Deterministic Almost-All Recovery for Square Four-Tensor Isomorphism

Jingchuan Ma[1], Yanhua Liu[1], and Qiaoyun Huang[1]

Fuzhou University Zhicheng College, Fuzhou 350002, China
kiciot@qq.com

Abstract. We give deterministic exact recovery for almost all square four-tensors over every finite field. For every $n \geq 2^{24}$ and prime power $q \geq 2$, an explicit $\mathrm{GL}_n(\mathbb{F}_q)^4$-invariant family contains a $1-O((\log n)^2/n)$ fraction of all q^{n^4} coefficient arrays, with an explicit failure bound uniform in q. Membership and isomorphism decision and recovery against every target are deterministic in $\mathrm{poly}(n, \log q)$ bit time, with four invertible factors over the original field. The source distribution is uniform without any conditioning on flattening ranks. The proof combines linked extension-field eigenmatrices, a submodule sieve under the joint law of a matrix and its inverse, and recovery from the kernels of singular original flattenings. Partial traces of matrix products provide a factorization-free algorithm in large characteristic. Complementary results give explicit positive source fractions in smaller dimensions, including a $1/q+O(q^{-2})$ family with deterministic recovery. The algorithms invert the coefficient-uniform square-tensor group action with the stated success probabilities. A deterministic exhaustive fallback gives total correctness and a polynomial running-time bound with high probability; no polynomial expected running time is claimed.

Keywords: Tensor isomorphism · Finite fields · Almost-all algorithms · Group-action cryptography · Matrix Kloosterman sums

1 Introduction

Tensor isomorphism asks whether changes of basis in the tensor factors carry one coefficient array to another. Its worst-case reductions connect matrix spaces, algebras, groups, and forms [13].

For square four-tensors with uniformly chosen coefficients, Li–Li–Qiao–Tao–Wang (LLQTW) [18, Theorem 1.3] give $\mathrm{poly}(n, q)$ recovery on a $1/\mathrm{poly}(q)$ source fraction. We give deterministic $\mathrm{poly}(n, \log q)$ exact decision and recovery against every target on a $1 - O((\log n)^2/n)$ fraction, for $n \geq 2^{24}$, uniformly over every finite field and without conditioning on flattening ranks.

Exact rejection requires a search that preserves every genuine transporter, not only verification of a candidate. Section 8.2 gives the resulting inversion guarantee for the coefficient-uniform group action.

A square four-tensor has four modes, each of dimension n. Write $k = \mathbb{F}_q$, where $q = p^e$, and let $V = (k^n)^{\otimes 4}$. The action is

$$T_{ijkl} = \sum_{a,b,c,d} (g_1)_{ia}(g_2)_{jb}(g_3)_{kc}(g_4)_{ld} A_{abcd}, \qquad g_i \in \mathrm{GL}_n(k). \tag{1}$$

Let U be uniform on *all* q^{n^4} arrays in V. Field elements are given in a valid polynomial-basis representation; arithmetic and equality in that representation have their usual deterministic polynomial bit costs. The input size is $n^4\lceil \log_2 q \rceil$ up to representation overhead. For $n \geq 2^{24}$, set

$$h = 2\lceil 4 \log_2 n \rceil, \qquad d_0 = 3 + \lceil \log_q(n^8) \rceil, \qquad D = \left\lfloor \frac{n}{100h} \right\rfloor, \tag{2}$$

and define the failure bound

$$\varepsilon(n,q) = \min\left\{1, \frac{d_0+1}{D+1} + 7n^{-4}\right\}. \tag{3}$$

Theorem 1 (Unconditional deterministic almost-all recovery). *For every integer $n \geq 2^{24}$ and every prime power $q \geq 2$, the explicit action-invariant family $\mathcal{K}(n,q)$ of Definition 3 satisfies*

$$\begin{aligned} \Pr_{A \leftarrow U}[A \notin \mathcal{K}(n,q)] &\leq \varepsilon(n,q) \\ &\leq \min\left\{1, \frac{6400(1+\log_2 n)^2}{n} + 7n^{-4}\right\}. \end{aligned} \tag{4}$$

There is one deterministic polynomial bit bound, uniform in n, q, for membership testing and for exact isomorphism decision/search from every $A \in \mathcal{K}(n,q)$ against every $T \in V$. On YES*, the algorithm returns four matrices in $\mathrm{GL}_n(k)$ satisfying* (1)*; otherwise it returns a correct* NO*. A source outside the family is reported as* SOURCE_FAIL*. One conservative bound is $O(n^{40}(1+\log_2 q)^{12})$ bit operations.*

The probability in Theorem 1 is over the source; every accepted source admits deterministic exact decision and recovery against every target. The term almost all refers to convergence of $\varepsilon(n,q)$ to zero as $n \to \infty$, uniformly over prime powers q. Since h, D are independent of q and d_0 is nonincreasing, $\varepsilon(n,q) \leq \varepsilon(n,2)$. Exact evaluation of (3) gives:

n	2^{24}	2^{28}	2^{32}	2^{40}
$\varepsilon(n,q) <$	0.225	0.0191	0.00155	9.43×10^{-6}

These strict upper bounds are rounded upwards; neither they nor the conservative runtime envelope are performance measurements.

Table 1 lists separate families with finite-parameter guarantees, including dimensions below the main cutoff. They have their own source-mass bounds, not an extension of (3); Section 7 states the results and their proofs.

Table 1. Lower bounds on the coefficient-uniform source fraction of each family. Only the first row uses $1-\varepsilon(n,q)$; the other rows give separate finite-parameter guarantees. Recognition and recovery take deterministic polynomial bit time (D) or Las Vegas expected polynomial bit time (LV); the LV algorithms are deterministic in characteristics $2, 3, 5$. The last row lists $\mathcal{S}_{\mathrm{det}} \subseteq \mathcal{S}_{\mathrm{CAS}}$, with $a(q) = \max\{75/(2048q), 1/q - 7/q^2\}$; recovery times follow that order.

Family	Parameters	Source fraction	Recognition / recovery
$\mathcal{K}$	$q \geq 2,\ n \geq 2^{24}$	$\geq 1-\varepsilon(n,q)$	D / D
$\mathcal{H}_{\mathrm{loc}}$	$q \geq 2,\ n \geq 4$	$\geq [1-(4n^4-n^2)/q]_+$	D / D
$\mathcal{R}_3$	$q \geq 5,\ n \geq 23$	$> 2^{-36}$	LV / LV
$\mathcal{R}_3$	$q \geq 5,\ n \geq 27$	$> 2^{-16}$	LV / LV
$\mathcal{R}_3$	$q = 4,\ n \geq 26$	$\geq c_4 > 2^{-15}$	D / D
$\mathcal{S}_{\mathrm{det}}, \mathcal{S}_{\mathrm{CAS}}$	$q \geq 5,\ n = 27$ or $n \geq 29$	$\geq a(q)$	D / (D, LV)

1.1 Mechanisms and proof structure

Set $m = n^2$ and use the three balanced flattenings

$$F_{ij,kl} = G_{ik,jl} = Q_{il,jk} = A_{ijkl}. \tag{5}$$

Each is marginally a uniform m-square matrix under U, but they are reshufflings of the same tensor. Whenever G is invertible, write

$$H = G^{-1}, \qquad W_{kl,ij} = H_{jl,ik}, \qquad X = FW. \tag{6}$$

The covariance of X supplies an extension-field eigenmatrix P on modes 12. Its linked image $Z = \mathrm{mat}(W \operatorname{vec} P)$ lies on modes 34 and retains the same scalar normalization. Two Frobenius ratios for each matrix reduce recovery to two simultaneous intertwiner systems. This two-sided construction does not need F^{-1}, and so it covers sources with singular F or Q.

Counting such eigenmatrices requires the joint law of H and H^{-1}. Matrix Kloosterman estimates give relative mixing for linearly independent contraction lists of length up to $\lfloor n/100 \rfloor$. We count embeddings of irreducible modules in a degree window and use an even subspace-lattice sieve of order $O(\log n)$. Its weights account for repeated primary factors. Polynomial zero bounds remove bad reshapes on both linked sides. This yields a conditional failure of $(d_0 + 1)/(D+1) + 9n^{-5}$ given G invertible. The window in (2) uses the available total-module-dimension budget $hD \leq \lfloor n/100 \rfloor$.

The singular-G region has positive mass at fixed q and requires a separate construction. Right and left kernel lines give a linear map $\mathrm{M}_n/kI \to \mathfrak{sl}_n$ and its trace adjoint. Quadratic and cubic cycle contractions recover modes $2, 3$; a large slice span recovers modes $1, 4$. Kernel reshapes may themselves be singular. The marked-kernel law reduces counting to a rectangular Hom space with an explicit factor for conditioning on rank. We control both the rectangular core and its extension columns to obtain a scalar centralizer for the full pair.

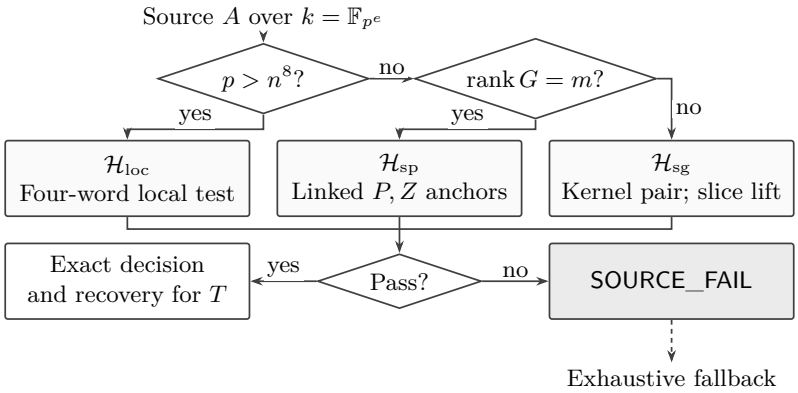


Fig. 1. Fast dispatcher for $\mathcal{K}(n, q)$ (Definition 3). Each leaf tests its full source predicate before any target is classified. Failure gives SOURCE_FAIL, not NO; the dashed arrow invokes the exhaustive fallback (Proposition 4) outside the fast family. There is no global full-rank gate on the original flattenings.

Finally, deterministic factorization has polynomial dependence on the characteristic p in the implementation used here. We employ it only for $p \le n^8$. For $p > n^8$, partial traces of X^2, Y^2, XY, YX give a factorization-free recovery predicate with source failure probability at most $(4n^4 - n^2)/q$. A characteristic- and parity-independent nonvanishing construction proves this estimate. These branches form the exhaustive dispatcher in Figure 1.

Proof roadmap. Definition 3 fixes the source predicates and parameters. For invertible G, Lemmas 3–4 yield Theorem 2 ($n \ge 1024$) and Corollary 1; the sieve with $hD \le \lfloor n/100 \rfloor$ and Lemma 5 then give Theorem 3. For singular G, Lemma 6, the marked law and projection (41)–(44), and Theorem 4 ($n \ge 4096$, $19n/20 \le s, t \le n$; Appendix C) give (54). Together with the rank tails, these bounds give (55) when $p \le n^8$. For $p > n^8$, Proposition 3 and Appendix D instead give $(4n^4 - n^2)/q < 4n^{-4}$, since $q \ge p$. Propositions 1–3 establish recovery; Lemma 7 and Section 6.2 establish the uniform bit bound. Section 6 combines these ingredients into Theorem 1 for $n \ge 2^{24}$.

1.2 Relation to earlier algorithms

LLQTW's flattening-kernel method [18, Theorem 1.3] is the closest prior algorithm in the same source model. Their Problem 1.6 asks for an average-case polynomial-time algorithm handling a $1/\Theta(q)$ source fraction. Theorem 1 exceeds this coverage for $n \ge 2^{24}$. At smaller dimensions, Theorem 6 resolves the $q \ge 5$, $n \in \{27\} \cup \{29, 30, \ldots\}$ regime with deterministic $\mathrm{poly}(n, \log q)$ recovery on a $\Theta(1/q)$ source fraction; Theorem 5 gives constant mass on the disjoint invertible-W branch. No all-parameter resolution is claimed.

Li–Qiao [19] and Brooksbank–Li–Qiao–Wilson [5] use random matrix-space structure to develop average-case algorithms for alternating matrix-space isometry. Related equivalence problems use different group actions and source laws. Grochow, Qiao, and Tang [14, Theorem 3.7] give deterministic $q^{O(n)}$ algorithms against arbitrary targets for almost all trilinear forms under a single simultaneous GL_n action. Narayanan–Qiao–Tang [21] use rank and corank invariants for matrix codes and alternating trilinear forms. Couvreur–Levrat [9] use dual-code contractions and hulls for matrix codes under independent left/right transformations. Graph-theoretic methods also occur in polynomial isomorphism [4]. Beullens [2] studies graph-theoretic heuristic algorithms for alternating trilinear equivalence under a single simultaneous factor, including concrete parameter analyses. These heuristic and experimental guarantees differ from our unconditional source bounds. The source spaces and actions differ from (1); worst-case TI reductions [13] do not identify their uniform measures. Our analytic inputs are the matrix Kloosterman estimates of Erdélyi–Tóth [11] and the purity and flag-cohomology results of Erdélyi–Sawin–Tóth [10]. Exact versions and the explicit dimension-dependent constant are specified in Section 3.

1.3 Notation and algorithmic conventions

Vectorization is row-major, so $(a \otimes b)\operatorname{vec} P = \operatorname{vec}(aPb^t)$. For $M \in \mathrm{M}_{n^2}(k)$, its *realignment* is

$$\mathscr{R}(M)_{(i,j),(k,l)} = M_{(i,k),(j,l)}.$$

This rearrangement converts Kronecker products into outer products [26]: $\mathscr{R}(B \otimes C) = \operatorname{vec}(B)\operatorname{vec}(C)^t$. Hence a nonzero M is a Kronecker product exactly when $\operatorname{rank}\mathscr{R}(M) = 1$. In that case, a nonzero pivot R_{ab} of $R = \mathscr{R}(M)$ gives $R = R_{:b}(R_{a:}/R_{ab})$; reshaping the two vectors recovers B, C. All operations take place in k.

All powers P^{q^j} are entrywise Frobenius powers. Matrix words such as X^2 use ordinary multiplication. Write $\mathrm{Cent}(C_1, \dots, C_s)$ for the simultaneous matrix centralizer. The trace pairing between $\mathrm{M}_n(k)/kI$ and $\mathfrak{sl}_n(k)$ is perfect in every characteristic: the trace pairing on M_n is nondegenerate and the annihilator of the trace-zero hyperplane is kI. No division by n is involved. The classes of the off-diagonal matrix units and E_{ii} for $0 \le i < n-1$ give a quotient basis in every characteristic. We distinguish SOURCE_FAIL, a statement about source recognition, from NO, an exact statement about a target for an accepted source.

2 Deterministic covariant recovery

We first give the pointwise algorithms. Their source probabilities are proved in Sections 4–6. We choose anchors from the source and consider their complete candidate images on the target.

2.1 Linked two-sided spectral anchors

For G invertible, use (6). If $T = (g_1, g_2, g_3, g_4)A$, put $L = g_1 \otimes g_2$ and $K = g_3 \otimes g_4$. Directly from the indices,

$$F_T = LF_AK^t, \quad W_T = K^{-t}W_AL^{-1}, \quad X_T = LX_AL^{-1}. \tag{7}$$

Indeed $G_T = (g_1 \otimes g_3)G_A(g_2 \otimes g_4)^t$, and its inverse reshuffles to the middle identity. Neither F nor W is inverted.

Definition 1 (Linked spectral predicate). *Fix a degree window $[d_0, D]$ with $d_0 \geq 3$. A source belongs to $\mathcal{H}_{\mathrm{sp}}$ if G is invertible and some irreducible factor f of χ_X in this window has the following properties. In $E = k[t]/f$, with $\alpha = t \bmod f$, the eigenspace $\ker(X - \alpha I)$ is a line. For a generator $v = \mathrm{vec}\, P$, put $z = Wv = \mathrm{vec}\, Z$ without independently normalizing Z. Require $P, Z \in \mathrm{GL}_n(E)$ and*

$$\begin{gathered} R_{P,j} = P^{q^j}P^{-1}, \quad R_{Z,j} = Z^{q^j}Z^{-1} \quad (j = 1, 2), \\ \mathrm{Cent}(R_{P,1}, R_{P,2}) = \mathrm{Cent}(R_{Z,1}, R_{Z,2}) = EI. \end{gathered} \tag{8}$$

Geometric multiplicity one suffices; no algebraic simplicity assumption is made. The property is invariant under changing the generator or the root of f: the first change scales both ratio pairs, and the second applies a common Frobenius automorphism.

Proposition 1 (Two-sided recovery). *The predicate in Definition 1 is action invariant and admits complete deterministic original-field decision/search in $\mathrm{poly}(p, n, \log q)$ bit time. It uses no inverse of F, Q, W.*

Proof. Recognize the source first and retain one passing f and its linked pair. For a target require G_T invertible and form the target anchors at the *same* α in the *same* E. The eigenspace, invertibility, and centralizer tests are necessary for an isomorphic target. Its chosen generator then has one common scale $c \in E^*$:

$$P_T = cg_1P_Ag_2^t, \qquad Z_T = cg_3^{-t}Z_Ag_4^{-1}. \tag{9}$$

Since the g_i are defined over k, the Frobenius ratios satisfy

$$R_{P_T,j} = \beta_j g_1 R_{P_A,j} g_1^{-1}, \quad R_{Z_T,j} = \beta_j g_3^{-t} R_{Z_A,j} g_3^t, \qquad \beta_j = c^{q^j - 1}. \tag{10}$$

For each j compare the two determinant ratios and, if they agree, list *all* roots in E of

$$b^n = \det R_{P_T,j} / \det R_{P_A,j} = \det R_{Z_T,j} / \det R_{Z_A,j}. \tag{11}$$

Each list has at most n distinct roots, also when $p \mid n$. For each of at most n^2 pairs (β_1, β_2) solve both systems

$$R_{P_T,j}D_1 = \beta_j D_1 R_{P_A,j}, \qquad R_{Z_T,j}D_3 = \beta_j D_3 R_{Z_A,j} \quad (j = 1, 2). \tag{12}$$

Retain only one-dimensional solution spaces with invertible generators. On the true pair the solution spaces identify with the source scalar centralizers, so $D_1 = sg_1$ and $D_3 = tg_3^{-t}$ for $s, t \neq 0$. Define

$$\widehat{g}_1 = D_1, \quad \widehat{g}_2 = (P_A^{-1} D_1^{-1} P_T)^t, \quad \widehat{g}_3 = D_3^{-t}, \quad \widehat{g}_4 = (Z_A^{-1} D_3^{-1} Z_T)^{-1}. \tag{13}$$

The four scales relative to the true factors are $(s, c/s, t^{-1}, t/c)$, whose product is one. All inverses in this formula are local n-square inverses. In particular the true branch reproduces T even with singular original F or Q.

To descend the candidate factors to k, divide each $\widehat{g}_i$ by its first nonzero coordinate and require every resulting entry to satisfy $x^q = x$. This tests precisely that its projective matrix line is defined over k. The true branch passes by the displayed scales. Let the four normalized k-matrices send A to B. A nonzero coefficient of B determines $\rho = T_j / B_j \in k$; require $\rho \neq 0$, absorb it into one factor, and check all original coefficients. Here $B \neq 0$ because G_A is invertible and the factors are invertible. The true branch passes, whereas a false branch can never produce an incorrect YES. Exhausting the complete lists therefore proves NO. The deterministic field operations and factorization costs are proved in Section 6.2.

2.2 Singular original flattenings

Let $1 \leq r = \operatorname{corank} G$ and choose nonzero right and left kernel *lines* kP, kV, with $G \operatorname{vec} P = 0$ and $(\operatorname{vec} V)^t G = 0$. Define

$$B([M]) = V^t \operatorname{mat}(G \operatorname{vec}(MP)), \qquad B : \mathrm{M}_n(k)/kI \longrightarrow \mathfrak{sl}_n(k). \tag{14}$$

This is well-defined: $B(I) = 0$ and $\operatorname{tr} B(M) = (\operatorname{vec} V)^t G \operatorname{vec}(MP) = 0$. Its trace adjoint is defined by $\operatorname{tr}(B^\dagger([N])M) = \operatorname{tr}(B([M])N)$ using the perfect quotient pairing. For the transpose permutation $S \operatorname{vec} M = \operatorname{vec}(M^t)$, its full coordinate matrix is SB^tS. Kernel reshapes P, V need not be invertible. For any such Hom map put $B_{ij} = B([E_{ij}])$ and

$$C_2(B) = \sum_{i,j} B_{ij} B_{ji}, \qquad C_3(B) = \sum_{i,j,l} B_{ij} B_{jl} B_{li}. \tag{15}$$

Definition 2 (Singular-source predicate). *For a nonnegative cutoff R, a source belongs to $\mathcal{H}_{\mathrm{sg}}$ when $1 \leq \operatorname{corank} G \leq R$, $\operatorname{rank} Q > m - n$, and some kernel-line pair has scalar common centralizer for both $(C_2(B), C_3(B))$ and $(C_2(B^\dagger), C_3(B^\dagger))$.*

Lemma 1 (Covariance of the dual contractions). *Set $h_1 = g_2^{-t}, h_2 = g_4^{-t}, h_3 = g_1^{-t}, h_4 = g_3^{-t}$. For matched representatives $h_1 P_A h_2^t = cP_T$ and $h_3 V_A h_4^t = dV_T$, let $\beta = (cd)^{-1}$. Then*

$$B_T = \beta \operatorname{Ad}(h_4) B_A \operatorname{Ad}(h_1)^{-1}, \qquad B_T^\dagger = \beta \operatorname{Ad}(h_1) B_A^\dagger \operatorname{Ad}(h_4)^{-1}, \tag{16}$$

$$C_j(B_T) = \beta^j h_4 C_j(B_A) h_4^{-1}, \qquad C_j(B_T^\dagger) = \beta^j h_1 C_j(B_A^\dagger) h_1^{-1} \quad (j = 2, 3). \tag{17}$$

Proof. Substitute $G_T = (h_3 \otimes h_4)^{-t} G_A (h_1 \otimes h_2)^{-1}$ in (14); the kernel equations cancel the right local factors without any reshape inverse. Taking the trace adjoint uses $(\mathrm{Ad}\, h)^\dagger = \mathrm{Ad}(h^{-1})$ and proves the second formula. The tensors $\sum E_{ij} \otimes E_{ji}$ and $\sum E_{ij} \otimes E_{jl} \otimes E_{li}$ are invariant under simultaneous conjugation: expanding the entries pairs each factor of h with one of h^{-1}. Applying B in each position and multiplying the outputs gives the last two formulas. Invertibility of B is never used.

Lemma 2 (Slice-span uniqueness). *If $\dim \mathcal{U} > n^2 - n$ for $\mathcal{U} \le \mathrm{M}_n(k)$ and $AM = MD$ for all $M \in \mathcal{U}$, then $A = D = \lambda I$.*

Proof. If A is nonscalar there is v such that v, Av are linearly independent. Otherwise all lines are invariant, and basis vectors and their pairwise sums force A scalar over any field. Choose $\ell(v) = 0, \ell(Av) = 1$. For $M = vw^t$, $\ell(AM - MD) = w^t$, so the Sylvester map has rank at least n. If $A = aI$ but $D \neq aI$, its rank is $n \operatorname{rank}(aI - D) \ge n$. Either case contradicts a kernel of dimension greater than $n^2 - n$.

Proposition 2 (Kernel-pair recovery). *For a source of corank r satisfying Definition 2, recognition and complete recovery against any target take deterministic time $q^{2r+1} \operatorname{poly}(n, \log q)$. All recovered factors lie in k, and no original flattening inverse is used.*

Proof. Enumerate source kernel-line pairs with first nonzero coordinate normalized to one, and retain one passing pair. On a target require the same rank of G and the slice gate. Enumerate every target kernel-line pair and every $\beta \in k^*$. For $j = 2, 3$ solve

$$C_j(B_T) D_v = \beta^j D_v C_j(B_A), \qquad C_j(B_T^\dagger) D_u = \beta^j D_u C_j(B_A^\dagger).$$

Retain only one-dimensional solution spaces with invertible generators. The true branch has $D_v = th_4$, $D_u = sh_1$. Apply D_u^t, D_v^t in modes $2, 3$ of T to obtain $T' = st(g_1, I, I, g_4)A$. For each j, k let $M_{jk} = (A_{ijkl})_{il}$ and $N_{jk} = (T'_{ijkl})_{il}$, and solve

$$LM_{jk} = N_{jk} J \quad \text{for all } j, k. \tag{18}$$

The true solution is $(L, J) = (stg_1, g_4^{-t})$. Relative to an invertible solution, any other solution gives a Sylvester pair on the slice span, whose dimension is $\operatorname{rank} Q > m - n$. Lemma 2 therefore makes the solution space a line. An invertible generator gives the candidate $(L, D_u^{-t}, D_v^{-t}, J^{-t})$. Scaling the generator scales L, J together and cancels in the action. Check every coefficient.

Every source line has a target image, both image lines occur in the enumeration, and their common scale occurs in k^*. The true branch survives all systems. Consequently exhaustion proves NO, separately from the coefficient check that proves soundness. There are $(q^r - 1)/(q - 1) < q^r$ lines per side, at most q^{2r} source pairs and q^{2r+1} target branches.

2.3 Recovery from partial traces in large characteristic

On $\mathcal{S}_{\text{full}} = \{\det F \det G \det Q \neq 0\}$ define, in addition to X,

$$Y_{ij,ab} = \sum_{k,l} A_{ijkl}(Q^{-1})_{bk,al}.$$

Both X, Y conjugate by $g_1 \otimes g_2$. For the ordered words $\mathcal{W} = (X^2, Y^2, XY, YX)$, let $\mathcal{C}_1$ take partial trace over mode 2 and $\mathcal{C}_2$ over mode 1. Define $\mathcal{H}_{\text{loc}} \subseteq \mathcal{S}_{\text{full}}$ by $\text{Cent}(\mathcal{C}_i) = kI$ for $i = 1, 2$. These are algebraic partial traces without division by n, and they conjugate by g_i.

Proposition 3 (Local recovery from matrix products). *For every $n \geq 4$ and every q, the family $\mathcal{H}_{\text{loc}}$ admits deterministic polynomial-bit recognition and complete recovery, with*

$$U(\mathcal{H}_{\text{loc}}) \geq [1 - (4n^4 - n^2)/q]_+. \tag{19}$$

Proof. For an accepted source $A \in \mathcal{H}_{\text{loc}}$, first test F_T, G_T, Q_T for invertibility and return NO if any test fails. Only then form $\mathcal{C}_i(T)$ and solve $D_i\mathcal{C}_i(A) = \mathcal{C}_i(T)D_i$ simultaneously over the four words, for $i = 1, 2$. Return NO unless each solution space is a line with an invertible generator. Indeed, if $D_{i,0}$ is invertible, the map $D_i \mapsto D_{i,0}^{-1}D_i$ identifies that space with $\text{Cent}(\mathcal{C}_i(A)) = kI$. A true target therefore yields $D_i = a_i g_i$. Set $\widehat{L} = D_1 \otimes D_2$ and $\widehat{K} = (F_A^{-1}\widehat{L}^{-1}F_T)^t$. Return NO unless its realignment has rank one and its two recovered local factors are invertible. Together with D_1, D_2, accept these factors only if they reproduce every tensor coefficient; otherwise return NO. Every true transporter survives these necessary tests and the final lift, so the rejection steps are complete. Source rejection, in contrast, is SOURCE_FAIL. Here F_A^{-1} is legitimate because $A \in \mathcal{S}_{\text{full}}$. Matrix arithmetic gives $O(n^6)$ field operations.

We describe the nonvanishing mechanism; Appendix D supplies the indexed expansions. Work over indeterminates in any characteristic. Take $A^0_{ijji} = 1$ for $i \neq j$, $A^0_{iiii} = t_i$, and perturb by $\epsilon B_{ijij} = \epsilon w_{ij}$ for $i \neq j$. The first two flattenings are invertible at $\epsilon = 0$, whereas the third is invertible over Laurent series. Both local X^2 and Y^2 matrices are diagonal along this curve, with leading diagonals

$$a_j = n + \epsilon^2 \sum_{h \neq j} w_{jh}w_{hj}(1 - (t_j t_h)^{-1})^2 + O(\epsilon^3),$$
$$b_j = \epsilon^{-2} \sum_{h \neq j} (w_{jh}w_{hj})^{-1} + O(\epsilon^{-1}).$$

The a_j are pairwise distinct. A transverse tensor entry $S_{ijil} = 1$ for distinct i, j, l produces a nonzero $(a_j - a_l)b'_{jl} - (b_j - b_l)a'_{jl}$ in *both* traced modes. A fourth index h makes this nonzero via a unique Laurent monomial in w_{jh}. Thus it works for every $n \geq 4$, including characteristic two. Choose independent transverse perturbations on a spanning tree. Eliminating off-diagonal commuting variables leaves a weighted reduced tree-incidence matrix on diagonal variables modulo scalars.

Its determinant is nonzero in every characteristic, proving scalar centralizers in both modes.

Each crossed-operator numerator has degree m, and each quadratic-word numerator has degree $2m$. Each mode has a nonzero $(m-1)$-row commutator minor of degree at most $2m(m-1)$. Their product times $\det F \det G \det Q$ has degree at most $4m(m-1)+3m = 4n^4 - n^2$. The formal extension is used only to prove that this polynomial over $\mathbb{F}_p$ is nonzero; the algorithm operates over k. Polynomial zero counting over k [22,27] proves (19).

Definition 3 (Recognized source family). *For $n \geq 2^{24}$ use h, d_0, D from* (2) *and set*

$$R = \max\{r \geq 0 : q^{r^2} \leq n^4\}. \tag{20}$$

If $p > n^8$, accept precisely $\mathcal{H}_{\mathrm{loc}}$. If $p \leq n^8$ and G is invertible, use $\mathcal{H}_{\mathrm{sp}}$ with this degree window. If $p \leq n^8$ and G is singular, use $\mathcal{H}_{\mathrm{sg}}$ with this cutoff. Their disjoint union is $\mathcal{K}(n,q)$. A rejected branch returns SOURCE_FAIL.

3 Mixing under the actual matrix/inverse law

All probabilities in this section condition only on G invertible. Then $H = G^{-1}$ is exactly uniform on $\mathrm{GL}_m(k)$, and F is the prescribed reshuffling of its own inverse. Write $P_a = \prod_{i=1}^{a}(1-q^{-i})$, with $P_0 = 1$. The elementary product bound

$$P_a \geq (1-2^{-1})(1-2^{-2})\left(1 - \sum_{j\geq 3} 2^{-j}\right) = 9/32 > 1/4 \tag{21}$$

holds for every $q \geq 2$.

3.1 Analytic inputs with explicit constants

Fix a nontrivial additive character ψ of k and put

$$\mathrm{Kl}_a(U,V) = \sum_{H \in \mathrm{GL}_a(k)} \psi(\mathrm{tr}(UH + VH^{-1})), \qquad C_a = 2^a a!.$$

We use the author versions arXiv:2109.00762v4 of [11] and arXiv:2210.07133v2 of [10] for theorem numbering. Theorems 1.9–1.10 of [11] give

$$\mathrm{Kl}_a(0,V) = (-1)^u q^{au - u(u+1)/2} |\,\mathrm{GL}_{a-u}(k)|, \qquad u = \mathrm{rank}\, V. \tag{22}$$

For both coefficients singular, put $u = \max(\mathrm{rank}\, U, \mathrm{rank}\, V)$. When $u > 0$, their bound is

$$|\,\mathrm{Kl}_a(U,V)| \leq 2q^{a^2 - \Delta_a(u)}, \qquad \Delta_a(u) = au - u^2 - \binom{\min(u, a-u)}{2}. \tag{23}$$

Lemma 3 (One invertible coefficient). *If at least one coefficient is invertible, then* $|\operatorname{Kl}_a(U,V)| \le C_a q^{3a^2/4}$ *over every finite field.*

Proof. Reduce by substitution to $K_a(D) = \operatorname{Kl}_a(D, I)$ for arbitrary D. For its nilpotent and invertible Fitting parts of dimensions b, c, respectively, Theorems 1.1–1.2 of [11] give

$$K_a(D) = q^{bc} K_b(D_0) K_c(D_\times), \qquad K_b(D_0) = (-1)^b q^{b(b-1)/2}.$$

On the invertible locus, the cohomological degree cutoff in Proposition 3.18 and Corollary 3.21 of [11], together with Theorem 1 and Corollary 9 of [10], gives

$$|K_c(D_\times)| \le 2^c c! q^{(3c^2 - (c \bmod 2))/4}.$$

Here the constant is explicit: the flag-cohomology direct summand has coefficient rank 2^c and total fixed-flag Betti number at most $c!$. Appendix B proves this count, including nonsplit matrices; it is not an extraction of a constant from the $\ll_c$ notation in Theorem 1.8 of [11]. Finally $bc + b(b-1)/2 + 3c^2/4 \le 3a^2/4$ and $2^c c! \le C_a$.

Lemma 4 (Rank-sensitive cancellation). *Over every k, the following bounds hold. If* $\operatorname{rank} U = s$, $\operatorname{rank} V \ge a - c$, *and* $a > 4s + c$, *then*

$$|\operatorname{Kl}_a(U,V)| \le 4^{a-s} q^{a^2 - (a-4s-c)(a+c+1)/2}. \tag{24}$$

If $b, c \ge 1$, $\operatorname{rank} U \ge a - b$, *and* $\operatorname{rank} V \ge a - c$, *then*

$$|\operatorname{Kl}_a(U,V)| \le C_a q^{a^2 - ((a-2b-c+1)_+)^2/4}. \tag{25}$$

Either coefficient orientation may be used.

The essential cancellation is summation over a complete Schur fibre before taking absolute values. Normalize $U = \operatorname{diag}(I_r, 0_b)$. After elimination of the nonsingular part of a bottom block, its singular sector has matrix $\left(\begin{smallmatrix} A' & B_0 \\ C_0 & 0 \end{smallmatrix}\right)$. The fibre of A' over its induced quotient map has difference space $\{B_0\alpha + \beta C_0\}$. Its inverse phase is constant; additive orthogonality forces *both* $MB_0 = 0$ and $C_0 M = 0$, where M is the remaining linear coefficient. Counting only these surviving borders closes the induction. Appendix A gives the full pivot argument, using the all-field rank-count prefactor 4 from (21).

3.2 Growing contraction lists

For $P, Z \in \mathrm{M}_n(k)$ define

$$S_P(H)_{lk} = \sum_{i,j} P_{ij} H_{jl,ik}, \qquad T_Z(G)_{ij} = \sum_{k,l} G_{ik,jl} Z_{kl}.$$

Thus $S_P(H)^t = \operatorname{mat}(W \operatorname{vec} P)$ and $T_Z(H^{-1}) = \operatorname{mat}(F \operatorname{vec} Z)$.

Theorem 2 (Growing-list mixing). *For $n \geq 1024$, $q \geq 2$, and $1 \leq L, K \leq \lfloor n/100 \rfloor$, fix linearly independent lists $P_1, \ldots, P_L$ and $Z_1, \ldots, Z_K$. For arbitrary output matrices R_i, T_j, uniformly for $H \in \mathrm{GL}_m(k)$,*

$$\Pr[S_{P_i}(H) = R_i^t,\ T_{Z_j}(H^{-1}) = T_j \text{ for all } i, j] = (1+\theta)q^{-(L+K)m}, \qquad |\theta| \leq \eta := q^{-m/4}. \tag{26}$$

No minimum rank is required in either span.

Proof. The trace identities are

$$\mathrm{tr}((P \otimes A)H) = \mathrm{tr}(AS_P(H)), \qquad \mathrm{tr}((B \otimes Z^t)G) = \mathrm{tr}(BT_Z(G)).$$

They give a Fourier expansion with coefficients

$$U = \sum_i P_i \otimes A_i, \qquad V = \sum_j B_j \otimes Z_j^t.$$

The zero frequency contributes the claimed main term. Both coefficient maps are injective. Choosing pivot entries for each linearly independent list gives a fixed injective restriction to an Ln-square or Kn-square matrix, of rank at most the original coefficient rank. By (21), for $1 \leq u \leq n/2$ each restricted rank count is at most $8q^{2Lnu}$ or $8q^{2Knu}$; hence at most $64q^{2(L+K)nu}$ pairs have maximum rank at most u. There is no union over pivot choices.

Put $r = L + K \leq n/50$ and $c_0 = r + 4 \leq n/40$. For each nonzero frequency pair set

$$u = \max\{\mathrm{rank}\, U, \mathrm{rank}\, V\}, \qquad s = \min\{\mathrm{rank}\, U, \mathrm{rank}\, V\},$$
$$c = m - u = \min\{\mathrm{corank}\, U, \mathrm{corank}\, V\}.$$

Thus c is the corank of a maximum-rank coefficient. There are q^{rm} frequency pairs. Divide the nonzero absolute Fourier sum by $|\mathrm{GL}_m(k)|$ and use the following exhaustive sectors:

Coefficient ranks	Total normalized contribution
$1 \leq u \leq n/2$	$1024q^{-m/2}$
Both singular, $n/2 < u \leq m - c_0$	$8q^{-3m}$
$1 \leq c < c_0, \quad s \leq m/8$	$4q^{-m^2/8}$
$1 \leq c < c_0, \quad s > m/8$	$4q^{-m^2/2048}$
At least one invertible coefficient	$4q^{-m^2/8}$

For the first sector, $\Delta_m(u) \geq (m - 3n/4)u$. The exponent $2(L+K)nu$ in the rank-profile count is at most $mu/25$, leaving a saving of at least $mu/2$; sum the geometric tail. For the second, concavity of $\Delta_m(u) - rm$ on either side of $m/2$ reduces to $u = n/2, m/2, m - c_0$. The respective lower bounds are $12n^3/25 - 3m/8 + n/4$, $m^2/8 - nm/50 + m/4$, and $4m - 3m/3200$, all at least $3m$.

In the third sector (24) saves at least $49m^2/200$; the frequency count and chart factor together contribute an exponent of at most $m^2/100$. In the fourth, the positive parameter in (25) is $s - 2c + 1 > 21m/200$. For $m \geq 2^{20}$,

$$C_m \leq (2m)^m \leq q^{m^2/1024}. \tag{27}$$

Indeed $(2m)^{1024} \leq 2^m$ holds at $m = 2^{20}$ and propagates since $(1 + 1/m)^{1024} \leq (1 - 1024/m)^{-1} < 2$. Also $rm \leq m^2/1024$, and $441/160000 - 2/1024 \geq 1/2048$. This proves the fourth entry. Lemma 3 and the same factorial budget prove the fifth. Every nonzero frequency belongs to one sector, so no further rank factor is required. The sum is at most $1044q^{-m/2} \leq q^{-m/4}$. Affine right-hand sides affect only phases of modulus one.

Corollary 1 (Actual-source module embeddings). *Assume $n \geq 1024$ and $q \geq 2$. For an invertible $s \times s$ matrix T, where $1 \leq s \leq \lfloor n/100 \rfloor$, let $\mathrm{Emb}_T(X)$ count injective maps $P : k^s \to k^m$ with $XP = PT$. Then*

$$\mathbb{E}[\mathrm{Emb}_T(X) \mid G \text{ invertible}] = (1 + \theta)a_s^2, \quad a_s = \prod_{j=0}^{s-1}(1 - q^{j-m}), \quad |\theta| \leq \eta. \tag{28}$$

Proof. If $XP = PT$, the linked $Z = WP$ is injective: $Zv = 0$ implies $PTv = 0$. For fixed injective P, Z, the two incidences $WP = Z, FZ = PT$ have probability $(1 \pm \eta)q^{-2sm}$ by Theorem 2. There are $q^{sm}a_s$ choices for each list; the alternatives Z for a fixed P are disjoint. Summing proves the result. Both factors a_s belong to the actual tensor law.

4 A growing-order irreducible-module sieve

Use the window $[d_0, D]$ and the even order h in (2). Let $I_q(d)$ count monic irreducibles of degree d, set $Q_f = q^{\deg f}$, and put

$$\mu = \sum_{d=d_0}^{D} \frac{I_q(d)}{q^d - 1}, \qquad \eta = q^{-m/4}. \tag{29}$$

Here h is the sieve order, not the inverse matrix H; ln denotes the natural logarithm below.

Theorem 3 (Two-sided anchor mass). *For every $n \geq 2^{24}$ and $q \geq 2$,*

$$\Pr[A \notin \mathcal{H}_{\mathrm{sp}} \mid G \text{ invertible}] \leq \frac{d_0 + 1}{D + 1} + 9n^{-5}. \tag{30}$$

4.1 A pointwise sieve including repeated primary factors

Throughout this section, f ranges over the monic irreducible polynomials with degree in $[d_0, D]$. Put $E_f = k[t]/(f)$ and

$$r_f = \dim_{E_f} \ker f(X) = \frac{\dim_k \ker f(X)}{\deg f},$$

where $t \bmod f$ acts as X on $\ker f(X)$. Thus r_f is the geometric multiplicity at any root of f. For $j \geq 0$ define

$$T_j(X) = \sum_{\sum_f b_f = j} \prod_f Q_f^{\binom{b_f}{2}} \begin{bmatrix} r_f \\ b_f \end{bmatrix}_{Q_f}. \tag{31}$$

The Gaussian coefficient is zero for $b_f > r_f$. The finite q-binomial identity gives

$$\prod_f \prod_{a=0}^{r_f - 1} (1 - Q_f^a z) = \sum_{j \geq 0} (-1)^j T_j(X) z^j.$$

If all r_f vanish, the even partial sum at $z = 1$ is one. Otherwise remove one factor $(1 - z)$; the remaining polynomial has alternating coefficients $(-1)^j B_j$ with $B_j \geq 0$. Then $T_j = B_j + B_{j-1}$ and telescoping gives

$$\mathbf{1}_{\{r_f = 0 \ \forall f\}} \leq \sum_{j=0}^{h} (-1)^j T_j(X) \qquad (h \text{ even}). \tag{32}$$

This bound holds for each source, however large its geometric multiplicities. Ordinary higher factorial moments would not suffice: three distinct lines in a two-dimensional $\mathbb{F}_2$-space give six ordered triples, although no injective map from $\mathbb{F}_2^3$ exists.

For a multiplicity vector $\boldsymbol{b}$, the semisimple module $M_{\boldsymbol{b}} = \bigoplus_f (k[t]/f)^{b_f}$ has dimension $s = \sum_f b_f \deg f \leq hD$ and automorphism group $\prod_f \mathrm{GL}_{b_f}(Q_f)$. Consequently, exactly on every source,

$$\prod_f \begin{bmatrix} r_f \\ b_f \end{bmatrix}_{Q_f} = \frac{\mathrm{Emb}_{M_{\boldsymbol{b}}}(X)}{\prod_f |\mathrm{GL}_{b_f}(Q_f)|}. \tag{33}$$

We show below that $hD \leq \lfloor n/100 \rfloor$, so Corollary 1 applies even to repeated factors. Using $|\mathrm{GL}_b(Q)| = Q^{\binom{b}{2}} \prod_{a=1}^{b} (Q^a - 1)$ gives comparison coefficients

$$c_j = \sum_{\sum_f b_f = j} \prod_f \left(\prod_{a=1}^{b_f} (Q_f^a - 1) \right)^{-1}.$$

Since $s \leq hD \leq n/100$, the finite-size correction satisfies $2q^{hD-m} \leq q^{-m/4} = \eta$. Together with $1 - 2q^{hD-m} \leq a_s^2 \leq 1$, Corollary 1 gives $|\mathbb{E}T_j - c_j| \leq 2\eta c_j$ for $1 \leq j \leq h$, and $\mathbb{E}T_0 = c_0 = 1$. The aggregate error includes every multiplicity vector.

4.2 Generating function and the truncation error

The absolutely convergent identity

$$\sum_{b \geq 0} \frac{z^b}{\prod_{a=1}^{b} (Q^a - 1)} = \prod_{a \geq 1} (1 + zQ^{-a}) \tag{34}$$

follows by writing the right side as $F(z) = (1 + z/Q)F(z/Q)$ and comparing coefficients. Thus c_j is the jth elementary symmetric function of the positive weights Q_f^{-a}, whose sum is μ. Hence

$$c_j \le \mu^j/j!, \quad (j+1)c_{j+1} \le \mu c_j, \quad \sum_j c_j \le e^{\mu},$$
$$\sum_j (-1)^j c_j = \prod_{f,a} (1 - Q_f^{-a}) \le e^{-\mu}.$$

For $h \ge \mu$, the alternating remainder is at most c_{h+1}. Taking expectations in (32) therefore yields

$$\Pr[\text{no window factor} \mid G \text{ invertible}] \le e^{-\mu} + \frac{\mu^{h+1}}{(h+1)!} + 2\eta e^{\mu}. \tag{35}$$

No independent-factor or Poisson law has been assumed.

For the uniform parameter bounds, take $2^r \le n < 2^{r+1}$, $r \ge 24$. Then $d_0 < 8r + 12$ and $h \le 8r + 10$. The integer inequality

$$100(8r + 10)(8r + 12) \le 2^r \tag{36}$$

holds at $r = 24$, since its left side is $4{,}120{,}800 < 2^{24}$. The successive left ratio is at most $(210 \cdot 212)/(202 \cdot 204) < 2$ and decreases with r. Thus (36) holds for every integer $r \ge 24$. It follows that $D \ge d_0$, while its definition gives $hD \le \lfloor n/100 \rfloor$. This proves a nonempty window and the required module budget for every integer n in range, not only powers of two.

From $q^d = \sum_{a \mid d} a I_q(a)$ and $\sum_{j=1}^{\lfloor d/2 \rfloor} q^j \le 2q^{d/2}$, we obtain $dI_q(d) \ge q^d - 2q^{d/2}$. Therefore

$$\begin{aligned} \mu &\ge \ln \frac{D+1}{d_0} - \varepsilon_{\mathrm{irr}}, \\ \varepsilon_{\mathrm{irr}} &:= 2 \sum_{d=d_0}^{D} \frac{q^{-d/2}}{d} \le \frac{2q^{-d_0/2}}{d_0(1 - q^{-1/2})} < \frac{3n^{-4}}{d_0}. \end{aligned} \tag{37}$$

Indeed $q^{d_0} \ge q^3 n^8$ and $2q^{-3/2}/(1 - q^{-1/2}) \le \sqrt{2} + 1 < 3$. The final bound in (37) is less than $1/(d_0 + 1) \le \ln(1 + 1/d_0)$, so

$$e^{-\mu} \le \frac{d_0 + 1}{D + 1}. \tag{38}$$

For the upper bound, $d_0 \ge 4$ and $I_q(d) \le q^d/d$ give

$$\mu \le \frac{16}{15} \sum_{d=4}^{D} \frac{1}{d} \le \frac{16}{15} \ln n, \qquad e^{\mu} \le n^2, \qquad h \ge 8\mu.$$

The last inequality follows from $h \ge 8 \log_2 n$ and $\ln 2 < 15/16$. Using $j! \ge (j/e)^j$ and $e < 3$, we have $\mu^{h+1}/(h+1)! \le (3/8)^{h+1} \le 2^{-h} \le n^{-8}$. Also $m/4 \ge 12n$ for $n \ge 48$ and $2^n \ge n$, whence $2\eta e^{\mu} \le 2n^{-10} \le n^{-8}$. Equation (35) is consequently at most $(d_0 + 1)/(D + 1) + 2n^{-8}$.

4.3 Removing bad linked anchors

Lemma 5 (Frobenius rigidity). *For $n, q \geq 2$, a nonzero polynomial over k of degree*

$$B = n + (n^2 - 1)(q^2 + n - 1) \leq n^3 q^2 \tag{39}$$

has nonzero locus consisting of invertible P for which

$$\mathrm{Cent}(P^q P^{-1}, P^{q^2} P^{-1}) = EI$$

over any extension E/k.

Proof. The map $\Phi(P) = P^q P^{-1}$ is a form of Lang's map on GL_n [17]. We recall a direct proof of its dominance for the degree estimate below. The differential

$$d\Phi = -\Phi(P)(dP)P^{-1}$$

is invertible on tangent coordinates. Algebraically, a polynomial relation of least positive degree, chosen irreducible over the perfect field k, has a nonzero partial derivative. The chain rule would make that derivative a smaller relation, a contradiction. For a generic matrix A, A, A^q have scalar centralizer: specialize $A = \mathrm{diag}(t_i) + J$, with J the cyclic shift. Then $A - A^q$ is diagonal with distinct entries, and commuting with J forces a diagonal commuting matrix to be scalar. Now $P^{q^2} P^{-1} = \Phi(P)^q \Phi(P)$. Use a nonzero maximal commutator minor on M_n/kI, clear P^{-1} with its adjugate, and multiply by $\det P$. This gives (39) without dividing by n.

For one f of degree d, identify a vector in E^m with its m-by-d coordinate matrix over k using the companion action of f. A nonzero eigenvector gives an injective module embedding. Its linked vector is also injective by the proof of Corollary 1. With $Q = q^d$, there are at most BQ^{m-1} bad candidates for either reshape. For each injective pair the required event is exactly $WP = Z, FZ = PT_f$, of probability at most $(1 + \eta)q^{-2dm}$. There are $q^{dm} a_d$ possibilities for the unrestricted side. Divide by $Q - 1$ for simultaneous scaling of the linked pair. The expected number of bad simple f-submodules is at most

$$\frac{2(1 + \eta)B}{Q(Q - 1)}.$$

Summing $I_q(d) \leq Q/d$ over the window bounds the probability of any bad linked submodule by $16Bq^{-d_0} \leq 8n^{-5}$. This is a union bound on actual incidences, not an independence assertion about P, Z.

Geometric multiplicity at least two gives an embedded E-two-space. Dividing its embedding count by $|\mathrm{GL}_2(Q)| \geq 3Q^4/8$ and summing gives at most $8q^{-3d_0} \leq \frac{n^{-5}}{4}$. The incidences here have module dimensions d or $2d$, both at most hD. These multiplicities were fully included in the sieve before being removed for recovery. Combining the zero event, bad linked loci, and this multiplicity loss gives

$$\frac{d_0 + 1}{D + 1} + 2n^{-8} + 8n^{-5} + \frac{n^{-5}}{4} \leq \frac{d_0 + 1}{D + 1} + 9n^{-5},$$

which is Theorem 3. The recognizer factors only χ_X and tests its factors; it never enumerates the sieve terms or all irreducible polynomials in the window.

5 Singular sources: marked kernels and rectangular maps

Here probabilities begin under the raw law U. We account for singular G without replacing it by an independently sampled inverse or requiring an invertible kernel reshape.

5.1 Rank tails and the exact marked law

Lemma 6 (Rank probabilities). *For a uniform a-square matrix over k,*

$$\Pr[\operatorname{corank} M = r] = q^{-r^2} \frac{P_a^2}{P_r^2 P_{a-r}}. \tag{40}$$

Moreover $\Pr[\operatorname{corank} M \geq b] \leq 4q^{-b^2}$.

Proof. There are $\begin{bmatrix} a \\ r \end{bmatrix}_q$ choices each for kernel and image, and $|\operatorname{GL}_{a-r}(k)|$ induced isomorphisms. Their product proves the equality. For the tail, union bound over b-dimensional kernel subspaces: there are at most $4q^{b(a-b)}$ such subspaces by (21), and annihilating a fixed one has probability q^{-ab}.

Thus $\Pr[\operatorname{corank} G > R] < 4n^{-4}$ for (20), and the raw failure probability of $\operatorname{rank} Q > m - n$ is at most $4q^{-n^2}$. These are marginal bounds; no independence of G, Q is used.

Fix a corank stratum $1 \leq r \leq R$. Mark independently, given G, a uniform nonzero right kernel vector u and left kernel vector v. Every rank-$(m-r)$ matrix has exactly $(q^r - 1)^2$ marks. The independent left/right GL_m action is transitive on nonzero pairs (u, v), so the marked vectors are independent uniform nonzero vectors in k^m. Conditional on them, G is uniform on the rank-$(m - r)$ part of

$$\mathcal{L}(u, v) = \{G : Gu = 0,\ v^t G = 0\} = \operatorname{Hom}(k^m / ku, v^\perp). \tag{41}$$

In quotient/image bases this is a full $(m - 1)$-square matrix space. Its rank probability from (40) is at least $q^{-(r-1)^2}/16$. Consequently every event E satisfies

$$\Pr[E \mid u, v, \operatorname{corank} G = r] \leq 16q^{(r-1)^2} \Pr_{\mathcal{L}(u,v)}[E] \leq 16n^4 \Pr_{\mathcal{L}(u,v)}[E]. \tag{42}$$

This is an explicit cost of conditioning on rank.

Let $S = \max\{s \geq 0 : q^{s^2} \leq n^{32}\}$. Since u, v are uniform nonzero vectors, Lemma 6 and a union bound give

$$\Pr[\operatorname{corank} \operatorname{mat} u > S \text{ or } \operatorname{corank} \operatorname{mat} v > S] < 16n^{-32}. \tag{43}$$

If every kernel pair fails the recognizer then this marked pair fails. Thus a failure upper bound for the marked construction bounds source failure, without charging for the recognizer's enumeration.

5.2 Exact projection to a rectangular Hom law

For fixed reshapes $P = \operatorname{mat} u, V = \operatorname{mat} v$ of ranks s, t, independent local transformations normalize them to $\operatorname{diag}(I_s, 0)$ and $\operatorname{diag}(I_t, 0)$. Formula (16) preserves the tests. Put $D_P = \{MP : M \in \mathrm{M}_n\}$ and let W_V be the matrices with their last $n - t$ rows zero and trace zero. Multiplication by V^t maps $v^\perp$ onto W_V. Restriction followed by this projection is a surjective linear map

$$\mathcal{L}(u, v) \longrightarrow \operatorname{Hom}(D_P/kP, W_V).$$

To prove surjectivity, lift a prescribed map through a linear section of the output projection, then extend it from D_P/kP to k^m/kP. All fibres have equal cardinality, so the pushed-forward law is exactly uniform. The dimensions are $ns - 1$ and $nt - 1$.

Writing $X_{ij} = B(E_{ij})$, we have $X_{ij} = 0$ for $j > s$; only indices at most s occur in nonzero cyclic products in (15). For those indices the exact law is

$$X_{ij} = \begin{pmatrix} Y_{ij} & Z_{ij} \\ 0 & 0 \end{pmatrix}, \quad Y_{ij} \in \mathfrak{sl}_t(k), \quad Z_{ij} \in \mathrm{M}_{t,n-t}(k), \quad \sum_{i=1}^{s} Y_{ii} = \sum_{i=1}^{s} Z_{ii} = 0. \tag{44}$$

The two arrays are independent uniform arrays subject to these relations. In particular Y represents a uniform map $\mathrm{M}_s/kI_s \to \mathfrak{sl}_t$. The trace-adjoint direction has the same marginal law with s, t exchanged. Indeed $G' = SG^tS$, $P' = V^t$, $V' = P^t$ give its map exactly $SB^tS = B^\dagger$. This is a bijection of slices, not independence of the two directions.

5.3 The rectangular core and its extension columns

Write the full covariants of (44) as

$$C_j = \begin{pmatrix} A_j & B_j \\ 0 & 0 \end{pmatrix} \quad (j = 2, 3), \qquad \begin{aligned} A_2 &= \sum Y_{ij}Y_{ji}, & B_2 &= \sum Y_{ij}Z_{ji}, \\ A_3 &= \sum Y_{ij}Y_{jl}Y_{li}, & B_3 &= \sum Y_{ij}Y_{jl}Z_{li}. \end{aligned} \tag{45}$$

Theorem 4 (Rectangular distribution and centralizers). *For $n \geq 4096$, $q \geq 2$, and $19n/20 \leq s, t \leq n$, the full n-square pair in* (45) *satisfies*

$$\Pr[\operatorname{Cent}(C_2, C_3) \neq kI_n] \leq q^{-n/4}. \tag{46}$$

Proof. A_2 takes values in $\mathcal{W}_2 = \mathrm{M}_t$, with $\mathcal{W}_2 = \mathfrak{sl}_t$ in characteristic two; A_3 takes values in $\mathcal{W}_3 = \mathrm{M}_t$, with $\mathcal{W}_3 = \mathfrak{sl}_t$ in characteristic three. Cyclically rotating the trace terms proves these restrictions: off-diagonal cycles cancel in pairs or triples, and the remaining traces are $\operatorname{tr}(Y_{ii})^2$ or $\operatorname{tr}(Y_{ii})^3 = 0$ in the respective characteristic.

Let u be uniform product measure on $\mathcal{W}_2 \times \mathcal{W}_3$, and put $h_0 = \lfloor (s-1)/2 \rfloor$. We first prove

$$\mathrm{TV}(\mathcal{L}(A_2, A_3), u) \leq \tau = \tfrac{1}{2} q^{t^2/2 - s(s-1)(t-2)} + \tfrac{1}{2} q^{t^2 - (5/2)th_0}. \tag{47}$$

Take Fourier frequencies U, V modulo scalars when the corresponding comparison space is trace zero. At a nontrivial quadratic-only frequency, condition on the diagonal blocks. Every independent unordered off-diagonal pair contributes $\operatorname{tr} U(YZ + ZY)$. The Jordan-block kernel calculation for $Y \mapsto UY + YU$ has rank at least $2t - 2$ for nonscalar U; restricting input and output trace pairings loses at most two ranks. For nonzero scalar U in odd characteristic the rank is at least $t^2 - 2$. It follows that

$$|\Phi(U, 0)| \le q^{-s(s-1)(t-2)}. \tag{48}$$

For a nontrivial cubic frequency set $\beta_V(Y, A) = (VY + YV)A + YAV$. The all-field bilinear-rank estimate is

$$\Pr_{Y,A \leftarrow \mathfrak{sl}_t} [\beta_V(Y, A) \bmod kI = D] \le q^{-5t/2} \quad (t \ge 1024), \tag{49}$$

uniformly for every scalar coset D. The rank mechanism is as follows. For $d = \min_\lambda \operatorname{rank}(V - \lambda I)$, completing a corank-$e$ input to an invertible one gives the saving $\lceil d/2 \rceil^2$. Near a scalar with $3\lambda \neq 0$, the main map $3\lambda YA$ gives saving $3t^2/4 - 3td + 2d^2$. In the remaining sector factor $V - \lambda I = UC$ and retain separately the ranks across $\operatorname{im} U$ and $\ker C$. Intersecting the incidence and corank events takes the minimum of their probability bounds, hence the maximum of their exponent savings. The resulting saving is at least $3dt - 19d^2/4 - \sqrt{2t}\, d^{3/2}$, minimized at the endpoint $d = 1$ in the needed range. After two trace restrictions, one scalar coset and all rank prefactors, it still exceeds $5t/2$. Appendix C gives the ranks, overlap calculation and uniform absorptions proving (49).

Choose one index 0 and disjoint sets I, J of size h_0. Condition on all blocks except Y_{ai}, Y_{i0}, Y_{0a} for $i \in I, a \in J$. Each triangle contains at most one J-to-I edge. Integration of Y_{ai} forces $\beta_V(Y_{i0}, Y_{0a}) + D_{ai} \in kI$; all other coefficients depend only on conditioned blocks, including the quadratic frequency. Keep a matching of h_0 equations. Their remaining input pairs are independent, giving $|\Phi(U, V)| \le q^{-5th_0/2}$ for nontrivial V. There are at most q^{t^2} quadratic-only frequencies and q^{2t^2} others. Parseval and TV $\le \frac{1}{2}\sqrt{\sum_{\xi \neq 0} |\Phi(\xi)|^2}$ prove (47).

For the independent comparison pair, the chance of a nonscalar common centralizer is at most $q^{-t/2}$ for $t \ge 16$; the chance of a common nonzero left or right kernel is at most $2q^{-t}$. Here is a direct count. A nonscalar commuting matrix either has a polynomial with nonzero proper kernel, or generates an embedded field of degree $d \ge 2$ dividing t. A common a-space costs $q^{-2a(t-a)}$, and counting spaces gives at most $4 \sum_{a=1}^{t-1} q^{-a(t-a)}$. Counting embedded fields and their matrix centralizers gives at most $4tq^{-t^2/2}$. Trace is surjective on each preserving algebra, including the extension-field centralizer, so fixed trace fibres do not change these probabilities. Their sum is at most $13q^{-(t-1)} \le q^{-t/2}$. For a fixed nonzero vector the two kernels cost q^{-2t}, and a union bound gives the stated $2q^{-t}$. Transfer this *union* of core bad events once using (47).

It remains to handle the $a = n - t$ extension columns. Conditional on Y, each column array z is uniform subject to $\sum_i z_{ii} = 0$ and maps by

$$L_Y(z) = \left(\sum_{i,j} Y_{ji} z_{ij},\ \sum_{i,j} (Y^2)_{ji} z_{ij} \right), \qquad (Y^2)_{ji} = \sum_l Y_{jl} Y_{li}. \tag{50}$$

An elementary exposure gives

$$\Pr[\operatorname{rank} L_Y < 2t] \le 3q^{-t(s-4)}. \tag{51}$$

To see this, a left annihilator (u, v) on the off-diagonal input columns satisfies $uY_{ba} + v(Y^2)_{ba} = 0$ for all $a \neq b$. For fixed $v \neq 0$, expose row 0. The exceptional event $vY_{0j} = 0$ for all $j \neq 0$ costs $q^{-t(s-1)}$. Otherwise choose the first j with $vY_{0j} \neq 0$. Leave Y_{ja}, $a \neq 0, j$, unexposed. In the $(0, a)$ equation this is the sole unexposed block, with nonzero coefficient vY_{0j}. The map $M \mapsto vY_{0j}M$ from $\mathfrak{sl}_t$ onto k^t is surjective, so the $s - 2$ equations cost $q^{-t(s-2)}$. Union over fewer than q^{2t} pairs gives $2q^{-t(s-4)}$. The $v = 0$ case contributes at most $q^{-t(s-2)}$, proving (51).

Assume the core centralizer is scalar, its joint kernels vanish, and L_Y is onto. Then $\delta(x) = (A_2 x, A_3 x)$ is injective and belongs to $\operatorname{im} L_Y$: use $z_{ij} = Y_{ij} x$. The extension columns are independent uniform in k^{2t} and hence in its t-dimensional quotient by $\operatorname{im} \delta$. Linear dependence of these images has probability at most

$$1 - \prod_{j=0}^{a-1} (1 - q^{j-t}) \le \frac{q^{a-t}}{q-1}. \tag{52}$$

The exact expression is zero when $a = 0$. For a commuting matrix $D = \left(\begin{smallmatrix} H & E \\ F & J \end{smallmatrix}\right)$, the joint left kernel forces $F = 0$, and then $H = \lambda I_t$. The upper-right equations modulo $\operatorname{im} \delta$ force $J = \lambda I_a$ by linear independence of those images. Injectivity of δ forces $E = 0$. Thus the full pair, not just the core, has scalar centralizer.

Its failure is bounded by

$$q^{-t/2} + 2q^{-t} + \tau + 3q^{-t(s-4)} + q^{n-2t}/(q-1). \tag{53}$$

For $s, t \ge 19n/20$, $h_0 \ge (s-2)/2$ gives both exponents in τ at most $-n^2/8$. The remaining terms each save at least $n/3$. Thus (53) is at most $8q^{-n/3} \le q^{-n/4}$. This proves (46).

For $n \ge 4096$, $S \le n/20$. Indeed on $2^r \le n < 2^{r+1}$, $r \ge 12$, the inequality $12800(r+1) \le 2^{2r}$ holds at 12 and propagates with right ratio 4 and left ratio less than 2. It implies $32 \log_q n \le n^2/400$. Apply Theorem 4 to both marginal directions of the marked law, using (42) for rank conditioning. Together with (43), the singular low-corank failure, apart from the raw Q gate, is at most

$$16n^{-32} + 32n^4 q^{-n/4}. \tag{54}$$

This is uniform at each corank $1 \le r \le R$. Averaging over those strata costs no factor R, since their probabilities sum to at most one.

6 Raw probability composition and uniform bit complexity

6.1 Proof of the headline probability bound

For $p \le n^8$, combine Theorem 3, the two marginal rank tails, and (54). Multiplying the conditional invertible-G loss by $\Pr[G \text{ invertible}] \le 1$ only improves it. Every probability now has the common raw denominator q^{n^4}:

$$\begin{aligned}\Pr_U[A \notin \mathcal{K}] \le \frac{d_0+1}{D+1} + 9n^{-5} + 4n^{-4} + 16n^{-32} \\ + 32n^4 q^{-n/4} + 4q^{-n^2}.\end{aligned} \tag{55}$$

If $q > n^4$, $R = 0$ and the low-corank terms are absent; the same upper bound is valid. For $n \ge 2^{24}$ and $q \ge 2$, both exponential terms in (55) are at most n^{-4}. Indeed $32n^8 \le 2^{n/4}$ follows, after taking fourth powers, from the seed $2^{20}1024^{32} \le 2^{1024}$ and the consecutive ratio $(1+1/n)^{32} \le (1-32/n)^{-1} < 2$ for $n \ge 1024$. It also implies $4n^4 \le 2^{n/8} \le 2^{n^2}$. Since $9/n + 16/n^{28} \le 1$, every term other than $(d_0+1)/(D+1)$ sums to at most $7n^{-4}$.

For $p > n^8$, the local bound (19) gives $(4n^4 - n^2)/q < 4n^{-4}$ directly under U, which is smaller than the same bound. The two characteristic cases are exhaustive, including unbounded extension degrees in the bounded-characteristic branch. Capping by one proves the first inequality of (4). For the second, put $x = \log_2 n$. Since $h \le 8x + 2$, $d_0 + 1 \le 8x + 5$, and $D + 1 > n/(100h)$,

$$\frac{d_0+1}{D+1} < \frac{100h(d_0+1)}{n} \le \frac{6400(1+\log_2 n)^2}{n}.$$

This proves the stated field-uniform envelope.

6.2 Field arithmetic and unsuccessful candidates

All predicates are action invariant by their covariance laws. Source recognition is completed before target classification, and Propositions 1, 2, and 3 establish complete decision/search on their accepted sources. The parameters h, d_0, D, R can be computed by integer-power comparisons and integer division, without approximate logarithms. We record the uniform cost, including rejected factors and candidate branches.

Lemma 7 (Deterministic factor and root lists). *Polynomials of degree a over $\mathbb{F}_{p^e}$ can be factored deterministically in time polynomial in $p, a, e, \log p$. All roots of a degree-b polynomial over a represented extension E can be listed with the same type of bound.*

Proof. After squarefree decomposition by gcds, derivatives, and inverse Frobenius, view $R = k[x]/f$ for squarefree f as an ae-dimensional $\mathbb{F}_p$-algebra. The fixed space of absolute p-Frobenius is $\mathbb{F}_p^r$, where r is the number of field factors.

Compute a basis by linear algebra. It separates the components: for any two distinct coordinate evaluations some basis element differs. Gcds with each basis element minus every $c \in \mathbb{F}_p$ therefore split all components. This is a deterministic version of the fixed-algebra factorization method [1,12]; its scalar scan is over p, not q. For roots in E, first take the gcd with $x^{|E|} - x$ by modular powering, then apply the same procedure. The remaining factors are linear and comprise all distinct roots, including roots of inseparable input polynomials.

In the spectral branch $p \leq n^8$. The characteristic polynomial has degree $m = n^2$, hence at most m source factors. Only factors of degree $d \leq D \leq n \leq m$ are tested. The field-operation bounds below are uniform for extension degree at most m. Each factor test uses a polynomial-size extension-field nullspace, local inverses, and two centralizer systems. Target search has at most n^2 paired scalar candidates. Each root list and failed candidate is covered by Lemma 7. Extension-field operations and projective descent have polynomial bit cost in $n, \log q$.

For the singular branch the rank cutoff gives

$$q^{2r} \leq q^{2r^2} \leq n^8, \qquad q^{2r+1} \leq q^{3r^2} \leq n^{12}.$$

Moreover $r \geq 1$ implies $q \leq n^4$, so enumerating $\beta \in k^*$ is polynomial here. Each candidate uses at most $O(n^{10})$ field operations for contractions and ordinary linear systems, a loose bound covering the trace adjoint and full coefficient check. This includes all failing kernel pairs and singular intertwiner generators. If the corank exceeds R, recognition stops after the rank computation.

In the large-characteristic branch no factoring or field-element scan occurs; the local procedure uses $O(n^6)$ field operations and stores $O(n^4)$ elements. Arithmetic in a valid polynomial-basis representation has polynomial bit cost in $\log q$. Combining these explicit loop bounds with the fixed-space construction gives the conservative common bound $O(n^{40}(1 + \log_2 q)^{12})$. Its absolute exponents allow for all source factor tests, extension arithmetic, target candidates, and verification. The growing-order sieve is used only in the proof of density and contributes no enumeration to the algorithm. This completes the proof of Theorem 1.

7 Complementary finite-parameter regimes

Theorem 1 covers every field but starts at a large dimension. The following families give explicit guarantees far below that cutoff. Their smaller-dimensional source-mass statements are not subsumed by the almost-all theorem. The recovery procedures use the covariance of (7); their density proofs need sharper short-list estimates at smaller dimensions. No containment in $\mathcal{K}$, or converse containment, is used.

7.1 Three rational spectral anchors

Put $\mathcal{B}_{\mathrm{inv}} = \{A : \det F \det G \det Q \det W \neq 0\}$.

Definition 4 (Three-anchor family). *A source in $\mathcal{B}_{\mathrm{inv}}$ belongs to $\mathcal{R}_3$ if $X = FW$ has three distinct nonzero eigenvalues $\lambda_1, \lambda_2, \lambda_3 \in k$, each with a one-dimensional geometric eigenspace, whose reshapes P_1, P_2, P_3 are invertible and satisfy* $\mathrm{Cent}(P_2 P_1^{-1}, P_3 P_1^{-1}) = kI$.

Algebraically repeated roots are allowed. Independent rescaling and permutation of the anchors preserve the predicate. An auxiliary minimum-rank condition used below is not part of recognition.

Theorem 5 (Constant density at finite dimensions). *For every prime power $q \geq 5$,*

$$U(\mathcal{R}_3) > 2^{-36} \quad (n \geq 23), \qquad U(\mathcal{R}_3) > 2^{-16} \quad (n \geq 27). \tag{56}$$

Over $\mathbb{F}_4$, for every $n \geq 26$,

$$U(\mathcal{R}_3) \geq c_4 := \frac{1228503645664423293814479}{264452523040700131966976000000} > 2^{-15}. \tag{57}$$

For $q \geq 5, n \geq 51$, also $U(\mathcal{R}_3) \geq c(q,n) \geq c(q) \geq 1/34 - 1/q$, with the explicit envelopes (93)–(94)*. Recognition and exact decision/search against every target take Las Vegas expected* $\mathrm{poly}(n, \log q)$ *bit time, deterministically in characteristics* $2, 3, 5$*. The pointwise algorithms apply for $n \geq 2$ on their promised sources.*

Appendix E gives the complete recovery algorithm and rational stabilizer analysis. Appendix G proves the explicit envelopes for $q \geq 5, n \geq 51$; Appendix K proves the lower-dimensional and field-four bounds under the same source law. The family $\mathcal{R}_3$ is empty over $\mathbb{F}_2$ and $\mathbb{F}_3$.

7.2 Kernel anchors of an inverse-derived matrix

Here the source is in $\mathcal{S}_{\mathrm{full}}$ and $\operatorname{rank} W = N = m - 1$, with invertible right and left kernel reshapes P_A, R_A. These are kernels of the rearranged inverse W, not kernels of the original singular G used in Section 2. Define

$$\mathcal{B}_A^W : \mathrm{M}_n(k)/kI \longrightarrow \mathfrak{sl}_n(k), \qquad [M] \longmapsto R_A^t \operatorname{mat}(W_A \operatorname{vec}(MP_A)). \tag{58}$$

The kernel equations make this map well-defined and trace zero; invertible reshapes and rank N make it an isomorphism. Write C_2, C_3 for its quadratic and cubic contractions, defined by applying (15) to $\mathcal{B}_A^W$. This notation is local to the inverse-derived construction.

Definition 5. *The family $\mathcal{S}_{\mathrm{CAS}} \subseteq \mathcal{S}_{\mathrm{full}}$ consists of the sources just described with* $\mathrm{Cent}(C_2, C_3) = kI$*. Put $\mathcal{S}_{\mathrm{det}} = \mathcal{S}_{\mathrm{CAS}}$ for $p \in \{2, 3, 5\}$, and $\mathcal{S}_{\mathrm{det}} = \{A \in \mathcal{S}_{\mathrm{CAS}} : \operatorname{tr} C_2(A) \operatorname{tr} C_3(A) \neq 0\}$ for $p \geq 7$.*

Rescaling kernel representatives scales $\mathcal{B}^W$ and the two homogeneous contractions; the predicates are action invariant and deterministically recognizable by polynomial-size linear algebra.

Theorem 6 (Inverse-derived corank-one families). *For $q \geq 5$ and $n \in \{27\} \cup \{29, 30, \ldots\}$,*

$$\max\left\{\frac{75}{2048q}, \frac{1}{q} - \frac{7}{q^2}\right\} \leq U(\mathcal{S}_{\mathrm{det}}) \leq U(\mathcal{S}_{\mathrm{CAS}}) \leq \frac{1}{q-1}, \qquad \mathcal{S}_{\mathrm{det}} \subseteq \mathcal{S}_{\mathrm{CAS}} \subseteq \mathcal{S}_{\mathrm{full}}. \tag{59}$$

Both masses are $1/q + O(q^{-2})$, uniformly in this dimension set. No source-mass claim at $n = 28$ is included. Recovery on $\mathcal{S}_{\mathrm{det}}$ is deterministic polynomial-bit; recovery on $\mathcal{S}_{\mathrm{CAS}}$ is Las Vegas expected polynomial-bit, deterministically for $p = 2, 3, 5$. These complete arbitrary-target procedures apply pointwise already for $n \geq 2$.

Appendix H supplies the matched-representative covariance, complete scalar recovery and raw-source normalization. Appendices I–J give the finite-band and analytic estimates for $n \geq 51$; Appendix K extends the bound to $n = 27, 29, \ldots, 50$ by the bilateral fibre estimate and exact certificates.

Corollary 2 (Disjoint finite-parameter unions). *For $q \geq 5$ and $n \in \{27\} \cup \{29, 30, \ldots\}$, both $\mathcal{R}_3 \cup \mathcal{S}_{\mathrm{det}}$ and $\mathcal{R}_3 \cup \mathcal{S}_{\mathrm{CAS}}$ have raw mass strictly greater than*

$$2^{-16} + \max\left\{\frac{75}{2048q}, \frac{1}{q} - \frac{7}{q^2}\right\},$$

and, when $n \geq 51$, at least $1/34 - 7/q^2$. Both admit complete Las Vegas expected polynomial-bit recovery, deterministically in characteristics $2, 3, 5$.

Proof. Dispatch by $\operatorname{rank} W$: it is m on $\mathcal{R}_3$ and $m - 1$ on $\mathcal{S}_{\mathrm{CAS}}$, so the families are pointwise disjoint. Add their lower bounds, and for the second use $(1/34 - 1/q) + (1/q - 7/q^2)$.

8 Total correctness and cryptographic consequences

8.1 A total solver with a high-probability time bound

Proposition 4 (Total extension of the fast algorithm). *Run the recognizer and selected recovery algorithm first. Return its* YES *witness or* NO *on an accepted source. On* SOURCE_FAIL*, enumerate $\mathrm{GL}_n(k)^4$ and test (1), returning a witness or* NO *after exhaustion. For dimensions outside Definition 3, use this finite enumeration directly. This is a deterministic total solver. For one polynomial P, uniform in n, q,*

$$\Pr_{A \leftarrow U}\left[\max_{T \in V} \mathrm{Time}(\mathrm{Solve}_{\mathrm{total}}(A, T)) > P(n, \log q)\right] \leq \varepsilon(n, q) \qquad (n \geq 2^{24}). \tag{60}$$

Proof. Enumeration of a finite group is complete. For every $A \in \mathcal{K}$, all targets terminate within the same polynomial bound of Theorem 1, before any exhaustive search is initialized. The event in (60) is therefore contained in $V \setminus \mathcal{K}$.

Equation (60) bounds the worst-target time with high probability, not its expectation: the exceptional sources may require expensive search.

8.2 Inversion of the coefficient-uniform group action

In the group-action inversion experiment [15, Section 2.3, Definitions 3–4], sample $A \leftarrow U$ and $g \leftarrow \mathrm{GL}_n(k)^4$, and reveal (A, gA). Run the polynomial-time recognizer and fast recovery, returning no witness on source rejection. Theorem 1 gives a valid transporter with probability at least $1 - \varepsilon(n, q)$ for $n \geq 2^{24}$. For each $A \in \mathcal{K}$, the guarantee holds for every g, not merely on average over targets. Any valid transporter solves the inversion problem; recovering the sampled g itself is unnecessary.

This is an asymptotic obstruction to using the coefficient-uniform square four-tensor action as a one-way problem, not practical cryptanalysis at the stated cutoff or a worst-case algorithm for arbitrary tensors. Alternating-trilinear and matrix-code signatures [25,8], including ALTEQ and MEDS [3,7], use different actions or formats. Worst-case reductions alone do not transfer the theorem to them.

The complementary finite-dimensional results also give strict polynomial inversion budgets. Put $a(q) = \max\{75/(2048q), 1/q - 7/q^2\}$ and $\rho_p = 1$ for $p \in \{2, 3, 5\}$, and $\rho_p = 1/2$ otherwise. Give the $\mathcal{H}_{\mathrm{loc}}$ and $\mathcal{S}_{\mathrm{det}}$ procedures their deterministic budgets. In the rational spectral procedure, for $p \geq 7$ stop after twice the pointwise expected bound T_0 in (77); use its full deterministic bound in the other three characteristics. Markov's inequality, applied separately to each fixed accepted source and target, retains at least a ρ_p fraction of its inversion probability. The disjointness in Corollary 2 therefore gives success strictly larger than $\rho_p 2^{-16} + a(q)$ on its dimension range, with a polynomial budget on every random tape. For $q \geq 5, n \geq 51$, the sharper bound is

$$\max\{[1 - (4n^4 - n^2)/q]_+,\ \rho_p c(q, n) + a(q)\}.$$

The maximum, rather than a sum, accounts for possible overlap with $\mathcal{H}_{\mathrm{loc}}$. These are randomized bounded-time inversion guarantees when $p \geq 7$; they do not alter the deterministic headline algorithm.

Dependence on the source distribution. For another distribution D on the same square tensor space,

$$\Pr_D[A \notin \mathcal{K}] \leq \min\{1,\ \varepsilon(n, q) + \mathrm{TV}(D, U)\}.$$

This is useful only with a proved comparison between D and U. For example, the nontrivial boundary-format sampler of [20, Sections 2–2.1.1] has side lengths satisfying $b_1 - 1 = \sum_{j=2}^{4}(b_j - 1)$ and is not square. After padding to $N = \max b_j$, some side has length $b_j < N$, so every balanced flattening has rank at most $b_j N < N^2$, also after invertible local changes of basis. Its padded distribution therefore has no mass on $\mathcal{R}_3 \subset \mathcal{S}_{\mathrm{full}}$. This observation distinguishes the source laws; it is not a claim that the padded sampler avoids the new family $\mathcal{K}$.

8.3 Evidence and reproducibility

The probability bounds are analytic. The verification package contains exact finite-field checks of marked-kernel normalization, contraction covariance and original-field recovery, together with submodule-sieve identities, sieve-parameter inequalities and the finite certificates for Appendices I and K. Run `python3 -B run_verification.py` with `--verify-manifest` to check integrity, or with `--output-dir replay` and `--compare-to results` to recompute the checks. The README gives the environment and coverage; the manifest and manuscript hash identify the accompanying version. These checks verify the stated finite computations, not asymptotic source density or a production implementation of the full algorithm.

Acknowledgments. The authors used ChatGPT for language and presentation editing, document preparation, auxiliary verification code, and the development and checking of the accompanying Lean 4 formalization.

References

1. Berlekamp, E.R.: Factoring polynomials over large finite fields. Mathematics of Computation **24**(111), 713–735 (1970). https://doi.org/10.1090/S0025-5718-1970-0276200-X
2. Beullens, W.: Graph-theoretic algorithms for the alternating trilinear form equivalence problem. In: Advances in Cryptology – CRYPTO 2023. LNCS, vol. 14083, pp. 101–126. Springer (2023). https://doi.org/10.1007/978-3-031-38548-3_4
3. Bläser, M., Duong, D.H., Narayanan, A.K., Plantard, T., Qiao, Y., Sipasseuth, A., Tang, G.: The ALTEQ Signature Scheme: Algorithm Specifications and Supporting Documentation. NIST-hosted first-round submission specification (2023), https://csrc.nist.gov/csrc/media/Projects/pqc-dig-sig/documents/round-1/spec-files/ALTEQ-Spec-web.pdf, document version 2023-06-01
4. Bouillaguet, C., Fouque, P.A., Véber, A.: Graph-theoretic algorithms for the "isomorphism of polynomials" problem. In: Advances in Cryptology – EUROCRYPT 2013. LNCS, vol. 7881, pp. 211–227. Springer (2013). https://doi.org/10.1007/978-3-642-38348-9_13
5. Brooksbank, P.A., Li, Y., Qiao, Y., Wilson, J.B.: Improved algorithms for alternating matrix space isometry: From theory to practice. In: 28th Annual European Symposium on Algorithms (ESA 2020). LIPIcs, vol. 173, pp. 26:1–26:15. Schloss Dagstuhl – Leibniz-Zentrum für Informatik (2020). https://doi.org/10.4230/LIPIcs.ESA.2020.26
6. Cantor, D.G., Zassenhaus, H.: A new algorithm for factoring polynomials over finite fields. Mathematics of Computation **36**(154), 587–592 (1981). https://doi.org/10.1090/S0025-5718-1981-0606517-5
7. Chou, T., Niederhagen, R., Persichetti, E., Ran, L., Randrianarisoa, T.H., Reijnders, K., Samardjiska, S., Trimoska, M.: Matrix Equivalence Digital Signature. NIST-hosted first-round submission specification (2023), https://csrc.nist.gov/csrc/media/Projects/pqc-dig-sig/documents/round-1/spec-files/MEDS-spec-web.pdf, version 1.0.1, 30 June 2023
8. Chou, T., Niederhagen, R., Persichetti, E., Randrianarisoa, T.H., Reijnders, K., Samardjiska, S., Trimoska, M.: Take your MEDS: Digital signatures from matrix code equivalence. In: Progress in Cryptology – AFRICACRYPT 2023. LNCS, vol. 14064, pp. 28–52. Springer (2023). https://doi.org/10.1007/978-3-031-37679-5_2
9. Couvreur, A., Levrat, C.: Highway to hull: An algorithm for solving the general matrix code equivalence problem. In: Advances in Cryptology – CRYPTO 2025. LNCS, vol. 16000, pp. 253–283. Springer (2025). https://doi.org/10.1007/978-3-032-01855-7_9
10. Erdélyi, M., Sawin, W., Tóth, Á.: The purity locus of matrix Kloosterman sums. Transactions of the American Mathematical Society **377**(6), 4117–4132 (2024). https://doi.org/10.1090/tran/9149, theorem numbering: arXiv:2210.07133v2
11. Erdélyi, M., Tóth, Á.: Matrix Kloosterman sums. Algebra & Number Theory **18**(12), 2247–2308 (2024). https://doi.org/10.2140/ant.2024.18.2247, result numbering: arXiv:2109.00762v4
12. von zur Gathen, J., Gerhard, J.: Modern Computer Algebra. Cambridge University Press, 3 edn. (2013). https://doi.org/10.1017/CBO9781139856065
13. Grochow, J.A., Qiao, Y.: On the complexity of isomorphism problems for tensors, groups, and polynomials I: Tensor isomorphism-completeness. SIAM Journal on Computing **52**(2), 568–617 (2023). https://doi.org/10.1137/21M1441110

14. Grochow, J.A., Qiao, Y., Tang, G.: Average-case algorithms for testing isomorphism of polynomials, algebras, and multilinear forms. Journal of Groups, Complexity, Cryptology **14**(1), 1:1–1:22 (2022). https://doi.org/10.46298/jgcc.2022.14.1.9431
15. Ji, Z., Qiao, Y., Song, F., Yun, A.: General linear group action on tensors: A candidate for post-quantum cryptography. In: Theory of Cryptography – TCC 2019. LNCS, vol. 11891, pp. 251–281. Springer (2019). https://doi.org/10.1007/978-3-030-36030-6_11, definition numbering: arXiv:1906.04330v1
16. Katz, N.M.: Gauss Sums, Kloosterman Sums, and Monodromy Groups, Annals of Mathematics Studies, vol. 116. Princeton University Press, Princeton (1988). https://doi.org/10.1515/9781400882120
17. Lang, S.: Algebraic groups over finite fields. American Journal of Mathematics **78**(3), 555–563 (1956). https://doi.org/10.2307/2372673
18. Li, T., Li, Y., Qiao, Y., Tao, D., Wang, Y.: On the average-case complexity landscape for Tensor-Isomorphism-complete problems over finite fields. arXiv:2604.00591v1 (2026). https://doi.org/10.48550/arXiv.2604.00591
19. Li, Y., Qiao, Y.: Linear algebraic analogues of the graph isomorphism problem and the Erdős–Rényi model. In: 2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS). pp. 463–474. IEEE (2017). https://doi.org/10.1109/FOCS.2017.49
20. Narayanan, A.K.: Strong keys for tensor isomorphism cryptography. In: 50th International Symposium on Mathematical Foundations of Computer Science (MFCS 2025). LIPIcs, vol. 345, pp. 78:1–78:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik (2025). https://doi.org/10.4230/LIPIcs.MFCS.2025.78
21. Narayanan, A.K., Qiao, Y., Tang, G.: Algorithms for matrix code and alternating trilinear form equivalences via new isomorphism invariants. In: Advances in Cryptology – EUROCRYPT 2024. LNCS, vol. 14653, pp. 160–187. Springer (2024). https://doi.org/10.1007/978-3-031-58734-4_6
22. Schwartz, J.T.: Fast probabilistic algorithms for verification of polynomial identities. Journal of the ACM **27**(4), 701–717 (1980). https://doi.org/10.1145/322217.322225
23. Shimomura, N.: A theorem on the fixed point set of a unipotent transformation on the flag manifold. Journal of the Mathematical Society of Japan **32**(1), 55–64 (1980). https://doi.org/10.2969/jmsj/03210055
24. Shimomura, N.: The fixed point subvarieties of unipotent transformations on the flag varieties. Journal of the Mathematical Society of Japan **37**(3), 537–556 (1985). https://doi.org/10.2969/jmsj/03730537
25. Tang, G., Duong, D.H., Joux, A., Plantard, T., Qiao, Y., Susilo, W.: Practical post-quantum signature schemes from isomorphism problems of trilinear forms. In: Advances in Cryptology – EUROCRYPT 2022. LNCS, vol. 13277, pp. 582–612. Springer (2022). https://doi.org/10.1007/978-3-031-07082-2_21
26. Van Loan, C.F., Pitsianis, N.: Approximation with Kronecker products. In: Linear Algebra for Large Scale and Real-Time Applications, NATO ASI Series E, vol. 232, pp. 293–314. Kluwer Academic Publishers, Dordrecht (1993). https://doi.org/10.1007/978-94-015-8196-7_17
27. Zippel, R.: Probabilistic algorithms for sparse polynomials. In: Symbolic and Algebraic Computation (EUROSAM 1979). LNCS, vol. 72, pp. 216–226. Springer (1979). https://doi.org/10.1007/3-540-09519-5_73

Supplementary material

Appendices A–D support the main almost-all theorem; Appendices E–L give the proofs, auxiliary algorithms, and quantitative certificates for the complementary finite-parameter results.

A Complete-fibre cancellation over every finite field

We prove Lemma 4, keeping the coefficient orientations and all-field constants explicit.

A.1 Asymmetric border pivot

Normalize $U = \operatorname{diag}(I_s, 0_{a-s})$ and write the summation matrix as $X = \left(\begin{smallmatrix} A & B \\ C & D \end{smallmatrix}\right)$ with A of size s. The case $s = 0$ is (22). Put $j = \operatorname{rank} A$ and $h = s - j$. Fix row and column reductions of A to $\operatorname{diag}(I_j, 0_h)$ as functions of A alone. Invertibility of X forces the last h rows B_0 and columns C_0 of the reduced borders to have rank h. Choose the first h bottom columns and rows giving invertible minors of B_0, C_0 in a fixed subset ordering. These choices depend only on A, B, C.

The resulting pivot P containing the first s rows and columns has size $b = s + h \le 2s$. After eliminating I_j, its remaining $2h$-block has zero upper-left block and invertible off-diagonal blocks. Its determinant is $(-1)^h \det B_{0,J} \det C_{0,I}$ up to the fixed nonzero reduction factors, and hence is independent of the bottom block D. For each h there are at most $\binom{a-s}{h}^2$ disjoint charts. Permute bottom rows and columns separately to place P first. The support of U stays inside this pivot.

Write the reordered matrix as $\left(\begin{smallmatrix} P & B' \\ C' & D' \end{smallmatrix}\right)$, put $J = D' - C'P^{-1}B'$, and let $\ell = a - b$. Fixing P, B', C' determines the whole chart condition; the only remaining condition is $J \in \mathrm{GL}_\ell$. Its phase is $\theta + \operatorname{tr}(V_0 J^{-1})$, with no J-linear term because U is supported in P. Here

$$V_0 = R_0 V' L_0, \qquad L_0 = \begin{pmatrix} -P^{-1}B' \\ I_\ell \end{pmatrix}, \quad R_0 = \begin{pmatrix} -C'P^{-1} & I_\ell \end{pmatrix}.$$

Both outer maps have rank ℓ, so $\operatorname{rank} V_0 \ge a - c - 2b = \ell - (b + c)$. For actual corank d of V_0, (22) gives magnitude $q^{\ell(\ell-1)/2 + d(d+1)/2} \prod_{i=1}^{d} (1 - q^{-i})$. Since $d \le b + c < \ell$, the saving is at least

$$\begin{aligned} \frac{\ell(\ell+1) - (b+c)(b+c+1)}{2} &= \frac{(a - 2b - c)(a + c + 1)}{2} \\ &\ge \frac{(a - 4s - c)(a + c + 1)}{2}. \end{aligned}$$

There are at most $q^{a^2 - \ell^2}$ outer assignments per chart. Sum the full residual character before taking absolute values, multiply by this count, and use $\sum_h \binom{a-s}{h}^2 \le 4^{a-s}$. This proves (24).

A.2 Two coefficient coranks

We use dimension a and corank bounds $b, c \geq 1$. Put $t = a - 2b - c + 1$. For $t \leq 0$ use the trivial bound; if one coefficient is invertible use Lemma 3, since $t_+ \leq a$. Replacing b by the actual positive corank of U strengthens the target. Normalize $U = \mathrm{diag}(I_r, 0_b)$, $r = a - b$, invert the summation variable, and write the phase as $\mathrm{tr}(UY^{-1} + VY)$. Let M be the upper-left r-square block of V.

Fix the bottom b-square block of Y of rank $b-j$ and reduce it to $\mathrm{diag}(I_{b-j}, 0_j)$ by bottom row and column operations, which preserve M. Eliminate this identity block. There are $q^{2r(b-j)}$ choices for the eliminated borders, and the remaining matrix is $\widetilde{Y} = \left(\begin{smallmatrix} A' & B_0 \\ C_0 & 0_j \end{smallmatrix}\right)$. For $j = 0$, the remaining inverse coefficient is I_r; counting all outer variables bounds the complete sector by $C_r q^{a^2 - r^2/4}$.

For $j \geq 1$, B_0, C_0 must both have rank j. With $i : \ker C_0 \hookrightarrow k^r$ and $\pi : k^r \to k^r / \operatorname{im} B_0$, the induced $E = \pi A' i$ has size $N = r - j$, and its invertibility is equivalent to that of $\widetilde{Y}$. The upper inverse block is $iE^{-1}\pi$, while the A'-fibre has dimension $2rj - j^2$ and difference space $\{B_0\alpha + \beta C_0\}$. Summing its entire linear character annihilates the summand unless

$$MB_0 = 0, \qquad C_0 M = 0.$$

These joint equations are retained together. In the surviving case $M = iL\pi$, $\operatorname{rank} L = \operatorname{rank} M \geq a - c - 2b$, and $J = \pi i$ has rank at least $N - j$. The left and right kernels of M have dimensions at most $b + c$, so at most $q^{2j(b+c)}$ pairs survive. The residual complete sum is $\mathrm{Kl}_N(L, J)$ times a unit-modulus phase and $q^{2rj - j^2}$.

By (21), the number of bottom blocks of rank $b - j$ is at most $4q^{b^2 - j^2}$. All counting exponents add to $a^2 - \Delta_j$, where $\Delta_j = 2j(t-1) + j^2$. Induction, with corank bounds $j, b + c - j$ and exchanged coefficients if needed, gives residual saving $((t - 2j)_+)^2/4$. An invertible residual coefficient uses Lemma 3 instead. A sector with $j > b + c$ cannot survive; also $t > 0$ implies $N \geq c \geq 1$ in every relevant sector. For integer $t, j \geq 1$,

$$\Delta_j + ((t - 2j)_+)^2/4 \geq t^2/4.$$

For $t \geq 2j$ the difference is $jt + 2j(j-1) \geq 0$; on $1 \leq t \leq 2j$ it is concave with nonnegative endpoint values $j^2 - 1/4$ and $4j^2 - 2j$. The $j = 0$ sector also obeys the desired saving since $r \geq t$. There are at most b singular sectors, each with $N \leq a - 2$. The constants close for $a \geq 3$ because

$$\frac{C_{a-1} + 4bC_{a-2}}{C_a} = \frac{1}{2a} + \frac{b}{a(a-1)} \leq \frac{1}{2a} + \frac{1}{a-1} \leq \frac{2}{3} < 1.$$

This proves (25) in all characteristics and for $q = 2, 3$ as well as larger fields.

B Total cohomology dimension

Fix an auxiliary prime $\ell \neq \mathrm{char}\, k$ and geometric $\overline{\mathbb{Q}}_\ell$-cohomology. For $D \in \mathrm{GL}_r(k)$ set $\mathcal{H}^j(D) = H^j_c(\mathrm{GL}_{r,\bar{k}}, \mathcal{L}_\psi(\mathrm{tr}(DX^{-1} + X)))$ and $\nu(t) = (3t^2 - (t \bmod 2))/2$; inversion identifies this phase with K_r.

Cohomological degree bound. For primary $\alpha I + J$ with $\alpha \neq 0$, Theorem 3.19, Lemma 3.20 and Corollary 3.21 of [11]'s arXiv v4 give $\mathcal{H}^j = 0$ for $j > \nu(t)$. Over a finite splitting extension, Proposition 3.18 decomposes the complex into nonzero-primary blocks of sizes t_a and affine factors of total dimension $\sum_{a<b} t_a t_b$. The degree is at most $\sum_a \nu(t_a) + 2\sum_{a<b} t_a t_b \leq \nu(r)$, since the difference is $\sum_{a<b} t_a t_b + (\sum_a (t_a \bmod 2) - (r \bmod 2))/2 \geq 0$. Geometric vanishing is invariant under finite extension.

Affine paving and rank of the coefficient system. Let $\mathcal{F}_D$ be the projective fixed complete-flag variety. A nilpotent J and $I + J$ fix the same flags. Shimomura's Schubert intersections are affine [23,24]; see Definition 4.1, Proposition 4.5, Theorem 4.7 and Corollary 4.8 of [24]. For type λ, distinct tableau labels increase along each Jordan row. Assigning labels to rows gives

$$\#\text{pieces} = t!/\prod_i \lambda_i!. \tag{61}$$

Equal-length rows remain distinct. Order Schubert cells by a linear extension of Bruhat order, lower cells first. Each prefix is a closed union of Schubert varieties; intersecting the closed fixed locus gives a closed filtration with these affine differences.

Over $\bar{k}$, generalized-eigenspace projectors are polynomials in D of locally constant ranks on tautological subspaces. Interleaving words therefore give open-and-closed products of primary flag varieties. For eigenspace dimensions r_α, the total cell count is $(r!/\prod_\alpha r_\alpha!)\prod_\alpha (r_\alpha!/\prod_i \lambda_{\alpha i}!) \leq r!$. The compact-support exact sequences give zero odd cohomology and one total Betti contribution per piece. Projectivity yields

$$\sum_j \dim H^j(\mathcal{F}_D, \overline{\mathbb{Q}}_\ell) \leq r!. \tag{62}$$

This is geometric: cells need not descend to k and no comparison with complex topology is used.

Put $\mathcal{L}_D = \bigotimes_{i=1}^r \lambda_i^* \mathcal{K}\!\updownarrow_2$. After forgetting Frobenius actions and Tate twists, the direct-summand statement of Corollary 9 of [10] gives

$$\dim \mathcal{H}^j(D) \leq \dim H^{j-r^2}(\mathcal{F}_D, \mathcal{L}_D).$$

The flag-quotient eigenvalue maps λ_i have finite image and are constant on geometric components. For invertible D their values lie in $\mathbb{G}_m$, where $\mathcal{K}\!\updownarrow_2$ is lisse of rank two [16, Theorem 4.1.1]. Thus $\mathcal{L}_D$ is geometrically constant of rank 2^r on each component. Summing over degrees, the cohomological shift only reindexes the sum, so

$$\sum_j \dim \mathcal{H}^j(D) \leq 2^r \sum_j \dim H^j(\mathcal{F}_D, \overline{\mathbb{Q}}_\ell) \leq 2^r r!$$

by (62). This dimension estimate does not use Frobenius weights or a sign convention for Tate twists. Independently, Theorem 1 of [10] makes the original complex $\mathcal{H}^*(D)$ pure of weight zero, since D is invertible. Its degree-j Frobenius

eigenvalues therefore have absolute value $q^{j/2}$. The trace formula and degree cutoff give $|K_r(D)| \le 2^r r! q^{\nu(r)/2}$, as Lemma 3 requires. Singular matrices use the nilpotent Fitting-factor formula [11, Theorem 1.2].

C All-field bilinear fibre estimates

We prove (49). The matrix dimension here is t, independent of the input-index size s in the rectangular array. Frequencies are modulo scalars when the corresponding comparison space is trace zero.

C.1 Full-matrix ranks and overlap costs

For nontrivial cubic frequency V, set $d = \min_{\lambda \in \bar{k}} \operatorname{rank}(V - \lambda I)$ and $\beta_V(Y, A) = (VY + YV)A + YAV$. Fix an invertible A_0. Multiplication by A_0^{-1} gives the Sylvester operator $Y \mapsto VY + Y(V + A_0 V A_0^{-1})$. Over a splitting field, each pair of Jordan blocks contributes at most the minimum of their sizes to its kernel. V has at most $t-d$ blocks per eigenvalue, so the kernel dimension is at most $t(t-d)$ and the rank at least td. Completing a corank-e input to A_0 changes the image by a subspace of dimension at most $2te$, using its two row-space constraints. Hence the zero-fibre probability for independent full matrices is at most

$$F_d \le 4(t+1) q^{-\lceil d/2 \rceil^2}, \qquad \min_{0 \le e \le t} \{e^2 + t(d-2e)_+\} = \lceil d/2 \rceil^2. \tag{63}$$

For odd $d = 2h+1$, compare $h^2 + t$ and $(h+1)^2$ and use $t \ge d$; even d is immediate. The corank probability is at most $4q^{-e^2}$.

If $d < t/2$, the eigenvalue with eigenspace dimension greater than $t/2$ is unique and Frobenius fixed, hence belongs to k. Write $V = \lambda I + R$. For $3\lambda \ne 0$, the main map $3\lambda YA$ has rank $t(t-e)$. The remaining image lies in a column-space-d plus row-space-$2d$ space, of dimension at most $3td - 2d^2$. Completing the square in the corank sum yields

$$S_d \le 4(t+1) q^{-3t^2/4 + 3td - 2d^2}. \tag{64}$$

In the residual case $3\lambda = 0$, $1 \le d < t/2$, factor $R = UC$, where U is injective and C surjective. For uniform full Y let

$$b = \operatorname{rank}(\operatorname{im} U \xrightarrow{Y} k^t / \operatorname{im} U), \quad c = \operatorname{rank}(\ker C \xrightarrow{Y} k^t \xrightarrow{C} k^d), \quad e = \operatorname{corank} Y.$$

The matrices $[U, YU]$ and $\binom{CY}{C}$ have ranks $d+b, d+c$. Sylvester's inequality gives $\operatorname{rank}(RY + YR) \ge b + c$. Restricting A to $AU = 0$ gives an image of dimension at least $(b+c)(t-d)$. Projecting modulo $\operatorname{im}[U, YU]$ gives the independent extra dimension $d(t - e - d - b)_+$, since $A \mapsto AU$ and C are surjective. Thus the full output rank is at least

$$r(b, c, e) = (b+c)(t-d) + d(t - e - d - b)_+. \tag{65}$$

Put $U_0 = \operatorname{im} U$ and $K_0 = \ker C$. A matrix in the (b,c) stratum admits subspaces $S_0 \le U_0$, $T_0 \le K_0$ of dimensions $d-b$ and $t-d-c$ with $Y(S_0) \subseteq U_0$ and $Y(T_0) \subseteq K_0$. For fixed S_0, T_0, these linear conditions have codimensions $(d-b)(t-d)$ and $d(t-d-c)$. In the dual of $\mathrm{End}(k^t)$, their equation spaces are $S_0 \otimes U_0^{\perp}$ and $T_0 \otimes K_0^{\perp}$, where $\perp$ denotes the annihilator in $(k^t)^*$. Their overlap has dimension

$$\dim(S_0 \cap T_0)\, \dim(U_0^{\perp} \cap K_0^{\perp}) \le (d-b)d.$$

The numbers of possible subspaces are at most $4q^{b(d-b)}$ and $4q^{c(t-d-c)}$, respectively, by (21). Subtracting these choice exponents from the joint codimension gives

$$(d-b)(t-d) + d(t-d-c) - d(d-b) - b(d-b) - c(t-d-c) = L_{bc}.$$

A union bound therefore gives $16q^{-L_{bc}}$, where

$$L_{bc} = (d-b)(t-2d-b) + (d-c)(t-d-c). \tag{66}$$

Intersecting with the corank event, take the minimum of the two probability bounds, equivalently the maximum of their exponent savings after bounding both prefactors by 16. The stratum probability is at most $16q^{-\max(L_{bc}, e^2)}$. Therefore

$$R_d \le 16(t+1) \sum_{b,c=0}^{d} q^{-\min_{0 \le e \le t}\{\max(L_{bc}, e^2) + r(b,c,e)\}}. \tag{67}$$

The exponents of the overlapping conditions are not added.

C.2 Uniform absorption in output dimension

Set $D_0 = \lfloor (3t+2)/10 \rfloor \le t/3$. For $d > D_0$, (63) saves at least $(3t+3)^2/400 \ge 41t/16$. For $d \le D_0$ and $3\lambda \ne 0$, (64) saves at least $3t^2/100 - 9t/25 + 2/25 \ge 41t/16$. After clearing denominators the differences have positive polynomials $9t^2 - 1007t + 9$ and $12t^2 - 1169t + 32$ for $t \ge 1024$.

For the residual case put $x = d-b, z = d-c$. Then $0 \le L = x(t-3d+x) + z(t-2d+z) \le 2dt$ and

$$r + L \ge 3dt - 4d^2 + x^2 - dx + z^2 - dz - de.$$

Using $\max(L, e^2) - de \ge L - d\sqrt{L} - d^2/4$ and $x^2 - dx, z^2 - dz \ge -d^2/4$, the saving is at least

$$E(d) = 3dt - 19d^2/4 - \sqrt{2t}\, d^{3/2}.$$

It is concave on $[1, t/3]$. At the second endpoint it is $t^2(17/36 - \sqrt{2/27}) \ge t^2/6 \ge 3t$; thus $E(d) \ge E(1) = 3t - 19/4 - \sqrt{2t}$. Since $(t+1)(d+1)^2 \le t^3$ here, the prefactor is at most $16t^3$. For all $q \ge 2, t \ge 1024$,

$$4q^3(t+1) \le q^{t/16}, \quad 16q^3t^3 \le q^{t/4}, \quad 19/4 + \sqrt{2t} \le t/4. \tag{68}$$

For the first two it suffices to check $q = 2, t = 1024$ and compare consecutive ratios: $(1 + 1/1025)^{16} < 2$ and $(1 + 1/1024)^{12} < 2$. For the third square $\sqrt{2t} \le (t - 19)/4$; the resulting quadratic is positive and increasing from 1024. These inequalities leave saving at least $5t/2$ in every sector.

For fixed first input, a shifted linear fibre is empty or has the same size as the zero fibre. Restricting Y, A each to trace zero costs at most q^2, and membership in a scalar coset costs at most q. These are exactly the q^3 costs in (68), proving (49). In characteristic three the cubic dual is modulo scalars and a nontrivial near-scalar class has $d \ge 1$. Otherwise a nonzero scalar frequency uses S_0, and the residual case has $\lambda = 0, d \ge 1$. Thus no characteristic or dual frequency is omitted.

C.3 The quadratic Sylvester bound

For completeness, for nonscalar U the operator $Y \mapsto UY + YU$ has rank at least $2t - 2$. Extend scalars to split U; rank does not change. If distinct eigenvalues occur, matching opposite eigenvalues bounds the kernel by $(t - 1)^2 + 1$: among decompositions with at least two eigenspaces, the sum of squared dimensions is largest at $(t - 1, 1)$, and opposite pairing cannot exceed it. If there is only one eigenvalue and twice it is nonzero, the operator is invertible. In the remaining single-eigenvalue case, the nonscalar nilpotent part has Jordan partition with at least one block of size two. The kernel dimension is the sum of squares of its column lengths, at most $(t-1)^2+1$. Replacing one nilpotent by its negative does not change the Jordan sizes. This proves the rank bound uniformly, including characteristic two. Restricting the input to trace zero and taking the trace-dual output modulo scalars loses at most two ranks, as used in (48).

D A parity-free local nonvanishing construction

We give the indexed construction for Proposition 3. Fix an arbitrary prime p and $n \ge 4$. All constructions take place over rational function fields over $\mathbb{F}_p$, followed by formal Laurent-series extensions. Their role is solely to establish nonzero numerator polynomials over $\mathbb{F}_p$.

D.1 Diagonal base curve

Take independent t_i, w_{ij} for $i \ne j$ and set

$$A^0_{ijji} = 1 \ (i \ne j), \qquad A^0_{iiii} = t_i, \qquad B_{ijij} = w_{ij} \ (i \ne j),$$

with all other entries zero. Put $A(\epsilon) = A^0 + \epsilon B$. $F(0), G(0)$ swap off-diagonal basis vectors and multiply e_{ii} by t_i, so they are invertible. The diagonal restriction of $Q(\epsilon)$ is the matrix C with diagonal t_i and all off-diagonal entries one. Its determinant has monomial $\prod_i t_i$ with coefficient one in every characteristic. Off the diagonal, $Q(\epsilon)e_{ij} = \epsilon w_{ji} e_{ji}$. Hence

$$Q(\epsilon)^{-1} = \epsilon^{-1} R + R_0, \qquad R e_{ij} = w_{ij}^{-1} e_{ji} \ (i \ne j),$$

where R is zero on the diagonal and $R_0 = C^{-1}$ there, zero elsewhere. All three flattenings are invertible over Laurent series.

Direct crossing gives $X(A^0) = I$. Both crossed operators preserve each off-diagonal unordered-pair subspace and each individual diagonal line. The partial traces of their words are therefore diagonal along this curve. The universal identities

$$\operatorname{ptr}_1 X = \operatorname{ptr}_2 X = \operatorname{ptr}_1 Y = \operatorname{ptr}_2 Y = nI \tag{69}$$

follow from a row of G or Q contracted against its corresponding inverse column, summed over the remaining index. They require no division by n. Let $r_{ij} = 1 - (t_i t_j)^{-1}$. The coefficient L of ϵ in X and the coefficient V of ϵ^{-1} in Y are

$$L_{ij,ji} = w_{ij} r_{ij}, \qquad V_{ij,ji} = w_{ji}^{-1} \quad (i \neq j),$$

and zero elsewhere. By (69), all higher Taylor coefficients of $X - I$ have zero partial traces. Thus for each traced mode the X^2 and Y^2 diagonals have the expansions

$$a_j(\epsilon) = n + \epsilon^2 c_j + O(\epsilon^3), \qquad c_j = \sum_{h \neq j} w_{jh} w_{hj} r_{jh}^2, \tag{70}$$

$$b_j(\epsilon) = \epsilon^{-2} d_j + O(\epsilon^{-1}), \qquad d_j = \sum_{h \neq j} (w_{jh} w_{hj})^{-1}. \tag{71}$$

These leading diagonals agree between modes because their products are symmetric. For $j \neq l$, choose $h \notin \{j, l\}$. The variable w_{jh} has a unique positive first power in $c_j - c_l$ and a unique negative first power in $d_j - d_l$. These differences are nonzero, so the local X^2 diagonals are pairwise distinct in every characteristic.

D.2 Transverse entries in both local modes

For pairwise distinct i, j, l, let $S_{iji} = 1$ be the sole nonzero tensor entry of S. Differentiate at zero in the perturbation $A(\epsilon) + zS$. At $\epsilon = 0$, the derivative of X has sole entry

$$(L_S)_{ij,li} = a_i^*, \qquad a_i^* = 1 - t_i^{-1}.$$

Indeed the direct term is $e_{ij} e_{li}^t$, and the inverse derivative $-G(0)^{-1} S_G G(0)^{-1}$ gives $-t_i^{-1} e_{ij} e_{li}^t$ after crossing. Identity (69) removes the partial traces of the remaining linear derivative terms. The leading X^2 derivative is $\epsilon \operatorname{ptr}(LL_S + L_S L)$, with sole (j, l) entry

$$u_1 = w_{ji} r_{ij} a_i^*, \qquad u_2 = w_{li} r_{il} a_i^*$$

in modes $1, 2$, respectively.

The leading derivative of Y is $\epsilon^{-2} Z$, where

$$Z = -F(0) \operatorname{cross}_Q(R S_Q R) = -(w_{il} w_{ij})^{-1} e_{ji} e_{il}^t.$$

Here $S_Q = e_{il}e_{ji}^t$ and $RS_QR = (w_{il}w_{ij})^{-1}e_{li}e_{ij}^t$, which checks the two transposes directly. Therefore the leading local Y^2 derivative is $\epsilon^{-3}\operatorname{ptr}(VZ+ZV)$ with sole (j,l) entries

$$v_1 = -(w_{ij}w_{il}w_{li})^{-1}, \qquad v_2 = -(w_{ij}w_{il}w_{ji})^{-1}.$$

For either mode, the transverse difference is

$$(a_j-a_l)(Y^2\text{-local})'_{jl} - (b_j-b_l)(X^2\text{-local})'_{jl} = \epsilon^{-1}\big((c_j-c_l)v_\nu - (d_j-d_l)u_\nu\big) + O(1). \tag{72}$$

Choose h distinct from i,j,l, possible for $n \geq 4$. The coefficient of the positive first power of w_{jh} in the bracket is exactly $v_\nu w_{hj} r_{jh}^2 \neq 0$. The d difference has only the negative first power, and u_ν, v_ν do not depend on w_{jh}. Thus (72) is nonzero in both modes, also in characteristic two.

D.3 Tree transversality and polynomial degree

Suppose $C(0) = \operatorname{diag}(a_i)$ with all a_i distinct and $D(0) = \operatorname{diag}(b_i)$. If $(a_j-a_l)D'_{jl} - (b_j-b_l)C'_{jl} \neq 0$ on a spanning tree, then the pair $C(\tau), D(\tau)$ has scalar centralizer over Laurent series. To prove this, take the commutator map on M_n/kI in the quotient basis specified in the introduction. Select its $n(n-1)$ off-diagonal rows from $[Z,C]$ and its $n-1$ tree-edge rows from $[Z,D]$. The off-diagonal block is invertible at $\tau = 0$. After its elimination the first coefficient on diagonal variables is a weighted reduced tree-incidence matrix, with nonzero weights equal to the transverse entries divided by $a_j - a_l$. Its reduced incidence determinant is ± 1. The selected maximal minor therefore has a nonzero coefficient at order τ^{n-1}, in every characteristic.

For each tree edge choose the tensor S above and an independent coefficient z_j. The corresponding transverse entry has a nonzero z_j coefficient, so it is nonzero over the joint rational function field. A finite product of these nonzero functions is nonzero, ensuring simultaneous transversality in both modes. Apply the lemma to the tensor $A(\epsilon) + \tau\sum_j z_j S_j$.

Finally the adjugate numerators of X, Y have degree m. Clear rows for X^2, Y^2, XY, YX by $(\det G)^2, (\det Q)^2, \det G \det Q$ as appropriate. These denominators are nonzero on the formal construction. The selected numerator minor in either mode has degree at most $2m(m-1)$ and is not identically zero over $\mathbb{F}_p$. Row selections may depend on p, but the degree bound does not. Multiply the two minors and $\det F \det G \det Q$ to obtain degree at most $4m(m-1)+3m = 4n^4 - n^2$. This is a genericity proof, not a claim that the displayed indeterminates can be specialized successfully in every small field. The polynomial zero bound over $\mathbb{F}_q$ gives exactly the raw assertion of Proposition 3, including its possibly vacuous small-q range.

Finite-field nonemptiness in odd dimension. For every odd n and every finite field, the 0/1 tensor indexed by $\mathbb{Z}/n\mathbb{Z}$ with $A_{ijkl} = 1$ exactly when $k = i+j$ and $l = i+2j$ has all three original flattenings permutation matrices. Indeed each

balanced pair of indices determines the other pair uniquely; only the pairing involving $2j$ uses the invertibility of 2 modulo n. This is an explicit $\mathcal{S}_{\mathrm{full}}$ witness over the original field, independently of the characteristic. It asserts no additional $\mathcal{H}_{\mathrm{loc}}$ membership.

E Three rational anchors: recovery and stabilizers

The spectral procedure analyzed here is pointwise valid for $A \in \mathcal{R}_3$ over $k = \mathbb{F}_q$, $n \geq 2$. Write $m = n^2$, $p = \operatorname{char} k$, $e = [k : \mathbb{F}_p]$, $\ell = \lceil \log_2 q \rceil$, and $d = \gcd(n, q-1)$. For $T = (\bigotimes_i g_i)A$, put $L = g_1 \otimes g_2$, $K = g_3 \otimes g_4$. The shared covariance (7) preserves rational eigenvalue labels. The partial-trace identities (69) hold even when $p \mid n$; they do not make the law of X uniform.

Scope of the procedure. This appendix specifies the three-anchor algorithm and its source-specific stabilizer and operation counts. The complementary selection used in Section 8.2 may first test $\mathcal{H}_{\mathrm{loc}}$, then select by $\operatorname{rank} W$ between $\mathcal{R}_3$ and $\mathcal{S}_{\mathrm{det}}$ (or $\mathcal{S}_{\mathrm{CAS}}$). All these families lie in $\mathcal{S}_{\mathrm{full}}$. This selection is separate from Definition 3.

Lemma 8 (Characteristic polynomial in $O(m^3)$ field operations). *For $M \in \mathrm{M}_m(k)$, its characteristic polynomial takes $O(m^3)$ classical field operations and $O(m^2)$ stored field entries, in every characteristic.*

Proof. Reduce the matrix to upper Hessenberg form by similarities, then compute successive leading characteristic polynomials by last-column expansion. Both stages use $O(m^3)$ field operations. The recurrence and in-place similarity check are given in Appendix L.

Finding all roots in k. The following routines use Cantor–Zassenhaus splitting and the absolute-Frobenius method [6,1,12]. For nonzero f of degree at most h, modular powering computes $f_{\mathrm{field}} = \gcd(f, x^q - x)$ in $O(h^2\ell)$ operations, retaining all distinct rational roots of repeated or inseparable inputs. For $p \geq 7$, recursively split using uniform residue polynomials and gcds with the residue or its $(q-1)/2$ power minus one. CRT evaluations at two distinct roots are independent uniform elements; opposite nonzero quadratic characters have probability $(q-1)^2/(2q^2) \geq 18/49 > 1/3$ and ensure a split. Conditionally on every history, each of at most $h-1$ successful splits therefore costs fewer than three $O(h^2\ell)$ trials in expectation. Product/root verification gives always-correct Las Vegas cost $R_{\mathrm{LV}}(h) = O(h^3\ell)$ pointwise, with almost-sure termination.

For $p \in \{2, 3, 5\}$, the p-power Frobenius map on the split algebra $k[x]/(f_{\mathrm{field}})$, of degree $a \leq h$, has fixed space $\mathbb{F}_p^a$. A basis separates CRT coordinates; gcd refinement by each basis element minus each of the p constants yields singleton factors. Building Frobenius on ae prime coordinates costs $O(a^3 e)$ operations in k, elimination $O((ae)^3)$ in $\mathbb{F}_p$, and refinement $O(a^3)$ in k. Separately formed polynomial-basis powers cost $O(e \log(e+1))$, giving deterministic

$$\begin{aligned} R_{\mathrm{fixed}}(h) = O(h^2\ell + h^3 e + e\log(e+1)) \quad & \text{operations in } k, \\ + O((he)^3) \quad & \text{operations in } \mathbb{F}_p. \end{aligned} \tag{73}$$

Prime-coordinate storage is $O((he)^2)$.

Recognition. Check invertibility of F, G, Q, W, form X, and list its nonzero rational labels. At each label solve the eigenspace, retaining only geometric dimension one and invertible reshapes, normalized by the first nonzero entry. Write $r \le \min(m, q-1)$ for the label count and $s \le r$ for the number of usable anchors. For anchors P, Q, R, the pair $U = QP^{-1}, V = RP^{-1}$ changes to U^{-1}, VU^{-1} when based at Q. Each pair recovers the other by inverses/products, preserving its centralizer. Swapping ratios proves invariance under all six permutations. Thus lexicographic $i < j < k$ tests retain the first passing ordered triple for any fixed candidate order: each set's earliest permutation is increasing. Exhausting $\binom{s}{3}$ tests is exactly rejection from $\mathcal{R}_3$. Store anchor inverses; for each base i, reuse all $P_j P_i^{-1}$ $(j > i)$, then recycle them. This costs $O(s^2 n^3)$ operations and $O(sn^2)$ temporary entries; retain the selected inverse for recovery.

Scale searches and original-tensor lift. Fix the accepted triple. Test the target conditions and the same three labeled eigenspaces/anchors, with no new spectral-root or triple search. A true witness has $g_1 P_{Ai} g_2^t = c_i P_{Ti}$, $c_i \neq 0$. For $i = 2, 3$, put $Q_{Ai} = P_{Ai} P_{A1}^{-1}$, $Q_{Ti} = P_{Ti} P_{T1}^{-1}$; then

$$Q_{Ti} = \alpha_i g_1 Q_{Ai} g_1^{-1}, \qquad \alpha_i = c_1/c_i, \qquad \alpha_i^n = \det Q_{Ti} / \det Q_{Ai}.$$

The power map on k^* has kernel size d, so each root set has size zero or d. List both completely, including inseparable cases. When $d = 1$, the root of $x^n = a$ is a^v, $nv \equiv 1 \pmod{q-1}$, obtained deterministically by extended gcd and powering. Three nonzero labels exclude modulus one. This shortcut affects only the scalar searches. For each of at most d^2 pairs solve

$$Q_{Ti} D = \alpha_i D Q_{Ai} \quad (i = 2, 3). \tag{74}$$

If D_0 is invertible, $D \mapsto D_0^{-1} D$ identifies this space with the scalar source centralizer. Reject any space other than a line with invertible generator D; otherwise form

$$\widehat{g}_2^t = P_{A1}^{-1} D^{-1} P_{T1}, \quad \widehat{L} = D \otimes \widehat{g}_2, \quad \widehat{K} = (F_A^{-1} \widehat{L}^{-1} F_T)^t. \tag{75}$$

Reuse P_{A1}^{-1}, F_A^{-1} and $\widehat{L}^{-1} = D^{-1} \otimes \widehat{g}_2^{-1}$. In a true branch $D = s g_1$, so $\widehat{L} = c_1^{-1} L$ and $\widehat{K} = c_1 K$. Rank-one realignment splits $\widehat{K}$ at a nonzero pivot; its invertibility forces invertible factors. Accept only after reproducing every original coefficient. Every true witness survives the conditions, labels, some scale pair, intertwiner line, and lift. Failure of a necessary condition or exhaustion of the pairs therefore proves exact nonisomorphism; errors and interruptions do not.

Rational stabilizer. Let $\mathcal{G} = \mathrm{GL}_n(k)^4$ and $Z = \{(a_1 I, a_2 I, a_3 I, a_4 I) : a_j \in k^*, \prod_j a_j = 1\}$.

Theorem 7. *For every $A \in \mathcal{R}_3$, $\mathrm{Stab}_{\mathcal{G}}(A)/Z$ embeds into $\mu_n(k)^2$. Here $\mu_n(k) = \{z \in k^* : z^n = 1\}$. Its order is at most d^2, and is one if $d = 1$. For each isomorphic target, successful scalar pairs biject with witness classes modulo Z, each containing $(q-1)^3$ tuples.*

Proof. A stabilizer g preserves each selected labeled eigenline: $g_1 P_i g_2^t = c_i(g) P_i$. Composition makes the unique c_i characters. For $Q_i = P_i P_1^{-1}$, the characters $\chi_i = c_i/c_1$ satisfy $g_1 Q_i g_1^{-1} = \chi_i Q_i$, hence $\chi_i^n = 1$ by determinants, also when $p \mid n$. If $\chi_2 = \chi_3 = 1$, the scalar common-centralizer condition forces $g_1 = aI$, the first anchor forces $g_2 = bI$, and invertible F_A forces $g_3 \otimes g_4 = (ab)^{-1} I_m$. Off-diagonal blocks force g_3 diagonal; one nonzero diagonal block makes g_4 scalar, and the remaining blocks make g_3 scalar. Thus $\ker \chi = Z$; the converse is immediate from the tensor action. Cyclicity of k^* gives $|\mu_n(k)| = d$.

For $T = g_0 A$, witnesses form $g_0 \operatorname{Stab}_{\mathcal{G}}(A)$. Equal scale pairs for g, h give $\chi_i(h^{-1}g) = (c_i(g)/c_i(h))/(c_1(g)/c_1(h)) = 1$, so their quotient lies in Z; multiplying by Z preserves the pair. Conversely (75) and (74) give $D P_{A1} \widehat{g}_2^t = P_{T1}$ and $D P_{Ai} \widehat{g}_2^t = \alpha_i^{-1} P_{Ti}$. Every successful branch therefore realizes its processed pair; true-branch completeness proves bijectivity. Finally $|Z| = (q-1)^3$.

This concerns groups of k-rational points, not group schemes or sampled secrets; repeated algebraic roots impose no additional restriction.

Operation and bit bounds. Let $\tau \le \binom{s}{3}$ count common-centralizer tests. Invertibility tests, inverses, X, and its characteristic polynomial cost $O(n^6)$ field operations. Source kernels cost $O(rn^6)$; common-centralizer/intertwiner elimination on $2n^2$-by-n^2 systems costs $O(n^6)$ per test. Target tests and three kernels cost $O(n^6)$. Per surviving pair, the lift costs $O(n^6)$, realignment $O(n^4)$, and verification $O(n^5)$: four n-by-n times n-by-n^3 mode contractions followed by coefficient comparison. Since $s \le r \le n^2$, caches are absorbed in

$$O\big((1 + r + \tau + d^2)n^6 + R(m) + 2R(n)\big), \tag{76}$$

with prime-field work in (73) separate. Non-root storage is $O(n^4)$ field entries. For $p \ge 7$, the pointwise expected cost, including rejection, is $O(n^{12} + n^6 \ell)$; a recognized source costs $O((1 + d^2)n^6 + 2R(n))$ per target. The bound includes every tested source triple; no distributional assumption on the number of failed tests is used here.

In the given polynomial-basis representation, multiplication costs $O(\ell^2)$ bits and exponentiation-based inversion $O(\ell^3)$. Charging $O(\ell^3)$ per k-operation and $O(1)$ per prime-field operation for $p \in \{2, 3, 5\}$, and using $e \le \ell$ and $e \log(e+1) = O(\ell^2)$, gives

$$T_0(n, \ell) = C_{\text{bit}} \big[n^{12} \ell^3 + n^6 \ell^4 + \ell^5 \big] \tag{77}$$

for an absolute arithmetic-model constant C_{bit}. This uniform pointwise expected bound covers coordinatewise binary-rejection sampling and is deterministic for $p \in \{2, 3, 5\}$. The $d = 1$ shortcut removes only scalar splittings; all eigenvalue labels remain rational. This proves Theorem 5's algorithmic claims independently of density.

F Mixing for linearly independent lists

Let $m = n^2$. For $P, Z \in \mathrm{M}_n(\mathbb{F}_q)$, define

$$S_P(H)_{lk} = \sum_{i,j} P_{ij} H_{jl,ik}, \qquad T_Z(G)_{ij} = \sum_{k,l} G_{ik,jl} Z_{kl}. \tag{78}$$

The flattening indices give

$$S_P(H)^t = \mathrm{mat}(W \operatorname{vec} P), \qquad T_Z(H^{-1}) = \mathrm{mat}(F \operatorname{vec} Z).$$

Write $P_d = |\operatorname{GL}_d(\mathbb{F}_q)| q^{-d^2} = \prod_{i=1}^{d}(1 - q^{-i})$. For $q \geq 5$, the product inequality gives $P_d \geq 3/4$.

Lemma 9 (Mixing for linearly independent contraction lists). *Suppose $q \geq 5$, $n \geq 51$, and $1 \leq L, K \leq 6$. Fix linearly independent matrices $P_1, \ldots, P_L$ and linearly independent matrices $Z_1, \ldots, Z_K$ in $\mathrm{M}_n(\mathbb{F}_q)$. For arbitrary R_i, T_j, let N count $H \in \mathrm{GL}_m(\mathbb{F}_q)$ satisfying*

$$S_{P_i}(H) = R_i^t \ (1 \leq i \leq L), \qquad T_{Z_j}(H^{-1}) = T_j \ (1 \leq j \leq K).$$

Then

$$\left| \frac{N}{q^{-(L+K)m} |\operatorname{GL}_m(\mathbb{F}_q)|} - 1 \right| \leq \eta_6, \qquad \eta_6 = 8q^{-m/2} \leq q^{-m/4}. \tag{79}$$

No rank condition is required on individual matrices or on nonzero combinations in either span.

Proof. The trace identities $\operatorname{tr}((P \otimes A)H) = \operatorname{tr}(AS_P(H))$ and $\operatorname{tr}((B \otimes Z^t)G) = \operatorname{tr}(BT_Z(G))$ give the exact expansion

$$\begin{aligned} N = q^{-(L+K)m} \sum_{\boldsymbol{A},\boldsymbol{B}} \psi \left(- \sum_i \operatorname{tr}(A_i R_i^t) - \sum_j \operatorname{tr}(B_j T_j) \right) \\ \cdot \operatorname{Kl}_m \left(\sum_i P_i \otimes A_i, \sum_j B_j \otimes Z_j^t \right). \end{aligned} \tag{80}$$

Both coefficient maps are injective: the unique joint zero frequency supplies the main term. Of at most q^{12m} frequency pairs, write the coefficients as U, V and bound the nonzero absolute sum divided by $|\operatorname{GL}_m|$; affine phases have modulus one.

An injective restriction. For each fixed linearly independent list choose one entry position per member with invertible coefficient matrix. Their n-by-n blocks determine the tuple. Restrict to their block rows I and columns J, retaining every chosen block, and pad to size $6n$. This injective restriction has rank at most $\operatorname{rank} U$ and depends only on the list, so no union is taken. The matrix-rank count gives, for $0 \leq u \leq n$,

$$\#\{\boldsymbol{A} : \operatorname{rank} U \leq u\} \leq \frac{4}{3} \sum_{r=0}^{u} q^{12nr - r^2} \leq \frac{4}{3}(u+1) q^{12nu - u^2}. \tag{81}$$

Swapping factors gives the same bound for V, with no separate A_i rank bounds.

Low maximum rank. For $1 \le u = \max(\operatorname{rank} U, \operatorname{rank} V) \le n$, both coefficients are singular and $u < m/2$. Equation (23) saves $\Delta_m(u) = mu - 3u^2/2 + u/2$. Multiplying the two counts (81) gives a contribution at rank u of at most $\frac{128}{27}(u+1)^2 q^{-(m-24n)u-u^2/2-u/2}$. Its exponent decreases for $n \ge 51$. Summing at most n ranks gives

$$\rho_{\text{low}} \le \frac{128}{27} n(n+1)^2 q^{-m+24n-1} \le q^{-m+25n-1} \le q^{-m/2}. \tag{82}$$

The prefactor is at most 5^n at $n = 51$; its consecutive ratio $(n+2)^2/[n(n+1)] < 5$ decreases thereafter. The final inequality is $n^2/2 - 25n + 1 \ge 0$.

Middle singular ranks. For $n < u \le m - 13$, the saving $\Delta_m(u)$ is concave on each side of $m/2$. Its endpoint values at $n, m/2, m-13$ are $n^3 - \frac{3}{2}n^2 + \frac{n}{2}$, $\frac{m^2}{8} + \frac{m}{4}$, $13m - 247$. The last is the smallest for $n \ge 51$: subtracting it from the first two gives respectively $n^3 - 29n^2/2 + n/2 + 247 > 0$ and $m^2/8 - 51m/4 + 247 > 0$. Thus the full frequency count gives

$$\rho_{\text{mid}} \le \frac{8}{3} q^{-m+247} \le \frac{8}{3} q^{-m/2}. \tag{83}$$

Near-full singular ranks. Let the maximum rank have corank $1 \le c \le 12$, and let s be the other rank. If $s \le m/5$, then $m > 4s + c$. Lemma 4 saves at least $\frac{(m/5-12)(m+13)}{2} = \frac{m^2}{10} - \frac{47m}{10} - 78$, because its saving decreases in s, c throughout this range. Its chart factor is at most $4^m \le q^m$. Hence

$$\rho_{\text{near,low}} \le \frac{4}{3} q^{-m^2/10+177m/10+78} \le \frac{4}{3} q^{-m/2}. \tag{84}$$

This includes $s = 0$. If $s > m/5$, orient Lemma 4 with positive coefficient coranks $a = c, b = m - s$. Its positive argument is $s - 2c + 1 > m/5 - 23$, so its saving is at least $(m/5 - 23)^2/4$.

We use the following factorial bound throughout both mixing estimates:

$$C_m = 2^m m! \le (2m)^m \le q^{m^2/384} \qquad (m \ge 2601). \tag{85}$$

The exact integer seed is $5202^{384} < 5^{2601}$. For $m \ge 2601$, $(1 + 1/m)^{384} \le \sum_{j \ge 0} (384/m)^j \le \frac{2601}{2217} < 5$, so $(2m)^{384} \le 5^m$ propagates at every integer step. It follows that

$$\rho_{\text{near,high}} \le \frac{4}{3} q^{-71m^2/9600+143m/10-529/4} \le \frac{4}{3} q^{-m/2}. \tag{86}$$

An invertible coefficient. Lemma 3, including singular products, and (85) give

$$\rho_{\text{inv}} \le \frac{4}{3} q^{-95m^2/384+12m} \le \frac{4}{3} q^{-m/2}. \tag{87}$$

The last three exponent comparisons require, for $m \ge 2601$, positivity of $\frac{m^2}{10} - \frac{91m}{5} - 78$, $\frac{71m^2}{9600} - \frac{74m}{5} + \frac{529}{4}$, $\frac{95m^2}{384} - \frac{25m}{2}$. Each quadratic and its derivative

are positive at 2601, and its leading coefficient is positive. The middle-rank comparisons above are checked in the same way at $n = 51$ or $m = 2601$. These are inequalities on the full parameter range.

These five disjoint sectors exhaust the nonzero frequencies; their full counts require no additional corank or orientation factor. The constants sum to $1 + 8/3 + 4/3 + 4/3 + 4/3 = 23/3 < 8$. Finally $8 \le q^{m/4}$ proves the second inequality in (79).

This includes $(L, K) = (3, 4), (4, 3)$ and unions of up to six linearly independent anchors with low-rank combinations. Dependent lists require removing their redundancies. Throughout, H^{-1} is the inverse of the same H.

G A constant raw source fraction

All moments use denominator q^{m^2} and the bijection $G \leftrightarrow H = G^{-1}$ on GL_m. For $q \ge 5$, the product $P_d = \prod_{j=1}^{d}(1-q^{-j})$ satisfies $P_d \ge (1-1/5)(1-1/25)(1-\sum_{j\ge 3} 5^{-j}) = 2376/3125 > 3/4$.

G.1 Auxiliary triples with scalar common centralizer

Define the auxiliary condition $\mathcal{T}$ on an ordered triple: it holds when its members are invertible, every nonzero scalar combination has rank at least $n - 2$, and its two ratios relative to the first member have scalar common centralizer. This implies linear independence and is invariant under individual nonzero scalings and common invertible left/right multiplication. The minimum-rank condition restricts the count, not recognition.

Lemma 10. *For $q \ge 5, n \ge 51$, a uniform triple in $\mathrm{GL}_n(k)^3$ has $\mathcal{T}$ with probability at least $1 - (5/2)q^{-7} > 1 - 1/1000$.*

Proof. The rank-r count is at most $(4/3)q^{2nr-r^2}$, hence the rank-at-most-$(n-3)$ count is at most $(4/3)q^{m-9}/(1-q^{-7})$. For fixed projective coefficients, condition on all but one matrix with nonzero coefficient. Translation is injective; divide by $|\,\mathrm{GL}_n\,|$ and sum over $q^2 + q + 1$ coefficient points:

$$\mathbb{P}(\text{rank condition fails}) \le \frac{16}{9}\frac{1 + q^{-1} + q^{-2}}{1 - q^{-7}}q^{-7} \le \frac{9}{4}q^{-7}.$$

The last inequality holds at five and improves with q.

Conditionally on the first matrix, its ratios are independent uniform members of GL_n. For a nonscalar similarity class $[S]$, put $c_S = |\,\mathrm{Cent}_{\mathrm{GL}_n}(S)|$. Union over its $|\,\mathrm{GL}_n\,|/c_S$ conjugates bounds the probability of both ratios commuting with some member by $c_S/|\,\mathrm{GL}_n\,|$. There are at most $2^{n-1}q^n$ classes: invariant-factor degrees compose n, and each composition permits at most q^n monic-polynomial lists, before divisibility restrictions. Over an algebraic closure, the centralizer dimension is the sum of squared column lengths of the primary Jordan partitions,

obtained by summing $\min(a,b)$ over same-eigenvalue blocks. For nonscalar S there are at least two lengths, summing to n, so $\dim \mathrm{Cent}(S) \le (n-1)^2+1 = m-2n+2$. Thus

$$\mathbb{P}(\text{centralizer fails}) \le \sum_{[S]\ \text{nonscalar}} \frac{c_S}{|\,\mathrm{GL}_n\,|} \le \frac{2^{n-1}q^{2-n}}{P_n} \le \frac{1}{4}q^{-7}.$$

The last bound holds at $(51,5)$ by $(4/3)2^{50}5^{-42} < 1/4$ and decreases in both parameters. Add the losses.

G.2 The third-flattening estimate

The Q contraction requires a different orientation from $(3,4)$ mixing.

Lemma 11. *Let $q \ge 5, n \ge 51$. Fix linearly independent triples P_i, Z_i, each span having minimum nonzero rank at least $n-2$, scalars λ_i, and $E \ne 0$. The number of $H \in \mathrm{GL}_m$ satisfying*

$$W\,\mathrm{vec}(P_i) = \mathrm{vec}(Z_i), \quad F\,\mathrm{vec}(Z_i) = \lambda_i\,\mathrm{vec}(P_i)\ (i=1,2,3), \qquad Q\,\mathrm{vec}(E) = 0$$

differs from $q^{-7m}|\,\mathrm{GL}_m\,|$ by relative error at most $10q^{-m/2} < 1/1000$.

Proof. Orthogonality has prefactor q^{-7m} and

$$\text{phase} = \psi\left(-\sum_i \mathrm{tr}(A_i Z_i^t) - \sum_i \lambda_i\,\mathrm{tr}(B_i P_i)\right),$$
$$U = \sum_i P_i \otimes A_i, \qquad V = \sum_i B_i \otimes Z_i^t + (E\otimes D)S.$$

Here S swaps factors: $E_{jk}D_{li}$ pairs with $G_{ik,jl}$ in the Q contraction. Scalar compression on the second factor gives $\mathrm{rank}\,U \ge n-2$ for a nonzero A tuple. The (j,i) block of V is $\sum_a (B_a)_{ji} Z_a^t + (\mathrm{col}_i\,D)(\mathrm{row}_j\,E)$. A nonzero B coefficient makes this block's rank at least $n-3$. Thus $\mathrm{rank}\,V < n-3$ forces $B=0$; $V=0$ further forces $D=0$. Both frequency maps are injective.

If $A=0$ and $1 \le r = \mathrm{rank}\,V \le n-4$, then $r = \mathrm{rank}\,E\,\mathrm{rank}\,D$ gives at most $(4/3)(n+1)q^{2nr}$ frequencies. By the one-zero formula each has relative magnitude at most $(4/3)q^{-rm+r(r-1)/2}$. For $r \ge n-3$, use all q^{4m} frequencies. Excluding zero, this sector contributes at most

$$\frac{16}{9}(n+1)(n-4)q^{-m+2n} + \frac{4}{3}q^{4m-(n-3)m+(n-3)(n-4)/2} \le \left(1+\frac{4}{3}\right)q^{-m/2}. \quad (88)$$

The first exponent decreases in r. Its prefactor is at most 5^n at 51 and thereafter: its consecutive ratio $(n+2)(n-3)/[(n+1)(n-4)]$ decreases below five. Thus $-m+3n \le -m/2$ suffices. The second exponent plus $m/2$ is $-n^3+8n^2-7n/2+6<0$.

For $A \neq 0$, use at most q^{7m} frequencies. Use Lemma 3 if a coefficient is invertible. For singular coefficients with $R = \max(\operatorname{rank} U, \operatorname{rank} V) \leq m-8$, one has $R \geq n-2$ and Erdélyi–Tóth saving at least $8m-92$. For maximum-rank corank $k \leq 7$, split the other rank at $m/5$. Below it, Lemma 4 saves $(m/5-7)(m+8)/2 = m^2/10 - 27m/10 - 28$ with chart factor at most q^m; above it, Lemma 4 saves $(m/5-13)^2/4$. The hypotheses $m/5-7>0$, $m/5-13>0$ and the factorial bound (85) hold for $m \geq 2601$. The four bounds have coefficients $4/3, 8/3, 4/3, 4/3$ and exponents

$$-95m^2/384 + 7m, \qquad -m + 92,$$
$$-m^2/10 + 107m/10 + 28, \qquad -71m^2/9600 + 83m/10 - 169/4.$$

Subtracting each exponent from $-m/2$ gives a polynomial positive and increasing at 2601, with nonnegative higher derivatives. The saving $8m-92$ uses the endpoint comparisons $n^3 - 23n^2/2 + 13n/2 + 85 > 0$ and $m^2/8 - 31m/4 + 92 > 0$, verified identically at 51 and 2601. Including (88), coefficients sum to $9 < 10$. The sector counts require no extra corank factor.

G.3 First moment and excluded events

Henceforth $q \geq 5, n \geq 51$. Let $\mathcal{N}_0$ count ordered distinct nonzero labels λ_i and projective eigenlines $[p_i]$ with both triples P_i and $Z_i = \operatorname{mat}(Wp_i)$ satisfying $\mathcal{T}$. Set $\mathcal{N}_0 = 0$ if $\det G = 0$. First-line representatives determine Z_i; rescaling preserves $\mathcal{T}$. Let $\mathcal{N}$ additionally require $\mathcal{B}_{\mathrm{inv}}$ and geometric dimension one at each label. Then $\mathcal{N} > 0$ implies $\mathcal{R}_3$. Put

$$L(q) = \frac{(q-2)(q-3)}{(q-1)^2}, \qquad B = P_m P_n^6 L(q),$$
$$\delta = \frac{5}{2} q^{-7}, \qquad e_6 = 8q^{-m/2}, \qquad e_Q = 10q^{-m/2}.$$

Mixing gives $q^{-6m}|\operatorname{GL}_m|(1 \pm e_6)$ sources for each pair of linearly independent lists. The $(q-1)(q-2)(q-3)$ labels, three projective first anchors and three vector intermediates supply six $|\operatorname{GL}_n|$ factors, three $(q-1)$ divisions and denominator q^{m^2}, yielding B. Lemma 10 gives

$$\mathbb{E}\mathcal{N}_0 \geq B(1-\delta)^2(1-e_6). \tag{89}$$

These are summed incidences, not independent source outputs. Dropping the centralizer and minimum-rank restrictions but retaining invertible anchors gives upper-bound factor B. Distinct nonzero labels force both actual lists to be linearly independent before requiring F, W to be invertible.

Loss from singular matrices. Put $f = 1-q^{3-m}$, $g = 1-q^{-m}$. An F-kernel direction lies outside $\langle Z_1, Z_2, Z_3 \rangle$, whose images are linearly independent. Summing $(3,4)$ mixing over $(q^m - q^3)/(q-1)$ projective directions gives $\mathbb{E}[\mathcal{N}_0 \, \#\mathbb{P}(\ker F)] \leq$

$Bf(1+e_6)/(q-1)$. The involution $J(F) = W$ satisfies $G(J(F)) = SG(F)^{-t}S$, hence $J^2(F) = F$, where S is the commutation matrix. It sends the selected lines to $[Wp_i]$ with intermediates $\lambda_i p_i$. Both $\mathcal{T}$ restrictions preserve $\mathcal{N}_0$ under this raw-source bijection, equating the weighted F and W losses. Summing Lemma 11 over $(q^m-1)/(q-1)$ projective directions bounds the Q loss by $Bg(1+e_Q)/(q-1)$. Projective kernel counts dominate singularity indicators, so a weighted union bound accounts for all three singularity events.

Geometric multiplicity. Let ℓ_i count projective lines in the λ_i eigenspace. Dimension at least two implies $\ell_i - 1 \geq q$. Distinct labels force an extra line outside the first-anchor span. Its W image lies outside the intermediate span: applying F to a dependence would put that eigenvector in the first span. Thus $(4,4)$ mixing applies despite the repeated fourth label. Summing over $(q^m - q^3)/(q-1)$ extra projective first anchors and $q^m - q^3$ vector intermediates gives $\mathbb{E}[\mathcal{N}_0 \sum_i (\ell_i - 1)] \leq 3Bf^2(1+e_6)/(q-1)$. Divide by q for geometric failures, allowing repeated algebraic roots. Subtracting the four losses, put

$$h(q,n) = (1-\delta)^2(1-e_6) - \frac{2f(1+e_6) + g(1+e_Q)}{q-1} - \frac{3f^2(1+e_6)}{q(q-1)}. \tag{90}$$

Nonnegativity gives

$$\mathbb{E}\mathcal{N} \geq B \max\{h(q,n), 0\}. \tag{91}$$

G.4 Second moment and overlap factors

Geometric dimension one makes triples sharing s labels share those lines. Their union has $R = 6 - s$ distinct labels, linearly independent first/intermediate lists, and $a_s = \binom{3}{s}^2 s! = 1, 9, 18, 6$ patterns. Write $(q-1)_R = \prod_{j=1}^{R}(q-j)$ for the falling factorial, zero for $R > q-1$. Per pattern, discard the other conditions and allow arbitrary intermediate vectors, giving at most $(q-1)_R(|\operatorname{GL}_n|/(q-1))^R q^{Rm}$ tuples. Mixing, without a minimum-rank condition on the union span, bounds each pattern by $(1+e_6)P_m P_n^R (q-1)_R/(q-1)^R$. Define $t_R(q) = (q-1)_R/(q-1)^R$, with $t_R = 0$ if $R > q-1$. Then

$$\mathbb{E}\mathcal{N}^2 \leq (1+e_6)P_m \sum_{s=0}^{3} a_s P_n^{6-s} t_{6-s}(q) \leq 34(1+e_6) < 35. \tag{92}$$

G.5 Explicit envelopes

Equations (91)–(92) give $\mathbb{P}(\mathcal{R}_3) \geq c_{\mathcal{R}_3}(q,n)$, where

$$c_{\mathcal{R}_3}(q,n) = \frac{P_m P_n^{12} L(q)^2 \max\{h(q,n),0\}^2}{(1+e_6)\sum_{s=0}^{3} a_s P_n^{6-s} t_{6-s}(q)}. \tag{93}$$

The denominator is positive since $t_3(q) > 0$. Fractional exponents denote positive real powers in every characteristic. For a bound independent of n, define

$$l(q) = (1-q^{-1})(1-q^{-2})\left(1 - \frac{1}{q^2(q-1)}\right), \qquad \epsilon(q) = q^{-8},$$
$$A(q) = \frac{3(q+1)}{q(q-1)}, \qquad D(q,x) = \sum_{s=0}^{3} a_s x^{3-s} t_{6-s}(q),$$
$$H(q) = (1-\delta(q))^2(1-\epsilon(q)) - A(q)(1+\epsilon(q)).$$

The product-tail bound gives $P_d \geq l(q)$ for $d \geq 1$. Also $e_6, e_Q \leq \epsilon$, since $10 \leq q^{m/2-8}$ for $n \geq 51$. Replacing f, g by one gives $h(q,n) \geq H(q)$. The positive factors in H increase and its subtracted term decreases with q; $H(5) > 9/100$ proves positivity throughout.

Cancelling P_n^3 leaves the product ratio $P_m P_n^9 / D(q, P_n)$. For $x > 0$,

$$\frac{d}{dx}\frac{x^9}{D(q,x)} = \frac{x^8 \sum_{s=0}^{3}(6+s)a_s x^{3-s} t_{6-s}(q)}{D(q,x)^2} > 0.$$

Substituting in the whole quotient gives

$$c_{\mathcal{R}_3}(q,n) \geq c_{\mathcal{R}_3}(q) := \frac{l(q)^{10} L(q)^2 H(q)^2}{(1+\epsilon(q))D(q,l(q))}. \tag{94}$$

For the limit, put $r = 1/(q-1)$ and use $D(q,l) \leq 34$. The three deficits defining l sum to r, so $l^{10} \geq 1 - 10r$; $L = 1 - 3r + 2r^2$ gives $L^2 \geq 1 - 6r$. Also

$$H \geq 1 - A - 2\delta - \epsilon(1+A), \qquad 2\delta + \epsilon(1+A) \leq 6q^{-7} \leq 4r - A.$$

Indeed $A \leq 9/10$, $1 + A \leq q$, and $4r - A = (q-3)/[q(q-1)] \geq 2/q^2 \geq 6/q^7$. Hence $H \geq 1 - 4r \geq 0$, $H^2 \geq 1 - 8r$, and $(1+\epsilon)^{-1} \geq 1 - r$. These four actual factors lie in $[0,1]$; their product's deficit is at most the sum of deficits, even if the resulting lower bound is negative. Therefore

$$\mathbb{P}(\mathcal{R}_3) \geq c_{\mathcal{R}_3}(q) \geq \frac{1 - 25/(q-1)}{34} \geq \frac{1}{34} - \frac{1}{q}, \tag{95}$$

where $25q \leq 34(q-1)$ proves the last inequality. Thus the actual mass has lower liminf at least $1/34$, uniformly over $n \geq 51$ as q tends to infinity along prime powers. Either envelope can be combined with the constant bound in Theorem 5.

q	Theoretical lower bound $c(q)$
5	0.000025661301
7	0.001176688007
8	0.002185371657
9	0.003303773701
16	0.010356763829
25	0.015734071280
64	0.023398915529
256	0.027824833019

H Inverse-derived kernel recovery and source counting

Use the W-kernel map and predicates of Section 7.2. Put $N = m - 1$ and use the quotient and trace pairing fixed in the main text.

Matched representatives and covariance. For matched representatives,

$$g_1 P_A g_2^t = cP_T, \quad g_3 R_A g_4^t = dR_T, \tag{96}$$

$$\mathcal{B}_T^W = \beta \operatorname{Ad}(g_4)\mathcal{B}_A^W \operatorname{Ad}(g_1)^{-1}, \quad \beta = (cd)^{-1}. \tag{97}$$

This covariance follows from that of W by direct substitution; its modes differ from (16). Apply the cycle contractions (15) to $\mathcal{B}_A^W$ and write their outputs as C_2^W, C_3^W. Within the inverse-derived analysis only, abbreviate these as C_2, C_3:

$$X_{ij} = \mathcal{B}_A^W([E_{ij}]), \quad C_2 = \sum_{i,j} X_{ij}X_{ji}, \quad C_3 = \sum_{i,j,l} X_{ij}X_{jl}X_{li}, \tag{98}$$

$$C_j(T) = \beta^j g_4 C_j(A) g_4^{-1} \quad (j = 2, 3). \tag{99}$$

Scalar recovery and lift. Adjoint action has determinant one on $\mathrm{M}_n(k)$ and on its quotient by kI and its trace-zero subspace. In the two invariant exact sequences, the scalar line and the trace quotient are trivial representations. Thus every true scalar obeys

$$\beta^N = a := \det \mathcal{B}_T^W / \det \mathcal{B}_A^W, \qquad u_j \beta^j = v_j \quad (j = 2, 3), \tag{100}$$

$$h = \gcd(z^N - a, u_2 z^2 - v_2, u_3 z^3 - v_3, z^q - z). \tag{101}$$

Here $u_j = \operatorname{tr} C_j(A)$ and $v_j = \operatorname{tr} C_j(T)$. The roots of h in k are precisely the solutions of (100). When both source traces are nonzero, a zero target trace causes rejection. If both target traces are also nonzero, the only candidate is $\beta = (v_3/u_3)/(v_2/u_2)$, which is checked against all equations in (100). This gives the $p \geq 7$ deterministic branch on $\mathcal{S}_{\det}$. When either source trace is zero, enumerate all roots of h using the procedure in Section E. Modular powering avoids a degree-q array; the p-power Frobenius map gives deterministic splitting for $p = 2, 3, 5$, and Las Vegas splitting covers $\mathcal{S}_{\mathrm{CAS}}$ in growing characteristic [1,6,12]. Inseparability loses no field root; zero polynomial filters impose no condition and incompatible constants reject.

For each candidate solve

$$\mathcal{I}_\beta = \{D : C_j(T)D = \beta^j D C_j(A),\ j = 2, 3\}. \tag{102}$$

As in (74), an invertible solution identifies this space with the scalar common centralizer at the source. Retain only an invertible line generator D, and form

$$\widehat{g}_3 = R_T D^{-t} R_A^{-1}, \quad \widehat{K} = \widehat{g}_3 \otimes D, \quad \widehat{L} = F_T \widehat{K}^{-t} F_A^{-1}. \tag{103}$$

On a true branch $D = tg_4$, so (96) gives $\widehat{g}_3 = (td)^{-1} g_3$, $\widehat{K} = d^{-1}K$ and $\widehat{L} = dL$. Rank-one realignment and invertibility therefore recover the other two factors,

exactly as in (75); verify all original coefficients before YES. Source rejection is SOURCE_FAIL. For an accepted source, test target flattenings before inversion, then its rank, kernel-reshape invertibility and common-centralizer tests. Every true witness survives these tests, one of at most N scalars, its intertwiner and lift. Exhaustion is exact NO; all other steps are polynomial linear algebra.

Counting correlated sources with prescribed kernel lines. Fix right and left kernel lines. Before rank and invertibility conditions, their linear slice induces the uniform distribution on $\mathrm{Hom}(\mathrm{M}_n(k)/kI, \mathfrak{sl}_n(k))$. Bound bad maps under this law and subtract them from the count satisfying those conditions, then sum over unique kernel-line pairs and normalize by all raw tensors. No uniformity after conditioning is asserted.

Write $P_t = \prod_{j=1}^{t}(1-q^{-j})$, $p_t = \prod_{j=2}^{t}(1-q^{-j})$ and $e = \mathrm{vec}(I)$. Start with the uniform (m^2-m)-dimensional slice $We = 0$. In crossed inverse coordinates $H = G_A^{-1}$ it is $\sum_a H[(a,l),(a,k)] = 0$. For nonzero Z, the additional constraint $F_A \,\mathrm{vec}(Z) = 0$ on H^{-1} has orthogonality sum $q^{-2m}\sum_{B,D}\mathrm{Kl}_m(I_n \otimes B, D \otimes Z^t)$. Its coefficient ranks are $n \,\mathrm{rank}\, B$ and $\mathrm{rank}\, D \,\mathrm{rank}\, Z$; the third flattening has the same ranks after rearrangement. Separate the zero pair, two singular coefficients, and an invertible coefficient. The single/singular-pair formulas [11, Theorems 1.9–1.10] and Lemma 3 give

$$\rho_{\mathrm{ss}} = 5(n+1)^2 q^{-n^2/8}, \quad \rho_{\mathrm{inv}} = \tfrac{4}{3}2^{n^2}(n^2)!\,q^{-n^4/4+2n^2}, \quad d_n = 2nq^{-n^3/2}. \tag{104}$$

Here d_n is the single-constraint error. For two singular coefficients, the larger rank $1 \le r \le m-n$ saves $\Delta(r) = mr - r^2 - \binom{\min(r,m-r)}{2}$. The outer count costs q^{2nr} for $r < n$, and q^{2m} otherwise. Concave-quadratic endpoint checks save $n^2/8$ for $n \ge 5$; there are at most $(n+1)^2$ rank pairs, with constant $2(4/3)^3 < 5$. The invertible sector has at most q^{2m} pairs and uses the all-product bound, including singular products. Single-constraint rank a contributes at most $(16/9)q^{-n^3a+n^2a^2/2+3na/2-a^2}$; convexity selects $a = 1$, whose exponent is at most $-n^3/2$.

In the original slice, rank N has probability p_m and its left kernel is a uniform projective line, with invertible-reshape fraction $P_n/(1-q^{-m})$. Put $s_m = (1-q^{-m})/(q-1)$. Subtract singular-H loss $1 - P_m + P_m d_n$ and each remaining flattening-singularity loss $P_m s_m(1+\rho_{\mathrm{ss}}+\rho_{\mathrm{inv}})$ in this same slice. The mass satisfying these conditions is at least $g^0 - E_g$, where

$$g^0 = \frac{P_m p_n}{1-q^{-m}} - (1-P_m) - 2P_m s_m, \qquad E_g = \rho_{\mathrm{ss}} + \rho_{\mathrm{inv}} + d_n; \tag{105}$$

$P_m, 2P_m s_m \le 1$ absorb the excess coefficients. There are $v_n = q^{m-1}p_n$ projective lines with invertible reshapes. Local actions preserve these conditions and are transitive on either kernel line, so each slice with both lines fixed contains at least $q^{N^2}(g^0 - E_g)/p_n$ maps satisfying the conditions. Before imposing them, the full slice induces the uniform distribution on $\mathrm{Hom}(\mathrm{M}_n(k)/kI, \mathfrak{sl}_n(k))$, with inverse

$$\mathrm{mat}(W\,\mathrm{vec}(Y)) = R^{-t}\mathcal{B}^W([YP^{-1}]). \tag{106}$$

Distribution of the quadratic and cubic covariants. Under the uniform law on $\mathrm{Hom}(\mathrm{M}_n(k)/kI, \mathfrak{sl}_n(k))$, the off-diagonal blocks X_{ij} and any $n-1$ diagonal blocks are independent and uniform in $\mathfrak{sl}_n(k)$. The remaining diagonal block is determined by $\sum_i X_{ii} = 0$. The quadratic and cubic covariants take values in $\mathcal{W}_2 = \mathrm{M}_n(k)$ (respectively $\mathfrak{sl}_n(k)$ if $p = 2$) and $\mathcal{W}_3 = \mathrm{M}_n(k)$ (respectively $\mathfrak{sl}_n(k)$ if $p = 3$), by pairing quadratic traces and rotating cubic traces. Fourier duals use the trace pairing, modulo scalars for a trace-zero output. A nontrivial quadratic frequency satisfies

$$|\Phi(U, 0)| \le q^{-n(n-1)(n-2)}. \tag{107}$$

Each of the $\binom{n}{2}$ independent off-diagonal pairs has bilinear rank at least $2n-4$: the full Jordan-module kernel bound is $n^2 - 2n + 2$, and trace restriction and scalar quotient each lose at most one rank.

For cubic frequencies choose 0 and disjoint s-sets I, J, $s = \lfloor (n-1)/2 \rfloor$. Expose everything except X_{ai}, X_{i0}, X_{0a} $(i \in I, a \in J)$. A triangle uses at most one J-to-I edge; averaging those edges forces $\beta_V(Y_i, A_a) + D_{ai} \in kI$, where $\beta_V(Y, A) = (VY + YV)A + YAV$ and D_{ai} is fixed. Keep a matching of s independent pairs. A shifted full-matrix fibre is no larger than the zero fibre; imposing two trace-zero conditions and a scalar coset costs q^3.

Lemma 12 (Bilinear fibre and joint distribution estimates). *Let $q \ge 5$. The following conclusions concern different laws and ranges.*

1. One bilinear fibre, $n \ge 106$. *For each nontrivial cubic frequency V and independent uniform $Y, A \in \mathfrak{sl}_n(k)$, with β_V as above,*

$$\max_{D \in \mathrm{M}_n(k)/kI} \Pr[\beta_V(Y, A) \bmod kI = D] \le q^{-5n/2}. \tag{108}$$

2. Joint distribution, $51 \le n \le 105$. *For uniform $\mathcal{B}^W \in \mathrm{Hom}(\mathrm{M}_n(k)/kI, \mathfrak{sl}_n(k))$ before additional source conditions, the joint (C_2, C_3) law has total-variation distance at most τ_{fin} from independent uniform elements of the comparison spaces $\mathcal{W}_2, \mathcal{W}_3$ defined above, where the finite-band bound is*

$$\tau_{\mathrm{fin}} \le (2q^4)^{-1}. \tag{109}$$

Proof. Appendix I proves the rank, tail, frequency and finite-band bounds, retaining overlapping incidence strata and handling the characteristic-three scalar quotient without an extra scalar factor.

Parseval and the matching argument give

$$\tau_{\mathrm{eff}} = \begin{cases} \tau_{\mathrm{fin}}, & 51 \le n \le 105, \\ \frac{1}{2}\sqrt{q^{-2n^3+7n^2-4n} + q^{2n^2-5ns}}, & n \ge 106. \end{cases} \tag{110}$$

For uniform independent outputs on the comparison spaces, centralizer failure is at most

$$E_c = q^2 \left(\frac{q}{q-1}\right)^n \left((2q)^n q^{-n^2/2} + \frac{2q^{4-2n}}{1 - 2/q}\right). \tag{111}$$

Union over nonscalar Z gives $q^2\sum_Z q^{-2(n^2-\dim \mathrm{Cent}(Z))}$, paying for trace restrictions. A rational-canonical type of centralizer dimension c has at most $(q/(q-1))^n q^{n^2-c}$ elements, with at most $(2q)^n$ types. For $c > n^2/2$, a unique primary column $n-t > n/2$ belongs to a field eigenvalue; removing it leaves at most $q(2q)^t$ types and codimension at least $2t(n-t)$. The geometric envelope gives the second term. Appendix J proves

$$E_g + \tau_{\mathrm{eff}} + E_c \le 7q^{-4}, \qquad \tau_{\mathrm{eff}} + E_c \le 4q^{-4} \quad (q \ge 5, n \ge 51). \tag{112}$$

Normalization by all raw tensors. For $p \ge 7$, the trace failure under independent uniform outputs is exactly $\delta_{\mathrm{tr}} = 2/q - 1/q^2$. Union it with centralizer failure, then apply the total-variation bound once; in small characteristic there is no trace loss. Subtract this bad set in each full slice from the count satisfying the rank and invertibility conditions. Rank-N matrices W have unique right and left kernel lines, so summing over v_n^2 pairs is disjoint. Crossed inversion/rearrangement is injective, and the raw denominator q^{m^2} gives

$$qU(\mathcal{S}_{\mathrm{det}}) \ge p_n g^0 - p_n E_g - p_n^2(\tau_{\mathrm{eff}} + E_c + \mathbf{1}_{p\ge 7}\delta_{\mathrm{tr}}), \tag{113}$$

since $N^2 + 2(m-1) - m^2 = -1$. No uniformity after conditioning is used. Here g^0 and E_g are defined in (105), and E_c in (111). The grouped-product estimate proved in Appendix J is $p_n g^0/q \ge L(q) = (q^3 - 7q^2 + 12q - 5)(q-2)/(q-1)^5$. Here $qL(q)$ increases on $q \ge 5$, with $5L(5) = 75/1024$. Equation (112) yields $75/(2048q)$ in small characteristic. For $p \ge 7$, use $q \ge 7$ and $7L(7) - 13/49 - 7/625 = 18830957/238140000 > 75/2048$. At $q = 5 + x$, the numerator of $L(q) - (1/q - 5/q^2)$ over $q^2(q-1)^5$ is $x^5 + 16x^4 + 101x^3 + 320x^2 + 526x + 375 > 0$. Including this trace loss in every characteristic therefore gives $U(\mathcal{S}_{\mathrm{det}}) \ge 1/q - 7/q^2 + (q^2-7)/q^5 \ge 1/q - 7/q^2$. Finally $A \mapsto W_A$ injects $\mathcal{S}_{\mathrm{CAS}}$ into rank-$(m-1)$ matrices, whose raw mass is $P_m^2/(qP_{m-1}P_1^2) \le 1/(q-1)$ by full-rank factorization modulo GL_{m-1}. Inclusion yields both families' mass bounds and leading term $1/q + O(q^{-2})$, uniformly for $n \ge 51$.

I Bilinear ranks and finite-parameter certificates

Here $q \ge 5$, $n \ge 5$, $s = \lfloor (n-1)/2 \rfloor$ and $d = \min_{\lambda \in \bar{k}} \mathrm{rank}(V - \lambda I)$; all rank minima are integral and coranks range over $0 \le e \le n$.

Fibres on the full matrix space. For a uniform square matrix, the corank probabilities are

$$\pi_e = q^{-e^2}\frac{\prod_{j=e+1}^{n}(1-q^{-j})^2}{\prod_{j=1}^{n-e}(1-q^{-j})} \le \tfrac{4}{3}q^{-e^2} \quad (0 \le e \le n). \tag{114}$$

The full-matrix rank argument of Appendix C, with $t = n$, gives rank at least $n(d-2e)_+$ on a corank-e input. This algebraic statement has no dimension cutoff. Thus

$$\min_e\{e^2 + n(d-2e)_+\} = \lceil d/2 \rceil^2. \tag{115}$$

For $d = 2h$ the minimum is h^2; for $d = 2h+1$, compare h^2+n with $(h+1)^2$ and use $n \geq d$.

If $d < n/2$ and $V = \lambda I + R$ with $3\lambda \neq 0$, the same argument gives rank at least $n(n-e) - 3nd + 2d^2$. Keeping the sharper corank prefactor in (114) gives zero-fibre bounds

$$F_d = \sum_{e=0}^{n} \pi_e q^{-n(d-2e)_+}, \quad S_d = \sum_{e=0}^{n} \pi_e q^{-[n(n-e)-3nd+2d^2]_+}, \tag{116}$$

$$F_d \leq \tfrac{4}{3}(n+1)q^{-\lceil d/2 \rceil^2}, \quad S_d \leq \tfrac{4}{3}(n+1)q^{-3n^2/4+3nd-2d^2}. \tag{117}$$

Completing the square without the positive part gives the second coarse bound; finite certificates retain it. A shifted linear fibre is empty or has the zero fibre's size, so averaging proves shifted-fibre domination. Two trace restrictions and a scalar coset cost q^3.

Overlapping incidence conditions. When $3\lambda = 0$, $1 \leq d < n/2$, factor $R = UC$, with U injective, C surjective; put $U_0 = \operatorname{im} U$, $K = \ker C$. For fixed Y let $a = \operatorname{rank}(U_0 \xrightarrow{Y} k^n/U_0)$, $b = \operatorname{rank}(K \xrightarrow{Y} k^n \xrightarrow{C} k^d)$ and $e = n - \operatorname{rank} Y$. Then $T_Y : A \mapsto \beta_R(Y, A)$ satisfies

$$\operatorname{rank} T_Y \geq r(a,b,e) = (a+b)(n-d) + d(n-e-d-a)_+. \tag{118}$$

This is (65) with (t,b,c) there replaced by (n,a,b) here. The shared subspace-incidence argument (66) also applies. For $q \geq 5$, each Gaussian coefficient has prefactor $4/3$, so the product is $16/9$ rather than the all-field 16. Retaining this sharper constant, set

$$L_{ab} = (d-a)(n-2d-a) + (d-b)(n-d-b), \tag{119}$$

$$M_{ab} = \min_{0 \leq e \leq n} \{\max(L_{ab}, e^2) + r(a,b,e)\},$$

$$R_d = \tfrac{16}{9}(n+1) \sum_{a,b=0}^{d} q^{-M_{ab}}. \tag{120}$$

Intersecting incidence and corank bounds gives $(16/9)q^{-\max(L_{ab},e^2)}$, also for $L_{ab} < 0$; their exponents do not add. The factor $n+1$ pays for coranks. Put $t = n-d-a$, $L_0 = \max(L_{ab}, 0)$. After removing $(a+b)(n-d)$, minimize the convex $f(e) = \max(L_0, e^2) + d(t-e)_+$. Its slopes across $\sqrt{L_0}, t$ are $-d$, then 0 or $2e-d$, then $2e$, giving $e_* = \min(t, \max(\sqrt{L_0}, d/2))$. Adjacent integers, including t in the flat case, suffice; comparisons $t^2 \leq L_0$, $4L_0 \geq d^2$ and integer square roots determine them.

One-pair bound in large dimensions. For $n \geq 106$ put $D = \lfloor (3n+2)/10 \rfloor$. For $d > D$, the saving in F_d is at least $(3n+3)^2/400 \geq 41n/16$ for $n \geq 112$: $9n^2 - 1007n + 9$ is 121 at 112 and increasing. For $106 \leq n \leq 111$, $D = (32, 32, 32, 32, 33, 33)$ and the saving is at least $289 \geq 41n/16$. For $d \leq D$ in the nonzero-scalar branch, the

decreasing quadratic $3n^2/4 - 3nd + 2d^2$ is at least $3n^2/100 - 9n/25 + 2/25 \geq 41n/16$, since $12n^2 - 1169n + 32$ is positive and increasing from 106.

For the residual branch $D \leq n/3$. Put $x = d-a$, $z = d-b$; then $0 \leq L = x(n-3d+x)+z(n-2d+z) \leq 2dn$ and $r+L \geq 3dn-4d^2+x^2-dx+z^2-dz-de$. Splitting at $e = \sqrt{L}$ and completing a square gives $\max(L, e^2) - de \geq L - d\sqrt{L} - d^2/4$. Together with $x^2 - dx, z^2 - dz \geq -d^2/4$, this proves

$$r + \max(L, e^2) \geq E(d) := 3dn - 19d^2/4 - \sqrt{2n}\, d^{3/2}. \tag{121}$$

Concavity gives $E(d) \geq E(1) = 3n - 19/4 - \sqrt{2n}$: the other endpoint satisfies $E(n/3) = n^2(17/36 - \sqrt{2/27}) \geq n^2/6 \geq 3n$, using $2/27 \leq (11/36)^2$. The stratum prefactor is at most n^3, since $(16/9)((3n+12)/10)^2(n+1) \leq n^3$; after clearing denominators its difference is $189n^3 - 324n^2 - 864n - 576$, positive after translation by 106. The remaining absorptions are $n^3 \leq q^{n/4}$, $31/4 + \sqrt{2n} \leq n/4$, and $q^3(4/3)(n+1) \leq q^{n/16}$. Endpoints are $106^{12} \leq 5^{106}$, $(106-31)^2 \geq 32 \cdot 106$, and $(4 \cdot 107)^{16} \leq 3^{16}5^{58}$. Propagation uses $(107/106)^{12} < 5$, the increasing polynomial $(n-31)^2 - 32n$, and $(108/107)^{16} < 5$; larger q strengthens them. They absorb all prefactors and the full q^3 cost, proving (108) in every branch.

Fourier bounds by rank. Set $b_R(d) = \min(1, q^3F_d, q^3R_d)$ for $1 \leq d < n/2$, $b_S(d) = \min(1, q^3F_d, q^3S_d)$ for $0 \leq d < n/2$, and $b_F(d) = \min(1, q^3F_d)$. The matching bounds each Fourier coefficient by the applicable b^s, uniformly in the quadratic frequency. Write $w_j = n^2 - \mathbf{1}_{p=j}$ $(j = 2, 3)$, $f = \lceil n/2 \rceil$. The rank-d count is $N_d = q^{2nd-d^2} \prod_{j=n-d+1}^{n}(1-q^{-j})^2 / \prod_{j=1}^{d}(1-q^{-j}) \leq (4/3)q^{2nd-d^2}$. For $d < n/2$, the minimizing scalar is unique and in k, since its $(n-d)$-dimensional eigenspace excludes a distinct Frobenius conjugate. Outside characteristic three, zero/nonzero scalar frequencies number at most $N_d, (q-1)N_d$. In characteristic three the dual is modulo scalars: each such class has a unique rank-d residual, since two would differ by a rank-n scalar with rank sum $2d < n$. No extra scalar factor occurs. The remaining $d \geq f$ frequencies are bounded by the whole cubic dual; F_d decreases termwise. Thus

$$\begin{aligned} H = {} & \sum_{d=1}^{s} N_d b_R(d)^{2s} + \mathbf{1}_{p\neq 3}(q-1)\sum_{d=0}^{s} N_d b_S(d)^{2s} \\ & + q^{w_3} b_F(f)^{2s}, \end{aligned} \tag{122}$$

$$\tau_{\text{fin}} = \min\{1, \tfrac{1}{2}\sqrt{(q^{w_2}-1)q^{-2n(n-1)(n-2)} + q^{w_2}H}\}. \tag{123}$$

The first sum excludes the trivial residual; $d = 0$ is retained only for nontrivial nonzero scalars. Parseval and (107) prove the total-variation bound for the covariants of a uniform linear map, on the stated comparison spaces and before additional source conditions.

Exact field-uniform certificate. For rational $C > 0$, let $\ell(C) = \min\{j \geq 0 : C \leq 5^j\}$; then $C \leq q^{\ell(C)}$ for every $q \geq 5$. Put $C_F = \frac{4}{3}(n+1)$, $C_R(d) = \frac{16}{9}(n+1)(d+1)^2$,

$m_F(d) = \lceil d/2 \rceil^2$, $m_S(d) = \min_e\{e^2 + [n(n-e) - 3nd + 2d^2]_+\}$ and $m_R(d) = \min_{0\le a,b\le d} M_{ab}$. Including the shift cost, set

$$\begin{aligned} k_S(d) &= \max(m_F(d), m_S(d)) - 3 - \ell(C_F), \\ k_R(d) &= \max(m_F(d) - 3 - \ell(C_F), m_R(d) - 3 - \ell(C_R(d))), \\ k_F &= m_F(f) - 3 - \ell(C_F). \end{aligned} \tag{124}$$

Each applicable b is at most q^{-k}. Enlarge both comparison-space dimensions to n^2, use $N_d \le q^{2nd-d^2+1}$ and $q - 1 \le q$. The $T = 2s + 3$ squared-Fourier terms are bounded by q^{-E}, where

$$\begin{aligned} E_Q &= 2n(n-1)(n-2) - n^2, & E_F &= 2sk_F - 2n^2, \\ E_R(d) &= 2sk_R(d) - n^2 - 2nd + d^2 - 1 & &(1 \le d \le s), \\ E_S(d) &= 2sk_S(d) - n^2 - 2nd + d^2 - 2 & &(0 \le d \le s). \end{aligned} \tag{125}$$

Proposition 5 (Finite-band certificate). *For $51 \le n \le 105$ and every prime power $q \ge 5$, $\tau_{\text{fin}} \le 1/(2q^4)$.*

Proof. For each integer n in the band and $0 \le d \le s$, compute m_F, m_S; for $d > 0$, minimize over all a, b by the adjacent-integer rule after (120). Repeated multiplication by five computes ℓ, then k, E, without a field grid, logarithms or approximate square roots. The exact comparisons give

$$\min\{E_Q, E_F, E_R(1), \dots, E_R(s), E_S(0), \dots, E_S(s)\} \ge 8 + \ell(T). \tag{126}$$

The minimum is 210 at $n = 54$, whereas $8 + \ell(T) \le 11$. Thus the whole squared-Fourier sum is at most $Tq^{-8-\ell(T)} \le q^{-8}$, proving the claim. All 55 dimensions and 1,289,546 residual strata are included; the prefactor inequality covers every $q \ge 5$.

The verification package recomputes all finite-band minima by exact integer arithmetic, independently enumerating the permitted integer e. The minimum values are stored as machine-readable output rather than as an additional printed table.

Individual lower bounds on source mass. For any applicable finite TV bound τ, the actual-source argument gives

$$\alpha_{\text{low}} = \frac{1}{q}[p_n g^0_{n,q} - p_n E_g - p_n^2(\tau + E_c + \mathbf{1}_{p\ge 7}(2/q - 1/q^2))]_+. \tag{127}$$

Use exact rationals or outward-rounded lower bounds, retaining products, factorials, floors and characteristic cases. Nonpositive bounds, failed hypotheses or insufficient precision do not imply zero mass. Below 51, certificates are individual and may be nonmonotone because of floors; they imply neither optimality nor practical runtime.

J Complementary source-mass bounds

Lemma 13. *For $q \geq 5, n \geq 51$,*

$$E_g + \tau_{\text{eff}} + E_c \leq 7q^{-4}, \qquad b := \tau_{\text{eff}} + E_c \leq 4q^{-4}. \tag{128}$$

For fixed q, $b(n,q) \to 0$.

Proof. For $m \geq 64$, $2^m m! \leq (2m)^m \leq q^{m^2/16}$: $(2m)^{16} \leq 5^m$ holds at 64 and propagates by $(65/64)^{16} < 5$. For $n \geq 106$, the middle bound below follows from (110). Use $(q/(q-1))^n (2q)^n \leq q^{2n}$, $(q/(q-1))^n \leq q^{n/2}$ and $2/(1-2/q) \leq q$ to obtain

$$E_g \leq 5(n+1)^2 q^{-n^2/8} + \tfrac{4}{3} q^{-3n^4/16+2n^2} + 2nq^{-n^3/2}, \tag{129}$$

$$\tau_{\text{eff}} \leq \tfrac{1}{2} q^{-n^3+7n^2/2-2n} + \tfrac{1}{2} q^{-n^2/4+5n/2}, \tag{130}$$

$$E_c \leq q^{2+2n-n^2/2} + q^{7-3n/2}. \tag{131}$$

At $n = 16$, the seven negative-exponent magnitudes are $(32, 11776, 2048, 3232, 24, 94, 17)$, with coefficients at most 5^{c_i} for $c_i = (5,1,3,0,0,0,0)$. Each magnitude starts above $c_i + 4$ and grows at least one per step: its excess over four and increment minus one have nonnegative coefficients after $n = 16 + t$. Prefactor ratios are below $4, 2$, so each term is at most 5^{-4} for $n \geq 16$. For $c(n) \geq 4$, $a(n) q^{-c(n)} \leq q^{-4} a(n) 5^{-(c(n)-4)} \leq q^{-4}$. Thus $E_g \leq 3q^{-4}$, $E_c \leq 2q^{-4}$, and the applicable TV tail is at most $2q^{-4}$; use (109) on $51 \leq n \leq 105$. All four majorants for b tend to zero for fixed q.

Put $P_\infty = \prod_{j \geq 1} (1 - q^{-j})$. Grouping the P_m coefficient in g^0, using $P_m \geq P_\infty$, $p_n \geq q P_\infty/(q-1)$ and $s_m \leq 1/(q-1)$, gives $g^0 \geq \delta_q = q P_\infty^2/(q-1) + (q-3)P_\infty/(q-1) - 1$. Substitute $P_\infty \geq (q-2)/(q-1)$: its numerator over $(q-1)^3$ is $q^3 - 7q^2 + 12q - 5 = 5 + 17x + 8x^2 + x^3 > 0$ at $q = 5 + x$. Multiplication by the lower bound for p_n is therefore valid, yielding $p_n g^0/q \geq L(q)$. Moreover $(qL(q))' = (4x^4 + 64x^3 + 369x^2 + 908x + 805)/(q-1)^6 > 0$. The accompanying checker regenerates all finite-band minima and the theoretical table. The displayed algebraic expansions prove the polynomial comparisons.

K Finite-dimensional contraction bounds

We prove the dimension ranges in Theorems 5 and 6 using the shared cancellation lemmas and the exact moment and full-Hom interfaces in Appendices G and H.

K.1 Finite contraction certificates and a stronger two-corank bound

Write $C_d = 2^d d!$ and $P_d = \prod_{j=1}^{d} (1 - q^{-j})$. For $q \geq 4$, coefficient ranks at least $d - a, d - b$, and $a, b \geq 1$, put

$$r = d - a, \qquad \kappa = d - 2a - b + 1, \qquad \sigma(d,a,b) = \min\{r^2/4, (\kappa_+)^2/2\}.$$

Then the two-corank estimate strengthens to

$$|\operatorname{Kl}_d(U,V)| \le C_d q^{d^2-\sigma(d,a,b)}.$$

To prove this, use the complete-affine-fibre cancellation in the two-corank lemma before taking absolute values. Its nonsingular lower-block sector saves $r^2/4$. In the lower-block corank-j sector, the residual core has size $r-j$, corank bounds $j, a+b-j$, and outer saving $\Delta_j = 2j(\kappa-1)+j^2$. If $1 \le \kappa < 2j$, the concave polynomial $\Delta_j - \kappa^2/2$ has nonnegative endpoint values on $[1,2j]$. If $\kappa \ge 2j$, then

$$\Delta_j + (\kappa-2j)^2/2 - \kappa^2/2 = 3j^2 - 2j \ge 0.$$

For the other core branch,

$$\Delta_j + (r-2j)^2/4 = r^2/4 + 2j(j-1) + j(2\kappa - r).$$

This is at least $r^2/4$ when $r \le 2\kappa$. When $r = 2\kappa + x$, $x \ge 0$, its lower bound $r^2/4 - (x+2)^2/8$ exceeds $\kappa^2/2$ by $(\kappa^2-1)/2 + (\kappa - 1/2)x + x^2/8 \ge 0$. The rank constraints in the complete-fibre reduction ensure legitimate smaller cores. Since $P_d \ge 2/3$ for $q \ge 4$, the induction constants close by

$$\frac{C_{d-1} + (3a/2)C_{d-2}}{C_d} \le \frac{1}{2d} + \frac{3}{8(d-1)} < 1 \qquad (d \ge 3).$$

Trivial and invertible-coefficient cases retain their original bounds. For finite calculations retain the entire first induction step: the $j=0$ term has coefficient C_r, and $j \ge 1$ has coefficient $(3/2)C_{r-j}$ with saving $\Delta_j + \sigma(r-j, j, a+b-j)$.

For linearly independent contraction lists of lengths L, K, the selected-block restriction is injective on each frequency space. Up to rank $u \le Ln$, the first has at most $2q^{2Lnu-u^2}$ coefficients, capped by q^{Lm}, and similarly for K. There are $q^{(L+K)m}$ frequency pairs. For each rank sector, take the minimum of the applicable upper bounds, including their prefactors. At maximum singular coefficient rank $u = m-c$, compare the singular-pair bound with the asymmetric/two-corank split: choose s_0 with $m > 4s_0 + c$; at smaller rank $s \le s_0$ use the actual chart factor $\sum_{h=0}^{s_0} \binom{m-s_0}{h}^2$, and above it use the unrolled bound at s_0+1, oriented with first corank c. The respective endpoint envelopes cover all smaller ranks. Sum low ranks using their profiles; above the retained cutoff use the complete frequency count once, without an extra rank multiplicity. The invertible sector uses $C_m q^{-m^2/4}$; normalization costs $P_m^{-1} \le 3/2$. The separate Q-gate has q^{7m} frequency pairs, and its one-zero sector uses the rank product $\operatorname{rank}(E)\operatorname{rank}(D)$.

The supplied exact certificate verifies relative errors at most 10^{-6} on these bands:

input	fields	dimensions
$(4,4)$	$q \ge 5$	23–50
$(6,6)$	$q \ge 7$	27–50
$(4,6), (5,5)$	$q \ge 4$	26–50
Q-gate	$q \ge 4$	26–50
Q-gate	$q \ge 5$	23–25

Inversion supplies $(6,4)$, and shorter lists follow by completion to linearly independent lists and summation over the added outputs. Each accepted error term is Dq^{-E} with $E \geq 0$ after including the full frequency count; its outward dyadic upper bound at the base field proves the entire stated field range. These bands specify 102 spectral and 28 third-flattening comparisons. The inequalities are in the field-size parameter, so evaluating their monotone envelopes does not require a grid of finite fields. For $n \geq 51$, the shared five-sector proof gives $\eta_6 \leq 9q^{-m/2}$ and $\eta_Q \leq 11q^{-m/2}$ for $q \geq 4$. Its factorial seed is $5202^{384} < 4^{2601}$, with the same propagation and nonnegative translated-polynomial comparisons as in the supplied certificate. Both tail errors are below 10^{-6}. Below, η_4 denotes the required $(4,4)$ relative error and η the maximum of the required field-four list errors; both are at most 10^{-6}.

K.2 Three rational anchors at dimensions 23 and 27

The auxiliary triple failure, before its coarse absorptions, is at most

$$\delta(q,n) = P_n^{-2}\frac{1+q^{-1}+q^{-2}}{1-q^{-7}}q^{-7} + P_n^{-1}2^{n-1}q^{2-n}.$$

It is at most $1/1000$ for $q \geq 5, n \geq 23$ and for $q = 4, n \geq 26$. Use the ordered-triple counts $\mathcal{N}_0, \mathcal{N}$ from Appendix G. The first moment and excluded-event argument needs only $(4,4)$ mixing and the separate Q gate. With $e = 1/1000$, put

$$H_0(q) = (1-e)^3 - \frac{3(1+e)}{q-1} - \frac{3(1+e)}{q(q-1)}.$$

The moment is at least $P_m P_n^6 L(q) H_0(q)$. Since $H_0(5) = 96102999/10^9 > 9/100$ and H_0 increases with q, this yields

$$\mathbb{E}\mathcal{N} \geq (3/4)^7(3/8)(9/100) > 1/256.$$

Let r count nonzero eigenvalues with one-dimensional eigenspaces and invertible right reshapes on the F, G, Q, W-invertible region, and set $r = 0$ elsewhere. The fourth factorial moment satisfies

$$\mathbb{E}(r)_4 \leq (1+\eta_4)P_m P_n^4 \frac{(q-1)_4}{(q-1)^4} \leq 1001/1000.$$

Since $\mathcal{N} \leq (r)_3$, for every integer $K > 3$,

$$\mathcal{N} \leq (K)_3 \mathbf{1}_{\mathcal{N}>0} + (r)_4/(K-3).$$

Taking $K = 345$ gives

$$U(\mathcal{R}_3) > \frac{1/256 - (1001/1000)/342}{345 \cdot 344 \cdot 343} = \frac{233}{9684783360000} > 2^{-36}.$$

For $n \geq 27$, the overlap second-moment proof applies with the certified errors from Appendix K. At $q = 5$ only overlap types with three or four distinct labels survive, so $(4, 4)$ suffices. At every prime power $q \geq 7$, use $(6, 6)$. Writing

$$\begin{aligned} A_5 &= (2376/3125)^7(3/8)(96102999/10^9), \\ B_5 &= (1001/1000)(96/125)^4\big((27/16)(96/125) + 9/4\big), \\ A_7 &= (3/4)^7(5/9)(425002999/10^9), \end{aligned}$$

we have $2^{16}A_5^2 > B_5$ and $2^{16}A_7^2 > 35$ by exact arithmetic. Cauchy–Schwarz gives the 2^{-16} bound. No prime power lies strictly between five and seven. This proves the bound for the family in Definition 4.

K.3 The mixed sieve at field four

For a fixed counted base triple, put $a_F = \dim\ker F$, $a_W = \dim\ker W$ and $d_i = \dim\ker(X - \lambda_i I) - 1$. Define

$$T_j = \sum_{a+b=j} q^{\binom{a}{2}+\binom{b}{2}} \begin{bmatrix} a_F \\ a \end{bmatrix}_q \begin{bmatrix} a_W \\ b \end{bmatrix}_q, \quad L_3 = T_0 - T_1 + T_2 - T_3, \quad U_2 = T_0 - T_1 + T_2.$$

Truncating the product of the two kernel polynomials gives $L_3 \leq \mathbf{1}_{a_F=a_W=0} \leq U_2$, with $U_2 \geq 0$. For $h_i = \begin{bmatrix} d_i \\ 1 \end{bmatrix}_q$,

$$\mathbf{1}_{a_F=a_W=d_1=d_2=d_3=0} \geq L_3 - U_2 \sum_i h_i.$$

The incidences have contraction lengths $(3+b+c, 3+a+c)$, with $a+b \leq 3, c = 0$, or $a + b \leq 2, c = 1$. The certified $(4, 6), (6, 4), (5, 5)$ bands therefore cover them.

Only for evaluating these incidence moments, take independent uniform full matrices constrained to carry the three fixed anchor pairs. With $t = m - 3$ and $K_t(a) = \begin{bmatrix} t \\ a \end{bmatrix}_q q^{-ta}$, the quotient-kernel moment is

$$M_t(a, b, c) = \left[\prod_{j=0}^{c-1}(1 - q^{j-t})\right]^2 \frac{q^{-c}}{|\mathrm{GL}_c(q)|} K_{t-c}(a)K_{t-c}(b), \qquad c = 0, 1.$$

Indeed choose the c-space and its injective intermediate image, then the a- and b-kernel spaces disjoint from those images. The two specified maps cost $q^{-t(a+b+2c)}$, and the free upper row costs q^{-c}. Counting bases and lifts proves the formula. Every actual incidence is between linearly independent contraction lists, so actual-law mixing transfers each nonnegative moment with the same relative error and common normalization P_m. There is no assumption of independent actual flattenings.

At $q = 4$, the limiting coefficients are $(1, 2/3, 7/45, 44/2835)$, giving

$$J_4 = 1991/5670, \qquad J_4 - 1/3 = 101/5670,$$

and the absolute coefficient sum is $13003/5670 < 3$. Finite products lose at most $18q^{-(m-5)}$, the mixing loss is at most 3η, and the separate Q gate costs $(1+\eta_Q)/(q-1)$. For $n \geq 26$ all three analytic errors are at most 10^{-6}. Thus, with $B = P_m P_n^6(2/9)$,

$$\begin{aligned}\mathbb{E}\mathcal{N} &\geq Bh_4,\\ h_4 &= (999/1000)^2(J_4 - 4\cdot 10^{-6}) - (1+10^{-6})/3\\ &= 89810616979/5250000000000.\end{aligned}$$

There are exactly three nonzero labels in $\mathbb{F}_4$. Their geometric eigenlines are unique. Changing the ratio base preserves the algebra generated by the two ratios, hence its centralizer; all other counting conditions are also permutation invariant. Consequently $\mathcal{N} \in \{0, 6\}$. Since $P_d \geq 705/1024$, this proves $U(\mathcal{R}_3) \geq (705/1024)^7 h_4/27 = c_4$. The family $\mathcal{R}_3$ is empty over $\mathbb{F}_2$ and $\mathbb{F}_3$.

K.4 The bilateral corank-one improvement

For $\beta_V(Y, A) = VYA + YVA + YAV$,

$$\beta_V(Y, A)^t = \beta_{V^t}(A^t, Y^t).$$

Put $d = \min_{\lambda\in\bar{k}} \operatorname{rank}(V - \lambda I)$. The fixed-input rank bounds of Appendix I are $r_F(e) = n(d-2e)_+$, and, when $d < n/2$ and $3\lambda \neq 0$,

$$r_S(e) = \max\{r_F(e), [n(n-e) - 3nd + 2d^2]_+\}.$$

For independent full matrices of coranks e, f, their zero-pair mass is bounded by each of $(16/9)q^{-e^2-f^2}$, $(4/3)q^{-e^2-r(e)}$, and $(4/3)q^{-f^2-r(f)}$. Taking their minimum and summing yields

$$\Pr[\beta_V(Y, A) = 0] \leq \frac{16}{9} \sum_{e,f=0}^{n} q^{-\max\{e^2+f^2, e^2+r(e), f^2+r(f)\}}.$$

This combines independent input-rank counts with two separate fibre bounds; it does not add overlapping fibre codimensions. A shifted fibre is dominated by the zero fibre after fixing one input. The two trace restrictions and scalar output coset therefore still cost q^3. Keep the residual overlapping-incidence bound of Appendix I when it is smaller, and retain the same matching size $\lfloor (n-1)/2 \rfloor$ and characteristic-three scalar quotient.

Substitution into Parseval gives $\tau \leq 3q^{-4}$ at $n = 27, 29, 30, 31$ and $\tau \leq (2q^4)^{-1}$ at $32 \leq n \leq 50$, over every $q \geq 5$. The certificates evaluate complete finite sums with integer exponents. For example, at $n = 30$ their field-uniform upper bound for the squared Fourier sum $\mathcal{S}$ in $\mathrm{TV} \leq \frac{1}{2}\sqrt{\mathcal{S}}$ is

$$q^8\mathcal{S} \leq 33.986337341613 < 36 \qquad (q \geq 5),$$

where the decimal is rounded upwards. This gives $\mathrm{TV} < 3q^{-4}$. Appendices I–J cover $n \geq 51$. The gate and centralizer bounds of Appendix J give $E_g \leq 3q^{-4}$

and $E_c \le 2q^{-4}$. Subtracting bad maps in the full Hom slice before conditioning, then summing over unique kernel-line pairs, yields

$$qU(\mathcal{S}_{\det}) \ge qL_{\det}(q) - 8q^{-4} - \mathbf{1}_{p\ge 7}(2/q - 1/q^2),$$

where $L_{\det}(q) = (q^3-7q^2+12q-5)(q-2)/(q-1)^5$. The monotonicity of $qL_{\det}(q)$ proved in Appendix J, together with

$$5L_{\det}(5) - 8/625 > 75/2048, \qquad 7L_{\det}(7) - 13/49 - 8/625 > 75/2048,$$

proves the first lower bound. Since $L_{\det}(q) \ge 1/q - 5/q^2$, the same inequality gives $U(\mathcal{S}_{\det}) \ge 1/q - 7/q^2 + (q^2-8)/q^5$. The upper bound $1/(q-1)$ and both recovery algorithms are established in Appendix H. At $n = 28$, the certified majorant fails in the sector $d = 10, e = f = 5$; no corank-one mass claim at this dimension follows from the proof.

L Characteristic-polynomial computation

This appendix supplies the elementary proof of Lemma 8, including the in-place similarity update.

Proof. Reduce M to upper Hessenberg form H by similarities. In each column $a = 1, \dots, m-2$, move a nonzero entry below the diagonal, if present, to row $a+1$ by simultaneously swapping its row and column. For $i > a+1$ with $H_{ia} \ne 0$, put $u = H_{ia}/H_{a+1,a}$ and apply

$$\mathrm{row}_i \leftarrow \mathrm{row}_i - u\,\mathrm{row}_{a+1}, \qquad \mathrm{col}_{a+1} \leftarrow \mathrm{col}_{a+1} + u\,\mathrm{col}_i\,.$$

Using the row-updated matrix in the column step gives EHE^{-1}, $E = I - uE_{i,a+1}$. Both rows vanish in earlier columns, preserving Hessenberg zeroes. The $O(m^2)$ eliminations cost $O(m)$ each and divide only by nonzero field pivots, never by an integer, trace, or dimension. Set $D_0 = 1$ and compute for $1 \le j \le m$:

$$D_j(x) = (x - H_{jj})D_{j-1}(x) - \sum_{i=1}^{j-1} H_{ij} \left(\prod_{t=i}^{j-1} H_{t+1,t} \right) D_{i-1}(x). \tag{132}$$

Last-column expansion of the leading j-by-j determinant forces the subdiagonal chain from i to j and leaves D_{i-1}; cofactor and off-diagonal signs give the minus sign. Hence $D_m = \det(xI - M)$. Reverse chain products and coefficient arrays cost $O(j^2)$ per stage, $O(m^3)$ overall, storing $O(m^2)$ coefficients. The recurrence is division-free, including zero subdiagonals, repeated/inseparable roots, and $p \mid m$.