1, \dots, y_k$ have operator norm at most one, then

$$\|x_1 \cdots x_k - y_1 \cdots y_k\| \le \sum_{i=1}^{k} \|x_i - y_i\|,$$

and similarly for tensor products. We use this for the unitary and contraction factors below, with the differences measured in Euclidean or Frobenius norm.

3 Spectral spaces along the diagonals

The simultaneous translations $T_h(a,b) = (a+h, b+h)$ preserve each diagonal $a - b = d$. We first collect the values of f on a diagonal into one vector. A positive operator will select a bounded-dimensional part of that vector which retains correlation with f. The operator is covariant under translation, so its selected spaces are transported unitarily along each diagonal.

We use the inner product $\langle g, h\rangle = \mathbb{E}_i \overline{g(i)} h(i)$ on $L^2(G)$, linear in the second argument. An array $A(i,j)$ denotes an integral kernel:

$$(Ag)(i) = \mathbb{E}_j A(i,j) g(j).$$

Thus its ordinary operator trace is $\mathbb{E}_i A(i,i)$, rather than $\sum_i A(i,i)$. Coordinate vectors and finite coordinate matrices have their usual, unnormalized Euclidean and Frobenius norms.

Lemma 3.1 (Coordinates on the diagonals)**.** *Suppose $|f| \le 1$, $\beta(f) \ge \delta > 0$, and $r = \lceil 8/\delta \rceil$. Fix a desired accuracy $0 < \zeta \le 1/2$. There exist $\delta/8 \le \lambda \le \delta/4$, a column $u(z) \in \mathbb{C}^r$, a row $\ell(z) \in \mathbb{C}^{1\times r}$, and unitary matrices $\sigma_h(z)$ such that*

$$\|u(z)\| \le 1, \qquad \|\ell(z)\| \le \lambda^{-1/2} \le \lambda^{-1}, \qquad \operatorname{Re} \mathbb{E}_z \overline{f(z)}\, \ell(z) u(z) \ge 3\delta/4. \tag{4}$$

They satisfy the exact identities

$$u(T_h z) = \sigma_h(z) u(z), \qquad \sigma_{h+k}(z) = \sigma_h(T_k z)\sigma_k(z). \tag{5}$$

There is an orthogonal projection $P_{\mathrm{act}}(d)$ onto the first $m(d) \le r$ coordinates, with $d = a - b$, for which

$$u = P_{\mathrm{act}} u, \qquad \ell = \ell P_{\mathrm{act}}.$$

The matrices σ_h preserve this active space and are the identity on its orthogonal complement.

The row and cocycle have step approximants $\widetilde{\ell}$ and $\widetilde{\sigma}_h$, with $\left\|\widetilde{\ell}\right\| \le \lambda^{-1}$ and $\widetilde{\sigma}_h$ unitary, such that

$$\left\|\ell - \widetilde{\ell}\right\|_{L^2_z} \le \zeta, \qquad \|\sigma - \widetilde{\sigma}\|_{L^2_{h,z}(\mathrm{F})} \le \zeta, \qquad \log L \le (C + \delta^{-1} + \zeta^{-1})^C. \tag{6}$$

Here L bounds the number of classes in each vertex partition. The partitions for the cocycle may depend on h.

In words: $u(z)$ records, in an orthonormal basis of the selected space on the diagonal through z, the coordinates of the part of the vector of values of f along that diagonal lying in the selected space; $\ell(z)$ is evaluation at the base point in these coordinates; and $\sigma_h(z)$ is the change of coordinates under translation by h.

The parameter ζ is fixed before the coordinates are chosen. No assertion is made that the same coordinates work with the displayed bound at every finer accuracy. This single prescribed accuracy is enough for the proof; the later parameters will be chosen so that all uses of these approximations request errors at least ζ.

Proof. For $t \in G$, write

$$\rho_t(a,b) = f(a,b)\overline{f(a+t, b+t)}, \qquad w_t(a,b) = \mathbb{E}_{c,s} \rho_t(c,b)\rho_t(a,s)\overline{\rho_t(c,s)}.$$

The function w_t is an average of products of one-bounded functions of a and of b. Moreover,

$$\mathbb{E}_{a,b} \overline{\rho_t(a,b)}\, w_t(a,b) = \|\rho_t\|_{\square}^4.$$

This is the conjugate of the rectangular average, which is real.

Set $v_z(i) = f(T_i z)$, and define

$$A_z(i,j) = w_{j-i}(T_i z). \tag{7}$$

To see its positivity, introduce vectors in $L^2(G^2)$:

$$\xi_{z,i}(c,s) = \overline{f(c+i, b+i)}\, \overline{f(a+i, s+i)}\, f(c+i, s+i). \tag{8}$$

Changing c, s by i in the inner product gives

$$A_z(i,j) = \langle \xi_{z,i}, \xi_{z,j} \rangle .$$

Thus A_z is positive and $\operatorname{Tr} A_z = \mathbb{E}_i \|\xi_{z,i}\|_2^2 \le 1$. Its spectrum is contained in $[0,1]$. Writing $t = j - i$, and then replacing z by $T_i z$, gives

$$\mathbb{E}_z \langle v_z, A_z v_z \rangle = \mathbb{E}_{t,z} \overline{\rho_t(z)} w_t(z) = \beta(f) \ge \delta. \tag{9}$$

For the moment take any $\lambda \in [\delta/8, \delta/4]$. Let P_z be the spectral projection of A_z onto the eigenvalues strictly greater than λ, and put $H_z = \operatorname{Ran} P_z$, $m_z = \dim H_z$. The trace bound implies $m_z \le \lambda^{-1} \le r$. Since $A_z \le \lambda I + P_z$,

$$\mathbb{E}_z \|P_z v_z\|_2^2 \ge \delta - \lambda \ge 3\delta/4. \tag{10}$$

The evaluation vectors $c_j(z) = P_z(\,\cdot\,, j)$ obey

$$\langle c_i(z), c_j(z) \rangle = P_z(i,j), \qquad \mathbb{E}_j c_j(z) c_j(z)^* = I_{H_z}. \tag{11}$$

Indeed $\langle c_j, g \rangle = g(j)$ for $g \in H_z$. Also $P_z \le \lambda^{-1} A_z$, so

$$\|c_j(z)\|_2^2 = P_z(j,j) \le \lambda^{-1} A_z(j,j) \le \lambda^{-1}. \tag{12}$$

In particular evaluation at any point has norm at most $\lambda^{-1/2}$, hence at most λ^{-1}.

Let $S_h g(i) = g(i+h)$. Formula (7) gives

$$A_{T_h z} = S_h A_z S_h^*,$$

and therefore $P_{T_h z} = S_h P_z S_h^*$. The shift S_h maps H_z unitarily onto $H_{T_h z}$, and m_z depends only on $d = a - b$. Choose an orthonormal basis of H_z, and let $\mathcal{U}_z : \mathbb{C}^r \to L^2(G)$ send the first m_z standard basis vectors to it and the other vectors to zero. Then

$$\mathcal{U}_z \mathcal{U}_z^* = P_z, \qquad \mathcal{U}_z^* \mathcal{U}_z = P_{\mathrm{act}}(d).$$

Define

$$u(z) = \mathcal{U}_z^* v_z, \qquad \ell(z) = \big((\mathcal{U}_z e_j)(0)\big)_{j=1}^{r},$$

and

$$\sigma_h(z) = \mathcal{U}_{T_h z}^* S_h \mathcal{U}_z + I - P_{\mathrm{act}}(d). \tag{13}$$

This is unitary: its first summand is the unitary change of coordinates on the active space, and its second summand is the identity on the complement. Composition of shifts proves the cocycle identity. The equality $v_{T_h z} = S_h v_z$ proves the other identity in (5).

The norm bounds follow from orthogonal projection and (12). Finally, $\ell(z) u(z) = (P_z v_z)(0)$. Translation covariance and averaging over a new uniform shift i show

$$\mathbb{E}_z \overline{f(z)}\, \ell(z) u(z) = \mathbb{E}_z \mathbb{E}_i \overline{v_z(i)} (P_z v_z)(i) = \mathbb{E}_z \|P_z v_z\|_2^2 .$$

This proves (4) for any choice of bases.

It remains to make that choice with controlled rectangular complexity. Lemma B.1 below chooses λ in the specified interval and bases for which (6) holds. None of the exact identities or estimates just proved depends on that choice. □

Fix these coordinates from now on. Comparing them to the coordinates at the point $(d, 0)$ on each diagonal gives

$$\Phi(a, b) = \sigma_b(d, 0), \qquad u_0(d) = u(d, 0). \tag{14}$$

The cocycle law yields

$$\Phi(T_h z) = \sigma_h(z)\Phi(z), \qquad u(z) = \Phi(z)u_0(d). \tag{15}$$

Both Φ and σ_h preserve the active space and fix its orthogonal complement.

The spectral construction gives more than a correlation with f. Every equivariant vector in the active space has a nontrivial directional cube average.

Lemma 3.2 (Nondegeneracy of the active space)**.** *Let $F(z) \in \operatorname{Ran} P_{\mathrm{act}}(d)$, with $\|F(z)\| \le 1$, satisfy $F(T_h z) = \sigma_h(z)F(z)$. Then $g(z) = \lambda\ell(z)F(z)$ is one-bounded and*

$$D(g)^2 \ge \lambda^3 \|F\|_{L^2_z}^2 . \tag{16}$$

Proof. On the diagonal through z, its vector of values is

$$v_z^g(h) = g(T_h z) = \lambda(\mathcal{U}_{T_h z}\sigma_h(z)F(z))(0) = \lambda(\mathcal{U}_z F(z))(h).$$

Thus $v_z^g \in H_z$ and $\|v_z^g\|_2 = \lambda \|F(z)\|$. Since $A_z \ge \lambda I$ on H_z, the same expansion as in (9) gives

$$\lambda^3 \|F\|_2^2 \le \mathbb{E}_z \langle v_z^g, A_z v_z^g\rangle = \mathbb{E}_{h,z} w_h(z)\overline{g(z)}g(T_h z) \le \mathbb{E}_h \left\| g\,\overline{T_h g}\right\|_{\square} .$$

For the last inequality, use that w_h is an average of bounded products of a function of a and a function of b, and apply rectangular Cauchy–Schwarz. Taking the fourth-power mean in h and using (3) bounds the last expression by $D(g)^2$. □

4 Separating equivariant matrix fields

Fix the coordinates supplied by Lemma 3.1, at an accuracy ζ which will be chosen in Section 7. Write $P(d) = P_{\mathrm{act}}(d)$. We use the normalized row $\lambda\ell$, which has norm at most one and a contraction step approximation of error at most ζ. The cocycle has a unitary step approximation with joint L^2 error at most ζ. Let $L \ge 2$ be a common class bound, so

$$\log L \le (C + \delta^{-1} + \zeta^{-1})^C. \tag{17}$$

The unitary field Φ from (14) satisfies

$$\Phi(T_h z) = \sigma_h(z)\Phi(z), \qquad u(z) = \Phi(z)u_0(d). \tag{18}$$

We will repeatedly use the following consequence of Lemma 3.2: for any active column of norm at most one satisfying $F(T_h z) = \sigma_h(z)F(z)$,

$$D(\lambda\ell F)^2 \ge \lambda^3 \|F\|_2^2 . \tag{19}$$

The purpose of this section is to construct a Hermitian field compatible with the same cocycle and with a polynomial lower bound for its spectral gaps on most diagonals.

The natural source of compatible fields is averaging along the diagonals: for any matrix field X, the field ΠX defined in Section 4.2 is automatically compatible with the cocycle, and on the diagonal d its eigenvalues are those of a single matrix $V_X(d)$. The difficulty is to choose X so that $V_X(d)$ has well separated eigenvalues. A repeated eigenvalue means that the compression of $V_X(d)$ to some two-dimensional plane is scalar, so we need fields which,

between them, see every plane. We first build an invariant tensor from f (Section 4.1), and use it to show that an explicit class of fields detects a nonscalar compression on any family of planes of positive measure (Lemma 4.1). We then select polynomially many fields that work for every plane on most diagonals (Lemma 4.2), and show that a random combination of them has polynomially separated eigenvalues (Lemma 4.3 and Corollary 4.4). Finally we record the complexity of the resulting field.

4.1 An invariant four-index tensor

For a rectangle $w = (a_0, a_1, b_0, b_1)$, write $z_{ij} = (a_i, b_j)$, $d_{ij} = a_i - b_j$. The corners $01, 10$ are conjugated in a rectangular product. Accordingly, let

$$G_w = \Phi(z_{00}) \otimes \overline{\Phi(z_{01})} \otimes \overline{\Phi(z_{10})} \otimes \Phi(z_{11}),$$

and let C_w be the analogous tensor product of the four rows $\lambda\ell(z_{ij})$, with the same conjugations. Let I_4 record the orbit under simultaneous translation of the four vertices. Define

$$\psi(w) = \mathbb{E}\big((C_w G_w)^* \mid I_4\big), \qquad F(w) = G_w \psi(w). \tag{20}$$

These are columns in $(\mathbb{C}^r)^{\otimes 4}$, of norm at most one. The tensor ψ is invariant under simultaneous translation. Both ψ and F are independent of any plane field that will be chosen later.

The covariance law gives the useful formula

$$F(w) = \mathbb{E}_h\big(\sigma_h(z_{00}) \otimes \overline{\sigma_h(z_{01})} \otimes \overline{\sigma_h(z_{10})} \otimes \sigma_h(z_{11})\big)^* C^*_{T_h w}. \tag{21}$$

In particular, if $\varepsilon > 0$ and $\zeta \le \varepsilon/32$, there is a partition $\mathcal{P}$ of G, with at most K classes, for which

$$F_0 = \mathbb{E}(F \mid \mathcal{P}(a_0), \mathcal{P}(a_1), \mathcal{P}(b_0), \mathcal{P}(b_1)), \quad \|F - F_0\|_2 \le \varepsilon, \quad \|F_0\| \le 1, \tag{22}$$

and

$$\log K \le C\varepsilon^{-2} \log(2L). \tag{23}$$

To prove this, replace the eight row and unitary factors in (21) by their prescribed step approximants. Telescoping and averaging in h cost at most $8\zeta \le \varepsilon/4$ in L^2. Lemma 2.2, applied in the Hilbert space of tensor-valued functions of the rectangle, supplies a sample of $O(\varepsilon^{-2})$ shifts with error at most $\varepsilon/4$. Refine all vertex partitions of the sampled factors, including their translates, into one partition $\mathcal{P}$ of G. The logarithm of its size is $O(\varepsilon^{-2}\log(2L))$. Conditional expectation of F onto these four labels improves the approximation and preserves its unit ball. This proves (22)–(23).

4.2 Detecting a nonscalar compression

For any complex matrix field X, put

$$(\Pi X)(z) = \mathbb{E}_h \sigma_h(z)^* X(T_h z)\sigma_h(z), \qquad V_X(d) = \Phi(z)^*(\Pi X)(z)\Phi(z).$$

Equivalently, $V_X(d)$ is the average of $\Phi(a,b)^* X(a,b)\Phi(a,b)$ over $a - b = d$. This follows by substituting (18) into the average. In particular V_X depends only on d, and

$$(\Pi X)(T_h z) = \sigma_h(z)(\Pi X)(z)\sigma_h(z)^*.$$

For a rank-two orthogonal projection R, define

$$D_R M = RMR - \tfrac{1}{2}\operatorname{Tr}(RMR)R.$$

This is the orthogonal projection onto the complex matrices supported on R and traceless there. On Hermitian matrices it is an orthogonal projection in the real Frobenius Hilbert space.

Fix $K \geq 2$; in the application, K is the class bound from (23). Let $\mathcal{X}_K$ contain the Hermitian real and imaginary parts of all fields

$$H(a,b) = h(a,b)v(a)w(b)E_{kl}, \tag{24}$$

where h is one-bounded and K^3-step, $|v|, |w| \leq 1$, and E_{kl} is a matrix unit. Here the Hermitian real and imaginary parts mean $(H + H^*)/2$ and $(H - H^*)/(2\mathrm{i})$. Each $X \in \mathcal{X}_K$ has $\|X(z)\|_{\mathrm{F}} \leq 1$. The constants E_{ll} are included. Every fixed member admits Hermitian contraction step approximants at all precisions e, with

$$L_X(e) \leq K^3(C/e)^4. \tag{25}$$

Indeed, retain the table h and round the two one-variable factors in the unit disk. The displayed count is deliberately generous. The definition of the class involves no plane field.

Lemma 4.1 (Polynomial detection)**.** *Fix $0 < q_0 < 1$, choose $0 < \kappa \leq (\lambda^3 q_0/(2r))^4$, and put*

$$\varepsilon = \frac{\kappa^2}{2^{20}(r+1)^3}, \qquad \nu = \frac{\kappa^2}{2^{20}(r+1)^{15}}.$$

Suppose $\zeta \leq \varepsilon/32$. Construct $\mathcal{P}, K, F_0$ from (22) at this ε, before choosing Q, R. If $Q \subseteq G$, $\mathbb{P}(Q) \geq q_0$, and $R(d)$ is a rank-two projection for every $d \in Q$, then some $X \in \mathcal{X}_K$ satisfies

$$\mathbb{E}_d \mathbf{1}_Q(d) \left\| D_{R(d)} V_X(d) \right\|_{\mathrm{F}} > \nu. \tag{26}$$

Proof. Suppose every test has the left side of (26) at most ν. Write $q = \mathbb{P}(Q)$, and extend R arbitrarily off Q.

Active mass and a nonzero tensor. If coordinate l is inactive on a particular diagonal, then $\Phi e_l = e_l$ there and $V_{E_{ll}} = E_{ll}$. For any plane R, the rank-one matrix $RE_{ll}R$ has trace $\|Re_l\|^2$, and its traceless part has Frobenius norm $\|Re_l\|^2/\sqrt{2}$. Since the constants E_{ll} belong to the test class, summing their bounds over the coordinates gives

$$r\nu \geq \sum_{l=1}^{r} \mathbb{E}\mathbf{1}_Q \mathbf{1}_{\{l > m(d)\}} \|D_R V_{E_{ll}}\|_{\mathrm{F}} = \frac{1}{\sqrt{2}} \mathbb{E}\mathbf{1}_Q \operatorname{Tr}(R(I-P)).$$

Since $\nu < 3q/(2\sqrt{2}r)$, it follows that

$$\mathbb{E}\mathbf{1}_Q \operatorname{Tr}(RP) \geq q/2.$$

Choose a standard coordinate vector v such that $\mathbb{E}\mathbf{1}_Q \|PRv\|^2 \geq q/(2r)$, and set

$$g(z) = \mathbf{1}_Q(d)\lambda\ell(z)\Phi(z)R(d)v.$$

The corresponding active column is $\mathbf{1}_Q \Phi P R v$. Equation (19) therefore gives

$$D(g)^8 \geq \left(\frac{\lambda^3 q}{2r}\right)^4 \geq \kappa.$$

The rectangular product of g is $C_w G_w\, u(w)$, where u is the tensor product of the four columns $\mathbf{1}_Q(d_{ij})R(d_{ij})v$, with the usual conjugations. It is orbit-invariant and has norm at most one.

It vanishes unless $d_{00} \in Q$, and its first index lies in $R(d_{00})$. Let R_0 act on that index, set $\xi = R_0\psi$, and write $Q_0 = \{w : d_{00} \in Q\}$. The rectangle identity (3) yields

$$\kappa \le \mathbb{E}|\psi^* u|^2 \le \mathbb{E}\mathbf{1}_{Q_0} \|\xi\|^2 . \tag{27}$$

A direct reduced-tensor pairing. For $i = 01$ or 10, arrange a four-index tensor v as an r^3-by-r matrix $Z_i(v)$, with index i as the column index, and put

$$\Gamma_i(v) = Z_i(v)Z_i(v)^*, \qquad \gamma_i = \Gamma_i(\psi), \qquad M_i = \Gamma_i(F_0).$$

If U_w is the tensor action of the three retained corner matrices, then

$$\Gamma_i(F) = U_w \gamma_i U_w^*.$$

Let P_R apply $D_{R(d_{00})}$ to every 00-block of a three-index matrix. Orthogonality of this projection gives the exact identity

$$\|P_R\gamma_i\|_{\mathrm{F}}^2 = \langle \Gamma_i(F), U_w(P_R\gamma_i)U_w^* \rangle .$$

The elementary estimates

$$\|\Gamma_i(v)\|_{\mathrm{F}} \le \|v\|^2 , \qquad \|\Gamma_i(v) - \Gamma_i(w)\|_{\mathrm{F}} \le (\|v\| + \|w\|)\, \|v - w\|$$

follow by expanding $ZZ^* - Z'Z'^* = (Z - Z')Z^* + Z'(Z - Z')^*$. Thus Γ_i is 2-Lipschitz on the unit ball. Also $\|P_R\gamma_i\|_{\mathrm{F}} \le 1$. Hence

$$\mathbb{E}\mathbf{1}_{Q_0} \|P_R\gamma_i\|_{\mathrm{F}}^2 \le 2\varepsilon + |\mathbb{E}\mathbf{1}_{Q_0} \langle M_i, U_w(P_R\gamma_i)U_w^* \rangle| . \tag{28}$$

The pairing retains the whole partition table in M_i. Its bound will depend on the number of matrix indices, but not on the number of entries in that table.

Expanding matrix indices, without expanding labels. For $i = 01$, parameterize the rectangle by

$$(a_0, a_1, b_0, b_1) = (x, y + s, y, y + t).$$

For fixed s, t, the other two retained corner matrices depend only on y. Write their prescribed conjugates as $A(y), C(y)$, so $U_w = \Phi(x, y) \otimes A(y) \otimes C(y)$. Let $\rho, \sigma, \alpha, \beta \in [r]^2$, and let

$$B_{\alpha\beta}(d) = D_{R(d)}\big((\gamma_i)_{(a,\alpha),(c,\beta)}\big)_{a,c}.$$

Because γ_i is orbit-invariant, this block depends only on $d = x - y$, once s, t are fixed. Moreover

$$D_R B_{\alpha\beta} = B_{\alpha\beta}, \qquad \|B_{\alpha\beta}\|_{\mathrm{F}} \le 1.$$

The pairing in (28) is the sum of at most $r^{10} \le r^{12}$ terms, indexed by $u, v, \rho, \sigma, \alpha, \beta$, whose integrands are

$$\begin{aligned} &\overline{(M_i)_{(u,\rho),(v,\sigma)}} A_{\rho_1,\alpha_1} C_{\rho_2,\alpha_2} \overline{A_{\sigma_1,\beta_1}}\, \overline{C_{\sigma_2,\beta_2}} \\ &\qquad \times \mathrm{Tr}(B_{\alpha\beta}(d)\Phi(x, y)^* E_{vu} \Phi(x, y)) . \end{aligned} \tag{29}$$

The first factor is a bounded table of $\mathcal{P}(x), \mathcal{P}(y), \mathcal{P}(y + s), \mathcal{P}(y + t)$. As a function of x, y, it has at most K classes in x and K^3 in y. The four other scalar entries combine into one bounded function of y. Thus their product with E_{vu} is precisely a field H of the form (24).

Average first over $x - y = d$ in each term. Since $B = RBR$ and $\mathrm{Tr}\, B = 0$, even for a non-Hermitian block,

$$\mathrm{Tr}(BV_H) = \mathrm{Tr}(BD_R V_H).$$

The Hermitian real and imaginary parts of H belong to $\mathcal{X}_K$. The contradictory assumption and Frobenius Cauchy–Schwarz therefore bound the integral of each term by 2ν. For $i = 10$, use $(a_0, a_1, b_0, b_1) = (x, x + s, y, x + t)$; the extra factors depend only on x, and the same conclusion holds. We have proved

$$\mathbb{E}\mathbf{1}_{Q_0} \|P_R\gamma_i\|_{\mathrm{F}}^2 \le 2\varepsilon + 2r^{12}\nu. \tag{30}$$

All test fields used in this calculation are independent of Q, R; only the paired blocks $B_{\alpha\beta}$ carry that dependence.

The two incompatible plane actions. Our parameter choices imply

$$\mathbb{E}\mathbf{1}_{Q_0} \|P_R\gamma_i\|_{\mathrm{F}} \le \frac{\kappa}{2^9(r+1)^{3/2}}. \tag{31}$$

Put $\rho_i = \Gamma_i(\xi)$. Its 00-blocks are the compressions RMR of those of γ_i. The nonscalar parts of these blocks are precisely $P_R\gamma_i$. Therefore, for any unitary A_0 acting on the plane at index 00,

$$\mathbb{E}\mathbf{1}_{Q_0} \|A_0\rho_i A_0^* - \rho_i\|_1 \le 2r^{3/2}\mathbb{E}\mathbf{1}_{Q_0} \|P_R\gamma_i\|_{\mathrm{F}} \le \kappa/256.$$

Here the trace norm is at most $r^{3/2}$ times the Frobenius norm on r^3-by-r^3 matrices.

We use Lemma C.2, proved in Appendix C:

$$\min_{D \text{ unitary}} \|UZ - ZD\|_{\mathrm{F}}^2 \le \|UZZ^*U^* - ZZ^*\|_1 .$$

For each $d \in Q$, fix an orthonormal basis of the plane $R(d)$, and let A and B be the sign change and the swap of the two basis vectors, so that $AB = -BA$; we use them on $R(d_{00})$. Applied to $Z_{01}(\xi)$, this inequality gives an action U_{01} on index 01 with

$$\mathbb{E}\mathbf{1}_{Q_0} \|(A_0 - U_{01})\xi\|^2 \le \kappa/256.$$

Applied to $Z_{10}(\xi)$, it gives an action V_{10} with the same bound for $B_0 - V_{10}$. These actions may depend on the rectangle; their indices are disjoint, so they commute pointwise. Commuting them past the plane actions and using the triangle inequality, then (27), gives

$$2\sqrt{\kappa} \le \|\mathbf{1}_{Q_0}(A_0B_0 - B_0A_0)\xi\|_2 \le 2\left(\sqrt{\kappa/256} + \sqrt{\kappa/256}\right) = \sqrt{\kappa}/4.$$

This contradiction proves the lemma. □

4.3 Selecting a polynomial number of tests

The detector can be converted into a short list that works for every plane on almost every diagonal. We use a bounded potential on the real Hilbert space of Hermitian matrices.

Lemma 4.2 (Covariance selection)**.** *Suppose every field in a class $\mathcal{X}$ is Hermitian and has Frobenius norm at most one. Assume that $\mathcal{X}$ has the detection property of Lemma 4.1 with parameters q_0, ν: whenever $Q \subseteq G$ has $\mathbb{P}(Q) \ge q_0$ and $R(d)$ is a rank-two projection for each $d \in Q$, some $X \in \mathcal{X}$ satisfies (26). One can select m fields $X_1, \dots, X_m$ from the class, where*

$$m \le \left\lceil 8r^2\nu^{-4} \right\rceil + 1,$$

such that, outside a set of d's of measure less than q_0, every rank-two projection R satisfies

$$\sum_{j=1}^{m} \left\| D_R V_{X_j}(d) \right\|_{\mathrm{F}}^2 > \nu^2. \tag{32}$$

Proof. Work in the real Hilbert space $\mathcal{H} = \mathrm{Herm}_r$, with Frobenius inner product; its dimension is r^2. For an existing list, let

$$v_j(d) = V_{X_j}(d), \qquad T(d) = \sum_j v_j(d) \otimes v_j(d).$$

All v_j have norm at most one. Set $t = \nu^2$, and use the bounded potential

$$\Psi(T) = \mathrm{Tr}_{\mathcal{H}}(T(T+tI)^{-1}) = r^2 - t\,\mathrm{Tr}_{\mathcal{H}}(T+tI)^{-1}, \qquad 0 \le \Psi(T) \le r^2.$$

The trace here is an operator trace on $\mathcal{H}$.

Let Q consist of the d's admitting a plane R with $\mathrm{Tr}_{\mathcal{H}}(D_R T(d)) \le t$. This trace is exactly $\sum_j \|D_R v_j(d)\|^2$. If $\mathbb{P}(Q) \ge q_0$, choose such a plane at each $d \in Q$, and use the detector to obtain a new field $v = V_X$ for which

$$\mathbb{E}\mathbf{1}_Q \|D_R v\| > \nu.$$

Fix $d \in Q$, and write $A = T + tI$, $a = v^* A^{-1} v$. The rank-one inverse identity

$$(A + vv^*)^{-1} = A^{-1} - \frac{A^{-1} v v^* A^{-1}}{1 + v^* A^{-1} v},$$

verified by multiplying by $A + vv^*$, gives

$$\Psi(T + v \otimes v) - \Psi(T) = \frac{t\, v^* A^{-2} v}{1 + a}. \tag{33}$$

If $D_R v \neq 0$, take $u = D_R v / \|D_R v\|$. Then $u^* A u \le 2t$, so Cauchy–Schwarz in the A-metric gives

$$a \ge \frac{|\langle u, v\rangle|^2}{u^* A u} \ge \frac{\|D_R v\|^2}{2t}.$$

Also $v^* A^{-2} v \ge a^2$, since $\|v\| \le 1$, and $a \le t^{-1}$. As $t \le 1$, (33) is at least

$$\frac{t^2 a^2}{2} \ge \frac{\|D_R v\|^4}{8}.$$

This remains valid when $D_R v = 0$. The increment is nonnegative at all other d's. Jensen or Hölder therefore gives

$$\mathbb{E}(\Psi(T + v \otimes v) - \Psi(T)) \ge \tfrac{1}{8}\mathbb{E}\mathbf{1}_Q \|D_R v\|^4 > \nu^4/8.$$

The potential is at most r^2, so after the asserted number of additions we must have $\mathbb{P}(Q) < q_0$. This is the conclusion. □

4.4 A polynomial bound for small eigenvalue gaps

For a Hermitian matrix of size at least two, mingap denotes the least difference between adjacent eigenvalues, in increasing order and counted with multiplicity. The discriminant is useful for identifying a measure-zero exceptional set, but its degree need not appear in a quantitative small-gap bound. Convex eigenvalue sums give the following bound.

Lemma 4.3 (Small-gap bound)**.** *Let $F_1, \dots, F_m$ be Hermitian r-by-r matrices, $\|F_i\|_{\mathrm{op}} \le 1$, with $r \ge 2$. Suppose that, for every rank-two orthogonal projection R,*

$$\sum_i \|D_R F_i\|_{\mathrm{F}}^2 \ge \nu^2.$$

For independent uniform $t_i \in [-1, 1]$, put $H(t) = \sum_i t_i F_i$. Then, for every $\epsilon > 0$,

$$\mathbb{P}_t(\mathrm{mingap}\, H(t) \le \epsilon) \le \frac{4mr^2\epsilon}{\nu^2}. \tag{34}$$

Proof. **Simple spectrum almost everywhere.** The real span of the F_i contains a simple-spectrum matrix. Here is an elementary proof. On the compact set of Hermitian matrices M with spectrum $1, \dots, r$, maximize $\sum_i (\operatorname{Tr}(MF_i))^2$, and put $H = \sum_i \operatorname{Tr}(MF_i) F_i$. Comparing the maximum with any other matrix M' in that set gives

$$2 \operatorname{Tr}((M - M')H) \geq \sum_i \big(\operatorname{Tr}((M' - M)F_i)\big)^2. \tag{35}$$

In particular, M maximizes the linear pairing with H. Take an orthonormal eigenbasis $u_1, \dots, u_r$ of M, with $Mu_j = ju_j$. On the plane spanned by u_j, u_k, $j < k$, replace the restriction of M by $jR + (k-j)ww^*$, where R is the orthogonal projection onto that plane and w is any unit vector in it; the resulting matrix M' again has spectrum $1, \dots, r$. Linear maximality says that u_k maximizes the Rayleigh quotient of the compressed H. Hence its off-diagonal entry is zero and its diagonal entries are ordered. Applying this to every pair shows that H is diagonal in this basis, say with $h_1 \leq \cdots \leq h_r$.

If $h_j = h_k$, some F_i has nonscalar compression to their plane, by the hypothesis of the lemma. Choose w there so that $w^* F_i w \neq u_k^* F_i u_k$, and use the same modification M'. The left side of (35) is zero and the right side is positive, a contradiction. Thus H has simple spectrum.

The discriminant of the characteristic polynomial of $\sum_i t_i F_i$ is consequently a nonzero polynomial in t. Its zero set has Lebesgue measure zero; this elementary fact follows by induction on the number of variables and Fubini's theorem. Thus the spectrum is simple almost everywhere. The discriminant is used only for this null-set assertion, and its degree does not enter any estimate below.

Convex functions controlling gap curvature. Write the ordered eigenvalues as $\lambda_1 \leq \cdots \leq \lambda_r$, and put

$$g_j = \lambda_{j+1} - \lambda_j, \qquad S_k = \text{sum of the } k \text{ largest eigenvalues}, \quad S_0 = 0,$$

$$W = \sum_{k=1}^{r} S_k = \sum_{j=1}^{r} j\lambda_j.$$

Each S_k is convex, since $S_k(H) = \max_{\operatorname{rank} P = k} \operatorname{Tr}(PH)$, with the maximum over orthogonal projections. To verify this formula, diagonalize H. The numbers P_{jj} belong to $[0,1]$ and sum to k, so their weighted eigenvalue sum is at most the sum of the k largest eigenvalues; the corresponding spectral projection attains it. On a line in any coefficient t_i, its second derivative is therefore a nonnegative Radon measure. The identity

$$g_j = 2S_{r-j} - S_{r-j-1} - S_{r-j+1}$$

implies, as measures on that line,

$$\sum_{j=1}^{r-1} |D^2 g_j| \leq 4 D^2 W. \tag{36}$$

The minimum–maximum formula $\lambda_j(H) = \min_{\dim V = j} \max_{v \in V,\, \|v\|=1} v^* H v$ gives $|\lambda_j(H + sF_i) - \lambda_j(H + tF_i)| \leq |s - t|$. Thus $|\lambda_j'| \leq 1$, and hence

$$\sum_j |g_j'| \leq 2r, \qquad |W'| \leq r(r+1)/2 \leq r^2. \tag{37}$$

In particular, the total mass of $D^2 W$ on a coefficient interval $(-1, 1)$ is at most $2r^2$.

Gradient mass where a gap is small. On a coefficient line set $h_j = \min(g_j, \epsilon)$. The gaps are Lipschitz differences of convex functions, so g_j' is locally of bounded variation. The product rule for the continuous Lipschitz function h_j and the BV function g_j' gives

$$\int_{-1}^{1} \mathbf{1}_{\{g_j<\epsilon\}} (g_j')^2 \, dt_i = [h_j g_j']_{-1}^{1} - \int_{(-1,1)} h_j \, d(g_j').$$

Interior one-sided derivatives are used at the endpoints. The chain rule on the level $g_j = \epsilon$ causes no extra term: a Lipschitz function has derivative zero almost everywhere on a level set. The same formula follows by smoothing the convex S_k and passing to the limit.

Sum over j. By (37), the absolute boundary contribution is at most $4r\epsilon$. By (36), the absolute measure contribution is at most $8r^2\epsilon$. Divide by two for uniform measure on the line, average the other coordinates, and sum over i. We obtain

$$\mathbb{E}_t \sum_j \mathbf{1}_{\{g_j<\epsilon\}} \sum_i (\partial_i g_j)^2 \le \epsilon m(2r + 4r^2). \tag{38}$$

Off-diagonal mass where a gap is small. Since the discriminant is a nonzero polynomial, for almost every choice of the other coefficients it does not vanish identically on a given coefficient line, and then it has only finitely many zeros on that line. We discard the null set of remaining lines; this does not affect any average in this proof. At a simple-spectrum point choose orthonormal eigenvectors u_j. Differentiating the eigenvalue equation gives

$$\partial_i \lambda_j = u_j^* F_i u_j, \qquad \partial_{ii} \lambda_j = 2 \sum_{k \ne j} \frac{|u_k^* F_i u_j|^2}{\lambda_j - \lambda_k}.$$

Indeed, the component of $\partial_i u_j$ along u_k is $(u_k^* F_i u_j)/(\lambda_j - \lambda_k)$. Consequently

$$\partial_{ii} W = 2 \sum_{a<b} (b-a) \frac{|u_a^* F_i u_b|^2}{\lambda_b - \lambda_a}. \tag{39}$$

Every term is nonnegative. Therefore

$$2 \sum_j \mathbf{1}_{\{g_j<\epsilon\}} |u_j^* F_i u_{j+1}|^2 \le \epsilon \, \partial_{ii} W.$$

The integral of the classical Hessian density is bounded by that of the full positive second-derivative measure. Its average on a coefficient line is at most r^2, so

$$2\mathbb{E}_t \sum_j \mathbf{1}_{\{g_j<\epsilon\}} \sum_i |u_j^* F_i u_{j+1}|^2 \le \epsilon m r^2. \tag{40}$$

For the plane spanned by u_j, u_{j+1}, the hypothesis says exactly

$$\nu^2 \le \frac{1}{2} \sum_i (\partial_i g_j)^2 + 2 \sum_i |u_j^* F_i u_{j+1}|^2.$$

Sum over the small gaps, integrate, and use (38)–(40):

$$\nu^2 \, \mathbb{E}_t \#\{j : g_j < \epsilon\} \le \epsilon m(r + 3r^2) \le 4\epsilon m r^2.$$

The probability of a small gap is at most the expected number of small gaps. This proves (34) with a strict threshold; apply that result to decreasing thresholds above ϵ to obtain the displayed non-strict version. □

Corollary 4.4 (A covariant field with polynomial gaps)**.** *Let $0 < \eta < 1$. Apply Lemma 4.1 with $q_0 = \eta/2$ and any admissible κ, and then Lemma 4.2. There are coefficients $t_i \in [-1, 1]$ such that*

$$X = \frac{1}{m}\sum_{i=1}^{m} t_i X_i, \qquad A_* = \Pi X$$

satisfy $\|X\|_{\mathrm{F}}, \|A_*\|_{\mathrm{F}} \le 1$, the covariance identity for A_*, and

$$\mathbb{P}_z(\operatorname{mingap} A_*(z) \le \gamma) < \eta, \qquad \gamma = \frac{\eta\nu^2}{16m^2r^2}.$$

In particular, if κ^{-1} is polynomial in δ^{-1}, η^{-1}, then so is γ^{-1}.

Proof. Outside an exceptional set of d's of measure less than $\eta/2$, the matrices $F_i(d) = V_{X_i}(d)$ satisfy the hypothesis of Lemma 4.3. They have Frobenius, hence operator, norm at most one. Apply that lemma with threshold $\epsilon = m\gamma = \eta\nu^2/(16mr^2)$. The expected measure of additional bad diagonals is at most $\eta/4$. Some choice of coefficients therefore has total bad measure less than $3\eta/4 < \eta$. The eigenvalues of $A_*(z)$ are those of $m^{-1}\sum_i t_i F_i(d)$. Finally, the displayed formulas for ν, m, γ involve only fixed powers of $\kappa^{-1}, r, \eta^{-1}$. □

4.5 Step approximations to the separated field

We next record the rectangular complexity of the chosen field. The point of keeping a complete partition table in the detector is that its size enters only through its logarithm in these estimates.

For the fixed combination in Corollary 4.4, approximate every X_i to error e using (25) and refine its m pairs of partitions. The resulting Hermitian contraction approximates X with L^2 error at most e, and

$$L_{\mathrm{sum}}(e) \le [K^3(C/e)^4]^m, \qquad \log L_{\mathrm{sum}}(e) \le m\big(3\log K + C\log(C/e)\big). \tag{41}$$

All fields X_i, X, A_* remain fixed as the requested precision varies.

To approximate

$$A_*(z) = \mathbb{E}_h \sigma_h(z)^* X(T_h z)\sigma_h(z),$$

suppose $e \ge 16\zeta$. Use the prescribed unitary step approximation to σ_h and a Hermitian contraction step approximation to X at accuracy $e/8$. Changing the three factors costs at most $2\zeta + e/8 \le e/4$ in L^2. Each new summand is Hermitian with Frobenius norm at most one. A suitable sample of $O(e^{-2})$ shifts costs at most another $e/4$. Refine the partitions and translated partitions for these shifts. The resulting Hermitian contraction has error at most e, with class bound satisfying

$$\log L_A(e) \le Ce^{-2}\left[m\big(\log K + \log(C/e)\big) + \log(2L)\right]. \tag{42}$$

Suppose $\kappa^{-1}, \eta^{-1}, e^{-1}, \zeta^{-1}$ are polynomial in δ^{-1}. Then $\varepsilon^{-1}, \nu^{-1}, m, \gamma^{-1}$ are polynomial, $\log K$ is polynomial by (23), and the right side of (42) is polynomial. This is the bound needed to control the eigenvectors.

5 Translation in a controlled eigenbasis

We now turn the separated Hermitian field into scalar phases. The covariance law transports each ordered eigenline to the corresponding eigenline at the translated point. Once a unit

vector has been chosen on each line, this transport is multiplication by a scalar of modulus one. The main issue is to make the choices compatible with rectangular step approximation.

Two elementary observations will be useful. The first provides a threshold which remains useful at every later approximation scale. The second controls the movement of a simple eigenline.

Lemma 5.1 (A threshold with little nearby mass)**.** *Let $Y_1, \dots, Y_N$ be real random variables, and let I be an interval of length $0 < l \le 1$. There exists $v \in I$ such that, simultaneously for every $s > 0$,*

$$\sum_{i=1}^{N} \mathbb{P}(|Y_i - v| \le s) \le \frac{32N}{l}\sqrt{s}.$$

Proof. Integrating on the two sides of a fixed real number y gives $\int_I |y - v|^{-1/2}\, dv \le 4\sqrt{l}$. Averaging in v and then in the random variables therefore gives a choice with

$$\sum_i \mathbb{E}|Y_i - v|^{-1/2} \le 4N/\sqrt{l}.$$

On the event $|Y_i - v| \le s$, the integrand is at least $s^{-1/2}$. Multiplication by $\sqrt{s}$, followed by summation, proves the assertion, with room in the stated constant. The same choice of v works for every s. □

Lemma 5.2 (Movement of an eigenline)**.** *Suppose that $H, H' \in \mathrm{Herm}_r$ have distinct eigenvalues, and that the gaps between consecutive eigenvalues of both matrices are at least $\gamma/2$. Let $P_j(H)$ denote the rank-one orthogonal projection onto the jth eigenline, with eigenvalues counted in increasing order. Then*

$$\left\|P_j(H) - P_j(H')\right\|_{\mathrm{op}} \le \frac{8}{\gamma}\left\|H - H'\right\|_{\mathrm{op}}.$$

Proof. Put $w = \left\|H - H'\right\|_{\mathrm{op}}$. The ordered eigenvalues of H and H' differ by at most w. To see this directly, the spectral theorem gives

$$\lambda_j(H) = \min_{\dim V = j} \max_{\substack{u \in V \\ \|u\| = 1}} u^* H u.$$

Indeed, the span of the first j eigenvectors gives the upper bound, and every j-dimensional space intersects the span of the last $r - j + 1$ eigenvectors nontrivially, giving the lower bound. Changing H to H' changes each Rayleigh quotient by at most w, proving the displacement assertion.

If $w < \gamma/8$, write $\xi = \lambda_j(H')$, and take a unit eigenvector u' for ξ. Every eigenvalue of H with index different from j is at distance at least $3\gamma/8$ from ξ. Since $\|(H - \xi I)u'\| \le w$, expansion in an eigenbasis of H gives

$$\|(I - P_j(H))u'\| \le \frac{8w}{3\gamma}.$$

For unit vectors u, u', the difference of the projections onto their spans has operator norm $\sqrt{1 - |u^* u'|^2} = \|(I - uu^*)u'\|$. This follows by writing the two projections in the plane spanned by u, u'; their difference has eigenvalues $\pm\sqrt{1 - |u^* u'|^2}$ on that plane and is zero outside it. This proves the desired estimate when $w < \gamma/8$. Otherwise, the norm of the difference of two orthogonal projections is at most one, which is at most $8w/\gamma$. □

The following proposition records both the exact scalarization and the approximation estimate we need. Its hypotheses refer to just one prescribed approximation of the row and the cocycle. In the application, the separated field and its approximants are supplied by Corollary 4.4 and (42).

Proposition 5.3 (Scalarization at a prescribed accuracy)**.** *Let $r \geq 2$, $0 < \gamma \leq 1$, and $0 < \eta < 1$. Suppose that $A_* : G^2 \to \mathrm{Herm}_r$ has Frobenius norm at most one and satisfies*

$$A_*(T_h z) = \sigma_h(z) A_*(z) \sigma_h(z)^*, \qquad \mathbb{P}_z(\mathrm{mingap}\, A_*(z) \leq \gamma) < \eta.$$

Here σ_h is a unitary cocycle, and $\Phi : G^2 \to U(r)$ satisfies $\Phi(T_h z) = \sigma_h(z)\Phi(z)$. Let $\mathcal{R}(z) = \lambda \ell(z)$ be a row of Euclidean norm at most one. Suppose that $\mathcal{R}$ and σ_h have, respectively, a contraction row approximant and unitary matrix approximants with errors

$$\|\mathcal{R} - \mathcal{R}'\|_{L^2_z} \leq \zeta, \qquad \|\sigma - S\|_{L^2_{h,z}(\mathrm{F})} \leq \zeta,$$

whose vertex partitions have at most $L \geq 2$ classes. The partitions for S_h may depend on h. Suppose also that, for $16\zeta \leq a \leq 1$, the field A_* has a rectangular step approximant of L^2 Frobenius error at most a and class bound

$$\log L_A(a) \leq C a^{-2} \left[m(\log K + \log(C/a)) + \log(2L) \right], \qquad K \geq 2, \quad m \geq 1. \tag{43}$$

There is a translation-invariant set $E \subseteq G^2$, with $\mathbb{P}(E^c) < \eta$, a unitary field U, phases $\phi_1, \dots, \phi_r$, and a unitary matrix $W(d)$ depending only on the diagonal $d = a - b$, such that

$$\Phi^0(z) = U(z)\, \mathrm{diag}(\phi_1(z), \dots, \phi_r(z)) W(d). \tag{44}$$

Here Φ^0 agrees with Φ on E and is the identity on E^c. These exact objects are fixed independently of the accuracy requested below.

For $0 < t \leq 1/2$, define

$$\theta(t) = \left(\frac{t\gamma}{10^7 (r+1)^4} \right)^8, \qquad d_*(t) = \frac{\gamma \theta(t) t}{10^4 (r+1)^3}, \qquad a(t) = \frac{t\, d_*(t)}{100(r+1)}. \tag{45}$$

Let $0 < e \leq 1/2$, and suppose that

$$\zeta \leq \min\left\{ \frac{e}{100}, \frac{a(e/100)}{16} \right\}. \tag{46}$$

Then each coefficient and each multiplicative increment

$$c_j(z) = \mathbf{1}_E(z) \mathcal{R}(z) U_j(z), \qquad \rho_{j,h}(z) = \phi_j(z) \overline{\phi_j(T_h z)}$$

have one-bounded rectangular step approximants c'_j and $\rho'_{j,h}$ such that

$$\left\| c_j - c'_j \right\|_{L^2_z} \leq e, \qquad \left\| \rho_j - \rho'_j \right\|_{L^2_{h,z}} \leq e. \tag{47}$$

A common bound for their number of classes per vertex is

$$K_*(e) = \max\{2, L L_A(a(e/100))^2\}. \tag{48}$$

In particular, $\log K_*(e)$ is bounded by a polynomial with absolute coefficients and degree in

$$r,\ \gamma^{-1},\ e^{-1},\ m,\ \log K,\ \log(2L).$$

Proof. **Choosing the good diagonals.** Write $g(z) = \operatorname{mingap} A_*(z)$. Unitary covariance makes g constant on each diagonal. Apply Lemma 5.1 with the interval $[\gamma/2, \gamma]$, and choose v_* such that

$$\mathbb{P}(|g - v_*| \le s) \le (64/\gamma)\sqrt{s} \qquad (s > 0). \tag{49}$$

Set $E = \{g > v_*\}$. This set is translation-invariant and $E^c \subseteq \{g \le \gamma\}$, so $\mathbb{P}(E^c) < \eta$. On E^c, replace A_*, Φ, σ_h by $D_0 = \operatorname{diag}(1, \dots, r), I, I$, respectively, and denote the resulting fields by A, Φ^0, σ_h^0. Since E is invariant, they retain the identities

$$A(T_h z) = \sigma_h^0(z) A(z) \sigma_h^0(z)^*, \qquad \Phi^0(T_h z) = \sigma_h^0(z) \Phi^0(z). \tag{50}$$

The modified σ^0 is still a cocycle. Every eigenvalue gap of A is at least $\gamma/2$.

Choosing unit vectors on the eigenlines. For the standard coordinate vectors e_i, the lengths $Y_{ji}(z) = \|P_j(A(z))e_i\|$ satisfy

$$\sum_{i=1}^{r} Y_{ji}(z)^2 = \operatorname{Tr} P_j(A(z)) = 1.$$

Use Lemma 5.1 for these r^2 real random variables on the interval $[1/(3\sqrt{r}), 1/(2\sqrt{r})]$. There is a common threshold v_p in that interval with

$$\sum_{j,i} \mathbb{P}(|Y_{ji} - v_p| \le s) \le 192 r^{5/2} \sqrt{s}. \tag{51}$$

For each j, z, choose the first index i for which $Y_{ji}(z) > v_p$, and define

$$U_j(z) = \frac{P_j(A(z))e_i}{\|P_j(A(z))e_i\|}. \tag{52}$$

Such an index exists because the largest of these lengths is at least $1/\sqrt{r}$. The columns U_j form a unitary matrix U.

The covariance in (50) preserves the ordering of the eigenlines. Consequently there are unique phases $s_{j,h}$ such that

$$\sigma_h^0(z) U_j(z) = s_{j,h}(z) U_j(T_h z), \qquad (T_h U)^* \sigma_h^0 U = \operatorname{diag}(s_{1,h}, \dots, s_{r,h}). \tag{53}$$

They obey $s_{j,h+k}(z) = s_{j,h}(T_k z) s_{j,k}(z)$, by the cocycle law for σ^0. Put

$$V(z) = U(z)^* \Phi^0(z), \qquad W(d) = V(d, 0), \qquad \phi_j(a, b) = s_{j,b}(a - b, 0).$$

The jth row of $V(T_h z)$ equals $s_{j,h}(z)$ times the jth row of $V(z)$. Taking the translation from $(d, 0)$ to (a, b) proves (44). The scalar cocycle identity also gives

$$\phi_j(T_h z) = s_{j,h}(z) \phi_j(z), \qquad \rho_{j,h}(z) = \overline{s_{j,h}(z)}. \tag{54}$$

This fixes all the exact objects in the proposition.

Approximating the eigenvectors. Fix $0 < t \le 1/2$ with $a(t) \ge 16\zeta$, and abbreviate the three quantities in (45) by θ, d_*, a. Take a rectangular step approximant A' to A_*, with L^2 Frobenius error at most a, on partitions with at most $L_A(a)$ classes each. Call a point z good if all of the following conditions hold:

$$\|A_*(z) - A'(z)\|_{\mathrm{op}} \le d_*, \qquad |g(z) - v_*| > \theta, \qquad |Y_{ji}(z) - v_p| > \theta \quad \text{for every } j, i.$$

By (49), (51), the union bound, and Chebyshev's inequality (using $\|\cdot\|_{\mathrm{op}} \le \|\cdot\|_{\mathrm{F}}$), the proportion of points which are not good is at most

$$b_t := \left(\frac{64}{\gamma} + 192r^{5/2}\right)\sqrt{\theta} + \frac{a^2}{d_*^2} \le \frac{t^2}{1000r}. \tag{55}$$

For completeness, the coefficient of $\sqrt{\theta}$ is at most $256(r+1)^3/\gamma$. Its contribution is therefore at most $256t^4\gamma^3/[10^{28}(r+1)^{13}]$, which is less than $t^2/(2000r)$. The other contribution is $t^2/[10^4(r+1)^2]$, also at most $t^2/(2000r)$.

Consider two good points z, z' in the same cell of the pair of partitions for A'. Their matrices A_* differ in operator norm by at most $2d_*$. The ordered eigenvalues therefore differ by at most $2d_*$, and the minimum consecutive gaps differ by at most $4d_* < \theta$. Since both gaps are more than θ from v_*, the two points have the same membership in E.

If both points lie outside E, their modified matrices A are equal. If both lie in E, Lemma 5.2 gives

$$\|P_j(A(z)) - P_j(A(z'))\|_{\mathrm{op}} \le 16d_*/\gamma < \theta.$$

Thus the comparison of every Y_{ji} with v_p has the same outcome at the two points. In particular, each column in (52) uses the same coordinate index at both points. For nonzero vectors v, w with norms at least $b > 0$,

$$\left\| \frac{v}{\|v\|} - \frac{w}{\|w\|} \right\| \le \frac{\|v-w\|}{\|v\|} + \frac{|\,\|w\| - \|v\|\,|}{\|v\|} \le \frac{2}{b}\,\|v-w\|\,.$$

Here $b = 1/(3\sqrt{r})$. Hence each corresponding pair of columns differs by at most $96\sqrt{r}\,d_*/\gamma$, and

$$\|U(z) - U(z')\|_{\mathrm{F}} \le 96rd_*/\gamma < t/10. \tag{56}$$

In every cell containing a good point, choose one such point and use its values of U and $\mathbf{1}_E$ throughout the cell. In a cell containing no good point, use I and 0. This gives a unitary step field U' and a $\{0,1\}$-valued step field e', on the same partitions as A'. At good points the indicator is correct and the unitary error is less than $t/10$. At all other points those errors are at most 1 and $2\sqrt{r}$. It follows from (55) that

$$\|\mathbf{1}_E - e'\|_{L^2_z} \le \sqrt{b_t} < t, \qquad \|U - U'\|_{L^2_z(\mathrm{F})} \le t/10 + 2\sqrt{r}\sqrt{b_t} < t. \tag{57}$$

No additional partition refinement was used.

The scalar approximants and their complexity. For the requested accuracy e, use the preceding construction with $t = e/100$. Condition (46) ensures both that $a(t) \ge 16\zeta$ and that $\zeta \le t$. Set

$$c'_j = e'\mathcal{R}'U'_j.$$

All factors are contractions, so $|c'_j| \le 1$. Changing one factor at a time, and using (57), gives $\left\|c_j - c'_j\right\|_2 \le 2t + \zeta \le 3t < e$.

For the increments, form the unitary matrix

$$Q_h = (T_hU')^*(e'S_h + (1-e')I)U'.$$

The middle factor is unitary because e' takes only the values zero and one. Its difference from σ_h^0 has operator norm at most

$$\|S_h - \sigma_h\|_{\mathrm{op}} + 2|e' - \mathbf{1}_E|.$$

Translation preserves uniform measure, so the two outside factors each contribute at most t to the joint $L^2_{h,z}$ operator norm error. Thus

$$\left\| Q - (T_h U)^* \sigma_h^0 U \right\|_{L^2_{h,z}(\mathrm{op})} \le \zeta + 4t \le 5t < e.$$

Take $\rho'_{j,h} = \overline{(Q_h)_{jj}}$. A diagonal entry of a unitary matrix has modulus at most one, and (53)–(54) now prove (47). In particular, $\mathbb{E}_h \left\| \rho_{j,h} - \rho'_{j,h} \right\|_{L^2_z} \le e$, by Cauchy–Schwarz.

For c'_j, refine the partitions for U', e' by those for $\mathcal{R}'$. This gives at most $LL_A(a)$ classes per vertex. For $\rho'_{j,h}$, refine the partitions for U', e', their translates by h, and the partitions for S_h. This gives at most $LL_A(a)^2$ classes per vertex, proving (48). Finally, the exact formulas in (45) give

$$a(t)^{-1} = 10^{62}(r+1)^{36} t^{-10} \gamma^{-9}. \tag{58}$$

Combining this identity with (43) yields

$$\log K_*(e) \le C a(e/100)^{-2} \left[m(\log K + \log(C/a(e/100))) + \log(2L) \right],$$

after increasing the absolute constant. This is the asserted polynomial bound. □

The construction used the row and cocycle approximations only at their prescribed accuracy ζ. The thresholds were chosen before any request for e, so the exact phases and coefficients remain fixed. Whenever the later scalar argument calls for one accuracy satisfying (46), the displayed class bound supplies it without requiring a new choice of coordinates.

6 Scalar phases on products of cosets

The preceding construction reduces the problem to scalar phases whose simultaneous increments have controlled rectangular complexity. We now explain how to recover a scalar phase from those increments. The useful local domains are products of two cosets of one subspace. Such a product meets each relevant diagonal in the same number of points, so a local approximation can be transported along the whole diagonal.

Recall from (2) that, in the reflected coordinates, a *test* is a function

$$u(a)v(b)w(a-b)e_p(B(a,b)),$$

with one-bounded u, v, w and a bilinear form B. The *cost* of an expansion $\sum_j \lambda_j T_j$ into tests is $\sum_j |\lambda_j|$. Costs multiply when expansions are multiplied, because a product of tests is again a test. Multiplication by a one-bounded function of $a - b$ does not increase cost. Simultaneous translation also preserves the class of tests:

$$B(a+h, b+h) = B(a,b) + B(a,h) + B(h,b) + B(h,h),$$

and the last three terms are absorbed into the one-variable factors.

A one-bounded K-step function has an expansion of cost at most K. To see this, take the classes P of its first partition and choose $a_P \in P$. Then

$$h(a,b) = \sum_P \mathbf{1}_P(a) h(a_P, b).$$

The same observation applies on a product of cosets: extending the two one-variable factors by zero gives tests on all of G^2.

6.1 A Fourier lemma for labelled coordinates

We first establish the Fourier estimate that allows us to pass from vertex labels to cosets. For a subspace $H \le G$, write

$$\mathcal{I}_H F(x_1, \dots, x_k) = \mathbb{E}_{h \in H} F(x_1 + h, \dots, x_k + h).$$

Each x_i may range over a different coset of H. An origin in each coset identifies its character group with H^*.

Lemma 6.1 (The spectrum seen by a partition)**.** *Let P be a label on H taking at most K values, let X be uniform on H, and put*

$$a_P(\xi)^2 = \mathbb{E}\left|\mathbb{E}(e_p(\xi(X)) \mid P)\right|^2, \qquad \xi \in H^*.$$

For every linearly independent set $\Lambda \subseteq H^*$,

$$\sum_{\xi \in \Lambda} a_P(\xi)^2 \le 2 \log K.$$

Proof. Fix a nonempty class C, of relative size θ, and let

$$z_\xi = \mathbb{E}(e_p(\xi(X)) \mid X \in C), \qquad M = \sum_{\xi \in \Lambda} |z_\xi|^2.$$

The real random variable

$$Y(X) = \sum_{\xi \in \Lambda} \mathrm{Re}\big(\overline{z_\xi} e_p(\xi(X))\big)$$

has conditional mean M on C. On all of H, its summands are independent and have mean zero: the independent linear forms $\xi(X)$ are jointly uniform.

If a mean-zero real random variable U takes values in $[-r, r]$, the chord bound for the exponential gives $\mathbb{E}e^U \le \cosh r \le e^{r^2/2}$. Applying this to the summands and using Jensen's inequality on C,

$$e^M \le \mathbb{E}(e^Y \mid X \in C) \le \theta^{-1} \mathbb{E} e^Y \le \theta^{-1} e^{M/2}.$$

Hence $M \le 2 \log(1/\theta)$. Averaging over the classes gives

$$\sum_{\xi \in \Lambda} a_P(\xi)^2 \le 2 \sum_C \theta_C \log(1/\theta_C) \le 2 \log K,$$

the last inequality following from concavity of the logarithm. □

Lemma 6.2 (Common translation and coset approximation)**.** *Let $C_1, \dots, C_k$ be cosets of a subspace $H \le G$, and let $|F| \le 1$ on $\prod_i C_i$. Suppose F is determined by one label $P_i(x_i)$ per coordinate, each with at most K values. For $0 < \epsilon < 1$, there is a subspace $H' \le H$ with*

$$\mathrm{codim}_H H' \le 2k^2 \epsilon^{-2} \log(2K)$$

such that averaging $\mathcal{I}_H F$ on each product of H'-cosets changes it by at most ϵ in L^2. The resulting approximation has modulus at most one. If F takes values in $[0, 1]$, so does the approximation.

Proof. For each i, span the frequencies satisfying $a_{P_i}(\xi) > \epsilon/\sqrt{k}$. A basis chosen from this set has at most $2k\epsilon^{-2} \log K$ elements by Lemma 6.1. Let $V \le H^*$ be the sum of these k spans, and put $H' = V^\perp$. Its codimension has the stated bound.

Expand F in characters in all k coordinates. The average $\mathcal{I}_H$ retains precisely the frequencies $(\xi_1, \dots, \xi_k)$ with $\sum_i \xi_i = 0$. Fix all frequencies except the i-th, and consider the partial Fourier coefficient

$$f_i(x_i) = \mathbb{E}_{x_j,\ j\neq i} F(x_1, \dots, x_k) \prod_{j\neq i} \overline{e_p(\xi_j(x_j))}.$$

This is constant on the classes of P_i. There is only one choice $\xi_i = -\sum_{j\neq i} \xi_j$ that survives common translation. If it lies outside V, averaging within the classes of P_i and applying Cauchy–Schwarz gives

$$\left|\widehat{F}(\xi_1, \dots, \xi_k)\right| \le a_{P_i}(\xi_i)\, \|f_i\|_2 \le \frac{\epsilon}{\sqrt{k}}\, \|f_i\|_2 .$$

Parseval in the other coordinates gives $\sum_{(\xi_j)_{j\neq i}} \|f_i\|_2^2 = \|F\|_2^2 \le 1$. Thus the squared Fourier mass surviving common translation but lying outside V in slot i is at most ϵ^2/k. Summing over the slots costs at most ϵ^2.

Averaging each coordinate over its H'-coset retains exactly the frequencies in V in every slot. This proves the error bound. Both averaging operations preserve the unit disk and the interval $[0, 1]$. □

6.2 The scalar approximation statement

For a phase $\phi : G^2 \to \mathbb{T}$, set

$$\rho_h(z) = \phi(z)\overline{\phi(T_h z)}.$$

Thus $\phi(z) = \rho_h(z)\phi(T_h z)$; the increments are used in this orientation throughout. For $h \in G$ we also write $\Delta_h g = g\,\overline{T_h g}$ for the simultaneous increment, so that $\rho_h = \Delta_h \phi$. When g is a function of several vertices (for instance of a rectangle), T_h translates every vertex simultaneously.

The following proposition says, roughly, that if every non-negligible restriction of $cm\phi$ to a set of diagonals has a cube average bounded below, then $cm\phi$ is close to a combination of tests whose cost is p raised to a power polynomial in the inverse accuracy and in the logarithms of the relevant class bounds.

Proposition 6.3 (Scalar approximation)**.** *Let $\phi : G^2 \to \mathbb{T}$, let*

$$|c(a, b)| \le \omega(a - b) \le 1, \qquad |m(d)| \le 1,$$

and let $0 < s < 1$. Suppose that, for some $0 < \alpha_0 \le 1$, every restriction $m' = m\mathbf{1}_E$, $E \subseteq G$, satisfies

$$\|\omega m'\|_2 \ge s/2 \quad \Longrightarrow \quad D(cm'\phi)^8 \ge \alpha_0. \tag{59}$$

Put

$$\alpha = \min\{\alpha_0, s^8/8^4\}, \qquad u = \alpha/8, \qquad t = (\alpha/2^{60})^{64}, \qquad e = t^{16}. \tag{60}$$

Suppose that c and the increments ρ_h have one-bounded step approximations c_e, $c_{s/4}$, and g_h satisfying

$$\|c - c_e\|_2 \le e, \qquad \left\|c - c_{s/4}\right\|_2 \le s/4, \qquad \mathbb{E}_h \|\rho_h - g_h\|_2 \le e,$$

where c_e and every g_h are K_e-step, $c_{s/4}$ is K_s-step, and $K_e, K_s \ge 2$. The partitions for g_h may depend on h. Finally put

$$Q = \lceil e^{-2} \rceil, \qquad S = K_e^{2Q}, \qquad J = \left\lceil (32e^{-2} + 144t^{-2}) \log(2S) \right\rceil + 3. \tag{61}$$

Then

$$J \le 10^3 e^{-4} \log(2K_e),$$

and $cm\phi$ has an L^2 approximation of error at most s by a finite linear combination of tests of cost at most

$$4K_s\alpha^{-1}p^{2J}SK_e. \tag{62}$$

In particular, for an absolute C, the logarithm to base p of this cost is at most

$$Ce^{-4}\big(\log(2K_e)+\log(2K_s)\big). \tag{63}$$

Only the two displayed approximation accuracies are required. In particular, the proposition does not require a family of underlying phases or coordinate choices that is uniform over all precisions.

We prove the proposition by constructing a local approximation on one product of cosets, transporting it along its diagonals, and repeating on the remaining diagonals. The hypothesis (59) ensures that every non-negligible remainder admits another local construction.

6.3 Choosing four cosets

The aim of this subsection is to find a product R of four cosets of a subspace of polynomial codimension on which the phase of the rectangle function A defined below is nearly determined by the increments ρ_h, and is approximated by a function of bounded complexity.

Fix a restriction m' with $\|\omega m'\|_2 \geq s/2$. For $w=(a_0,a_1,b_0,b_1)$, write

$$\Xi(w)=\Box\phi(w), \qquad \Omega(w)=\mathcal{I}_G\Box(c\phi)(w), \qquad A(w)=\Xi(w)\overline{\Omega(w)}.$$

The second factor is invariant under simultaneous translation, so

$$|A|=|\Omega|, \qquad A=\mathbb{E}_h\Box(\rho_h\,\overline{T_hc}), \qquad \frac{A}{T_hA}=\Box\rho_h \quad \text{where } A\neq 0.$$

The sample-average formula follows by multiplying $\phi(z)\overline{(c\phi)(T_hz)}=\rho_h(z)\overline{c(T_hz)}$ at the four corners. In particular, $|A|$ is invariant.

The rectangle identity gives

$$D(cm'\phi)^8=\mathbb{E}|A(w)|^2\prod_{i,j=0}^{1}|m'(a_i-b_j)|^2.$$

Put $w_0(a_0,b_0)=|m'(a_0-b_0)|^2$. Since all the factors have modulus at most one,

$$\mathbb{E}w_0\mathbf{1}_{\{|A|\geq u\}}\geq\alpha-u^2. \tag{64}$$

In the sample-average formula for A, replace ρ_h by g_h and c by c_e. Telescoping the eight factors and averaging in h gives L^2 error at most $8e$. A suitable average of Q independent shifts adds error at most $Q^{-1/2}\leq e$. We obtain a function A', of modulus at most one, with

$$\|A-A'\|_2\leq 9e. \tag{65}$$

For each sampled shift, a vertex needs one label from the partition for g_h and one from the translated partition for c_e. Consequently A' is determined by at most $S=K_e^{2Q}$ labels per vertex. The two a-vertices have the same partition, as do the two b-vertices.

Apply Lemma 6.2 with $k=4$ to $|A'|^2$, at accuracy e. This gives $H\leq G$, of codimension at most $32e^{-2}\log(2S)$, and a function $\mathcal{L}$, constant on products of H-cosets, with $0\leq\mathcal{L}\leq 1$ and

$$\big\|\mathcal{L}-|A|^2\big\|_2\leq\big\|\mathcal{L}-\mathcal{I}_G|A'|^2\big\|_2+\big\|\mathcal{I}_G(|A'|^2-|A|^2)\big\|_2\leq 19e.$$

Let $U = \{\mathcal{L} \geq u^2/2\}$. Chebyshev's inequality, applied at the two thresholds $u^2/2$ and $u^2/4$, gives

$$\mathbb{E} w_0 \mathbf{1}_U \geq \alpha - u^2 - (38e/u^2)^2 \geq 3\alpha/4, \tag{66}$$

$$\mathbb{P}(U \cap \{|A| < u/2\}) \leq (76e/u^2)^2. \tag{67}$$

The numerical inequality follows from (60); in the estimates below it is enough to use the much weaker consequence

$$t \leq 2^{-200}\alpha^8, \qquad e = t^{16}.$$

We must also control the errors in the increments at the points used to transport the eventual local approximation. Define

$$Z(z') = \mathbb{E}_h \left| g_h(T_{-h}z') - \rho_h(T_{-h}z') \right|^2.$$

Since the difference of two one-bounded functions has modulus at most two,

$$\mathbb{E} Z \leq 2\mathbb{E}_h \left\| g_h - \rho_h \right\|_2 \leq 2e.$$

Consider the products of four H-cosets contained in U on which the average of w_0 is at least $\alpha/2$. Their union has measure at least $\alpha/4$, by (66). On that union the average of $\mathbf{1}_{\{|A|<u/2\}} + |A - A'|^2$ is at most

$$4\alpha^{-1}((76e/u^2)^2 + 81e^2),$$

and the average of $Z(a_0, b_0)$ is at most $8e/\alpha$. Markov's inequality, with factor three for each of the two quantities, therefore supplies

$$R = X_0 \times X_1 \times Y_0 \times Y_1$$

with X_i, Y_j cosets of H, such that

$$\mathbb{E}_R w_0 \geq \alpha/2, \tag{68}$$

$$\mathbb{E}_R(\mathbf{1}_{\{|A|<u/2\}} + |A - A'|^2) \leq 12\alpha^{-1}((76e/u^2)^2 + 81e^2) =: E_0, \tag{69}$$

$$\mathbb{E}_{X_0 \times Y_0} Z \leq 24e/\alpha. \tag{70}$$

Here and below norms and averages with a set as subscript use its normalized counting measure.

On R, let $\psi = A/|A|$ and $\psi' = A'/|A'|$, assigning the value one at zeros. Every $h \in H$ preserves R, and the zero set of A is invariant. Normalization changes the distance between A, A' by at most $4/u$ where $|A| \geq u/2$. Thus, for every $h \in H$,

$$\left\| \psi - \psi' \right\|_{2,R} \leq (2 + 4/u)\sqrt{E_0} \leq t, \qquad \left\| \Delta_h \psi - \Box \rho_h \right\|_{2,R} \leq 2\sqrt{E_0} \leq t. \tag{71}$$

These numerical bounds follow, for example, from $E_0 \leq 2^{29} e^2 \alpha^{-5}$ and (60).

6.4 Removing the dependence on four vertices

We next factor ψ, up to a small error and an invariant factor, into one phase on each edge. This step is the reason for the second application of the Fourier lemma.

Write $y = (b_0, b_1)$, and introduce independent $a_0' \in X_0, a_1' \in X_1$. For a phase P on R, put

$$\mathcal{R}_a P = \frac{P(a_0, a_1, y) P(a_0', a_1', y)}{P(a_0, a_1', y) P(a_0', a_1, y)}.$$

Define $\mathcal{R}_b$ similarly, using independent $b_0' \in Y_0, b_1' \in Y_1$. If P is a product of a function of a_0 and a function of a_1, with y fixed, then $\mathcal{R}_a P = 1$. In particular,

$$\mathcal{R}_a(\Box\rho_h) = \mathcal{R}_b(\Box\rho_h) = 1.$$

Let simultaneous translation now act on all six arguments of a quotient. By (71) and telescoping,

$$\|\Delta_h(\mathcal{R}_a\psi) - 1\|_2 \le 4t, \qquad \|\mathcal{R}_a\psi - \mathcal{R}_a\psi'\|_2 \le 4t.$$

Averaging the first estimate over $h \in H$, then using the second, gives

$$\|\mathcal{R}_a\psi - \mathcal{I}_H(\mathcal{R}_a\psi')\|_2 \le 8t.$$

The same statements hold for $\mathcal{R}_b$. Each of the six arguments of either quotient of ψ' has at most S labels. Apply Lemma 6.2 to the two quotients at precision t, and round the resulting functions to the nearest point of $\mathbb{T}$. Rounding at most doubles their distance to any phase. We obtain phases Θ_a, Θ_b such that

$$\|\mathcal{R}_a\psi - \Theta_a\|_2 \le 18t, \qquad \|\mathcal{R}_b\psi - \Theta_b\|_2 \le 18t. \tag{72}$$

Each is constant on products of cosets of a subspace of H, of codimension at most $72t^{-2}\log(2S)$. Intersect the two subspaces and call the result H'. Then

$$[G : H'] \le p^J.$$

In particular, changing even one argument of Θ_a or Θ_b by an element of H' leaves its value unchanged.

We now fix a_0', a_1', b_0', b_1'. The elementary identities underlying the factorization are worth recording. Set

$$L_0(y) = \psi(a_0, a_1', y), \quad L_1(y) = \psi(a_0', a_1, y), \quad L_*(y) = \psi(a_0', a_1', y),$$

and, for any of these three functions, set

$$(\mathcal{P}L)(y) = \frac{L(b_0, b_1')L(b_0', b_1)}{L(b_0', b_1')}.$$

Then $\psi = (\mathcal{R}_a\psi)L_0L_1/L_*$ and $L = (\mathcal{R}_bL)(\mathcal{P}L)$. Combining them gives

$$\psi = (\mathcal{R}_a\psi)\frac{(\mathcal{R}_bL_0)(\mathcal{R}_bL_1)}{\mathcal{R}_bL_*}\frac{(\mathcal{P}L_0)(\mathcal{P}L_1)}{\mathcal{P}L_*}. \tag{73}$$

The first four quotients are approximated by Θ_a and three evaluations of Θ_b. The final fraction consists of nine evaluations of ψ.

Choose all four fixed points together. Before fixing, each of the four quotient errors has squared mean at most $(18t)^2$, and each of the nine evaluations of $\psi - \psi'$ has squared mean at most t^2. The expected sum of these thirteen squared errors, averaged over whichever vertices remain free, is at most

$$4(18t)^2 + 9t^2 < (50t)^2.$$

Some choice of the fixed points has this bound. For that choice, each individual error has L^2 norm at most $50t$.

Replace the first four quotients in (73) by their approximants and call their product Θ_*. It is invariant under simultaneous H'-translation of the four free vertices. The last fraction is a product of four edge phases. One explicit choice is

$$\begin{aligned}
d_{00}(a_0,b_0) &= \psi(a_0,a_1',b_0,b_1'),\\
d_{01}(a_0,b_1) &= \frac{\psi(a_0,a_1',b_0',b_1)}{\psi(a_0,a_1',b_0',b_1')},\\
d_{10}(a_1,b_0) &= \frac{\psi(a_0',a_1,b_0,b_1')}{\psi(a_0',a_1',b_0,b_1')},\\
d_{11}(a_1,b_1) &= \frac{\psi(a_0',a_1,b_0',b_1)\psi(a_0',a_1',b_0',b_1')}{\psi(a_0',a_1,b_0',b_1')\psi(a_0',a_1',b_0',b_1)}.
\end{aligned}$$

It follows that

$$\Big\| \psi - \Theta_* \prod_{i,j} d_{ij} \Big\|_{2,R} \le 200t. \tag{74}$$

Moreover

$$\widetilde{d}_{00}(a_0,b_0) = \psi'(a_0,a_1',b_0,b_1')$$

is an S-step phase on $X_0 \times Y_0$, and

$$\big\| d_{00} - \widetilde{d}_{00} \big\|_{2,X_0\times Y_0} \le 50t. \tag{75}$$

6.5 A corrected phase with cube average close to one

For $i,j \in \{0,1\}$, put

$$g_{ij}(a_i,b_j) = \frac{\mathcal{C}^{\,i+j}\phi(a_i,b_j)}{d_{ij}(a_i,b_j)},$$

where $\mathcal{C}$ denotes complex conjugation. Then $\prod_{i,j} g_{ij} = \Xi / \prod_{i,j} d_{ij}$. For $h \in H'$, the invariant factor Θ_* disappears under Δ_h. Equations (71) and (74), applied at a point and its translate, therefore give

$$\Big\| \prod_{i,j} \Delta_h g_{ij} - 1 \Big\|_{2,R} \le 401t. \tag{76}$$

Indeed, $\Delta_h \Xi = \Box \rho_h$, and the left side is the distance between $\Delta_h \Xi$ and $\Delta_h \prod d_{ij}$, after division by a phase.

The real part of the mixed rectangle average in (76) is at least $1 - 401t$. Apply the mixed rectangle Cauchy–Schwarz inequality, conjugating the inputs at corners $01, 10$ so that their rectangular product is the product displayed there. All four rectangle norms are at most one. Consequently

$$\|\Delta_h g_{00}\|_{\square(X_0,Y_0)} \ge 1 - 401t.$$

We also need the two rectangular increments to lie in H'. Partition X_0 and Y_0 into H'-cosets. For any function F on $X_0 \times Y_0$, condition its rectangle average on the cosets P_0, P_1, Q_0, Q_1 of its four vertices. Mixed rectangle Cauchy–Schwarz and the arithmetic–geometric mean inequality give the bound

$$\prod_{i,j} \|F\|_{\square(P_i,Q_j)} \le \frac{1}{4} \sum_{i,j} \|F\|^4_{\square(P_i,Q_j)}\,.$$

Averaging over the cosets therefore yields

$$\|F\|_{\square(X_0,Y_0)}^4 \le \mathbb{E}_{C'} \|F\|_{\square(C')}^4 ,$$

where C' is a uniformly chosen product of two H'-cosets inside $X_0 \times Y_0$. Apply this with $F = \Delta_h g_{00}$, then average over $h \in H'$. The rectangle identity on each C' gives

$$\mathbb{E}_{C'} D_{C'}(g_{00})^8 \ge (1 - 401t)^4 \ge 1 - 1604t. \tag{77}$$

Here $D_{C'}$ uses increments in H' and normalized measure on C'.

By (68), the products with $\mathbb{E}_{C'} w_0 \ge \alpha/4$ constitute at least an $\alpha/4$ fraction. The average over all products of

$$1 - D_{C'}(g_{00})^8 + \left\| d_{00} - \widetilde{d}_{00} \right\|_{2,C'}^2$$

is at most $1604t + (50t)^2 \le 1605t$. The average of $\mathbb{E}_{C'} Z$ is at most $24e/\alpha$, by (70). Passing to the products with the required w_0-mass and applying Markov's inequality with factor three to each quantity supplies one product

$$C' = (a_* + H') \times (b_* + H')$$

for which

$$\mathbb{E}_{C'} w_0 \ge \alpha/4, \qquad \mathbb{E}_{C'} Z \le 300e/\alpha^2, \tag{78}$$

$$1 - D_{C'}(g_{00})^8 + \left\| d_{00} - \widetilde{d}_{00} \right\|_{2,C'}^2 \le 2 \cdot 10^4 t/\alpha. \tag{79}$$

Corollary A.4 now gives a test phase s_0 on C' with

$$\|g_{00} - s_0\|_{2,C'}^2 \le 8192 \cdot 2 \cdot 10^4 t/\alpha.$$

Its bilinear form and one-variable factors extend to a test on G^2, as described in that corollary. Put

$$v_0 = \mathbf{1}_{C'} \widetilde{d}_{00} s_0.$$

It has modulus at most one, is supported on C', and has a test expansion of cost at most S. Since $\phi = d_{00} g_{00}$ on $X_0 \times Y_0$,

$$\|\phi - v_0\|_{2,C'}^2 \le 10^9 t/\alpha. \tag{80}$$

The estimate follows by changing the two factors and using (79); the chosen constant bounds $2 \cdot 8192 \cdot 2 \cdot 10^4 + 4 \cdot 10^4$.

6.6 Transport along diagonals and iteration

Write $q = |H'|/|G|$, so $q \ge p^{-J}$. The simultaneous translates of C' fill

$$V = \{(a, b) : a - b \in a_* - b_* + H'\},$$

which has density q in G^2. For $z = (a, b) \in V$, exactly the shifts $h \in a_* - a + H'$ take z into C'. Their proportion is q. No shift takes a point outside V into C'. Define on all of G^2

$$v_1(z) = q^{-1} \mathbb{E}_h g_h(z) v_0(T_h z). \tag{81}$$

On V, this is an ordinary average over the shifts taking z into C'; off V, it is zero. In particular $|v_1| \le 1$. Translation preserves test cost, and multiplication by g_h has cost at most K_e. Hence (81) has a test expansion of cost at most

$$q^{-1} K_e S.$$

To estimate the error, choose z uniformly in V, and then choose h uniformly among the shifts taking z into C'. Under the change of variables $z' = T_h z$, this is the uniform product law on $C' \times G$. Indeed the inverse is $(z', h) \mapsto (T_{-h} z', h)$, and all admissible pairs have equal weight. For each pair, the identity $\phi = \rho_h T_h \phi$ gives

$$\phi(z) - g_h(z) v_0(z') = \rho_h(z)(\phi(z') - v_0(z')) + (\rho_h(z) - g_h(z)) v_0(z').$$

Jensen's inequality, (80), and (78) therefore give

$$\|\phi - v_1\|_{2,V}^2 \le 2 \|\phi - v_0\|_{2,C'}^2 + 2\mathbb{E}_{C'} Z \le 2 \cdot 10^9 t/\alpha + 600e/\alpha^2. \tag{82}$$

The difference $a - b$ is uniform on $a_* - b_* + H'$ for uniform points of either V or C'. Thus

$$\mathbb{E}_V |m'|^2 = \mathbb{E}_{C'} w_0 \ge \alpha/4. \tag{83}$$

Start with $m' = m$. At each step for which $\|\omega m'\|_2 \ge s/2$, carry out the construction above, insert the term $m' v_1$, and replace m' by $m' \mathbf{1}_{V^c}$. The inserted terms have disjoint supports: a later weight vanishes on every previously removed diagonal. Multiplication by $m'(d)$ does not increase their test costs.

By (83), the squared L^2 norm of the remaining weight decreases by at least $q\alpha/4$ at a step. It is initially at most one. The sum of the densities q of all chosen sets is therefore at most $4/\alpha$, and, since each $q \ge p^{-J}$, the number of steps is at most $4p^J/\alpha$. In particular the procedure terminates, with a final weight satisfying $\|\omega m'\|_2 < s/2$.

Let v be the sum of the inserted terms. It is one-bounded. On the removed diagonals the squared errors add, and (82) bounds their total by

$$\frac{4}{\alpha}\left(\frac{2 \cdot 10^9 t}{\alpha} + \frac{600e}{\alpha^2}\right) \le s^2/16.$$

For the final inequality, use $t \le 2^{-200}\alpha^8$, $e \le t$, and $\alpha \le s^8/8^4$. Thus cv approximates $cm\phi$ on the removed diagonals with L^2 error at most $s/4$. On the remaining diagonals $v = 0$, and the error is at most

$$\|cm'\phi\|_2 \le \|\omega m'\|_2 < s/2.$$

Finally replace c in cv by $c_{s/4}$. This costs at most another $s/4$, because $|v| \le 1$. The total error is at most s.

Each step has cost at most $p^J K_e S$, there are at most $4p^J/\alpha$ steps, and multiplication by $c_{s/4}$ costs at most K_s. This proves (62). Also $Q \le 2e^{-2}$ and

$$\log(2S) \le 2Q \log(2K_e) \le 4e^{-2} \log(2K_e), \qquad t^{-2} \le e^{-2}.$$

These estimates give $J \le 10^3 e^{-4} \log(2K_e)$. Taking logarithms in (62), using $p \ge 3$, and absorbing $\log(\alpha^{-1})$ into Ce^{-4}, gives (63) and completes the proof of Proposition 6.3.

7 Proof of the inverse theorem

We now assemble the constructions and choose their parameters. The order of this choice is useful: every accuracy required of the initial coordinates can be determined from δ alone, before those coordinates or their partitions are constructed.

7.1 Choosing the accuracy before the coordinates

Assume $\beta(f) \geq \delta$, and set

$$r = \lceil 8/\delta \rceil, \qquad \lambda_- = \delta/8, \qquad s = \frac{\lambda_- \delta}{100r}, \qquad \eta = s^2, \qquad q_0 = \eta/2. \tag{84}$$

For the matrix detector and the spectral gap, take

$$\kappa = \left(\frac{\lambda_-^3 q_0}{2r} \right)^4, \qquad \varepsilon = \frac{\kappa^2}{2^{20}(r+1)^3}, \qquad \nu = \frac{\kappa^2}{2^{20}(r+1)^{15}}, \tag{85}$$

$$m_{\max} = \left\lceil 8r^2 \nu^{-4} \right\rceil + 1, \qquad \gamma = \frac{\eta \nu^2}{16 m_{\max}^2 r^2}. \tag{86}$$

For the scalar approximation, take

$$\alpha_0 = \left(\frac{\lambda_-^3 s^2}{4} \right)^4, \qquad \alpha = \min\{\alpha_0, s^8/8^4\}, \qquad e = (\alpha/2^{60})^{1024}. \tag{87}$$

In particular, $e \leq s/4$. The eigenvector construction uses the matrix accuracy

$$a_0 = \left(10^{62}(r+1)^{36}(e/100)^{-10}\gamma^{-9} \right)^{-1}. \tag{88}$$

Finally prescribe the accuracy of the initial coordinates to be

$$\zeta = \min\{1/2,\ \varepsilon/32,\ e/100,\ a_0/16\}. \tag{89}$$

Table 1 summarizes the roles of these parameters. All these choices precede the application of Lemma 3.1.

Parameter	Role	Where used
r, λ_-	dimension bound and spectral cutoff	Lemma 3.1
s	L^2 accuracy for each scalar summand	Proposition 6.3
$\eta = s^2$, $q_0 = \eta/2$	measure of discarded diagonals	Corollary 4.4
κ, ε, ν	detector parameters	Lemma 4.1
$m_{\max}$, γ	number of selected fields; spectral gap	Lemma 4.2, Corollary 4.4
α_0, α, e	robustness and accuracy for scalar phases	Proposition 6.3
a_0	matrix accuracy for the eigenvectors	Proposition 5.3
ζ	accuracy of the initial coordinates	Lemma 3.1

Table 1: The parameters of Section 7. Each reciprocal is at most $(C/\delta)^C$.

Every reciprocal in (84)–(89), as well as r and $m_{\max}$, is bounded by $(C/\delta)^C$. This follows simply by substituting the displayed formulas: there are a fixed number of products, minima, and fixed powers, and $r \leq 9/\delta$.

Apply Lemma 3.1 at this ζ. It produces the exact data with some $\lambda \in [\lambda_-, \delta/4]$ and

$$\log L \leq (C/\delta)^C. \tag{90}$$

The value of κ chosen above is admissible in Lemma 4.1, because $\lambda \geq \lambda_-$. The condition $\zeta \leq \varepsilon/32$ permits its tensor approximation. Lemma 4.2 gives a list of at most $m_{\max}$ tests. Corollary 4.4 then gives the field A_* with

$$\mathbb{P}(\operatorname{mingap} A_* \leq \gamma) < \eta.$$

Indeed its stated gap, using the actual list length m, is at least the γ in (86).

The partition size K in the detector satisfies $\log K \le C\varepsilon^{-2}\log(2L) \le (C/\delta)^C$. Equation (42) applies at a_0, since $a_0 \ge 16\zeta$, and gives

$$\log L_A(a_0) \le (C/\delta)^C.$$

The construction of Section 5, at scalar accuracy e, therefore gives an invariant set E, with $\mathbb{P}(E^c) < \eta$, and a factorization

$$\Phi(z) = U(z)\,\mathrm{diag}(\phi_1(z), \dots, \phi_r(z))W(d) \qquad (z \in E). \tag{91}$$

It also gives one-bounded step approximations to

$$c_j = \mathbf{1}_E \lambda \ell U_j, \qquad \rho_{j,h} = \phi_j\,\overline{T_h \phi_j},$$

with L^2 errors at most e, averaged in h for the increments, and a common class bound $K_* \ge 2$ satisfying

$$\log K_* \le (C/\delta)^C. \tag{92}$$

The same coefficient approximation is accurate to $s/4$, because $e \le s/4$. This verifies all approximation requirements of Proposition 6.3 using only the one prescribed initial accuracy.

7.2 The scalar hypotheses and the total cost

Put

$$m_j(d) = (W(d)u_0(d))_j.$$

Since W is unitary and $\|u_0\| \le 1$, each m_j is one-bounded. Equation (91) gives

$$\mathbf{1}_E \lambda \ell u = \sum_{j=1}^{r} c_j(z) m_j(d) \phi_j(z). \tag{93}$$

The eigenvectors need not lie in the active coordinate space. Their active components supply the appropriate envelopes. Define

$$\omega_j(d) = \mathbf{1}_E \left\| P(d) U_j(z) \right\|.$$

This is a well-defined function of d. On E, simultaneous translation multiplies U_j by the cocycle and a scalar phase; the cocycle commutes with $P(d)$. Thus the displayed norm is constant along each diagonal. Moreover $0 \le \omega_j \le 1$ and $|c_j| \le \omega_j$, since $\ell = \ell P$ and $\|\lambda \ell\| \le 1$.

For every one-bounded function $v(d)$, the column

$$F(z) = \mathbf{1}_E v(d) P(d) U_j(z) \phi_j(z)$$

is active, has norm at most one, and satisfies $F(T_h z) = \sigma_h(z) F(z)$. On E this follows from the factorization; off E both sides are zero, because E is invariant. Lemma 3.2 yields

$$D(c_j v \phi_j)^2 \ge \lambda^3 \left\| \omega_j v \right\|_2^2. \tag{94}$$

In particular, if v is any restriction of m_j to a set of diagonals and $\|\omega_j v\|_2 \ge s/2$, then

$$D(c_j v \phi_j)^8 \ge (\lambda^3 s^2/4)^4 \ge \alpha_0.$$

This is exactly the robust nondegeneracy hypothesis of Proposition 6.3.

Apply that proposition to each summand of (93), with $K_e = K_s = K_*$. Each summand has an approximation of L^2 error at most s by a finite linear combination of tests, of total absolute coefficient at most

$$J_* = 4K_*^2\alpha^{-1}p^{2J}S, \qquad S = K_*^{2^{\lceil e^{-2}\rceil}}, \qquad J \le 10^3 e^{-4}\log(2K_*). \tag{95}$$

Every quantity in this expression has already been bounded. Since $p \ge 3$, equations (87) and (92) imply

$$\log_p J_* \le (C/\delta)^C. \tag{96}$$

In particular, the appearance of large partitions creates just one exponential: their logarithms are polynomial, and their powers in (95) are also polynomial.

7.3 Recovering the correlation

Let $\mathcal{S}$ be the sum of the r approximations just obtained. Its coefficient cost is at most rJ_*. Using $\|\lambda\ell u\|_\infty \le 1$, $\mathbb{P}(E^c) < \eta$, and the L^2 bounds for the individual summands, we have

$$\|\lambda\ell u - \mathcal{S}\|_1 \le \eta + rs \le 2rs = \lambda_-\delta/50 \le \lambda\delta/4.$$

Here $\eta = s^2 \le rs$. The initial correlation (4), multiplied by λ, now gives

$$\operatorname{Re}\mathbb{E}\overline{f}\,\mathcal{S} \ge \lambda\delta/2.$$

Every test in the expansion has absolute correlation with $\overline{f}$ at most $\tau(f)$, since the class is closed under conjugation. Therefore

$$\tau(f) \ge \frac{\lambda\delta}{2rJ_*}.$$

Equation (96) absorbs the polynomial prefactor and proves $\tau(f) \ge p^{-(C/\delta)^C}$. Returning from (a,b) to (x,y) gives the stated test in the original coordinates. The phase and alternating-form reductions in the introduction complete Theorem 1.1. □

Remark 7.1 (Where the quantitative bound is determined). The three estimates in Section 4 keep the matrix stage polynomial: the detector loses only powers of the matrix dimension, the covariance potential selects polynomially many fields, and the small-gap bound is linear in the requested gap threshold. All subsequent accuracies are fixed powers of these parameters. The scalar Fourier argument contributes a subspace of polynomial codimension, and hence a factor p to a polynomial power. These are the only ingredients needed for the single-exponential form of the final answer.

A Stability when the cube average is near one

This section identifies the functions that almost attain equality in the directional cube average. We work first in the coordinates (x,y), with directions $(h,0)$, $(0,k)$, and $(t,-t)$. All averages and L^2 norms are normalized.

Theorem A.1 (Near equality). *Let $G = \mathbb{F}_p^n$, where p is odd, and let $f : G^2 \to \mathbb{C}$ be one-bounded. There are phases $u, v, w : G \to \mathbb{T}$ and a bilinear form $B : G \times G \to \mathbb{F}_p$ such that the phase*

$$s(x,y) = u(x)v(y)w(x+y)e_p(B(x,y))$$

satisfies

$$\operatorname{Re}\mathbb{E}_{x,y} f(x,y)\overline{s(x,y)} \ge 1 - 4096(1-\beta(f)). \tag{97}$$

Consequently

$$\tau(f) \geq 1 - 4096(1 - \beta(f)), \qquad \|f - s\|_2^2 \leq 8192(1 - \beta(f)).$$

The relevant algebraic identity is the scalar two-cocycle equation

$$F(a, b)F(a + b, c) = F(b, c)F(a, b + c). \tag{98}$$

For example, both a bilinear phase and a function $v(a + b)/(v(a)v(b))$ satisfy this equation. If s is a test phase and X, Y are fixed, its normalized rectangle

$$\frac{s(X + a, Y + b)s(X, Y)}{s(X + a, Y)s(X, Y + b)}$$

is a product of precisely these two kinds of cocycle. The proof will first obtain an approximate cocycle from a nearly extremal function, then correct it and read off its factors.

A.1 Correcting an approximate scalar cocycle

For any phase $F : G^2 \to \mathbb{T}$, define its multiplicative defect by

$$(\partial F)(a, b, c) = \frac{F(b, c)F(a, b + c)}{F(a + b, c)F(a, b)}.$$

Thus $\partial F = 1$ is equivalent to (98).

Lemma A.2 (Cocycle correction)**.** *Let G be a finite abelian group, and let $F : G^2 \to \mathbb{T}$. If*

$$q^2 = \mathbb{E}_{a,b,c}|1 - \partial F(a, b, c)|^2, \qquad q \leq 1/32,$$

then there is an exact cocycle $F_0 : G^2 \to \mathbb{T}$ with

$$\|F - F_0\|_2 \leq 44q.$$

Proof. **Weighted shifts.** Put $N = |G|$. For matrices indexed by G, use the normalized Hilbert–Schmidt norm

$$\|A\|_{\mathrm{nHS}}^2 = N^{-1} \operatorname{Tr}(A^* A).$$

Every unitary has norm one in this normalization, and multiplication on either side by a unitary preserves the norm. For each $a \in G$, define a weighted translation by

$$(U_a v)(x) = F(x, a)v(x + a).$$

This is a unitary operator on $\mathbb{C}^G$. The two operators $U_a U_b$ and $F(a, b)U_{a+b}$ have their only nonzero entry in row x at column $x + a + b$. The quotient of those entries is $(\partial F(x, a, b))^{-1}$. Therefore

$$\left(\mathbb{E}_{a,b} \|U_a U_b - F(a, b)U_{a+b}\|_{\mathrm{nHS}}^2\right)^{1/2} = q. \tag{99}$$

From an average defect to a uniform projective defect. For unitaries A, B, define

$$d(A, B) = \min_{\zeta \in \mathbb{T}} \|A - \zeta B\|_{\mathrm{nHS}}.$$

This is a metric after identifying scalar multiples. The triangle inequality follows by multiplying the two optimizing phases, and the metric is invariant under unitary multiplication on either side. It measures how close two operators are when their scalar phases may be adjusted.

For a fixed increment a, the translated quotients

$$L_a(t) = U_{a+t}U_t^{-1}$$

give alternative candidates for U_a. Inserting U_{t-s} on the right of both factors in $U_{a+s}U_s^{-1}$ gives

$$\begin{aligned} d(L_a(s), L_a(t)) \le\ & d(U_{a+s}U_{t-s}, U_{a+t}) \\ & + d(U_sU_{t-s}, U_t). \end{aligned}$$

When s, t are independent and uniform, each of $(a+s, t-s)$ and $(s, t-s)$ is uniform on G^2. By (99) and the L^2 triangle inequality,

$$\left(\mathbb{E}_{s,t} d(L_a(s), L_a(t))^2\right)^{1/2} \le 2q.$$

For each a, choose s_a such that

$$\left(\mathbb{E}_t d(L_a(s_a), L_a(t))^2\right)^{1/2} \le 2q.$$

Choose the scalar multiple V_a of $L_a(s_a)$ closest to U_a. For every t,

$$\|V_a - U_a\|_{\mathrm{nHS}} = d(V_a, U_a) \le d(V_a, L_a(t)) + d(U_{a+t}, U_aU_t).$$

The first term has root mean square at most $2q$ for each a; the second has root mean square at most q after averaging a, t. It follows that

$$\left(\mathbb{E}_a \|V_a - U_a\|_{\mathrm{nHS}}^2\right)^{1/2} \le 3q. \tag{100}$$

For fixed a, b, the translated quotients satisfy the exact identity

$$L_a(b+t)L_b(t) = L_{a+b}(t).$$

Compare its three terms to V_a, V_b, V_{a+b} in the metric d, and average over t. Each replacement has mean at most $2q$. Thus

$$d(V_aV_b, V_{a+b}) \le 6q \qquad \text{for every } a, b.$$

Choose an optimizing scalar $\lambda(a, b) \in \mathbb{T}$. Then

$$\|V_aV_b - \lambda(a,b)V_{a+b}\|_{\mathrm{nHS}} \le 6q \qquad \text{for every } a, b. \tag{101}$$

The phases λ remain close to the original F:

$$\|\lambda - F\|_2 \le 16q. \tag{102}$$

To see this, compare $\lambda(a,b)V_{a+b}$ to $F(a,b)V_{a+b}$ through V_aV_b, U_aU_b, and $F(a,b)U_{a+b}$. The four successive errors have root mean square at most $6q$, $6q$, q, and $3q$, respectively, by (99), (100), and (101). The distance between the first and last terms is $|\lambda(a,b) - F(a,b)|$, proving (102).

Correcting the remaining scalar defect. Compute $V_aV_bV_c$ using either order of multiplication. Two applications of (101) in each order show that

$$|1 - \partial\lambda(a,b,c)| \le 24q \le 3/4 \qquad \text{for every } a, b, c.$$

Let $r(a,b,c) \in (-\pi,\pi]$ be the argument of $\partial\lambda(a,b,c)$. For $0 \le u \le 3/4$,

$$2\arcsin(u/2) \le 7u/6;$$

for instance, the derivative of the left side is at most $8/\sqrt{55} < 7/6$ on that interval. Consequently

$$\|r\|_\infty \le 28q. \tag{103}$$

The cancellation identity $\partial^2\lambda = 1$ means that

$$\begin{aligned} R(a,b,c,t) = r(b,c,t) - r(a+b,c,t) + r(a,b+c,t) \\ - r(a,b,c+t) + r(a,b,c) \end{aligned}$$

is an integer multiple of 2π. This can also be checked directly by substituting $e^{ir} = \partial\lambda$: every factor of λ cancels. But $|R| \le 140q < 2\pi$, so $R = 0$. Define the real two-variable function

$$w(a,b) = -\mathbb{E}_t r(a,b,t).$$

Averaging the identity $R = 0$ in t, and using that $c + t$ is uniform, gives

$$r(a,b,c) = w(b,c) - w(a+b,c) + w(a,b+c) - w(a,b).$$

Hence

$$F_0(a,b) = \lambda(a,b)e^{-iw(a,b)}$$

satisfies $\partial F_0 = 1$. By (103), $\|F_0 - \lambda\|_2 \le 28q$. Combining this with (102) gives

$$\|F - F_0\|_2 \le \|F - \lambda\|_2 + \|\lambda - F_0\|_2 \le 16q + 28q = 44q. \qquad \square$$

A.2 The form of an exact cocycle

Lemma A.3 (Scalar cocycles on a vector space)**.** *Let $G = \mathbb{F}_p^n$, where p is odd. Every phase $F_0 : G^2 \to \mathbb{T}$ satisfying (98) has the form*

$$F_0(a,b) = c\, e_p(B(a,b))\, \frac{v(a+b)}{v(a)v(b)}, \tag{104}$$

where $c \in \mathbb{T}$, $v : G \to \mathbb{T}$, and $B : G \times G \to \mathbb{F}_p$ is bilinear. The form B may be chosen alternating.

Proof. Taking the middle argument to be zero in the cocycle identity gives $F_0(a,0) = F_0(0,b) = c$, independent of a, b. Put $\sigma = F_0/c$, so that $\sigma(a,0) = \sigma(0,a) = 1$. On $\mathbb{T} \times G$, define

$$(z,a)(w,b) = (zw\sigma(a,b), a+b).$$

Associativity is precisely the cocycle equation. The identity is $(1,0)$, and the inverse of (z,a) is

$$(z^{-1}\sigma(a,-a)^{-1}, -a).$$

This is both a left and a right inverse because the cocycle equation at $(a,-a,a)$ gives $\sigma(a,-a) = \sigma(-a,a)$. Thus we have a group whose elements $(z,0)$ are central.

Set

$$\kappa(a,b) = \frac{\sigma(a,b)}{\sigma(b,a)}.$$

If P, Q have second coordinates a, b, then

$$PQ = (\kappa(a,b), 0)QP.$$

For a third element of second coordinate c, move it past Q and then past P. Since the scalar factors are central, comparison with moving it past PQ at once gives

$$\kappa(a+b,c) = \kappa(a,c)\kappa(b,c).$$

Also $\kappa(b,a) = \kappa(a,b)^{-1}$, so the same multiplicativity holds in the second variable. Since $pa = 0$, every $\kappa(a,b)$ is a p-th root of unity. We may therefore write

$$\kappa(a,b) = e_p(A(a,b)),$$

where A is additive in both variables, hence bilinear over $\mathbb{F}_p$. The equality $\kappa(a,a) = 1$ makes A alternating.

Because p is odd, $B = A/2$ is defined, and the bilinear phase $e_p(B(a,b))$ has alternating quotient $\kappa(a,b)$. Dividing σ by this phase gives a symmetric normalized cocycle σ_1. Repeat the group construction with σ_1. The resulting group is abelian.

Choose a basis $e_1, \dots, e_n$ of G. The p-th power of $(1, e_i)$ is $(z_i, 0)$ for some $z_i \in \mathbb{T}$. Choose $u_i \in \mathbb{T}$ with $u_i^p = z_i^{-1}$. Then $(u_i, e_i)^p = (1, 0)$. The adjusted lifts commute, so their products define a homomorphism from G to this group, of the form

$$a \longmapsto (v(a), a).$$

The exponents of the basis lifts can be read modulo p, so this map is well-defined. Its homomorphism property says

$$v(a)v(b)\sigma_1(a,b) = v(a+b).$$

Restoring $e_p(B)$ and c gives (104). □

A.3 From a near-extremal cube to an approximate cocycle

Proof of Theorem A.1. Put $\epsilon = 1 - \beta(f)$. If $\epsilon \geq 1/4096$, choose a constant phase s for which $\operatorname{Re} \mathbb{E} f\overline{s} \geq 0$; this already proves (97). We may therefore assume $\epsilon < 1/4096$.

Write $f = |f|g$, where g is a phase everywhere, assigning any phase at a zero of f. On a random directional cube, let C be the alternating product of the eight values of g, and let m be the product of the eight corresponding values of $|f|$. The cube product of f is mC. Since $0 \leq m \leq 1$ and $\operatorname{Re} C \leq 1$,

$$\mathbb{E}(1 - |f|) \leq \mathbb{E}(1 - m) \leq \epsilon, \qquad \mathbb{E}(1 - \operatorname{Re} C) \leq \mathbb{E}(1 - \operatorname{Re}(mC)) + \mathbb{E}(1 - m) \leq 2\epsilon.$$

For each choice of X, Y, form

$$F(a,b) = \frac{g(X+a, Y+b)g(X,Y)}{g(X+a,Y)g(X,Y+b)}.$$

Direct cancellation gives

$$\partial F(a,b,c) = \Delta_{(a,0)}\Delta_{(0,c)}\Delta_{(b,-b)}g(X, Y+b). \tag{105}$$

Indeed, the four factors in its numerator are the values at

$$(X, Y+b), \quad (X+a, Y+b+c), \quad (X+a+b, Y), \quad (X+b, Y+c),$$

and the denominator contains the other four cube vertices. When X, Y, a, b, c are uniform and independent, so are $X, Y+b, a, c, b$. Since $|1-C|^2 = 2(1-\operatorname{Re} C)$, averaging (105) over X, Y gives a choice of those two points for which

$$q^2 := \mathbb{E}_{a,b,c}|1-\partial F(a,b,c)|^2 \le 4\epsilon, \qquad q < 1/32.$$

By Lemma A.2, there is an exact cocycle F_0 with $\|F-F_0\|_2 \le 44q$. Lemma A.3 describes its factors.

Undo the rectangle normalization by defining

$$s(X+a, Y+b) = F_0(a,b)\frac{g(X+a,Y)g(X,Y+b)}{g(X,Y)}.$$

All the normalizing factors are phases, so

$$\|g-s\|_2 = \|F-F_0\|_2 \le 44q.$$

Moreover, s is a test phase. To check this explicitly, substitute (104), put $x = X+a, y = Y+b$, and note that $v(a+b)$ becomes a function of $x+y$. The bilinear term expands as

$$B(x-X, y-Y) = B(x,y) - B(x,Y) - B(X,y) + B(X,Y).$$

The last three terms, and the axis values of g, are absorbed into the one-variable phases and a constant.

For two phases, $\operatorname{Re}\mathbb{E} g\overline{s} = 1 - \frac{1}{2}\|g-s\|_2^2$. Consequently

$$\operatorname{Re}\mathbb{E} g\overline{s} \ge 1 - \tfrac{1}{2}(44q)^2 \ge 1 - 3872\epsilon.$$

Since $\mathbb{E}|f-g| = \mathbb{E}(1-|f|) \le \epsilon$, we obtain

$$\operatorname{Re}\mathbb{E} f\overline{s} \ge 1 - 3873\epsilon \ge 1 - 4096\epsilon.$$

The class of test phases is closed under conjugation, giving the asserted bound for $\tau(f)$. Finally,

$$\|f-s\|_2^2 = \mathbb{E}|f|^2 + 1 - 2\operatorname{Re}\mathbb{E} f\overline{s} \le 2\big(1 - \operatorname{Re}\mathbb{E} f\overline{s}\big) \le 8192\epsilon.$$

This argument also proves the last inequality in the previously discarded case $\epsilon \ge 1/4096$. □

Corollary A.4 (Restriction to a product of cosets)**.** *Let $H \le G$, and let*

$$C = (a_0 + H) \times (b_0 + H)$$

be a product of cosets in the reflected coordinates. For a one-bounded $g : C \to \mathbb{C}$, let $D_C(g)^8$ denote its cube average with both base coordinates uniform on their cosets and all three increments in H, in the horizontal, vertical, and simultaneous directions. There is a global test phase

$$s(a,b) = u(a)v(b)w(a-b)e_p(B(a,b))$$

such that

$$\|g-s\|_{2,C}^2 \le 8192(1 - D_C(g)^8).$$

Proof. Use the coordinates $x = a - a_0$, $y = b_0 - b$ on H^2. The simultaneous direction becomes $(h, -h)$, and the horizontal and vertical directions remain the corresponding coordinate directions, after an immaterial change of sign in a uniform increment. Theorem A.1 on H gives the required approximation there.

Extend its bilinear form to a form $\widetilde{B}$ on $G \times G$ by completing a basis of H to a basis of G. Substituting $x = a - a_0, y = b_0 - b$ gives the global bilinear term $B(a,b) = -\widetilde{B}(a,b)$, together with one-variable phases and a constant. The third factor is a function of $a-b$ on $a_0 - b_0 + H$. Extend the resulting three one-variable phases arbitrarily with modulus one outside their respective cosets. This produces a global test whose restriction has the stated error. □

B Choosing coordinates at a prescribed accuracy

We prove the remaining approximation assertion in Lemma 3.1. Only one accuracy is prescribed. This permits a finite random construction: take a moderate sample of projection columns, form Gaussian linear combinations, and orthonormalize them wherever they are well conditioned. The exceptional set is small enough that arbitrary bases may be used there.

Lemma B.1 (Coordinates at a prescribed accuracy)**.** *For the kernels A_z in (7) and any $0 < \zeta \le 1/2$, one can choose $\lambda \in [\delta/8, \delta/4]$ and the bases of H_z so that the row ℓ and the unitaries σ_h in (13) have the step approximations (6). In fact the cocycle error can be at most ζ in L^2_z for each h separately.*

Proof. We make the polynomial dependences explicit enough to keep track of the number of exponentials. Put

$$B = 1 + 8/\delta, \qquad V = 10^6(r+1)(B+1)\zeta^{-1}, \qquad v = \frac{\delta^2 V^{-12}}{1024}.$$

All constants and powers used below are absolute. Their size is immaterial; the point is that there are only polynomially many scalar quantities to approximate, and only polynomial inverse accuracies are required.

A cutoff avoiding the spectrum at one scale

Average λ uniformly over $[\delta/8, \delta/4]$. An eigenvalue lying within v of such a λ is greater than $\delta/16$; the trace bound permits at most $16/\delta$ such eigenvalues at any z. Each eigenvalue excludes a set of cutoffs of length at most $2v$. Consequently

$$\mathbb{E}_\lambda \mathbb{P}_z\big(\operatorname{spec}(A_z) \cap [\lambda - v, \lambda + v] \neq \varnothing\big) \le 256 v \delta^{-2} \le V^{-12}/4.$$

Choose one λ for which the exceptional measure is at most V^{-12}. Let χ vanish on $[0, \lambda - v]$, be linear from zero to one on $[\lambda - v, \lambda + v]$, and equal one thereafter. Outside this exceptional set,

$$\chi(A_z) = P_z. \tag{106}$$

A finite Gaussian choice of bases

Recall $c_j(z) = P_z(\,\cdot\,, j)$, the identities (11), and the bound $\|c_j(z)\|_2 \le B$. Take $J = \lceil V^{10} \rceil$ independent uniform indices $k_1, \dots, k_J \in G$. Independently take standard complex Gaussian variables γ_{jl}, $1 \le j \le J$, $1 \le l \le r$, with density $\pi^{-1} e^{-|w|^2}$ on $\mathbb{C}$. Set

$$\beta_{jl} = J^{-1/2} \gamma_{jl}, \qquad w_l(z) = \sum_{j=1}^{J} \beta_{jl} c_{k_j}(z).$$

Let $W(z) : \mathbb{C}^r \to H_z$ have columns $w_l(z)$, and let $W_0(z)$ consist of the first m_z columns.

For a fixed z, the empirical covariance satisfies

$$\mathbb{E} \left\| J^{-1} \sum_{j=1}^{J} c_{k_j}(z) c_{k_j}(z)^* - I_{H_z} \right\|_{\mathrm{HS}}^2 \le B^4/J.$$

Here the Hilbert–Schmidt norm is computed on H_z, with no normalization by its dimension. This estimate is the variance identity for a sample mean; each summand has Hilbert–Schmidt

norm at most B^2. Except with probability at most $4B^4/J$, the covariance is at least $\frac{1}{2}I_{H_z}$. Conditional on these indices, the columns of $W_0(z)$ are independent Gaussian vectors with that covariance.

Let $a_* = V^{-10}$. If $W_0(z)$ has least singular value less than a_*, there is a unit vector of coefficients whose image has norm less than a_*. One coefficient has magnitude at least $r^{-1/2}$; the corresponding column is within $\sqrt{r}\,a_*$ of the span of the other columns. Conditional on the other columns, its component in a unit normal direction is a complex Gaussian of variance at least $1/2$. Its density is at most $2/\pi$, so its probability of lying in a disk of radius $\sqrt{r}\,a_*$ is at most $2ra_*^2$. A union bound gives

$$\mathbb{P}\big(s_{\min}(W_0(z)) < a_*\big) \le 4B^4/J + 2r^2a_*^2 \le 6V^{-6}.$$

At rank zero there is no exceptional event. Averaging in z and applying Markov shows that, with probability at least $1 - 6V^{-3}$, the set

$$\mathcal{B} = \{z : m_z > 0,\ s_{\min}(W_0(z)) < a_*\}$$

has measure at most V^{-3}.

Also $\mathbb{E}|\gamma_{jl}| \le 1$, so $\mathbb{E}\sum_j |\beta_{jl}| \le \sqrt{J}$. With $U = 100r\sqrt{J}$, the probability that any column violates $\sum_j |\beta_{jl}| \le U$ is at most $1/100$. Fix one realization with both conclusions. We will use the crude bound $U \le V^7$. Outside $\mathcal{B}$, define the orthonormal basis by the polar formula

$$\mathcal{U}_z = W_0(z)\big(W_0(z)^* W_0(z)\big)^{-1/2}, \tag{107}$$

and append zero columns through index r. On $\mathcal{B}$, choose any orthonormal basis; at rank zero take $\mathcal{U}_z = 0$. These are exact bases everywhere, so all covariance and cocycle identities in Section 3 remain exact.

Recovering projection entries from a finite matrix

Although P_z acts on $L^2(G)$, we can approximate all its kernel entries simultaneously using polynomially many entries of A_z. Define

$$\mathcal{B}_z g = \mathbb{E}_i \xi_{z,i} g(i), \qquad K_z = \mathcal{B}_z \mathcal{B}_z^* = \mathbb{E}_i \xi_{z,i} \xi_{z,i}^*.$$

Then $\mathcal{B}_z^* \mathcal{B}_z = A_z$. Put

$$\psi(s) = \chi(s)/s, \qquad \phi(s) = \chi(s)/s^2 \quad (0 \le s \le 1),$$

with both functions zero near zero. Since $\lambda - v \ge \delta/16$, their suprema and Lipschitz constants are at most

$$L_0 = 10^4 \delta^{-2} v^{-1} \le V^{20}.$$

The singular-value decomposition gives $\chi(\mathcal{B}_z^* \mathcal{B}_z) = \mathcal{B}_z^* \psi(\mathcal{B}_z \mathcal{B}_z^*) \mathcal{B}_z$. In kernel notation,

$$\chi(A_z)(i,j) = \langle \xi_{z,i}, \psi(K_z) \xi_{z,j} \rangle\,. \tag{108}$$

We use an elementary Hilbert–Schmidt Lipschitz estimate: if D, D' are self-adjoint and q is Lipschitz on their spectra, then

$$\|q(D) - q(D')\|_{\mathrm{HS}} \le \mathrm{Lip}(q)\,\|D - D'\|_{\mathrm{HS}}\,. \tag{109}$$

Indeed, in an eigenbasis of D on the left and an eigenbasis of D' on the right, the corresponding entries are $(q(s) - q(t))\,\langle u, v\rangle$ and $(s - t)\,\langle u, v\rangle$. Squaring and summing proves the estimate.

Take $M = \lceil L_0^2 V^{260} \rceil$ independent uniform indices $n_1, \dots, n_M$, and set

$$K'_z = M^{-1} \sum_{\ell=1}^{M} \xi_{z,n_\ell} \xi_{z,n_\ell}^*.$$

Since $\|\xi_{z,i}\| \le 1$, the expected squared Hilbert–Schmidt error is at most $1/M$. Fix indices for which this is true after averaging over z. Outside a set of measure at most V^{-20},

$$\|K_z - K'_z\|_{\mathrm{HS}} \le V^{-120}/L_0.$$

By (109), replacing K_z by K'_z in (108) costs at most V^{-120}, simultaneously for all i, j.

There is a finite formula for this replacement. Let $T : \mathbb{C}^M \to L^2(G^2)$ have columns $\xi_{z,n_\ell}/\sqrt{M}$, and put

$$D = T^*T = \big(A_z(n_\ell, n_{\ell'})/M\big)_{\ell,\ell'}, \qquad a_i = T^*\xi_{z,i} = \big(A_z(n_\ell, i)/\sqrt{M}\big)_\ell.$$

Then $TT^* = K'_z$, $\|a_i\| \le 1$, and another application of the singular-value decomposition yields

$$\langle \xi_{z,i}, \psi(K'_z)\xi_{z,j} \rangle = a_i^* \phi(D) a_j. \tag{110}$$

Together with (106), this approximates $P_z(i, j)$ to error V^{-120} off a set of measure at most $V^{-12} + V^{-20}$.

We record its dependence on the kernel entries. If each required entry of A_z changes by at most τ_0 in real and imaginary parts between two points, then D changes by at most $2\tau_0$ in Hilbert–Schmidt norm and each a_i by at most $2\tau_0$ in Euclidean norm. Using (109) for ϕ and changing the three factors in (110), its value changes by at most $6L_0\tau_0$.

Recovering the row and the change of coordinates

Before orthonormalizing, the needed arrays are

$$t_z = \big(w_l(z)(0)\big)_{l=1}^{r}, \qquad G_z = W(z)^*W(z), \qquad C_{h,z} = W(T_h z)^* S_h W(z).$$

Their entries are fixed linear combinations of, respectively,

$$P_z(0, k_j), \qquad P_z(k_j, k_{j'}), \qquad P_z(k_j + h, k_{j'}). \tag{111}$$

For the last identity, note that $c_k(T_h z) = S_h c_{k+h}(z)$. The sums of the absolute coefficients in these expressions are at most U, U^2, and U^2. In particular

$$\|t_z\| \le \sqrt{r}\, U B^2 \le V^{10}, \qquad \|C_{h,z}\|_{\mathrm{op}} \le r U^2 B^2 \le V^{17}.$$

On the complement of $\mathcal{B}$, the nonzero eigenvalues of G_z are at least $a_*^2 = V^{-20}$, because $WW^* \ge W_0 W_0^* \ge a_*^2 I_{H_z}$. Thus sufficiently close Gram matrices at two such points have the same rank. At their common rank m, restrict to the first m coordinates and write

$$\Theta_z = (W_0(z)^* W_0(z))^{-1/2}.$$

The row is $t_z \Theta_z$ on those coordinates, and the active block of the transition is

$$\Theta_{T_h z} C_{h,z} \Theta_z.$$

The inactive block is the identity.

For positive $D, D' \geq a_*^2 I$, the integral representation of the inverse square root and the resolvent identity give

$$\left\| D^{-1/2} - D'^{-1/2} \right\|_{\mathrm{op}} \leq \frac{\|D - D'\|_{\mathrm{op}}}{2a_*^3}. \tag{112}$$

Indeed,

$$D^{-1/2} = \frac{1}{\pi} \int_0^\infty t^{-1/2} (D + tI)^{-1} \, dt, \qquad \frac{1}{\pi} \int_0^\infty \frac{t^{-1/2}}{(t + a_*^2)^2} \, dt = \frac{1}{2a_*^3}.$$

The integral identity follows first for positive scalars by the substitution $t = a_*^2 u$, and then for matrices by diagonalization.

Set

$$\tau_0 = \frac{V^{-140}}{100 L_0}.$$

Consider two points outside all the exceptional sets above. For a transition require the same at both translated points. If the real and imaginary parts of all the kernel entries used in the finite formulas differ by at most τ_0, then the projection entries in (111) differ by at most $3V^{-120}$. Hence the arrays t, G, C differ in their respective Euclidean or Frobenius norms by at most V^{-103}. Their ranks agree, since $V^{-103} < a_*^2$. We have $\|\Theta_z\|_{\mathrm{op}} \leq V^{10}$, and (112) gives $\|\Theta_z - \Theta_{z'}\|_{\mathrm{op}} \leq V^{-73}$. Changing one factor at a time in $t\Theta$ and $\Theta_h C \Theta$, and using the bounds just obtained, shows that both the row and the full transition matrix vary by at most V^{-40} in the required norm. For the transition, conversion from operator norm to Frobenius norm costs at most $\sqrt{r}$, which is absorbed by the stated bound. These estimates also cover rank zero, where the row is zero and the transition is the identity.

The rectangular partitions

For a fixed h, all indices needed in the finite formulas belong to

$$n_\ell, \ k_j, \ k_j + h, \ 0.$$

At most

$$N_{\mathrm{ent}} = 10(M + 2J + 1)^2 \leq V^{605}$$

real scalar kernel entries at z or at $T_h z$ suffice. Every such complex kernel entry is an average of products of one-bounded functions of the two separate vertices, by (7). Lemma 2.3 therefore approximates it to L^2 error d_1 with $\log L \leq C d_1^{-2} \log(C/d_1)$. Choose

$$d_1 = \frac{\tau_0}{10\sqrt{N_{\mathrm{ent}}}\, V^2}.$$

Refine the partitions over all N_{ent} entries. By Chebyshev and a union bound, the proportion of points at which any entry has error greater than $\tau_0/4$ is at most

$$16 N_{\mathrm{ent}} d_1^2 / \tau_0^2 \leq V^{-4}.$$

Two remaining points in the same cell have all required kernel entries within τ_0. Including the exceptional sets for the basis and projection at both z and $T_h z$, the total exceptional measure is at most

$$2(V^{-3} + V^{-12} + V^{-20}) + V^{-4} \leq 5V^{-3}.$$

On a cell containing a nonexceptional point, take the actual row or transition matrix at one such point as the approximating value. On any other cell use the zero row or the identity matrix. This preserves the row bound and unitarity. The error is at most V^{-40} at

nonexceptional points; elsewhere it is at most $2\lambda^{-1}$ for the row and $2\sqrt{r}$ for a unitary matrix. Consequently either squared L^2 error is at most

$$V^{-80} + 20\max(\lambda^{-2}, r)V^{-3} \le \zeta^2.$$

For the row, use only the entries at z, so that its partitions are independent of h. For each transition the same estimates hold separately, with partitions allowed to depend on h.

Finally, $J \le V^{11}$, $L_0 \le V^{20}$, $M \le V^{301}$, $\tau_0^{-1} \le V^{162}$, and $d_1^{-1} \le V^{470}$. Thus the logarithm of the total number of classes is bounded by

$$CN_{\mathrm{ent}} d_1^{-2} \log(C/d_1) \le V^{2000}.$$

Since V is polynomial in δ^{-1}, ζ^{-1}, this is the claimed bound (6). □

C Two elementary matrix inequalities

The arguments below use the ordinary, unnormalized Frobenius norm $\|M\|_{\mathrm{F}}^2 = \mathrm{Tr}(M^*M)$, and the trace norm $\|M\|_1$, the sum of the singular values of M. The following estimate will allow us to transfer an approximate symmetry of row inner products to an approximate symmetry of the rows themselves. We include the underlying square-root inequality.

Lemma C.1 (Square roots of positive matrices)**.** *If M, N are positive semidefinite matrices of the same size, then*

$$\left\| M^{1/2} - N^{1/2} \right\|_{\mathrm{F}}^2 \le \|M - N\|_1 .$$

Proof. Put $A = M^{1/2}$, $B = N^{1/2}$, and $S = A - B$. In an orthonormal eigenbasis of S, with eigenvalues s_j, positivity of A, B gives

$$(A + B)_{jj} \ge |A_{jj} - B_{jj}| = |s_j|.$$

Consequently

$$\mathrm{Tr}(S^2) \le \mathrm{Tr}\big(|S|(A + B)\big).$$

The identity $M - N = AS + SB$, together with cyclicity of trace, gives

$$\mathrm{Tr}\big(|S|(A + B)\big) = \mathrm{Tr}\big(\mathrm{sign}(S)(M - N)\big) \le \|M - N\|_1 .$$

For the last inequality, diagonalize $M - N$ and use that every diagonal entry of the Hermitian contraction $\mathrm{sign}(S)$ lies in $[-1, 1]$. Since $\mathrm{Tr}(S^2) = \|A - B\|_{\mathrm{F}}^2$, this proves the claim. □

Lemma C.2 (Transferring a symmetry to a factor)**.** *Let Z be an m-by-r complex matrix, let $\rho = ZZ^*$, and let U be an m-by-m unitary matrix. Then some r-by-r unitary matrix D satisfies*

$$\|UZ - ZD\|_{\mathrm{F}}^2 \le \|U\rho U^* - \rho\|_1 .$$

Proof. Expanding the square and maximizing the resulting trace pairing over unitary D gives

$$\min_D \|UZ - ZD\|_{\mathrm{F}}^2 = 2\,\mathrm{Tr}\,\rho - 2\,\|Z^*UZ\|_1 . \tag{113}$$

Indeed, for any square matrix A, singular-value decomposition shows $\max_{D\ \mathrm{unitary}} \mathrm{Re}\,\mathrm{Tr}(AD) = \|A\|_1$.

Write $Z = \rho^{1/2} V$ in polar form. The restriction of the partial isometry V is an isometry from $(\ker Z)^\perp$ onto $\operatorname{Ran} \rho$. The matrix $\rho^{1/2} U \rho^{1/2}$ is supported on $\operatorname{Ran} \rho$ on both sides. Thus multiplication by V^* and V preserves its nonzero singular values, and

$$\|Z^* U Z\|_1 = \left\|\rho^{1/2} U \rho^{1/2}\right\|_1 .$$

Set $\rho' = U \rho U^*$. Since $(\rho')^{1/2} = U \rho^{1/2} U^*$, right multiplication by U^* shows that the last trace norm is $\left\|\rho^{1/2} (\rho')^{1/2}\right\|_1$. It is at least the real part of the trace of that product. Since $\operatorname{Tr} \rho' = \operatorname{Tr} \rho$, equation (113) is therefore at most

$$\operatorname{Tr} \rho + \operatorname{Tr} \rho' - 2 \operatorname{Re} \operatorname{Tr}(\rho^{1/2} (\rho')^{1/2}) = \left\|\rho^{1/2} - (\rho')^{1/2}\right\|_{\mathrm{F}}^2 .$$

Apply Lemma C.1. □

When Z is obtained by putting one index of a tensor into its columns, right multiplication by D acts on that tensor index as D^{T}. The transpose is again unitary. Thus Lemma C.2 can be used directly to obtain an action on an omitted tensor index.

References

[1] T. Austin, *Partial difference equations over compact Abelian groups, I: modules of solutions*, arXiv:1305.7269, version 5 (2014). https://arxiv.org/abs/1305.7269.

[2] B. Green, *100 Open Problems*, https://people.maths.ox.ac.uk/greenbj/papers/open-problems.pdf.

[3] L. Milićević, *An inverse theorem for certain directional Gowers uniformity norms*, arXiv:2103.06354 (2021). https://arxiv.org/abs/2103.06354.

[4] S. Peluse, *Finite field models in arithmetic combinatorics—twenty years on*, arXiv:2312.08100 (2023), especially Problem 4.5. https://arxiv.org/abs/2312.08100.