Eight-point fiber certificates and binary-input list recovery of folded Reed–Solomon codes

Scott Schubr
Luminary Oddities, Countryside, Illinois, USA
scottschubr@luminaryoddities.com
ORCID: 0009-0009-7357-4439

October 18, 2026

Abstract

We prove that folded Reed–Solomon codes with primitive-power evaluation blocks and two allowed symbols at each coordinate have output list size at most seven at every asymptotic rate $R < 2/3 - 8\rho/9$, where $0 \le \rho < 3/4$ is the disagreement fraction. The result holds in the stronger average-radius sense, with explicit finite-parameter folding and field-size conditions. The proof assigns one nonnegative subspace certificate to each configuration of eight messages, controlling the two largest fibers of every linear quotient, including partial captures. Every eight-point configuration over any field has certificate cost at most nine, and the Boolean three-cube attains nine. Combining this inequality with classical folded-Wronskian rank bounds gives the input-two, output-seven case of the Chen–Zhang tradeoff. Retaining partial-overlap windows in the folded determinant gives a quadratic intersection budget and sharper integer finite-folding profiles. These certify, for example, $n = 10$, $k = 120$ and $\rho = 1/4$ at folding 29 over $\mathbb{F}_{512}$. The same certificate applies to multiplicity codes in characteristic at least k; its derivative-Wronskian ramification budget certifies folding 28 over $\mathbb{F}_{331}$ for the same n, k, ρ. Binary input refers to the two allowed symbols, independently of the field characteristic.

1 Introduction

A list-recovery input specifies a set of allowed symbols at each coordinate of a code. The objective is to bound the number of codewords that disagree with these sets on at most a prescribed fraction of coordinates. We consider folded Reed–Solomon codes and input sets of size at most two. Our main result gives output cap seven on the positive-error rate curve

$$R < \frac{2}{3} - \frac{8\rho}{9}.$$

This is the input-two, output-seven case of the tradeoff conjectured by Chen and Zhang [1, Conjecture 3.6]. The proof also gives the stronger average-radius statement and explicit finite-folding bounds.

1.1 The coding result

We use the folded alphabet $\Sigma = \mathbb{F}_q^s$. For a table $A = (A_0, \dots, A_{n-1})$, where $A_j \subseteq \Sigma$ and $|A_j| \le 2$, define the agreement of a codeword c by

$$\operatorname{agr}(c, A) = \frac{1}{n} \sum_{j=0}^{n-1} \mathbf{1}_{c_j \in A_j}.$$

Ordinary $(\rho, 2, 7)$ list recovery means that at most seven codewords have agreement at least $1 - \rho$ with any such table. We prove average-radius $(\rho, 2, 7)$ list recovery: every eight distinct codewords $c_1, \dots, c_8$ satisfy

$$\frac{1}{8} \sum_{i=1}^{8} \operatorname{agr}(c_i, A) < 1 - \rho \qquad \text{for every table } A.$$

The latter implies the former. Throughout, *binary input* refers to the size two of each allowed set. The code alphabet need not be binary, and the field may have any characteristic allowed by the stated parameters.

Theorem 1.1 (Binary-input list recovery)**.** *Let $n \ge 1$, $s \ge 7$, and $1 \le k \le sn$ be integers. Let $q > sn$ be a prime power and γ a generator of $\mathbb{F}_q^\times$. Encode each polynomial $f \in \mathbb{F}_q[X]$ of degree less than k as*

$$\operatorname{Enc}(f)_j = \big(f(\gamma^{js}), \dots, f(\gamma^{js+s-1})\big), \qquad 0 \le j < n.$$

If $0 \le \rho < 3/4$ and

$$\frac{k-1}{n(s-6)} < \frac{2}{3} - \frac{8\rho}{9}, \tag{1}$$

then this code is average-radius $(\rho, 2, 7)$ list recoverable.

The rate is $R = k/(sn)$. For every fixed $\rho < 3/4$ and positive slack below $2/3 - 8\rho/9$, a sufficiently large fixed folding parameter s makes (1) hold. For example, with $\rho = 1/4$, the asymptotic sufficient rate is $R < 4/9$. The generic entropy sufficient bound in [2], in the same normalization, is approximately 0.39624 at this error fraction. This comparison concerns those two sufficient bounds; it does not assert optimality at every finite parameter.

The finite profiles in Proposition 2.4 exploit the dimensions of the actual certificate supports. The cyclic folded-Wronskian budget in Lemma 2.5 then retains partial-overlap windows and gives a stronger integer receiver. At $n = 10$, $s = 29$, $k = 120$, $q = 512$ and $\rho = 1/4$, it bounds mean capture by $59/10 < 6 = 8(1 - \rho)$. The same message-degree and error parameters work at multiplicity 28 in characteristic 331. Section 2 gives the complete arithmetic and the distinct hypotheses for the two code families. These are list-size guarantees; algorithms for locating the codewords are studied in, for example, [5].

1.2 Why an affine-configuration inequality suffices

Consider eight candidate messages as points in the vector space of polynomials of degree less than k. At a fixed folded coordinate, evaluation is linear. Its kernel partitions the eight messages into fibers, one for each folded symbol. A table allowing two symbols can capture at most the total population of the two largest fibers.

It remains to control these populations simultaneously at all coordinates. A bound chosen separately for each coordinate need not interact with the global polynomial rank budget. We instead attach to the message configuration a fixed collection of nonnegative subspace weights. The

same weights then apply to every evaluation kernel. Summing the resulting inequalities reduces the agreement count to the established folded-Wronskian bounds for intersections with these subspaces.

Here is the geometric statement. Let S be a finite set of distinct points over a field, and let $V = \mathrm{dir}(S)$ be the span of its pairwise differences. For $K \le V$, write $b_2(S/K)$ for the sum of the two largest populations of affine K-cosets meeting S; the second population is zero if there is only one such coset.

Definition 1.2. A certificate for S is a finite family of subspaces $W \le V$ and rational weights $y_W \ge 0$ such that

$$b_2(S/K) \le 2 + \sum_W y_W \dim(K \cap W) \qquad (K \le V). \tag{2}$$

Its cost is $\sum_W y_W \dim W$. The family and weights may depend on S, but are fixed independently of K. Let $\mathcal{C}(S)$ be the infimum of these costs.

Theorem 1.3 (Eight-point certificate)**.** *Every set of eight distinct points over any field has a certificate of cost at most nine. For affine dimensions* $1, \dots, 7$*, respectively, sufficient costs are*

$$6, \quad 8, \quad 9, \quad 9, \quad 9, \quad 8, \quad 7.$$

For the Boolean three-cube, $\mathcal{C}(S) = 9$*.*

To see the numerical connection, the folded-Wronskian budget bounds the mean intersection dimension of each certificate support W by

$$\frac{1}{n}\sum_{j=0}^{n-1} \dim(W \cap K_j) \le r \dim W, \qquad r = \frac{k-1}{n(s-6)},$$

where K_j is the kernel of the jth folded evaluation map. The certificate therefore bounds the mean number of captured messages by $2 + 9r$. Eight messages with average agreement at least $1 - \rho$ would require at least $8(1-\rho)$, contradicting (1). Section 2 gives the complete argument and the dimension-sensitive finite-parameter refinements.

The cube lower bound proves optimality of nine for the stated class of nonnegative certificates; individual configurations can have smaller costs. The finite-folding coding inequalities are sufficient conditions.

1.3 Prior work and scope

Chen and Zhang [1, Conjecture 3.6 and Theorem 3.7] proposed the relevant power-of-input-list tradeoff and proved the input-two, output-three case. Brakensiek, Chen, Dhar and Zhang [2, Theorem 1.6] proved the general zero-error case. Their October 11, 2026 version retains a gap for the finer positive-error tradeoff [2, Remark 1.7 and Section 1.3]. The present result establishes the input-two, output-seven case for the primitive-power evaluation model specified above.

Brakensiek–Chen–Dhar–Zhang also attribute the binary-input zero-error case to Chen, Cheraghchi, Si, and Zhang, *Evidence Toward a Conjecture in List Recoverability of Folded Reed–Solomon Codes* (2025). The conjectured curve, Boolean-cube obstruction and polynomial rank-budget framework are prior ingredients. Our contribution is the arbitrary-field eight-point certificate analysis and its positive-error consequence. The folded-Wronskian transfer uses the established subspace-design framework [3, 4]. In particular, the October 2026 account in [2, Remark 1.7] identifies the

input-two, output-three result as the known positive-error special case beyond its zero-error theorem. The present certificate controls partial captures with the same weights at every quotient; this is the additional requirement used in the positive-error transfer.

Recent work of Goyal, Guruswami, Sun and Wootters [6] gives threshold equivalences for optimal nested subspace-design constructions, including multiplicity codes. The numerical input supplied here is the cost-nine inequality for every eight-message configuration. Proposition 2.6 records its explicit Wronskian application to the second polynomial code family.

The proof separates according to affine dimension. In dimensions at most four, large fibers force incidences among lines and planes. In dimensions five through seven, the affine dependency space has dimension at most two; its supports provide the additional subspaces needed in the certificate. These two arguments cover all configurations, rather than only Boolean cubes.

Organization. Section 2 proves the coding transfer from Theorem 1.3 and records the finite-parameter refinements. The affine tools and the small- and large-dimension arguments then prove Theorem 1.3.

2 From certificates to folded Reed–Solomon codes

We give the standard polynomial rank-budget argument explicitly to fix its characteristic, order and finite-folding conventions. This is established subspace-design machinery [3, 4].

Lemma 2.1 (Folded-Wronskian budget)**.** *Under the hypotheses of Theorem 1.1, let* $\mathcal{P}_k = \{f \in \mathbb{F}_q[X] : \deg f < k\}$ *and let* $K_j \le \mathcal{P}_k$ *be the kernel of the* j*th folded evaluation block. For* $W \le \mathcal{P}_k$ *with* $1 \le u = \dim W \le s$*,*

$$\sum_{j=0}^{n-1} \dim(W \cap K_j) \le \frac{u(k-1) - u(u-1)/2}{s-u+1} \le \frac{u(k-1)}{s-u+1}.$$

Proof. Choose a basis $g_1, \dots, g_u$ of distinct degrees $d_i < k$ by degree elimination. The polynomial

$$D(X) = \det\big(g_b(\gamma^a X)\big)_{0 \le a < u,\, 1 \le b \le u}$$

is nonzero: its coefficient in degree $\sum_b d_b$ is a product of the nonzero leading coefficients and the Vandermonde determinant in the distinct elements γ^{d_b}. They are distinct because γ has order $q-1 \ge sn$ and $0 \le d_b < k \le sn$. Because the degrees are distinct, $\deg D = \sum_b d_b \le u(k-1) - u(u-1)/2$.

Put $h_j = \dim(W \cap K_j)$. Change basis so the first h_j polynomials vanish on block j. For every $0 \le t \le s-u$, all entries in those columns vanish at $X = \gamma^{js+t}$. Each such column is divisible by $X - \gamma^{js+t}$, so the determinant has root multiplicity at least h_j. A basis change multiplies D by a nonzero scalar. All these roots, across j, t, are distinct by the same order condition. Summing multiplicities gives $(s-u+1)\sum_j h_j \le \deg D \le u(k-1) - u(u-1)/2$. □

Proof of Theorem 1.1. All sn evaluation points are distinct and $k \le sn$, so the encoding is injective. Suppose eight distinct messages formed a forbidden average-radius list. Their affine-direction space $V \le \mathcal{P}_k$ has dimension at most seven. At block j, the evaluation fibers are affine cosets of $K_j \cap V$, so the number of messages whose folded symbol is in the two-symbol table is at most $b_2(S/(K_j \cap V))$.

Apply Theorem 1.3. For each certificate support $W \le V$, its dimension is at most seven, and Lemma 2.1 gives

$$\frac{1}{n}\sum_j \dim(W \cap K_j) \le \frac{k-1}{n(s-6)} \dim W =: r \dim W.$$

The mean number of captured messages is therefore at most $2 + 9r$. Average agreement at least $1 - \rho$ would require it to be at least $8(1 - \rho)$, giving $r \geq 2/3 - 8\rho/9$, contrary to (1). If each of eight messages individually has agreement at least $1 - \rho$, its average does too, proving ordinary list recovery. □

Corollary 2.2 (Asymptotic form)**.** *Fix $0 \leq \rho < 3/4$, set $B = 2/3 - 8\rho/9$, and take $0 < \eta < B$. Under the remaining hypotheses of Theorem 1.1, rate $R = k/(sn) \leq B - \eta$ and folding $s \geq 7 + \lceil 12/\eta \rceil$ suffice.*

Proof. Here $R \leq B \leq 2/3$, and

$$r < \frac{Rs}{s-6} = R + \frac{6R}{s-6} < R + \eta \leq B.$$

The integer and field constraints still apply. The theorem does not include equality $r = B$. □

Remark 2.3. Sharpness concerns the stated nonnegative-certificate problem. The asymptotic target and its coding obstruction originate in [1]. Our proof establishes a sufficient inequality through a particular class of nonnegative certificates.

2.1 Finite-folding profiles

For nonnegative support budgets $c_1, \ldots, c_{\min(7,k)}$, put $t_d = \max_{1 \leq u \leq d} c_u/u$ and define, when the index is available,

$$\begin{aligned}
&C_1 = 6c_1, && C_2 = 8t_2, \quad C_3 = 9t_3, \quad C_4 = 9t_4,\\
&C_5 = \max\big\{c_5 + 2c_1, c_5 + 2c_2, c_5 + c_3, c_5 + c_4,\\
&\qquad\qquad 3c_5/2, 3c_5/2 + c_1/2, 3c_5/2 + c_2/2\big\},\\
&C_6 = c_6 + \max_{1 \leq u \leq 6} \frac{c_u}{\max(1, u-1)}, && C_7 = c_7.
\end{aligned}$$

Proposition 2.4 (Support-budget transfer)**.** *Suppose a linear polynomial encoding has kernels $K_0, \ldots, K_{n-1}$, is injective on $\mathcal{P}_k$, and satisfies*

$$\sum_j \dim(W \cap K_j) \leq c_u \qquad (\dim W = u,\ 1 \leq u \leq \min(7, k)).$$

If $0 \leq \rho < 3/4$ and $\max_{1 \leq d \leq \min(7,k)} C_d < n(6 - 8\rho)$, then it is average-radius $(\rho, 2, 7)$ list recoverable.

Proof. For a fixed eight-message configuration, summing its certificate over all blocks bounds total capture by $2n + \sum_W y_W c_{\dim W}$. For affine dimensions one through four, the cost bounds and $c_u \leq u t_d$ give $C_1, \ldots, C_4$. The proper-core certificates of Lemma 5.1 contribute $c_5 + 2c_u$ for $u = 1, 2$ or $c_5 + c_u$ for $u = 3, 4$. Its full-core certificates contribute $3c_5/2$, $3c_5/2 + c_u/2$ for $u = 1, 2$, or $c_5 + 2c_2$. These are exactly the branches of C_5. Lemma 5.2 gives the ambient weight one and the support weight $1/\max(1, u-1)$ in dimension six, and ambient weight one in dimension seven. These give C_6, C_7. The strict inequality therefore excludes total capture at least $8n(1 - \rho)$. No monotonicity of the budgets in u is required. □

The folded-Wronskian bound of Lemma 2.1 already supplies the budgets

$$b_u = \frac{u(k-1) - u(u-1)/2}{s - u + 1}.$$

Their numerators increase and their positive denominators decrease for $u \le k$, so b_u is nondecreasing. Substitution in the preceding profiles recovers the simpler dimension-sensitive bounds. We next improve the budgets themselves by retaining windows that cross the ends of an evaluation block.

Lemma 2.5 (Cyclic folded-Wronskian budget)**.** *Under the hypotheses of Theorem 1.1, let* $1 \le u \le \min(s,k)$*,* $\dim W = u$*, and* $h_j = \dim(W \cap K_j)$*. Then*

$$\sum_{j=0}^{n-1} h_j(s-u+h_j) \le D_u^{\mathrm{fold}} := u(k-1) - u(u-1)/2.$$

Proof. Use the folded matrix $M(X) = (g_b(\gamma^a X))_{0\le a<u,\,1\le b\le u}$ from Lemma 2.1. Its nonzero determinant $D(X)$ has degree at most D_u^{fold}. Write $L = q-1$. At $X = \gamma^t$, the rows evaluate at the u distinct cyclic exponents $t, t+1, \dots, t+u-1$ modulo L. Partition them among the disjoint coded blocks $B_j = \{js, \dots, js+s-1\}$ and any gap. If $r_j(t)$ rows belong to block j, their rank is at most $\min(r_j(t), u-h_j)$, since the full block map on W has kernel dimension h_j. Gap rows have rank at most their count. Writing $x_+ = \max\{x, 0\}$, we obtain

$$u - \operatorname{rank} M(\gamma^t) \ge \sum_j (h_j + r_j(t) - u)_+,$$

For any polynomial square matrix, its determinant vanishes to order at least its nullity at the evaluation point: constant invertible row and column operations make its last nullity-many rows zero at that point, so each such row is divisible by the root factor. This argument holds in every characteristic. It pays all block deficits at a shared root before any roots are summed.

If $n = 1$, then $k \le s$ makes the sole block map injective and $h_0 = 0$. Suppose $n \ge 2$. We have $L \ge sn \ge 2s \ge s+u$. For a fixed length-s cyclic block, the sliding length-u window has overlap u at $s-u+1$ starts and overlap r at two starts for each $1 \le r < u$. Thus, for $0 \le h \le u$,

$$\sum_{t \bmod L} (h + r_j(t) - u)_+ = h(s-u+1) + 2\sum_{a=1}^{h-1} a = h(s-u+h).$$

Sum determinant valuations at the L distinct roots γ^t, interchange the finite nonnegative sums, and use the degree bound. □

For each integer $H \ge 0$, write $m = \lfloor H/n \rfloor$ and $r = H - nm$, so $0 \le r < n$. For a numerator $D_u \ge 0$, define an integer budget

$$c_u(D_u) = \max\left\{ H \in \mathbb{Z} : \begin{array}{l} 0 \le H \le n\min(u, \max(k-s,0)), \quad H = nm+r,\ 0 \le r < n, \\ (s-u)H + nm^2 + r(2m+1) \le D_u \end{array} \right\}. \tag{3}$$

The set contains zero. This is a finite, explicit computation. For $H = \sum h_j$, integer convexity gives $\sum h_j^2 \ge nm^2 + r(2m+1)$: transferring one unit between two entries differing by at least two decreases their square sum. Also $h_j \le \min(u, \max(k-s,0))$, by divisibility by the block's degree-s vanishing polynomial. Lemma 2.5 therefore supplies the support budgets $c_u = c_u(D_u^{\mathrm{fold}})$ in Proposition 2.4.

For $n = 10$, $s = 29$, $k = 120$, $q = 512$, and $\rho = 1/4$, these budgets and the seven profiles are respectively

$$(4, 8, 12, 17, 22, 27, 32), \qquad (24, 32, 36, 153/4, 39, 35, 32).$$

Their maximum is $39 < 40 = n(6-8\rho)$, giving mean capture at most $2 + 39/10 = 59/10 < 6$. The primitive-order requirement holds since $511 \ge 290$. This certifies folding 29 for these parameters. The coarser budgets b_u certify folding 30 in the earlier example. The accompanying exact-arithmetic program reproduces both comparisons.

2.2 Hasse-derivative evaluation blocks

The same certificate applies to multiplicity codes by the classical derivative Wronskian in its large-characteristic range [3, 1].

Proposition 2.6 (Multiplicity-code receiver)**.** *Let $\mathbb{F}_q$ have characteristic p, let $n \geq 1$, $s \geq 7$, $1 \leq k \leq \min(p, sn)$, and choose distinct $\alpha_1, \dots, \alpha_n \in \mathbb{F}_q$. Encode $\deg f < k$ by the Hasse-derivative blocks $(f^{[0]}(\alpha_j), \dots, f^{[s-1]}(\alpha_j))$. For $0 \leq \rho < 3/4$, inequality (1) suffices under the present hypotheses for average-radius $(\rho, 2, 7)$ list recovery. More sharply, Proposition 2.4 applies with $c_u = c_u(D_u^{\mathrm{mult}})$, where $D_u^{\mathrm{mult}} = u(k-u)$ and c_u is given by (3).*

Proof. The block kernel consists of polynomials divisible by $(X - \alpha_j)^s$. Injectivity follows from $k \leq sn$. For $\dim W = u \leq \min(s, k)$, its ordinary derivative Wronskian is nonzero and has degree at most $u(k-u)$: eliminate leading degrees and use the Vandermonde determinant of their falling factorials. All degrees are distinct and below $k \leq p$. At a fixed α_j, Taylor elimination gives distinct orders $0 \leq \nu_1 < \cdots < \nu_u < k$, exactly $h_j = \dim(W \cap K_j)$ of which are at least s. The leading Wronskian coefficient is again a nonzero Vandermonde, so its valuation is $\sum_i \nu_i - u(u-1)/2$. The low orders are at least $0, \dots, u - h_j - 1$ and the high ones at least $s, \dots, s + h_j - 1$. Hence the valuation is at least $h_j(s - u + h_j)$. Summing at distinct evaluation points gives

$$\sum_j h_j(s - u + h_j) \leq u(k-u).$$

Integer convexity gives the displayed budgets. The coarse inequality $\sum_j h_j \leq u(k-1)/(s-u+1)$ also follows, so the proof of Theorem 1.1 gives its simpler sufficient condition. □

At $n = 10$, $s = 28$, $k = 120$, $p = q = 331$ and $\rho = 1/4$, the multiplicity budgets and profiles are

$$(4, 8, 13, 17, 22, 27, 32), \qquad (24, 32, 39, 39, 39, 35, 32).$$

Again the maximum is $39 < 40$. This folding-28 example uses the derivative characteristic hypothesis; the folded-29 example above instead uses a primitive element of the characteristic-two field $\mathbb{F}_{512}$.

3 Affine tools

Lemma 3.1 (Selected-fiber rank bound)**.** *Let $S = \{v_1, \dots, v_m\}$ have affine dimension d. On $E_0 = \{c \in \mathbb{F}^m : \sum_i c_i = 0\}$ define $T(c) = \sum_i c_i v_i$, and put $\Lambda = \ker T$. Then $\dim \Lambda = m - 1 - d$. If $K < V$, choose the two largest nonempty K-fibers A, B and set $c = |A| + |B|$. Let $E \leq E_0$ consist of coefficient vectors supported on $A \cup B$ with separate coefficient sums zero on A and B. With $\delta = \dim(E \cap \Lambda)$,*

$$b_2(S/K) - 2 \leq \dim K + \delta. \tag{4}$$

Proof. A proper K cannot contain all pairwise differences, so there are at least two nonempty fibers. We have $\dim E = c - 2$ and $T(E) \leq K$. Rank–nullity gives $c - 2 = \dim T(E) + \delta \leq \dim K + \delta$. □

Lemma 3.2 (Adjoining affine coloops)**.** *Let T consist of $m \geq 3$ distinct points with direction space U. Adjoin j points so that S has affine dimension $\dim U + j$; equivalently, their homogenized vectors are independent modulo the homogeneous span of T. For $K \leq \mathrm{dir}(S)$,*

$$b_2(S/K) \leq b_2(T/(K \cap U)) + \dim K - \dim(K \cap U) + \mathbf{1}_{U \subseteq K}. \tag{5}$$

Consequently, if a certificate for the modified quantity $b_2(T/H) + \mathbf{1}_{H=U}$ has ambient coefficient at least one and cost D, then S has an ordinary certificate of cost at most $D + j$.

Proof. Translate a core point to zero. The j outside points are then linearly independent in V/U. The case $K = V$ is immediate. Otherwise select two nonempty largest fibers, containing s outside points and c_T core points, and let $b \in \{0, 1, 2\}$ count the fibers meeting the core. The space of coefficients with zero sum on each selected fiber has image modulo U of dimension $s - (2 - b)$. Indeed a fiber not meeting the core imposes one independent zero-sum condition on its outside coefficients; a core coefficient absorbs that condition in a fiber meeting the core. Independence in V/U gives

$$s - (2 - b) \le \dim K - \dim(K \cap U).$$

If $U \not\subseteq K$, the core occupies at least two fibers, and $c_T + 2 - b \le b_2(T/(K \cap U))$: when only one selected fiber meets it, a second nonempty core fiber contributes at least one; when neither does, the two core fibers contribute at least two. If $U \subseteq K$, the same inequality holds with an additional one on the right. Adding proves (5). For the last assertion, add weight one on V and subtract one from the core ambient coefficient. The new coefficients are nonnegative and the cost increases by $\dim V - \dim U = j$. □

Lemma 3.3 (Modified planar certificate)**.** *For $m \ge 3$ points T spanning an affine plane with direction U, there are $a \ge 1$, $w \ge 0$ rational and a line $L_1 \le U$ with*

$$b_2(T/H) + \mathbf{1}_{H=U} \le 2 + a \dim H + w \dim(H \cap L_1), \qquad 2a + w \le m.$$

Thus a planar core together with j affine coloops has ordinary certificate cost at most $m + j$.

Proof. For a direction line L, set $m_L = b_2(T/L) - 2$. Positive values occur only in pair-difference directions, so there are finitely many. The two largest values $m_1 \ge m_2$ in distinct directions are positive, because an affine spanning triple supplies two directions with capture at least three. Two unions of two parallel lines in different directions intersect in at most four points. Therefore $m_1 + m_2 \le m$. Let L_1 maximize m_L and set

$$\lambda = \max\left\{1, \frac{m-1}{m_1 + m_2}\right\}, \quad a = \lambda m_2, \quad w = \lambda(m_1 - m_2).$$

Then $a \ge 1$, $w \ge 0$, and $2a + w = \max\{m - 1, m_1 + m_2\} \le m$. For $H = L_1$ the increment $a + w \ge m_1$; for every other line it is $a \ge m_2$. For $H = U$ the required increment is $m - 1$, and for $H = 0$ the bound is exact. Apply Lemma 3.2. □

Lemma 3.4 (Two directions)**.** *Two distinct parallel L-lines and two distinct parallel M-lines, with $L \ne M$, have at most four cross-intersections. If there are at least three, all four lines lie in one affine plane. If there are at least two, their affine span has dimension at most three.*

Proof. Each line pair has at most one intersection. Three edges of the bipartite 2-by-2 intersection graph connect its four vertices. An intersecting nonparallel pair determines a plane, and the parallel mates reached by those edges also lie in it. With two intersections, either two matched pairs span parallel planes of direction $L + M$, whose union has affine dimension at most three, or three lines lie in a plane and adjoining the fourth increases dimension by at most one. □

4 Eight points in affine dimensions at most four

Throughout, captures refer to $b_2(S/K)$. For dimension one, weight six on V works: the zero kernel captures two and V captures eight. Dimension two follows from Lemma 3.3, discarding its nonnegative indicator. Dimension seven will follow from affine independence in Section 5.

Lemma 4.1 (Dimension three)**.** *Eight distinct points of affine dimension three admit cost at most nine.*

Proof. If every line direction captures at most five, take $y_V = 3$. Lines get right-hand side five and planes eight, so all kernels are covered. Otherwise choose a line direction L with capture $c_L \geq 6$. It is at most seven: two parallel lines have affine span at most two. For any plane P not containing L, its two selected cosets capture at most four points on the selected two L-lines and at most $8 - c_L$ outsiders. Hence

$$b_2(S/P) \leq 12 - c_L. \tag{6}$$

If $c_L = 7$ and another direction M captured at least five, the selected subsets would meet in at least four points. There are at most four cross-intersections, so Lemma 3.4 would make all eight points coplanar, a contradiction. Thus other lines capture at most four. Weights $y_V = 2$, $y_L = 3$ work: line L gets seven, other lines four, planes containing L get nine, and transverse planes get six, enough by (6). Cost is nine.

If $c_L = 6$, no other line can capture six: four common selected points would again make all eight points coplanar. If another direction captures five, the selected subsets meet in three or four points. Three intersections put their eight-point union in a plane, impossible. Four intersections give seven coplanar points plus one affine coloop, covered by Lemma 3.3 at cost at most eight. Otherwise all other lines capture at most four. Take $y_V = 2$, $y_L = 2$, of cost eight. Planes containing L get eight and transverse planes get six, again covering (6). Zero and full kernels satisfy every displayed certificate directly. □

Lemma 4.2 (Modified seven-point core)**.** *Seven distinct points T of affine dimension three, with direction U, admit a certificate for $b_2(T/H) + \mathbf{1}_{H=U}$ of cost at most eight and with coefficient on U at least one.*

Proof. If every line captures at most four, weight $5/2$ on U works. Otherwise a line L captures five or six. If it captures six, another line capturing at least five would have four common selected points, putting all seven points in a plane. Thus all other lines capture at most four. Use $y_U = 2$, $y_L = 2$, of cost eight. A plane avoiding L captures at most five and gets six; a plane containing L gets eight. The modified full-kernel value is eight and is covered.

Suppose the largest line capture is five. If another direction captures five, the selected subsets meet in three or four points. Three would make all seven coplanar. Four give six coplanar points and one affine coloop. Lemma 3.3 supplies an ordinary certificate on U with ambient coefficient one and cost between six and seven: the modified planar core has cost at least $m - 1 = 5$ and adjoining the coloop adds one. Thus its full-kernel right-hand side is at least eight, covering the extra indicator as well. If L is the only direction capturing five, take $y_U = 2$, $y_L = 1$. Planes avoiding L capture at most $4 + 2 = 6$; those containing it get seven. The modified full-kernel value eight is covered by nine. All other ranks are immediate. □

Lemma 4.3 (Dimension four)**.** *Eight distinct points of affine dimension four admit cost at most nine.*

Proof. If a line direction captures at least six points, six points on two parallel lines span at most a plane. The remaining two points force that plane to have dimension two and are independent affine coloops, since the total affine dimension is four. Lemma 3.3 gives cost at most eight. Line captures seven or eight are incompatible with dimension four. Similarly, if a plane direction captures at least seven points, seven lie in an affine space of dimension at most three; dimension four forces a

seven-point three-dimensional core and one coloop. Lemmas 4.2 and 3.2 give cost at most nine. A plane cannot capture all eight.

We may assume line captures at most five and plane captures at most six. If no line captures five, $y_V = 9/4$ works: ranks one, two and three get 17/4, 13/2, 35/4, respectively. If just one line L captures five, $y_V = 2$, $y_L = 1$ works.

It remains to consider two different directions capturing five. Their selected subsets A, B meet in $t \in \{2, 3, 4\}$ points. If $t = 2$, their union is all eight and Lemma 3.4 bounds its affine dimension by three. If $t = 3$, seven points lie in a plane and the last point raises dimension to at most three. Both contradict the hypothesis. If $t = 4$, their union is six planar points; dimension four forces two independent affine coloops, already covered. This exhausts the cases; zero and full kernels satisfy the stated weights. □

5 Dependency supports in larger affine dimension

Use the notation of Lemma 3.1 with $m = 8$. Distinctness excludes a nonzero affine dependency supported on fewer than three points. We will repeatedly use the following observation. Let C be the exact support of a nonzero dependency, and suppose the dependency space restricted to C is one-dimensional. Then its direction space U_C has dimension $|C| - 2$. If that dependency belongs to E, all points of C lie in $b \le 2$ selected fibers, with coefficient sum zero on each. The coefficient space on C with these zero-sum conditions has dimension $|C| - b$; its image has dimension $|C| - b - 1$ and lies in $K \cap U_C$. This image is nonzero: for $|C| = 3$, a split into two nonempty groups would create a singleton with nonzero coefficient, so $b = 1$; for $|C| \ge 4$, the dimension is at least $|C| - 3 \ge 1$.

Lemma 5.1 (Dimension five)**.** *Every eight-point set of affine dimension five admits cost at most nine.*

Proof. Here $\dim \Lambda = 2$. Let C be the union of the supports of all dependencies, put $t = |C|$ and $U = \mathrm{dir}(C)$. All dependencies are supported on C, so $u = \dim U = t - 3$. Distinctness implies $t \ge 4$. For proper K, put $h = \dim(K \cap U)$ and retain $\delta = \dim(E \cap \Lambda)$.

First suppose C is proper, so $u \le 4$. If $\delta > 0$, then $h \ge 1$: otherwise projection modulo K is injective on the affine span of C, so the two selected fibers contain at most two core points, too few to support a dependency. If $\delta = 2$, all dependencies belong to E, so all of C lies in the two selected fibers. If $b \le 2$ of these meet C, the coefficient space on C with zero sums on the groups has image dimension $t - b - 2$. Thus

$$h \ge t - b - 2 \ge u - 1.$$

For $u = 3, 4$ these observations give $\delta \le h$; take $y_V = y_U = 1$, of cost $5 + u \le 9$. For $u = 1, 2$ they give $\delta \le 2h$; take $y_V = 1$, $y_U = 2$, of cost $5 + 2u \le 9$. Lemma 3.1 proves the inequalities for proper kernels, and the full kernel is covered because the cost is at least seven.

Now suppose $C = S$. Any two independent dependencies have supports whose union is all eight points, because they span Λ. Call a one-dimensional dependency subspace small if its support has size at most four. There are three possibilities: no small line; one small line of support size three or four; or two small lines with disjoint four-point supports. Indeed two different small lines must have union of supports all eight, hence both supports have size four and are disjoint. A combination using both has support eight, so there is no third small line.

For a proper K of dimension $k \ge 2$, we have $\delta \le k/2$. This is immediate for $\delta \le 1$. If $\delta = 2$, all eight points occur in the selected fibers, so $\dim E = 6$ and $\dim T(E) = 4$, forcing $k \ge 4$. Therefore

$$b_2(S/K) \le 2 + \tfrac{3}{2} \dim K \qquad (k \ge 2). \tag{7}$$

For a line kernel, $\delta = 2$ is impossible. If $\delta = 1$, Lemma 3.1 gives capture at most four, so the dependency in E spans a small line.

With no small line, weight $3/2$ on V works, of cost $15/2$. With one small support C_0, take $y_V = 3/2$, $y_{U_{C_0}} = 1/2$, of cost at most $17/2$. The observation preceding the lemma supplies a nonzero intersection for any exceptional line kernel, and (7) handles the larger proper kernels.

With two small lines, use their disjoint four-point supports C_1, C_2 and $U_i = \operatorname{dir}(C_i)$, each of dimension two. Take $y_V = y_{U_1} = y_{U_2} = 1$, of cost nine. If $\delta = 2$, both dependencies belong to E, giving nonzero intersections with both U_i. If $\delta = 1$ and its line is one of these two, one nonzero intersection suffices. Otherwise its generator uses both supports and hence all eight points. Then $\dim E = 6$ and $\dim T(E) = 5$, forcing $K = V$, contrary to propriety. Thus in every proper case

$$\delta \le \dim(K \cap U_1) + \dim(K \cap U_2),$$

as required by Lemma 3.1. In all three full-core cases, zero and full kernels satisfy the displayed certificates directly. □

Lemma 5.2 (Dimensions six and seven)**.** *In affine dimension six the sufficient cost is at most eight; in dimension seven it is seven.*

Proof. In dimension six the dependency space has dimension one. Let its nonzero support C have size $t \ge 3$, so $U = \operatorname{dir}(C)$ has dimension $t - 2$. If $\delta = 1$ for a proper kernel, the observation at the start of the section gives

$$\dim(K \cap U) \ge \max\{1, t - 3\} =: a.$$

Thus $y_V = 1$, $y_U = 1/a$ works and has cost $6 + (t-2)/a \le 8$. For $\delta = 0$ ambient weight one suffices. The full kernel is covered since the cost is at least six.

In dimension seven the points are affinely independent, so $\Lambda = 0$. Lemma 3.1 gives $b_2(S/K) \le 2 + \dim K$ for proper K. Weight one on V also covers the full kernel and has cost seven. □

Proof of Theorem 1.3. The upper bounds follow from the planar and low-dimensional arguments and Lemmas 5.1–5.2. All displayed weights are rational, nonnegative, finitely supported and fixed for S.

For sharpness let $S = \{0,1\}^3$ in a three-dimensional space. Let H_i be its three coordinate hyperplanes. Then $b_2(S/H_i) = 8$. For every W, the coordinate functionals restricted to W span W^*, so at least $\dim W$ restrictions are nonzero. Consequently

$$\sum_{i=1}^{3} \dim(W \cap H_i) \le 2 \dim W.$$

Summing any certificate's inequalities at these three kernels gives $18 \le \sum_W y_W \sum_i \dim(W \cap H_i) \le 2 \sum_W y_W \dim W$. Every certificate therefore costs at least nine, proving equality. □

Statements and Declarations

Funding. This research was self-funded. No external funding or material support was received.

Competing interests. The author owns Luminary Oddities and its Research Archive, where an author preprint of this work is available.

Data availability. All mathematical arguments supporting the results are contained in this article. No external dataset is required for the proofs. The accompanying code verifies the finite-parameter arithmetic example.

Author contribution. Scott Schubr directed the research and preparation of the manuscript and is responsible for its mathematical content and attribution.

Preprint availability. An earlier author preprint is available in the Luminary Research Archive, edition 2.

AI assistance. Generative AI tools, including ChatGPT and Codex, assisted research exploration, literature discovery, proof development, checking and manuscript preparation under the author's direction. The named author is responsible for the mathematical content and attribution.

References

[1] Y. Chen and Z. Zhang, *Explicit Folded Reed–Solomon and Multiplicity Codes Achieve Relaxed Generalized Singleton Bounds*, Proceedings of STOC 2025, pp. 1–12. `https://doi.org/10.1145/3717823.3718114`. arXiv:2408.15925v3, April 14, 2025; first version August 28, 2024.

[2] J. Brakensiek, Y. Chen, M. Dhar and Z. Zhang, *Combinatorial Bounds for List Recovery via Discrete Brascamp–Lieb Inequalities*, Proceedings of STOC 2026, pp. 365–376. `https://doi.org/10.1145/3798129.3800756`. Full version: arXiv:2510.13775v2, October 11, 2026; first version October 15, 2025.

[3] V. Guruswami and S. Kopparty, *Explicit subspace designs*, Combinatorica 36 (2016), 161–185. https://doi.org/10.1007/s00493-014-3169-1.

[4] J. Brakensiek, Y. Chen, M. Dhar and Z. Zhang, *From Random to Explicit via Subspace Designs With Applications to Local Properties and Matroids*, Proceedings of STOC 2026, pp. 619–630. `https://doi.org/10.1145/3798129.3800778`. arXiv:2510.13777v1, October 15, 2025.

[5] R. Goyal and V. Guruswami, *Structure Theorems (and Fast Algorithms) for List Recovery of Subspace-Design Codes*, arXiv:2512.08017v1, 2025.

[6] R. Goyal, V. Guruswami, Y. Sun and M. Wootters, *Improved LCL Equivalences and Curve Decoding Alphabet Sizes*, arXiv:2610.38862v1, October 28, 2026.