Sharp Minimum-Distance Tails for BAA Codes across Finite Fields

Majid Khabbazian

University of Alberta, Canada

Abstract. Fast randomized linear codes used in pseudorandom correlation generators and code-based proof systems require quantitative guarantees on the probability of sampling a code with insufficient distance. For chosen-block Block–Accumulate–Accumulate (BAA) codes, prior work provides finite-length upper bounds, sparse failure mechanisms, and permutation obstructions that persist under every nonzero coordinate scaling. We determine the global minimum-distance lower tail across field sizes. For a fixed-length constituent of minimum distance $d \geq 2$, sufficiently small fixed $\delta > 0$, and $Q = q - 1$, we prove, uniformly over finite fields and constituent coefficients,

$$\mathbb{P}[d_{\min} \leq \lfloor \delta N \rfloor] = \Theta\Big(N^{1-d} + N^{1-\lceil d/2 \rceil} Q^{-\lfloor d/2 \rfloor}\Big).$$

A support-dependent weighted-gap estimate controls every active-block count and every cancellation layer. It identifies linear field growth as the exact order needed for the N^{1-d} law. Certified extensions reach distance 0.30 for $[16, 8, 9]_q$ MDS blocks and 0.60 for $[16, 4, 13]_q$ MDS blocks over all $q \geq 2^{20}$; binary instances reach their stated useful thresholds. For the two MDS families, when $Q/N \to \infty$, we further prove that the known marked one-block coefficient is the leading constant of the actual failure probability. Grouping full-block gap events and bounding their dependence gives near-exact two-sided finite certificates at the published large-field benchmarks. These show that the existing numerical guarantees are essentially optimal for the unchanged ensemble.

Keywords: Random linear codes · BAA codes · Minimum distance · Pseudorandom correlation generators · Sampling failure

1 Introduction

Fast linear encoding is a central local operation in pseudorandom correlation generators (PCGs) and code-based proof systems. In the first setting, a structured linear map expands shares of sparse noise; distance rules out low-support linear tests but does not itself establish the underlying decoding assumption. In the second, code distance supplies separation for proximity and soundness arguments, while encoding contributes to prover cost [7,8,3,2]. These applications motivate linear-time encoders together with quantitative, rather than only asymptotic, distance guarantees.

A randomized encoder introduces a distinct question. Even if most sampled codes have linear distance, how likely is a sample to fall below the distance used by the application? For output length N and a prescribed relative threshold δ, the relevant quantity is

$$p_{N,q}(\delta) = \mathbb{P}[d_{\min}(C_N) \le \lfloor \delta N \rfloor].$$

An expected low-weight count upper-bounds this probability, but may count many words, supports, or cancellation traces generated by the same bad sample. Conversely, finding one rare failure pattern does not show that all other support sizes are negligible. Sharp sampling guarantees require both sides.

We study the chosen-block BAA ensemble of Peceny and Rindal [8]. A fixed short $[m, \sigma, d]_q$ constituent is used independently on M message blocks, followed by two independent permutation–scaling–prefix-sum stages. With $N = mM$, its code is

$$C_N = C_0^{\oplus M} \Pi_1 D_1 A \Pi_2 D_2 A. \tag{1}$$

The diagonal entries are independent uniform elements of $\mathbb{F}_q^*$; they are trivial over $\mathbb{F}_2$. The construction, its implementation advantages, and the relevance of its distance to cryptography are prior work [7,8]. We leave the encoder and sampling model unchanged.

The starting point. Peceny and Rindal give the exact binary first moment and a nonbinary support/shortening/trace upper bound. They identify the binary one-block scale $N^{1-\lceil d/2 \rceil}$, the large-field sparse scale N^{1-d}, and scaling-independent permutation obstructions. Their finite certificates give failure below $2^{-137.56}$ and $2^{-209.60}$ at their two large-field benchmark configurations, including all support sizes and surplus cancellations [8, Sections 4–7 and Appendices C, D, H]. Khabbazian proves a sharp growing-field tail for repetition-based RAA, identifies the paired-cancellation obstruction below linear field growth, and formulates a field-interpolation conjecture [6]. These results already provide the relevant sparse mechanisms. The remaining questions addressed here are a uniform global law across fields and the relation between the dominant marked count and the true failure probability.

1.1 Our results

A sharp law across field sizes. For every fixed $m \ge d \ge 2$, there is an explicit $\delta_*(m, d) > 0$ such that, for every fixed $0 < \delta \le \delta_*$,

$$p_{N,q}(\delta) = \Theta_{m,d,\delta}\Big(N^{1-d} + N^{1-\lceil d/2 \rceil}(q-1)^{-\lfloor d/2 \rfloor}\Big). \tag{2}$$

The constants are independent of the field and of the constituent's coefficients. The statement includes fixed fields, fields growing sublinearly with N, and linear or faster growth. In particular, $q - 1 = \Omega(N)$ is necessary and sufficient, in asymptotic order, for the constant-factor N^{1-d} law in this distance regime. For a repetition constituent, the result establishes the interpolation proposed in [6] at

a conservative positive distance threshold; it does not establish that conjecture throughout its larger proposed distance range.

The binary specialization gives the sharp order $N^{1-\lceil d/2 \rceil}$ and shows that local distance $d \geq 3$ is exactly what is needed for some positive relative-distance threshold to have vanishing failure. Prior BAA work already controls growing sublinear active-block counts at its finite large-field benchmarks [8, Appendix H.2]. The new increment here is a global, two-sided asymptotic theorem uniform across field sizes and constituents, rather than all-block-count coverage at those finite benchmarks alone.

Useful thresholds with smaller fields. An MDS-specific analysis proves the same interpolation at all fixed distances up to 0.30 for $[16, 8, 9]_q$ blocks and 0.60 for $[16, 4, 13]_q$ blocks, uniformly over $q \geq 2^{20}$. Thus the sharp law is not confined to an asymptotically tiny distance, nor to distances below 1/2. A separate binary spectral extension proves sharp tails for the four constituents in Table 1. The distance values themselves are chosen below the transitions studied in [8]; our claims are global two-sided tail theorems at the displayed exact thresholds, not new numerical estimates of those transitions.

From a marked count to an actual probability. For the two MDS parameter sets, and at their proved thresholds, we identify the leading constant when $(q-1)/N \to \infty$:

$$p_{N,q}(\delta) \sim \frac{d!}{m} \binom{m}{d}^2 \left(\frac{\delta}{1-\delta} \right)^d N^{1-d}. \tag{3}$$

The rate-1/4 marked coefficient is explicit in [8, Equation (41)]; the corresponding rate-1/2 marked contribution is evaluated in [8, Appendix D.2]. Our contribution is to prove that the coefficient describes the actual event probability, by grouping witnesses before counting and controlling their dependence. At the published benchmark lengths, the same argument confines $-\log_2 p$ to intervals of widths less than 0.000680 and 0.000286 bits. The existing certificates are therefore nearly tight for the unchanged ensemble. The tiny accompanying upper-bound improvement is a consequence, not the principal result.

1.2 Technical overview

Why field size and shortening dimension must be handled together. A support of size s need not contain just one codeword direction. If its shortening has dimension r, it contains at most $(q^r-1)/(q-1)$ projective directions. A prescribed total of $a+b$ zero returns across the accumulators has probability at most

$$\chi_q(r, a+b) = \min \left\{ 1, \frac{q^r-1}{q-1} (q-1)^{-a-b} \right\}.$$

This is the projective bound from [8], based on the cancellation framework of [6]. Simply paying the full projective count before summing traces can be wasteful, especially when q varies. The useful local invariant is the support excess $s-r+1 \geq (d-1)t+1$ for a support meeting t constituent blocks [8, Proposition 5.4].

A support-dependent weighted-gap estimate. The RAA analysis already uses weighted generating functions and an auxiliary field size [6, Section 6]. We adapt that strategy to arbitrary shortening dimensions and obtain the last-stage inequality

$$\sum_b \binom{h-b}{b} u^b g^{(2)}_{N,L}(h,b) \le \left(\delta + 2\sqrt{\delta u}\right)^h.$$

Here $g^{(2)}_{N,L}$ is the exact prescribed-trace low-weight probability. Unlike a fixed-trace half-distance estimate, this bound remains useful for $\delta > 1/2$ when u is small. Combining it with the known first-gap transform produces one support bound involving a matching polynomial. We choose $u = \max\{(q-1)^{-1}, cs/N\}$ after fixing the support size. When the second term wins, support excess yields the N^{1-d} contribution; when the actual inverse field wins, parity yields the second term of (2). Explicit geometric sums cover every t, rather than summing constants valid only for fixed t.

Grouping whole blocks instead of individual traces. The leading-constant argument exposes the first-permuted positions of all m coordinates of one constituent block. Its first accumulated word is constant on the resulting m intervals. Call an interval good if the second permutation maps it completely into a terminal segment of length L. If at least $d = m - \sigma + 1$ intervals are good, at most $\sigma - 1$ interval states must be set to zero. A nonzero solution exists in the σ-dimensional state space, for every scaling choice. This is a quantitative forward-gap realization of the coefficient-independent obstruction phenomenon already established in [8, Section 6], not a claim that such obstructions are new.

The full-block event has exact factorial moments and an inclusion–exclusion formula. A two-block overlap bound makes its union probability asymptotic to its first moment. In the upper bound, it covers every one-block no-surplus word with no second-stage return. All other one-block no-surplus terms cost an additional power of N; positive surplus costs $N/(q-1)$ relative to the leading scale. This proves (3) without treating a trace count as a probability or requiring a randomly transformed local code to remain MDS.

1.3 Relation to prior work and scope

Block codes followed by interleaved accumulators predate their cryptographic uses [4,9]. The original cryptographic Block–Accumulate work introduces an implementation-oriented family with independently sampled local maps and discusses reusing a suitable local code [7]. Peceny–Rindal develop that chosen-block specialization, its exact enumerators, nonbinary shortenings, generic-rank obstruction, and complete finite-benchmark certificates, including growing sublinear support counts [8]. We use those results explicitly as inputs. In particular, we do not claim the sparse exponents, the existence of scalar-independent bad permutations, or the rate-1/4 marked coefficient as discoveries of this work.

Large-field RAA distance analysis appears in [1]; sharp scaled and unscaled tails and the necessary linear-field order appear in [6]. Our nonbinary result

concerns the independently scaled ensemble only. Its field-uniform upper bound completes the interpolation at the stated thresholds and accommodates supports of differing shortening dimensions. Our large-field probability equivalent additionally resolves the clustering issue for the dominant rare events of the two certified MDS families.

2 Ensemble and main theorems

Let q be a prime power, $Q = q-1$, and let $C_0 \subseteq \mathbb{F}_q^m$ be a linear $[m, \sigma, d]_q$ code. Set $N = mM$ and $K = \sigma M$. Throughout, m and d are fixed; the field and constituent may vary with N but are selected before the global randomness. Sample mutually independent uniform permutations Π_1, Π_2 and diagonal matrices D_1, D_2 with independent uniform entries in $\mathbb{F}_q^*$. Define C_N by (1), with the zero-initialized, unterminated accumulator

$$(xA)_j = \sum_{i=1}^{j} x_i.$$

This action fixes the matrix convention: $A_{ij} = 1$ for $i \le j$ in row-vector notation. All maps after the constituent encoder are invertible, so the rate is always σ/m. No interleaver expurgation or terminal-state conditioning is imposed. For $0 < \delta < 1$, put

$$L = \lfloor \delta N \rfloor, \qquad p_{N,q}(\delta) = \mathbb{P}[d_{\min}(C_N) \le L],$$
$$a_d = \lceil d/2 \rceil, \qquad b_d = \lfloor d/2 \rfloor, \qquad \nu_d = a_d - 1.$$

All logarithms are natural unless the base is displayed. We write $\binom{u}{v} = 0$ for integer $u \ge 0$ and $v \notin \{0, \dots, u\}$, and $(u)_k = u(u-1)\cdots(u-k+1)$ for falling factorials.

Theorem 1 (Field-uniform sharp tail). *Fix $m \ge d \ge 2$ and set*

$$\delta_*(m, d) = \min\left\{ \frac{1}{144}, \frac{1}{36}\left(6 \cdot 128^m\right)^{-4/d} \right\}. \tag{4}$$

For every fixed $0 < \delta \le \delta_*(m, d)$ there are $c, C, N_0 > 0$, depending only on m, d, δ, such that for every multiple $N = mM \ge N_0$, every prime power q, and every $[m, \sigma, d]_q$ constituent,

$$c\Psi_{N,Q,d} \le p_{N,q}(\delta) \le C\Psi_{N,Q,d}, \qquad \Psi_{N,Q,d} = N^{1-d} + N^{1-a_d}Q^{-b_d}. \tag{5}$$

The explicit threshold is conservative. The proof in Section 4 is entirely analytic; no numerical certificate is used.

Corollary 1 (The field-size transition). *Under theorem 1,*

$$p_{N,q}(\delta) \asymp N^{1-d}\left[1 + (N/Q)^{b_d}\right].$$

Thus a sequence admits a constant-factor $\Theta(N^{1-d})$ law exactly when $Q = \Omega(N)$. If $Q = \Theta(N^\beta)$ with $0 \le \beta \le 1$, the order is $N^{1-a_d-\beta b_d}$; if $\beta \ge 1$, it is N^{1-d}. For $d = 2$, $p = \Theta(N^{-1} + Q^{-1})$. For fixed nonzero binary constituents, some fixed positive distance threshold has $p \to 0$ if and only if $d \ge 3$.

Table 1. Sharp probability orders at exact fixed thresholds. The MDS rows require $q \geq 2^{20}$; the binary rows use the spectra in [8].

Field	Constituent	δ_0	Sharp order
$\mathbb{F}_2$	Golay $[24, 12, 8]$	$1/10$	N^{-3}
$\mathbb{F}_2$	RM$(2, 5)$, $[32, 16, 8]$	$1/10$	N^{-3}
$\mathbb{F}_2$	BCH subcode $[32, 8, 12]$	$1/5$	N^{-5}
$\mathbb{F}_2$	Extended QR $[168, 84, 24]$	$11/100$	N^{-11}
$\mathbb{F}_q$	MDS $[16, 8, 9]$	$3/10$	$N^{-8} + N^{-4}Q^{-4}$
$\mathbb{F}_q$	MDS $[16, 4, 13]$	$3/5$	$N^{-12} + N^{-6}Q^{-6}$

The first claims follow by factoring Ψ. The binary necessity for $d = 1, 2$ is proved with the lower bounds in Section 4.3. No monotonicity of the actual failure probability as a function of q is assumed.

Theorem 2 (Certified useful thresholds). *The laws in Table 1 hold for every fixed $0 < \delta \leq \delta_0$ and along multiples of the constituent length. In the MDS rows, constants are uniform over $q \geq 2^{20}$ and over all constituent coefficients. In the binary rows, the named constituent is fixed.*

The proof combines analytic uniform sparse bounds with certified dense inequalities; Section 5 gives the reduction, and Sections A, B and D give the full details. The two MDS rows apply in particular to Reed–Solomon constituents. The field floor is a sufficient choice, not a claim of optimality.

Theorem 3 (Leading constant of the true probability). *Fix $(m, \sigma, d) = (16, 8, 9)$ and $0 < \delta \leq 3/10$, or $(m, \sigma, d) = (16, 4, 13)$ and $0 < \delta \leq 3/5$. For $Q \geq N$,*

$$p_{N,q}(\delta) = C_{m,d,\delta} N^{1-d} \left[1 + O_{\sigma,\delta}\left(N^{-1} + N/Q\right)\right], \tag{6}$$

uniformly over the fields and MDS coefficients, where

$$C_{m,d,\delta} = \frac{d!}{m} \binom{m}{d}^2 \left(\frac{\delta}{1-\delta}\right)^d. \tag{7}$$

In particular, $p_{N,q}(\delta) \sim C_{m,d,\delta} N^{1-d}$ whenever $Q/N \to \infty$.

The error in (6) need not vanish when $Q \asymp N$. The theorem is a probability equivalent only under the stated superlinear-field condition. Its proof appears in Section 6.

Theorem 4 (Near-exact finite certificates). *Let $K = 2^{20}$, $q \geq 2^{127}$, and $\mathsf{b}_{N,q}(\delta) = -\log_2 p_{N,q}(\delta)$. For every MDS constituent with the parameters indicated,*

$$137.5682716678 < \mathsf{b}_{2^{21},q}(3/10) < 137.5689511846, \qquad [16, 8, 9]_q, \tag{8}$$

$$209.6008685317 < \mathsf{b}_{2^{22},q}(3/5) < 209.6011545289, \qquad [16, 4, 13]_q. \tag{9}$$

The decimal endpoints are exact rationals. The left bit endpoints correspond to upper bounds on failure probability; the right bit endpoints correspond to lower bounds.

Section 7 explains the comparison and Section C proves the finite reductions. The executable certificates accompany the paper. A fixed large field at a benchmark length is not being treated as a superlinear field sequence at arbitrarily large lengths.

3 Support geometry and cancellation bounds

We state the exact quantities from [8,6] that the analysis uses. For an exact nonzero outer support $S \subseteq [N]$, write

$$s = |S|, \qquad r = \dim C_0^{\oplus M}[S], \qquad t = \#\{\text{constituent blocks met by } S\}.$$

Here $C[S]$ denotes the shortening whose words vanish outside S. Every nonempty local intersection of an exact support has a nonzero shortening. The Singleton shortening inequality [8, Proposition 5.4] gives

$$dt \le s \le mt, \qquad r \le s - (d-1)t, \qquad E := s - r + 1 \ge (d-1)t + 1. \tag{10}$$

There are at most $\binom{M}{t}(2^m)^t$ such supports. For an MDS constituent, $r = s - (d-1)t$ whenever all nonempty local intersections have size at least d.

An accumulator trace records the active indices where the prefix state is zero. Its first index is nonzero and its zero indices are nonadjacent. The number of size-a traces on s active positions and the prescribed-trace weight law are, respectively,

$$D(s,a) = \binom{s-a}{a}, \qquad 0 \le a \le \lfloor s/2 \rfloor, \tag{11}$$

$$g_N(s,a,h) = \frac{\binom{h-1}{s-a-1}\binom{N-h}{a}}{\binom{N}{s}}, \qquad h \ge 1. \tag{12}$$

For the second stage, define

$$g^{(2)}_{N,L}(h,b) = \sum_{j=1}^{L} g_N(h,b,j) = \mathbb{P}[X \ge h-b], \qquad X \sim \mathrm{Hyp}(N,L,h). \tag{13}$$

These are the exact gap and order-statistic laws of [6, Lemmas 3.5–3.6], used in [8, Section 5]. The hypergeometric identity is an equality in distribution; it does not identify every prescribed-trace event pointwise with suffix containment. We will also use the first-gap transform [6, Lemma 3.7]:

$$\sum_{h=1}^{N} g_N(s,a,h)\rho^h \le \left(\kappa \frac{s}{N}\right)^{s-a}, \qquad \kappa = \frac{\rho}{1-\rho}, \quad 0 < \rho < 1. \tag{14}$$

Proposition 1 (Support/trace bound, [8]). *Let $p(S)$ be the probability that some word with exact outer support S has final weight at most L. With*

$$\chi_q(r,c) = \min\left\{1, \frac{q^r-1}{Q}Q^{-c}\right\}, \tag{15}$$

one has

$$p(S) \le \sum_{a,h,b} D(s,a)g_N(s,a,h)D(h,b)g^{(2)}_{N,L}(h,b)\chi_q(r,a+b). \tag{16}$$

The trace sizes have the ranges in (11)*, $1 \le h \le N$, and impossible terms are zero. Summing over exact supports bounds $p_{N,q}(\delta)$.*

This is [8, Theorem 5.3]; at $q = 2$ the same upper bound is valid with $\chi = 1$, possibly counting unrealizable binary traces. The exposure order is essential: a first trace and its unordered input support determine the exact intermediate support before the second support geometry is averaged. The scalar bound is uniform in both active orders and both unordered supports. Thus a witness cannot select a new support after its geometry has been averaged. We inherit this conditioning from the cited master bound rather than introduce a new independence assumption.

For $0 \le z < 1$ and integer $k \ge 0$, write

$$S_k(z) = \sum_{t\ge1} t^k z^t, \qquad T_k(z) = \sum_{t\ge1} t^k z^{t-1}. \tag{17}$$

These are rational functions; for example $T_0(z) = (1-z)^{-1}$ and $T_{k+1}(z) = (zT_k(z))'$.

4 The field-uniform sharp law

The new upper bound sums trace choices together with a tunable field penalty before summing support sizes. We first derive the transform and then prove a uniform active-block estimate. The matching witnesses extend the known RAA mechanisms to arbitrary constituents and include the distance-two boundary case.

4.1 A weighted transform for arbitrary relative distance

Lemma 1 (Projective tilt and weighted last-stage tail). *If $Q \ge Q_0 \ge 1$ and $Q^{-1} \le u \le 1$, then*

$$\chi_q(r,c) \le (1+Q_0^{-1})^r u^{c-r+1} \le 2^r u^{c-r+1}. \tag{18}$$

For every $0 < \delta < 1$, $L = \lfloor \delta N \rfloor$, $u > 0$, and $1 \le h \le N$,

$$\sum_{b=0}^{\lfloor h/2 \rfloor} D(h,b)u^b g^{(2)}_{N,L}(h,b) \le \rho_\delta(u)^h, \qquad \rho_\delta(u) = \delta + 2\sqrt{\delta u}. \tag{19}$$

Proof. The projective count is at most $Q^{r-1}(1+Q^{-1})^r$. If $c \geq r$, use $Q^{-(c-r+1)} \leq u^{c-r+1}$; if $c \leq r-1$, the right side of (18) is at least one. This proves the first claim.

By (13), success for the second stage requires some specified $h-b$ of the h sampled labels to fall in a set of size L. Each specified choice has probability $(L)_{h-b}/(N)_{h-b} \leq \delta^{h-b}$, so

$$g^{(2)}_{N,L}(h,b) \leq \binom{h}{b}\delta^{h-b}.$$

Moreover $D(h,b)\binom{h}{b} = h!/[b!^2(h-2b)!]$. The left side of (19) is therefore at most

$$\sum_{b=0}^{\lfloor h/2 \rfloor} \frac{h!}{b!^2(h-2b)!}\delta^{h-b}u^b.$$

These are exactly the terms having equal last two exponents in the multinomial expansion of $(\delta + \sqrt{\delta u} + \sqrt{\delta u})^h$. All omitted terms are nonnegative. □

Define the matching polynomial

$$Z_s(v) = \sum_{a=0}^{\lfloor s/2 \rfloor} D(s,a)v^a.$$

It satisfies $Z_0 = Z_1 = 1$ and $Z_s = Z_{s-1} + vZ_{s-2}$.

Proposition 2 (Support moment bound). *Let $x = s/N$. If $Q^{-1} \leq u \leq 1$ and $\rho = \rho_\delta(u) < 1$, set $\kappa = \rho/(1-\rho)$. Then*

$$p(S) \leq 2^r u^{1-r}(\kappa x)^s Z_s\Big(\frac{u}{\kappa x}\Big). \tag{20}$$

When $Q \geq Q_0 \geq 1$, the factor 2^r may be replaced by $(1+Q_0^{-1})^r$.

Proof. Insert (18) into (16) and factor out $2^r u^{1-r}$. Sum the second trace with (19), and apply (14) to each first trace. The resulting polynomial is

$$\sum_a D(s,a)u^a(\kappa x)^{s-a}.$$

Factoring out $(\kappa x)^s$ gives (20). □

The last-stage estimate, not a bound on one fixed trace, is what permits a base approaching δ as $u \downarrow 0$. The support-dependent choice below refines the auxiliary-field approach of [6, Section 6] for supports with arbitrary shortening dimensions.

4.2 Summing all supports without an intermediate-weight gap

Assume $d \geq 2$ and $0 < \delta \leq 1/144$. For each support choose

$$u = \max\{Q^{-1}, \delta x\}, \qquad x = s/N. \tag{21}$$

Then $u \leq 1$ and $\rho_\delta(u) \leq \delta + 2\sqrt{\delta} < 1/2$. Partition supports according to which term attains the maximum; a tie can be assigned to either class.

Support-size tilt: $u = \delta x$. Here $\rho = \delta(1+2\sqrt{x})$ and $\delta \le \kappa \le 4\delta$. Since $Z_s(\delta/\kappa) \le 2^s$, (20) implies

$$p(S) \le 2^{r+s} 4^s (\delta x)^{s-r+1} \le 16^s (\delta x)^E. \tag{22}$$

For t active blocks, $s \le mt$, $x \le t/M$, and $E \ge (d-1)t+1$. Count supports using (10) and $\binom{M}{t} \le (3M/t)^t$. Their class-A contribution satisfies

$$P_t^{(A)} \le \frac{\delta t}{M} \left[\theta_A (t/M)^{d-2}\right]^t, \qquad \theta_A = 3 \cdot 32^m \delta^{d-1}. \tag{23}$$

Because $(d-2)t \ge d-2$ and $t/M \le 1$, summing over every t gives

$$\sum_t P_t^{(A)} \le \delta M^{1-d} S_{d-1}(\theta_A) \quad \text{if } \theta_A < 1. \tag{24}$$

Actual inverse field: $u = Q^{-1} \ge \delta x$. In this case

$$\kappa \le 6\sqrt{\delta} \le 1, \qquad Q\kappa x \le 6. \tag{25}$$

Indeed, $\kappa \le 2\rho \le 2\delta + 4\sqrt{\delta}$ and $Qx\rho = \delta Qx + 2x\sqrt{\delta Q} \le 1 + 2\sqrt{x} \le 3$. Let $j_s = \lfloor s/2 \rfloor$, $k_s = \lceil s/2 \rceil$, and $k_t = \lceil dt/2 \rceil$. For $v = (Q\kappa x)^{-1} \ge 1/6$, $Z_s(v) \le 2^s 6^{j_s} v^{j_s}$. Therefore (20) and $r \le s - (d-1)t$ give

$$p(S) \le 2^{r+s} 6^{j_s} (\kappa x)^{k_s} Q^{k_s - (d-1)t - 1}.$$

Replacing k_s by k_t costs at most $6^{k_s - k_t}$, by (25). Since $2^{r+s} 6^{j_s + k_s - k_t} \le 64^s$, support counting yields

$$P_t^{(B)} \le \theta_B^t (t/M)^{A_t} Q^{-B_t}, \qquad \theta_B = 3 \cdot 128^m (6\sqrt{\delta})^{d/2}, \tag{26}$$

where

$$A_t = \lceil dt/2 \rceil - t, \qquad B_t = (d-1)t + 1 - \lceil dt/2 \rceil. \tag{27}$$

The required bookkeeping is uniform in t:

$$\nu_d \le A_t \le \nu_d t, \qquad B_t \ge b_d, \qquad A_t + B_t = (d-2)t + 1. \tag{28}$$

For $d = 2b_d$ these follow from $A_t = (b_d - 1)t$, $B_t = (b_d - 1)t + 1$. For $d = 2b_d + 1$, they follow from $A_t = (b_d - 1)t + \lceil t/2 \rceil$ and $B_t = b_d t + 1 - \lceil t/2 \rceil$. Thus

$$\sum_t P_t^{(B)} \le M^{-\nu_d} Q^{-b_d} S_{\nu_d}(\theta_B) \quad \text{if } \theta_B < 1. \tag{29}$$

Proposition 3 (Finite all-field upper bound). *If $\delta \le 1/144$ and $\theta_A, \theta_B < 1$, then for every $N = mM$,*

$$p_{N,q}(\delta) \le \delta M^{1-d} S_{d-1}(\theta_A) + M^{-\nu_d} Q^{-b_d} S_{\nu_d}(\theta_B). \tag{30}$$

At $\delta \le \delta_*(m,d)$ the two series parameters are at most $1/2$, and the right side is at most

$$\delta\, 2^{d-1} (d-1)!\, M^{1-d} + 2^{\nu_d} \nu_d!\, M^{-\nu_d} Q^{-b_d}. \tag{31}$$

Proof. The two classes exhaust all exact supports; apply (24) and (29). The threshold (4) gives $\theta_B \le 1/2$, and $\theta_A/\theta_B = 4^{-m}6^{-d/2}\delta^{3d/4-1} \le 1$. For $k \ge 0$, $t^k \le k!\binom{t+k-1}{k}$ gives $S_k(1/2) \le 2^k k!$, proving (31). If $L = 0$, the failure probability is zero. □

Every bound preceding the series summation holds at each integer block count. In particular, (30) does not extrapolate a fixed-t asymptotic formula into a growing-t range.

4.3 Matching witnesses and their dependence

Choose one minimum-weight constituent word and place a copy in each of the M possible active blocks. Its M outer supports are disjoint. The two mechanisms below are the constituent-code versions of the boundary and paired-scalar witnesses of [6, Section 7]; the large-field sparse BAA scale is also identified in [8, Appendix D.1]. We include the probability and overlap bounds needed for a uniform law, including $d = 2$.

Boundary concentration. Let I be the final d positions and H the final L positions. Require a selected block's chosen support to map to I under Π_1, and require $\Pi_2(I) \subseteq H$. For every scaling, the first accumulated word is supported in I and the final one in H. It is nonzero by invertibility. The first events for different blocks are mutually exclusive; hence

$$p_{N,q}(\delta) \ge \frac{M\binom{L}{d}}{\binom{N}{d}^2} \sim \frac{d!\delta^d}{m} N^{1-d}. \tag{32}$$

Paired cancellations. Put $a = a_d$, $b = b_d$. For even $d = 2b$, require b disjoint adjacent pairs in the first-permuted support. For odd $d = 2b + 1$, require b adjacent pairs in $[N-1]$ and a singleton at N. There are $\binom{N-a}{b}$ possible supports in either case. In each pair, require the two scaled nonzero values to sum to zero. The pair equations involve disjoint fresh multipliers, so their probability is Q^{-b}. The first accumulator then has exactly a nonzero coordinates. Sending them into H under Π_2 forces failure. For the complete event F_i in block i,

$$\mathbb{P}(F_i) = \frac{\binom{N-a}{b}}{\binom{N}{d}} Q^{-b} \frac{\binom{L}{a}}{\binom{N}{a}} \sim \frac{d!}{b!}\delta^a N^{-a} Q^{-b}. \tag{33}$$

The asymptotic constants are independent of Q. For odd d, the F_i are mutually exclusive because each requires its own support at position N. For even d, set $X = \sum_i 1_{F_i}$, $\mu = \mathbb{E}X$, and $v_N = \binom{L}{a}/\binom{N}{a}$. Conditioning on one first support leaves the other uniform among the d-subsets of the remaining coordinates. Ignoring both second-permutation conditions and using the disjoint first-layer scalars gives, for $N \ge 2d$,

$$\mathbb{E}X^2 \le \mu + \frac{\binom{N}{d}}{\binom{N-d}{d}} \frac{\mu^2}{v_N^2}. \tag{34}$$

Indeed, at most $\binom{N-a}{a}$ supports in the whole path realize the second block's pair configuration, even before requiring avoidance of the first block. This proves the displayed overlap bound. Now $v_N \to \delta^a > 0$, the binomial ratio is bounded, and $\mu = \Theta(N^{1-a}Q^{-a})$ is uniformly bounded. Cauchy–Schwarz gives $\mathbb{P}[X > 0] \geq \mu^2/\mathbb{E}X^2 \geq c\mu$. In either parity,

$$p_{N,q}(\delta) \geq c_{m,d,\delta} N^{1-a_d} Q^{-b_d}. \tag{35}$$

The second-moment step matters at $d = 2$ and fixed Q, where the overlap need not be lower order than the first moment.

Proof (Completion of theorem 1). The upper bound is proposition 3, with $M = N/m$. The lower bounds (32) and (35) hold for every fixed $\delta > 0$. The maximum of two positive multiples of the two terms in Ψ bounds a positive multiple of their sum. □

For the binary threshold in corollary 1, $d \geq 3$ follows from the theorem. At $d = 2$, (34) gives a positive limiting lower bound for every fixed $\delta > 0$ and $Q = 1$. At $d = 1$, the boundary event in (32) gives the positive limit δ/m. Therefore neither $d = 1$ nor $d = 2$ permits a vanishing binary failure probability at any fixed positive threshold.

5 Extending the law to useful distance thresholds

The threshold in theorem 1 uses only (m, d) and deliberately forgets most constituent information. To reach the distances in Table 1, we use a uniform sparse estimate up to a fixed fraction of the blocks and a certified dense exponent on the entire complement. The underlying spectrum and support/trace exponents come from [8, Appendices C and D]; the additional step is a complete field-uniform summation with finite margins.

5.1 MDS constituents

For an MDS constituent, the exact shortening relation and support polynomial are

$$r = s - (d-1)t, \qquad U(z) = \sum_{\ell=d}^{m} \binom{m}{\ell} z^{\ell}. \tag{36}$$

There are at most $\binom{M}{t}[z^s]U(z)^t$ exact supports of size s meeting t blocks. We need only this upper count, not that every admissible set occurs as an exact support.

Choose $\lambda_0, c, \bar{\kappa}, Z > 0$ so that

$$\delta + 2\sqrt{\delta c \lambda_0} \leq \frac{\bar{\kappa}}{1+\bar{\kappa}} < 1, \qquad cZ^2 - \bar{\kappa}Z - \bar{\kappa} \geq 0. \tag{37}$$

For $t \le \lambda_0 M$, use $u = \max\{Q^{-1}, cx\}$ in proposition 2. When $u = cx$, the trace polynomial

$$F_s(\kappa, c) = \sum_a D(s,a) c^a \kappa^{s-a}$$

satisfies $F_s = \kappa F_{s-1} + c\kappa F_{s-2}$ and $F_s \le (cZ)^s$ by (37). The exact shortening relation consequently gives the pointwise bound

$$P_t^{(A)} \le c \frac{t}{M} \left[J_A (t/M)^{d-2} \right]^t, \qquad J_A = 3 \cdot 2^\sigma c^{d-1} U(Z). \tag{38}$$

When $u = Q^{-1}$, the inequality $cx \le Q^{-1}$ constrains the active-block fraction itself. Combining it with the parity relations (28) produces a second geometric series at the scale $M^{-\nu_d} Q^{-b_d}$. Proposition 7 proves, with explicit constants,

$$\sum_{t \le \lambda_0 M} P_t \le C_A M^{1-d} + C_B M^{-\nu_d} Q^{-b_d}. \tag{39}$$

The constants are independent of Q above the chosen field floor. Neither class discards intermediate weights or positive cancellation surplus.

For dense supports, let p_ℓ be the fraction of blocks with local support size ℓ, for $\ell \in \{0, d, \dots, m\}$. Write $\alpha = s/N$, $\beta = r/N$, $\eta = h/N$, $x = a/N$, $y = b/N$, and $\zeta = j/N$. The finite support/trace exponent $\mathcal{F}_{Q_0,\delta}$ is recalled precisely in Eq. (77). In particular it retains $\zeta \le \delta$ as a variable; it does not assume the final-weight maximum lies at $\delta > 1/2$. If that exponent is at most $-\varepsilon$ over every feasible tuple with $1 - p_0 \ge \lambda_0$, the complete dense contribution is at most

$$(M+1)^{\sigma+1} (N+1)^6 \exp\left[- \left(\varepsilon - \frac{\sigma}{m Q_0} \right) N \right]. \tag{40}$$

Lemma 7 accounts for the shifted binomials, type counts, projective prefactor, and the floor in L.

For both practical families use $\lambda_0 = 2/125$ and $Q_0 = Q_* = 2^{20} - 1$, with

σ	d	δ_0	c	$\bar{\kappa}$	Z	ε
8	9	3/10	1/4	3/5	16/5	1/1000
4	13	3/5	1/10	2	21	1/5000

All sparse checks are rational inequalities. Eighteen and twenty-one rational certificate leaves, respectively, prove the dense inequalities on the complete domains. Their acceptance test is stated in Section D; it does not rely on optimality of numerical saddle points.

5.2 Binary constituents

The binary endpoint admits a sharper, simpler sparse bound than the general projective estimate. Let $W_0(z) = \sum_w A_w z^w$ be the local spectrum, $J = 2^\sigma - 1$,

and let $\mu_{N,t}$ be the expected number of final low-weight words with t active outer blocks. For $\delta < 1/2$, put

$$\rho = 2\sqrt{\delta(1-\delta)}, \quad \kappa = \frac{\rho}{1-\rho}, \quad \theta = eJ\kappa^{d/2}, \quad H_\kappa = \max\{1, \sqrt{\kappa}\}.$$

The binary kernel and the first-gap transform give, for a prescribed outer word of weight w, the bound $(\kappa w/N)^{\lceil w/2 \rceil}$. Consequently, whenever $\kappa t/M \le 1$,

$$\mu_{N,t} \le H_\kappa \theta^t (t/M)^{\lceil (d-2)t/2 \rceil}. \tag{41}$$

For any sufficiently small fixed $\varepsilon > 0$, this is summable over all $t \le \varepsilon M$, with a tail starting at any fixed T of order $N^{-\lceil (d-2)T/2 \rceil}$. This uniform statement is proved in lemma 6.

The complement uses the known binary spectral exponent [8, Appendix C.3]. The complete rational certificates establish negative margins at the four thresholds in Table 1, over domains beginning at an explicit positive outer-weight fraction. Their 116 leaves and all sparse checks are specified in Sections A and D. The spectra are inputs from the cited paper; our calculation does not independently certify the QR constituent.

Proof (Proof of theorem 2). For the MDS rows, theorem 6 combines (39) and (40) with the parameters above. The certified margin satisfies $\varepsilon - \sigma/(mQ_*) \ge \varepsilon/2 > 0$ in both rows. This yields the desired upper orders uniformly for $Q \ge Q_*$. For the binary rows, apply theorem 5 with the data in Table 2. The lower bounds in Section 4.3 give the matching orders in all rows. A smaller fixed positive δ is covered by monotonicity of the upper event and by its own threshold-dependent lower-bound constant. □

Finite guarantees at linear-size fields. The MDS reduction is finite, not only asymptotic. At $K = 2^{20}$, exact-rational evaluation of Eq. (79) gives

$$\begin{aligned} &[16, 8, 9]_q : \;\; N = 2^{21},\ \delta = 3/10,\ q \ge N : p_{N,q}(\delta) < 2^{-83}, \\ &[16, 4, 13]_q : N = 2^{22},\ \delta = 3/5,\ q \ge N : \;\; p_{N,q}(\delta) < 2^{-111}. \end{aligned} \tag{42}$$

These are conservative sampling bounds, not optimized field-size thresholds. At much larger fields, retaining the exact one-block geometry gives the far tighter bounds developed next.

6 The leading constant of the actual failure probability

Throughout this section, C_0 is an $[m, \sigma, d]_q$ MDS constituent, so $d = m - \sigma + 1$. We first construct a geometric event for a whole constituent block, calculate its exact probability, and bound dependence between blocks. We then show that it covers the dominant trace class. The final asymptotic equivalent uses the two certified MDS families of theorem 2; the finite geometric lower bound itself is more general.

6.1 Full-block gaps and a dimension deficit

Fix one outer block and expose the images under Π_1 of all its m coordinate positions, including coordinates at which a subsequently chosen word may be zero. Write

$$p_1 < \cdots < p_m, \qquad p_{m+1} = N+1, \qquad I_i = \{p_i, \ldots, p_{i+1}-1\}, \quad \Delta_i = |I_i|.$$

Every first accumulated word confined to this block is constant on each I_i and zero before p_1. Its interval states form a code $V \subseteq \mathbb{F}_q^m$ of dimension σ, since the local permutation, nonzero scaling, and local prefix map are invertible.

Let $H = \{N-L+1, \ldots, N\}$ and $R = \Pi_2^{-1}(H)$. Conditional on Π_1, R is an independent uniform L-subset of $[N]$. Define

$$Y = |\{i \in [m] : I_i \subseteq R\}|, \qquad \mathcal{G} = \{Y \geq d\}. \tag{43}$$

An interval counted by Y will be called good.

Lemma 2 (Full-block gap witness). *The event $\mathcal{G}$ forces a nonzero final word of weight at most L, for every choice of the two layers of nonzero scalings. The implication requires only dimension σ and $d = m-\sigma+1$, not an MDS hypothesis on V.*

Proof. There are at most $m - d = \sigma - 1$ nongood intervals. Requiring their states to be zero imposes at most $\sigma - 1$ homogeneous linear equations on V. A nonzero solution exists. Its physical first accumulated support is a union of good intervals, and Π_2 sends it into H. The second scaling preserves support, and the last prefix sum remains supported in H. Invertibility ensures that the final word is nonzero. □

This is a concrete rank-deficiency witness within the structural-obstruction picture of [8, Section 6]. It is useful here because its full-block probability can be calculated without counting codeword directions or assuming a generic choice of scalings.

6.2 Exact moments and inter-block dependence

For $1 \leq k \leq m$, define

$$F_k(N,L) = \sum_{h=k}^{L} \frac{\binom{h-1}{k-1}\binom{N-h}{m-k}}{\binom{N}{m}} \frac{\binom{L}{h}}{\binom{N}{h}}. \tag{44}$$

Lemma 3 (Full-gap factorial moments). *For any fixed k-subset of the interval indices, $F_k(N,L)$ is the probability that all selected intervals are good. Therefore*

$$\begin{aligned} \mathbb{E}\binom{Y}{k} &= \binom{m}{k} F_k(N,L), \\ \pi_{N,L} := \mathbb{P}(\mathcal{G}) &= \sum_{k=d}^{m} (-1)^{k-d} \binom{k-1}{d-1} \binom{m}{k} F_k(N,L). \end{aligned} \tag{45}$$

Proof. The vector $(p_1 - 1, \Delta_1 - 1, \ldots, \Delta_m - 1)$ is a uniform weak composition of $N - m$. For a fixed k-subset of gaps, their sum has probability $\binom{h-1}{k-1}\binom{N-h}{m-k}/\binom{N}{m}$ of equaling h, by the known gap law applied to the full block. The union of those intervals is contained in a uniform L-set with probability $\binom{L}{h}/\binom{N}{h}$. This proves the first identity. For integer $0 \le y \le m$,

$$1_{\{y \ge d\}} = \sum_{k=d}^{m} (-1)^{k-d} \binom{k-1}{d-1} \binom{y}{k}.$$

For $y \ge d$, substitute $\binom{k-1}{d-1}\binom{y}{k} = (d/k)\binom{y}{d}\binom{y-d}{k-d}$ and integrate the binomial expansion of $(1-x)^{y-d}$ against x^{d-1} on $[0,1]$. For $y < d$ both sides vanish. Taking expectations proves (45). □

The alternating formula counts a block with more than d good intervals once, instead of once for every selected d-subset. It also groups the outer supports and trace descriptions that can yield the same block event.

Let $\mathcal{G}_i$ denote the event for block i. For $N \ge 2m$ and $0 < z < 1$, put

$$D_N^{\mathrm{disj}} = \frac{\binom{N-m}{m}}{\binom{N}{m}}, \qquad V_{N,d}(z) = \frac{\binom{N-d}{m-d}}{\binom{N}{m}} \left(\frac{z}{1-z}\right)^d. \tag{46}$$

Lemma 4 (Two-block overlap). *For distinct blocks i, j and $z^2 \ge L/N$,*

$$\mathbb{P}(\mathcal{G}_i \cap \mathcal{G}_j) \le \binom{m}{d}^2 \frac{V_{N,d}(z)^2}{D_N^{\mathrm{disj}}}. \tag{47}$$

Proof. Select d full gaps in each block and let their unions be U, V, of sizes h_1, h_2. Conditional on Π_1, the probability that both lie in R is

$$\frac{(L)_{|U \cup V|}}{(N)_{|U \cup V|}} \le (L/N)^{|U \cup V|} \le z^{h_1 + h_2},$$

because $|U \cup V| \ge (h_1 + h_2)/2$. If the union cannot fit, its probability is zero and the same upper bound holds. The two first-permuted full block supports have the law of independent uniform m-subsets conditioned to be disjoint. Removing that conditioning from a nonnegative numerator bounds

$$\mathbb{E}[z^{h_1+h_2} \mid \text{disjoint}] \le \frac{\mathbb{E}[z^{h_1}]\mathbb{E}[z^{h_2}]}{D_N^{\mathrm{disj}}}.$$

For one selected d-subset of gaps, the composition law gives

$$\mathbb{E} z^h \le \frac{\binom{N-d}{m-d}}{\binom{N}{m}} \sum_{h \ge d} \binom{h-1}{d-1} z^h = V_{N,d}(z).$$

Union-bound over the $\binom{m}{d}^2$ pairs of gap-index subsets. No independence between the actual block events is assumed. □

Proposition 4 (Finite lower bound on the actual probability). *For every full-rank $[m,\sigma]_q$ constituent, put $d = m-\sigma+1$. For $N = mM \geq 2m$, $1 \leq L < N$, and $0 < z < 1$ with $z^2 \geq L/N$,*

$$p_{N,q}(L/N) \geq M\pi_{N,L} - \mathcal{E}_N(z), \qquad \mathcal{E}_N(z) = \binom{M}{2}\binom{m}{d}^2 \frac{V_{N,d}(z)^2}{D_N^{\mathrm{disj}}}. \tag{48}$$

The bound remains valid when the nonzero scalings are chosen as functions of the permutations.

Proof. Each $\mathcal{G}_i$ forces failure for every scaling by lemma 2. Apply the first two Bonferroni terms, (45), and lemma 4. The only randomness needed for the lower bound is that of the two uniform permutations. □

6.3 A geometric cover for the dominant upper-bound class

Return to MDS constituents. Call a trace pair no-surplus if $a + b \leq r - 1$; then $\chi_q(r, a + b) = 1$. A positive-surplus pair has $a + b \geq r$.

Lemma 5 (Cover of the leading one-block class). *A bad word confined to one block with $b = 0$ and $a \leq r - 1$ implies the full-block event $\mathcal{G}$.*

Proof. The first trace has $s - a$ nonzero active-prefix gaps. Since $r = s - d + 1$, $s - a \geq d$. Each contains at least one full-block interval on which the state is nonzero, and distinct active-prefix gaps provide distinct intervals. When $b = 0$, the final prefix sum is nonzero from its first active input position to N. Final weight at most L therefore requires every second-stage input coordinate to lie in H. All the selected full intervals are consequently good. □

Only this direction is needed; not every realization of $\mathcal{G}$ must belong to the covered trace class. Retain the one-block no-surplus terms having $b \geq 1$ as the correction

$$\begin{aligned}\mathcal{C}_{N,L} = M \sum_{s=d}^{m} \binom{m}{s} \sum_{a=0}^{\min(\lfloor s/2 \rfloor, r-1)} \sum_{h=1}^{N} D(s,a) g_N(s,a,h) \\ \cdot \sum_{b=1}^{\min(r-1-a, \lfloor h/2 \rfloor)} D(h,b) g_{N,L}^{(2)}(h,b), \qquad r = s - d + 1.\end{aligned} \tag{49}$$

If $\mathcal{S}_N$ bounds all positive-surplus one-block terms, $\mathcal{B}_N$ all terms using $2 \leq t \leq T$ blocks, and $\mathcal{D}_N$ all terms using $t > T$ blocks, the exact class cover yields

$$p_{N,q}(\delta) \leq M\pi_{N,L} + \mathcal{C}_{N,L} + \mathcal{S}_N + \mathcal{B}_N + \mathcal{D}_N. \tag{50}$$

The geometric event can overlap the retained classes without invalidating this upper bound. Its role is to replace the multiplicity within the one-block, no-surplus, $b = 0$ class.

6.4 Extracting the leading constant

Proof (Proof of theorem 3). For fixed m, k, the full-gap formula has the expansion

$$F_k(N, L) = (m)_k N^{-k} \left(\frac{\delta}{1-\delta}\right)^k [1 + O_{m,\delta}(N^{-1})]. \tag{51}$$

Here is a uniform justification for summing the errors. Expanding finitely many falling-factorial factors gives

$$\frac{\binom{N-h}{m-k}}{\binom{N}{m}} = (m)_k N^{-k} \left[1 + O_m\left(\frac{h+1}{N}\right)\right]$$

on feasible terms, with the analogous absolute error bound at omitted endpoints. Also,

$$0 \le \delta^h - \frac{\binom{L}{h}}{\binom{N}{h}} \le \frac{C_\delta (h+1)^2}{N} \delta^h.$$

For $h \le L$, write the ratio as a product of h factors at most δ, and use $1 - \prod_i (1-u_i) \le \sum_i u_i$; the denominators satisfy $N - i \ge (1-\delta)N$. For $h > L$ the ratio is zero and the same inequality holds after enlarging C_δ, for large N. All polynomially weighted sums against $\binom{h-1}{k-1}\delta^h$ converge. These estimates prove (51), including its relative $O(N^{-1})$ error.

The $k = d$ term in (45) now gives

$$M\pi_{N,L} = C_{m,d,\delta} N^{1-d} [1 + O_{m,\delta}(N^{-1})], \tag{52}$$

while every $k > d$ term is lower order. The two copies of $\binom{m}{d}$ in $C_{m,d,\delta}$ come from selecting the good gaps and from $(m)_d = d!\binom{m}{d}$ in their placement probability. For any fixed $\sqrt{\delta} < z < 1$, $V_{N,d}(z) = O(N^{-d})$ and $D_N^{\mathrm{disj}} \to 1$, so the subtraction in (48) is $O(N^{2-2d})$. This proves the required lower equivalent.

For the upper bound, every term of (49) has $s - a \ge d + b \ge d + 1$. For $N \ge 2m$,

$$g_N(s, a, h) \le C_m N^{-(s-a)} \binom{h-1}{s-a-1},$$
$$D(h, b) g_{N,L}^{(2)}(h, b) \le \binom{h-b}{b} \binom{h}{b} \delta^{h-b}.$$

Since s, a, b range over fixed finite sets, these polynomially weighted geometric series give $\mathcal{C}_{N,L} = O(N^{-d})$ after multiplication by M. The complete remainder bounds proved in lemma 8 with $T = \lfloor N/1000 \rfloor$ give

$$\mathcal{S}_N = O(N^{2-d}/Q), \qquad \mathcal{B}_N = O(N^{3-2d}), \qquad \mathcal{D}_N = e^{-\Omega(N)}. \tag{53}$$

Those bounds use the practical parameters in Section 5. They apply for $Q \ge N$, since $cd > 1$ in both rows and hence $Q^{-1} \le cd/N$; their dense field floor holds for all sufficiently large N. For smaller fixed δ, the same parameter checks and

dense certificates remain valid by monotonicity, while the exact geometric terms use that actual δ. Insert (52) and (53) into (50). Divide the upper and lower bounds by $C_{m,d,\delta}N^{1-d}$. Since $d \geq 3$, every geometric error is $O(N^{-1})$ or smaller, and the field error is $O(N/Q)$. □

Thus the known marked coefficient is not an artificial leading multiplicity that can be discarded by replacing a first moment with an event count. The grouped geometric event attains it. This conclusion is specific to the large-field MDS probability law above: we do not identify the leading constant of the true binary failure probability.

7 Near-exact certification at cryptographic lengths

We evaluate (48) and (50) at the unchanged large-field benchmark parameters of [8, Theorem 7.1]:

Constituent	Rate	K	N	δ	L
$[16, 8, 9]_q$	1/2	2^{20}	2^{21}	3/10	629145
$[16, 4, 13]_q$	1/4	2^{20}	2^{22}	3/5	2516582

Throughout this comparison, $q \geq 2^{127}$. No extra accumulator, different constituent length, or altered permutation distribution is introduced.

Complete finite upper and lower bounds. The lower bound uses the exact $\pi_{N,L}$ and the two-block subtraction, with rational $z = 11/20$ and $31/40$ for the two rows. For the upper bound, Section C supplies explicit values of all four retained terms in (50). The finite sums use $h \leq 512$ and an outward geometric bound on every positive tail beyond that cutoff. All sparse block counts and positive-surplus layers are included; dense supports use the same 39 rational certificate leaves as the MDS interpolation theorem.

Proof (Proof of theorem 4). Proposition 4 and (50) sandwich the actual event probability. Lemma 8 bounds the remaining support/trace classes, and Section C.2 gives rational enclosures for all direct sums and tails. The supplied checker verifies both full dense coverings, the sparse conditions, tail ratios, and the strict rational-logarithm comparisons in (8) and (9). These are finite arithmetic certificates of the specified bounds, not empirical estimates from sampled codes. □

What the comparison shows. The source paper's full-precision certified lower bounds on $-\log_2 p$ are 137.5682486336 and 209.6008601171 [8, Appendix H.5]. Our left endpoints improve them by only about 0.0000230342 and 0.0000084146 bits. The substantive conclusion is the opposite-direction probability bound: no valid analysis of these unchanged ensembles can certify $-\log_2 p$ as large as 137.5689511846 or 209.6011545289, respectively. Even a 0.01-bit improvement over the published rounded values 137.56 and 209.60 is ruled out.

This comparison goes beyond precision in the old calculation. The source paper already observes that its complete trace upper bound is nearly equal to a directly summed positive partial count, while explicitly noting that this does not lower-bound the true failure probability [8, Appendix H.5]. The full-block event provides that missing lower bound. Enlarging the field or resampling only the scalings cannot eliminate the event, because it forces failure for every scaling choice.

Verification boundary. The general theorem is analytic. The useful-threshold results additionally use certified inequalities for explicit entropy functions, and the finite benchmark bounds additionally use outward rational summation. Section D specifies how to reconstruct and check each domain, including all coefficient residuals in the dual bounds. The artifact requires only Python's standard library. Finite regression tests supplement the proofs; they do not establish an all-length theorem. No proof-assistant formalization of the new analytic reductions is claimed.

8 Consequences and limitations

The field-uniform law separates two sources of sampling failure. The N^{1-d} term is geometric and remains under arbitrary nonzero scalings. The second term records the cost of pair cancellations and dominates below the linear-field scale. In the proved regime, increasing the field helps until that geometric scale is reached; beyond it, field growth no longer improves the polynomial order. For the two practical MDS families with $Q/N \to \infty$, theorem 3 further gives

$$-\log_2 p_{N,q}(\delta) = (d-1)\log_2 N - \log_2 C_{m,d,\delta} + o(1).$$

This quantifies the sampling-failure cost of the chosen distance and local redundancy. It is not a formula for the total security of a PCG or proof system: any application must still justify its own reduction, hardness assumption, and treatment of the code-sampling event.

The upper bounds assume independent uniform permutations and independent nonzero scalings; they do not cover implemented pseudorandom interleavers, shared diagonal layers, or expurgation. The geometric lower bound also holds when scalings are chosen after observing the permutations. The general theorem uses the conservative range (4); useful thresholds require the specified spectra or MDS parameters and certified margins.

We do not certify the optimized transitions or structural-container concentration studied in [8]. Nor do the binary first-moment constants of Section A.3 identify true probability constants. The near-exact finite bounds instead show that materially better sampling guarantees at the same distance require changing parameters or the sampling ensemble. Extending the distance ranges, obtaining exact constants at $Q \asymp N$, and analyzing additional accumulator layers remain separate questions.

Artifact availability. The companion artifact, release v1.0, contains the exact certificate inputs, checkers, regression tests, and recorded verification outputs [5]. Download the named release asset `baa_eprint_v1_artifact.zip`. Use Python 3.11 or later in the extracted artifact directory. Run `python verify_all.py` for all eight mathematical checks. Run `python verify_manifest.py` to check file integrity. Section D specifies the computational trust boundary and the archive checksum. The analytic theorem does not depend on a numerical certificate; the useful-threshold and finite-benchmark results do.

AI assistance and verification status. OpenAI's ChatGPT was used in the preparation of this work, including proof exploration and verification scripts. The author is responsible for all results.

References

1. Akhiani, P., Zhang, Y.: Distance of RAA codes over large finite fields (with applications in zkSNARKs and PCGs). Cryptology ePrint Archive, Paper 2026/524 (2026). https://eprint.iacr.org/2026/524
2. Block, A.R., Fang, Z., Katz, J., Thaler, J., Waldner, H., Zhang, Y.: Field-agnostic SNARKs from expand-accumulate codes. In: Advances in Cryptology—CRYPTO 2024. LNCS, vol. 14929, pp. 276–307. Springer (2024)
3. Brehm, M., Chen, B., Fisch, B., Resch, N., Rothblum, R.D., Zeilberger, H.: Blaze: Fast SNARKs from interleaved RAA codes. In: Advances in Cryptology—EUROCRYPT 2025. pp. 123–152. Springer (2025). Also Cryptology ePrint Archive, Paper 2024/1609
4. Graell i Amat, A., Le Bidan, R.: Minimum distance and convergence analysis of Hamming–accumulate–accumulate codes. arXiv:0905.4545 (2009)
5. Khabbazian, M.: Artifact for "Sharp Minimum-Distance Tails for BAA Codes across Finite Fields." Companion verification release v1.0 (2026). Archive identity and SHA-256 are specified in Section D. https://github.com/mkhabbazian/baa-sharp-tails/releases/tag/v1.0-eprint
6. Khabbazian, M.: Sharp minimum-distance lower tails for RAA codes. Cryptology ePrint Archive, Paper 2026/1075 (2026). https://eprint.iacr.org/2026/1075
7. Kolesnikov, V., Peceny, S., Rachuri, R., Raghuraman, S., Rindal, P., Shah, H.: Block-accumulate codes: Accelerated linear codes for PCGs and ZK. In: Advances in Cryptology—CRYPTO 2026, Part VIII. LNCS, vol. 16807, pp. 397–429. Springer (2026). Also Cryptology ePrint Archive, Paper 2025/1828
8. Peceny, S., Rindal, P.: Chosen-block BAA codes: Fast, near the Gilbert–Varshamov bound, and field-agnostic. Cryptology ePrint Archive, Paper 2026/1903 (2026). https://eprint.iacr.org/2026/1903
9. Rosnes, E., Graell i Amat, A.: On the analysis of weighted nonbinary repeat multiple-accumulate codes. arXiv:1101.5966 (2011)

A Binary refinements and certified thresholds

This appendix proves the binary part of theorem 2, gives the stronger analytic binary threshold, and records the first-moment constant. The input is the exact binary kernel and first moment from [8, Section 4], not a new enumerator.

A.1 A uniform active-block estimate

Let $W_0(z) = \sum_w A_w z^w$, $J = 2^\sigma - 1$, and $\mu_{N,t}$ be as in Section 5. Put $a(w) = \lceil w/2 \rceil$, $b(w) = \lfloor w/2 \rfloor$, and $\nu = \nu_d$. The exact single-accumulator transition is

$$P_N(w,h) = \frac{\binom{h-1}{a(w)-1}\binom{N-h}{b(w)}}{\binom{N}{w}} \quad (h \geq 1).$$

Its cumulative tail $Q_N(h,L) = \sum_{j=1}^{L} P_N(h,j)$ satisfies

$$Q_N(h,L) = \mathbb{P}[\mathrm{Hyp}(N,L,h) \geq \lceil h/2 \rceil] \leq \rho^h, \qquad \rho = 2\sqrt{\delta(1-\delta)} < 1, \tag{54}$$

when $\delta < 1/2$. This is the binary instance of the prescribed-gap law and hypergeometric Chernoff estimate in [6, Lemmas 3.6 and 5.1]. For a fixed outer word of weight w, write

$$F_N(w,\delta) = \sum_h P_N(w,h) Q_N(h,L) \leq (\kappa w/N)^{a(w)}, \qquad \kappa = \rho/(1-\rho). \tag{55}$$

The inequality follows directly from (14) with $a = b(w)$. Independence of the second interleaver holds conditional on the entire first output, so composition of the transitions is valid.

Set $c = (d-2)/2$, $g(t) = \lceil ct \rceil$, $\theta = eJ\kappa^{d/2}$, and $H_\kappa = \max(1, \sqrt{\kappa})$. There are $\binom{M}{t} J^t$ words using t active blocks, with $dt \leq w \leq mt$. If $\kappa t/M \leq 1$, (55) gives

$$\mu_{N,t} \leq \binom{M}{t} J^t (\kappa t/M)^{\lceil dt/2 \rceil} \leq H_\kappa \theta^t (t/M)^{g(t)}. \tag{56}$$

For the second inequality use $\binom{M}{t} \leq (eM/t)^t$ and $\lceil dt/2 \rceil = dt/2 + e_t$, where $e_t \in \{0, 1/2\}$. The factor κ^{e_t} is at most H_κ; it can be omitted for even d or $\kappa \leq 1$.

Lemma 6 (All sparse binary block counts). *Fix $d \geq 3$ and $0 < \delta < 1/2$. Choose $\varepsilon \in (0,1]$ with*

$$\kappa\varepsilon \leq 1, \qquad q_* := \theta\varepsilon^{(d-2)/2} < 1. \tag{57}$$

For every fixed integer $T \geq 1$,

$$\sum_{T \leq t \leq \varepsilon M} \mu_{N,t} = O_{C_0,\delta,\varepsilon,T}(N^{-g(T)}). \tag{58}$$

Such an $\varepsilon > 0$ exists for every fixed $\delta < 1/2$.

Proof. For $t \geq T$, $g(t) \geq g(T)$ and $t/M \leq \varepsilon$. Thus the right side of (56) is at most

$$H_\kappa M^{-g(T)} t^{g(T)} \theta^t \varepsilon^{g(t)-g(T)} \leq H_\kappa M^{-g(T)} \varepsilon^{-g(T)} t^{g(T)} q_*^t.$$

Sum the convergent series over all $t \geq T$. Existence follows from $d > 2$. □

Proposition 5 (Explicit analytic binary threshold). *If $d \geq 3$, $0 < \delta < 1/2$, and $\theta < 1$, then*

$$p_{N,2}(\delta) \leq \mu_N(\delta) \leq M^{-\nu} \sum_{t \geq 1} t^\nu \theta^t \leq \frac{\nu! \theta}{(1-\theta)^{\nu+1}} M^{-\nu}. \tag{59}$$

Together with (35)*, this gives $p_{N,2} = \Theta(N^{-\nu})$. A sufficient positive threshold is*

$$\delta \leq \frac{1 - \sqrt{1-\rho_*^2}}{2}, \qquad \rho_* = \frac{1}{1 + (2eJ)^{2/d}}. \tag{60}$$

Proof. The condition $\theta < 1$ implies $\kappa < 1$, so (56) applies to all $t \leq M$ without H_κ. Since $g(t) \geq \nu$, sum $\theta^t (t/M)^{g(t)} \leq M^{-\nu} t^\nu \theta^t$. The last inequality in (59) follows from $t^\nu \leq \nu! \binom{t+\nu-1}{\nu}$ and its generating series. At (60), $\theta \leq 1/2$. □

A.2 Joining the sparse estimate to a dense spectral margin

Write $H(u) = -u \log u - (1-u)\log(1-u)$ and $f(A,B) = AH(B/A)$ on $0 \leq B \leq A$, with $f(0,0) = 0$. The known binary exponents [8, Appendix C.3] are

$$\Gamma(u,v) = f(v, u/2) + f(1-v, u/2) - H(u), \tag{61}$$

$$\mathcal{A}_0(x) = \inf_{z>0} \{m^{-1} \log W_0(z) - x \log z\},$$

$$\Phi_\delta(x,y) = \mathcal{A}_0(x) + \Gamma(x,y) + \Gamma(y,\delta). \tag{62}$$

The transition domain is $u/2 \leq v \leq 1 - u/2$. Define

$$\mathcal{D}_{\delta,x_0} = \{(x,y) : x_0 \leq x \leq 1,\ x/2 \leq y \leq 1 - x/2,\ y \leq 2\delta\}. \tag{63}$$

Theorem 5 (Global binary spectral criterion). *Let $d \geq 3$ and $0 < \delta < 1/2$. Suppose* (57) *holds with $\varepsilon < 1$, put $x_0 = d\varepsilon/m$, and assume*

$$\sup_{(x,y) \in \mathcal{D}_{\delta,x_0}} \Phi_\delta(x,y) \leq -\eta_0 < 0. \tag{64}$$

Then $\sum_{t > \varepsilon M} \mu_{N,t} = e^{-\Omega(N)}$ and $p_{N,2}(\delta) = \Theta(N^{-\nu})$.

Proof. Any word with $t > \varepsilon M$ has outer weight $w \geq dt > x_0 N$. Bound this contribution by the part of the exact first moment

$$\sum_{w \geq x_0 N} [z^w] W_0(z)^M \sum_{h=1}^{N} P_N(w,h) \sum_{j=1}^{L} P_N(h,j).$$

Table 2. Binary certificate parameters. The margins hold throughout $\mathcal{D}_{\delta_0,x_0}$.

Constituent	δ_0	x_0	$\bar{\kappa}$	η_0	Leaves
Golay	1/10	1/1000	3/2	1/4000	23
RM$(2,5)$	1/10	1/1000	3/2	1/4000	21
BCH subcode	1/5	1/1000	4	1/4000	28
Extended QR	11/100	1/10000	2	1/200000	44

It has at most N^3 nonzero summands. Along any sequence of such summands, extract a subsequence with $(w/N, h/N, j/N) \to (x, y, v)$. Finite feasibility gives $x \geq x_0$, $x/2 \leq y \leq 1 - x/2$, and $y/2 \leq v \leq \delta < 1/2$. The exact coefficient bound $[z^w]W_0(z)^M \leq W_0(z)^M z^{-w}$ for every $z > 0$, followed by a limit superior and the infimum over z, gives exponent at most $\mathcal{A}_0(x)$. The finite binomial inequalities

$$\frac{e^{nH(k/n)}}{n+1} \leq \binom{n}{k} \leq e^{nH(k/n)} \tag{65}$$

and continuity of the entropy perspective on its closed feasible domain give transition exponents at most $\Gamma(x, y)$ and $\Gamma(y, v)$. The shifts by one and floor/ceiling terms vanish after division by N; this argument also applies to binomials at their endpoints.

On the feasible part below 1/2, $\Gamma(u, v)$ is nondecreasing in v: its interior derivative is $\log[(1 - u/(2(1 - v)))/(1 - u/(2v))] \geq 0$, and continuity covers boundaries. Thus $\Gamma(y, v) \leq \Gamma(y, \delta)$ and the limiting exponent is at most $-\eta_0$. If no uniform bound $e^{-\eta_0 N/2}$ held for individual summands at large N, a violating sequence and a compact subsequence would contradict this conclusion. Summing at most N^3 terms proves exponential decay. Add lemma 6 and (35). □

For the spectra in [8, Appendices A–C], use the data in Table 2, $\varepsilon = mx_0/d$, and $\bar{\theta} = 3J\bar{\kappa}^{d/2}$. The checker verifies the exact rational inequalities

$$4\delta_0(1 - \delta_0) \leq \left(\frac{\bar{\kappa}}{1 + \bar{\kappa}}\right)^2, \qquad \bar{\kappa}\varepsilon \leq 1, \qquad \bar{\theta}\varepsilon^{(d-2)/2} < \tfrac{1}{2}. \tag{66}$$

All four local distances are even. These inequalities imply (57); the exact certificate procedure in Section D proves (64). This completes the binary part of theorem 2.

A.3 The leading binary first moment

This refinement concerns the expected number of low-weight words, not its identification with the true probability. It also records a leading two-block contribution when $d = 3$. For fixed h, put

$$\beta_h(\delta) = \mathbb{P}[\mathrm{Bin}(h, \delta) \geq \lceil h/2 \rceil], \qquad K_w(\delta) = \frac{w!}{b(w)!} \sum_{h \geq a(w)} \binom{h-1}{a(w)-1} \beta_h(\delta).$$

The series converges for $\delta < 1/2$, by the binomial Chernoff bound $\beta_h \le \rho^h$. The exact kernel and the fixed-size hypergeometric limit give

$$F_N(w,\delta) \sim K_w(\delta) N^{-a(w)}. \tag{67}$$

Indeed, $N^{a(w)} P_N(w,h) \to [w!/b(w)!]\binom{h-1}{a(w)-1}$ and $Q_N(h,L) \to \beta_h(\delta)$. The summand is bounded by $w^{a(w)}\binom{h-1}{a(w)-1}\rho^h$, independently of N, using the falling-factorial ratio in the first-gap transform. Dominated convergence proves (67).

Proposition 6 (First-moment constant and leading block counts). *Under proposition 5 or theorem 5,*

$$\begin{aligned} \mu_N(\delta) &\sim C_{C_0}(\delta) N^{-\nu}, \\ C_{C_0}(\delta) &= \frac{1}{m} \sum_{\substack{d \le w \le m \\ a(w)=a(d)}} A_w K_w(\delta) + 1_{\{d=3\}} \frac{A_3^2}{2m^2} K_6(\delta). \end{aligned} \tag{68}$$

For $d \ge 4$, words meeting at least two blocks contribute $o(N^{-\nu})$ to the first moment. For $d = 3$, words meeting at least three blocks contribute $o(N^{-1})$.

Proof. Apply lemma 6 and, if needed, the exponentially small dense bound. For $d \ge 4$, $g(2) = d - 2 > \nu$, leaving only $\mu_{N,1} = M \sum_w A_w F_N(w,\delta)$. Use (67) in this finite sum. For $d = 3$, $g(3) = 2 > \nu = 1$; the two-block expression is

$$\mu_{N,2} = \binom{M}{2} \sum_w [z^w](W_0(z) - 1)^2 F_N(w,\delta).$$

Its smallest weight is six, with coefficient A_3^2, producing the displayed term by (67). All higher weights have $a(w) \ge 4$ and contribute $O(N^{-2})$. □

For even d, only local weight d appears in the first sum; for odd d, weights d and $d+1$ may both appear. Overlaps between low-weight words inside a single binary block are not controlled by this proposition, so no leading constant for $p_{N,2}$ is inferred from it.

B Uniform MDS bounds at the practical thresholds

Fix an $[m, \sigma, d]_q$ MDS constituent with $d \ge 3$, set $h_0 = d - 1 = m - \sigma$, and use the support polynomial U in (36). Write $U_0 = U(1)$. This appendix proves the finite bounds underlying the MDS rows of theorem 2 and (42).

B.1 Every sparse block count at every field above a floor

Choose $Q_0 \ge 1$, $0 < \lambda_0 < 1$, and $c, \bar{\kappa}, Z > 0$ such that

$$0 < c\lambda_0 \le 1, \qquad Q_0^{-1} \le c\lambda_0, \tag{69}$$

and (37) holds. Define

$$J_A = 3 \cdot 2^{\sigma} c^{h_0} U(Z), \qquad \vartheta_A = J_A \lambda_0^{h_0 - 1}, \tag{70}$$

$$W = \max\{1, \bar{\kappa}/c\}, \qquad D_0 = \max\{1, m/(cd)\},$$

$$J_B = 3U_0\big(4W \max\{1, \bar{\kappa}\}\big)^m, \qquad \vartheta_B = J_B D_0^{\nu_d} Q_0^{-(d-2)}. \tag{71}$$

Assume $\vartheta_A, \vartheta_B < 1$.

Proposition 7 (Uniform sparse MDS bound). *For every $Q \geq Q_0$, the sum of support-event upper bounds over supports meeting at most $\lambda_0 M$ blocks is at most*

$$cJ_A T_{h_0}(\vartheta_A) M^{-h_0} + J_B T_{\nu_d}(\vartheta_B) M^{-\nu_d} Q^{-b_d}. \tag{72}$$

Proof. Let $x = s/N \leq t/M \leq \lambda_0$ and choose $u = \max\{Q^{-1}, cx\}$. The parameter conditions imply $u \leq 1$, $\rho_\delta(u) < 1$, and $\kappa \leq \bar{\kappa}$ in proposition 2. If $u = cx$, (38) and $t/M \leq \lambda_0$ give

$$P_t^{(A)} \leq cJ_A M^{-h_0} t^{h_0} \vartheta_A^{t-1}. \tag{73}$$

For clarity, its trace-polynomial step uses $F_0 = 1$, $F_1 = \kappa$ and $F_s = \kappa F_{s-1} + c\kappa F_{s-2}$. The number $\Lambda = cZ$ satisfies $\Lambda^2 \geq \bar{\kappa}\Lambda + c\bar{\kappa}$ and $\Lambda > \bar{\kappa}$, so $F_s \leq (cZ)^s$ by induction. With $r = s - h_0 t \leq \sigma t$, the support bound becomes $2^{\sigma t} c^{h_0 t + 1} Z^s x^{h_0 t + 1}$. Summing Z^s over a fixed set of active blocks gives at most $U(Z)^t$, and $\binom{M}{t} \leq (3M/t)^t$ proves (38) and (73).

If $u = Q^{-1} \geq cx$, then $Q\kappa x \leq \bar{\kappa}/c \leq W$. For $v = (Q\kappa x)^{-1} \geq 1/W$, $Z_s(v) \leq 2^s W^{\lfloor s/2 \rfloor} v^{\lfloor s/2 \rfloor}$. Substitute into (20), use $r = s - h_0 t$, and replace $\lceil s/2 \rceil$ by $k_t = \lceil dt/2 \rceil$, paying at most $W^{\lceil s/2 \rceil - k_t}$. This yields

$$p(S) \leq (4W)^s (\kappa x)^{k_t} Q^{k_t - h_0 t - 1}.$$

There are at most $\binom{M}{t} U_0^t$ supports. Use $s \leq mt$, $x \leq t/M$, and $\kappa^{k_t} \leq \max(1, \bar{\kappa})^{mt}$ to get

$$P_t^{(B)} \leq J_B^t (t/M)^{A_t} Q^{-B_t}, \tag{74}$$

with A_t, B_t from (27). If this class contains any support, then $dt/N \leq x \leq 1/(cQ)$ and $t/M \leq D_0/Q$. By (28),

$$0 \leq A_t - \nu_d \leq \nu_d(t-1), \qquad (A_t - \nu_d) + (B_t - b_d) = (d-2)(t-1).$$

Consequently

$$\begin{aligned}(t/M)^{A_t} Q^{-B_t} &\leq M^{-\nu_d} Q^{-b_d} t^{\nu_d} D_0^{A_t - \nu_d} Q^{-(d-2)(t-1)} \\ &\leq M^{-\nu_d} Q^{-b_d} t^{\nu_d} \big(D_0^{\nu_d} Q_0^{-(d-2)}\big)^{t-1}.\end{aligned}$$

Thus $P_t^{(B)} \leq J_B M^{-\nu_d} Q^{-b_d} t^{\nu_d} \vartheta_B^{t-1}$. Sum this and (73) over all $t \geq 1$, obtaining (72). □

B.2 A finite dense reduction with every prefactor retained

The dense exponent is the support/trace exponent of [8, Appendices D and H], evaluated with a fixed field floor. We state its domain and finite prefactor because the use here is uniform in both N and q. For a probability vector $p = (p_0, p_d, \dots, p_m)$, define

$$\lambda = 1 - p_0, \qquad \alpha = \frac{1}{m}\sum_{\ell=d}^{m} \ell p_\ell, \qquad \beta = \frac{1}{m}\sum_{\ell=d}^{m} (\ell - h_0) p_\ell. \tag{75}$$

The support-type entropy is

$$\mathcal{S}(p) = \frac{1}{m}\left[-\sum_{\ell} p_\ell \log p_\ell + \sum_{\ell=d}^{m} p_\ell \log \binom{m}{\ell} \right].$$

Use $\eta = h/N$, $x = a/N$, $y = b/N$, and $\zeta = j/N$. The domain $\mathcal{D}(\lambda_0, \delta)$ consists of the probability-simplex constraints, (75), $\lambda \geq \lambda_0$, $0 \leq \zeta \leq \delta$, all scalar variables in $[0, 1]$, and

$$\begin{aligned} &2x \leq \alpha, \qquad x + \eta \leq 1, \qquad \alpha - x \leq \eta, \\ &2y \leq \eta, \qquad y + \zeta \leq 1, \qquad \eta - y \leq \zeta. \end{aligned} \tag{76}$$

Include a variable $v \in [0, 1]$ with $v \geq x + y - \beta$. The choice $v = (x + y - \beta)_+$ is always feasible. With H and f as in Section A, define

$$\begin{aligned} \mathcal{F}_{Q_0,\delta} = \mathcal{S}(p) &+ f(\alpha - x, x) + f(1 - \eta, x) + f(\eta, \alpha - x) \\ &+ f(\eta - y, y) + f(1 - \zeta, y) + f(\zeta, \eta - y) \\ &- H(\alpha) - H(\eta) - v \log Q_0. \end{aligned} \tag{77}$$

Lemma 7 (Complete finite dense bound). *If $\mathcal{F}_{Q_0,\delta} \leq -\varepsilon$ on $\mathcal{D}(\lambda_0, \delta)$ and $\kappa_0 = \varepsilon - \sigma/(mQ_0) > 0$, then for all $Q \geq Q_0$, the total master-bound contribution with $t > \lambda_0 M$ is at most*

$$(M + 1)^{\sigma+1} (N + 1)^6 e^{-\kappa_0 N}. \tag{78}$$

Proof. A type with $k_\ell = Mp_\ell$ blocks of each size has at most

$$\frac{M!}{\prod_\ell k_\ell!} \prod_{\ell=d}^{m} \binom{m}{\ell}^{k_\ell} \leq e^{N\mathcal{S}(p)}$$

supports. Increase $\binom{h-1}{s-a-1}$ to $\binom{h}{s-a}$ and the second-stage analogue in the exact master sum. Their respective ratios to the replacements are $(s - a)/h \leq 1$ and $(h - b)/j \leq 1$. Apply (65) to the six numerator binomials and two denominators. The exponent is the six perspective terms minus $H(\alpha) + H(\eta)$, with prefactor $(N + 1)^2$. The projective factor has the uniform envelope

$$\begin{aligned} \chi_q(r, a + b) &\leq (1 + Q_0^{-1})^r Q_0^{-(a+b-r+1)_+} \\ &\leq e^{\sigma N/(mQ_0)} Q_0^{-N(x+y-\beta)_+}. \end{aligned}$$

Hence each type/trace term is at most $(N+1)^2 e^{-\kappa_0 N}$. There are at most $(M+1)^{\sigma+1}$ types and $(N+1)^4$ integer choices of (a,h,b,j). Every actual tuple with $j \le \lfloor \delta N \rfloor$ lies in the stated domain. Multiplication gives (78). □

Theorem 6 (Finite MDS interpolation criterion). *Suppose* (69), (37), $\vartheta_A, \vartheta_B < 1$, *and the hypotheses of lemma 7 hold. Then for every* $Q \ge Q_0$,

$$\begin{aligned} p_{N,q}(\delta) \le{}& cJ_A T_{h_0}(\vartheta_A) M^{-h_0} + J_B T_{\nu_d}(\vartheta_B) M^{-\nu_d} Q^{-b_d} \\ &+ (M+1)^{\sigma+1} (N+1)^6 e^{-\kappa_0 N}. \end{aligned} \tag{79}$$

In particular, $p_{N,q}(\delta) = \Theta(\Psi_{N,Q,d})$, *uniformly above the field floor and in the MDS coefficient choices.*

Proof. Split at $t \le \lambda_0 M$ versus $t > \lambda_0 M$. Apply proposition 7 and lemma 7. The exponential term is $O(N^{1-d})$, uniformly in Q. The other terms have the required orders, and Section 4.3 supplies the matching lower bound. □

B.3 Instantiating and evaluating the criterion

Use the parameters of Section 5, $\lambda_0 = 2/125$, $Q_0 = 2^{20} - 1$, and $m = 16$. The square-root comparison in (37) is checked by

$$\frac{\bar\kappa}{1+\bar\kappa} > \delta_0, \qquad \left(\frac{\bar\kappa}{1+\bar\kappa} - \delta_0\right)^2 \ge 4\delta_0 c \lambda_0.$$

All other sparse checks are rational. Direct substitution gives

Parameters	ϑ_A	ϑ_B
$[16,8,9]$, $\delta_0 = 3/10$	$< 3 \cdot 10^{-5}$	$< 8 \cdot 10^{-19}$
$[16,4,13]$, $\delta_0 = 3/5$	$< 3 \cdot 10^{-9}$	$< 8 \cdot 10^{-22}$

The verifier proves, on the complete domains,

$$\begin{aligned} &\mathcal{F}_{Q_*,3/10} \le -1/1000, \ [16,8,9], \\ &\mathcal{F}_{Q_*,3/5} \le -1/5000, \ \ [16,4,13]. \end{aligned} \tag{80}$$

The respective trees have 18 and 21 leaves. The exact verification rule and data locations are given in Section D. Since $\varepsilon - \sigma/(mQ_*) \ge \varepsilon/2$, these comparisons establish every hypothesis of theorem 6.

For the finite evaluations, all sparse quantities in (79) are rational. The power series T_k can be evaluated by the Eulerian recurrence or by $T_{k+1} = (zT_k)'$. The dense term has the rational upper bound

$$(M+1)^{\sigma+1}(N+1)^6 2^{-\lfloor 5\varepsilon N/7 \rfloor}, \tag{81}$$

because $\kappa_0 \ge \varepsilon/2$ and $\log 2 < 7/10$. At $Q = N-1$ and the two benchmark lengths, the checker compares the resulting rational sums with 2^{-83} and 2^{-111}, respectively. These inequalities remain valid for every prime-power $q \ge N$ because the only remaining actual-field dependence in (79) is Q^{-b_d}. This proves (42) without assuming monotonicity of the actual ensemble probability.

C Complete finite two-sided bounds

This appendix proves the finite remainder estimates used in Sections 6 and 7, then specifies the outward evaluation of (48) and (50). Keep the practical MDS parameters of Section 5, put $\lambda_0 = 2/125$, and take $T = \lfloor N/1000 \rfloor$.

C.1 All support and cancellation classes

For a positive-surplus term, let $e = a + b - r + 1 \geq 1$. For $Q^{-1} \leq u \leq 1$, the projective bound improves to

$$\chi_q(r, a+b) \leq 2^r Q^{-e} \leq \frac{2^r}{Qu} u^e. \tag{82}$$

The second inequality uses $Qu \geq 1$ and $e \geq 1$. Therefore a support's positive-surplus contribution is at most $1/(Qu)$ times the moment bound (20).

For $t \leq T$, take $u = cs/N$. At the benchmark lengths and for all sufficiently large N with $Q \geq N$, this is permitted: $Q^{-1} \leq cd/N \leq u \leq c\lambda_0 < 1$, since $cd > 1$ in both rows. Now all trace classes, rather than only class A of the maximum-tilt split, are bounded by (38). Define

$$J = 3 \cdot 2^\sigma c^{d-1} U(Z), \quad \theta = J\lambda_0^{d-2}, \quad P_t = c\frac{t}{M}\left[J(t/M)^{d-2}\right]^t. \tag{83}$$

For $2 \leq t < T$,

$$\frac{P_{t+1}}{P_t} = J\left(\frac{t+1}{M}\right)^{d-2}(1+1/t)^{(d-2)t+1} \leq 2 \cdot 3^{d-2}\theta =: R_{\mathrm{blk}} < 1. \tag{84}$$

The inequality uses $(1+1/t)^t < e < 3$ and $(t+1)/M \leq \lambda_0$; the final strict comparison is checked at the displayed rational parameters.

Lemma 8 (Complete refinement remainders). *At either benchmark, and for all sufficiently large N with $Q \geq N$, the following choices are valid in* (50)*:*

$$\mathcal{S}_N = \frac{N}{Qcd}P_1, \qquad \mathcal{B}_N = \frac{P_2}{1 - R_{\mathrm{blk}}}, \tag{85}$$

$$\mathcal{D}_N = (M+1)^{\sigma+1}(N+1)^6 2^{-\lfloor 5\varepsilon_\sigma N/7 \rfloor}, \qquad \varepsilon_8 = \frac{1}{1000}, \quad \varepsilon_4 = \frac{1}{5000}. \tag{86}$$

In particular they have the asymptotic orders in (53).

Proof. For one active block, (82) gives a factor at most $N/(Qcd)$ relative to the complete P_1 bound. This covers every positive-surplus layer and every intermediate weight. For $2 \leq t \leq T$, sum (84) as a geometric tail. The dense estimate follows from lemma 7 and (81). Every actual $t > T$ satisfies $t/M > 16/1000 = \lambda_0$, so no block count is omitted. Finally $P_1 = O(N^{1-d})$, $P_2 = O(N^{3-2d})$, and the polynomial dense prefactor is dominated by its exponential decay. At smaller fixed δ the same bounds remain valid by monotonicity. □

C.2 Direct finite sums and complete tails

The two exact expressions to evaluate are the good-gap probability (45) and the no-surplus correction (49). The checker directly sums their positive summands through $h = 512$. The cutoff is computational only; every omitted positive tail is bounded.

For the kth moment summand in (44), write its value as $\tau_{k,h}$. Inside its positive support,

$$\frac{\tau_{k,h+1}}{\tau_{k,h}} = \frac{h}{h-k+1}\frac{N-h-m+k}{N-h}\frac{L-h}{N-h}. \tag{87}$$

Each factor is nonincreasing in h. To evaluate the correction, expand (13) as

$$g^{(2)}_{N,L}(h,b) = \binom{N}{h}^{-1}\sum_{v=0}^{b}\binom{L}{h-v}\binom{N-L}{v}.$$

For fixed s, a, b, v, the resulting positive summand including multiplicity is

$$T_{s,a,b,v,h} = \frac{M\binom{m}{s}D(s,a)\binom{N-h}{a}\binom{h-1}{s-a-1}\binom{h-b}{b}\binom{L}{h-v}\binom{N-L}{v}}{\binom{N}{s}\binom{N}{h}}. \tag{88}$$

Put $k = s - a$. Its consecutive ratio is

$$\frac{T_{s,a,b,v,h+1}}{T_{s,a,b,v,h}} = \frac{N-h-a}{N-h}\frac{h}{h-k+1}\frac{h+1-b}{h+1-2b}\frac{L-h+v}{N-h}\frac{h+1}{h+1-v}. \tag{89}$$

All five factors are nonincreasing on the positive tail; in particular, $L + v < N$ at the benchmark parameters. At $h = 513$ all denominators are positive. The checker verifies that the initial ratio r_{513} is in $[0, 1)$ for each summand family. Its entire tail from 513 onward is then at most the first omitted term divided by $1 - r_{513}$. At the end of a positive support the sequence becomes zero and cannot reappear.

The direct partial sum of F_k is a lower bound, and adding its geometric tail gives an upper bound. Assemble the alternating expression (45) using the upper endpoint for a negative coefficient when forming a lower bound, and conversely when forming an upper bound. This gives a rigorous interval for $M\pi_{N,L}$; it is not an alternating truncation with an assumed error sign. The correction is a positive sum and is enclosed without sign changes.

C.3 Certified numerical comparisons

Every positive rational summand is rounded outward to a multiple of 2^{-512} by integer division with remainder. Tail bounds are kept as exact fractions. The direct computations contain 64,064 insertions at rate 1/2 and 11,964 at rate 1/4, including the old no-surplus partial sum used for a separate comparison. Logarithms used to check decimal bit thresholds are enclosed by the rational method in Section D.

The regional outputs are below. Entries are lower bounds on $-\log_2$ of the corresponding upper-bounding quantity, rounded down.

Quantity	$\sigma = 8$	$\sigma = 4$
$M\pi_{N,L}$, upper enclosure	137.568951184	209.601154528
$\mathcal{C}_{N,L}$	148.620581094	221.901490012
$\mathcal{S}_N$	218.488442136	287.634941942
$\mathcal{B}_N$	188.442127929	320.189638706
$\mathcal{D}_N$	1217.999896810	376.999970419
$\mathcal{E}_N(z)$	248.923202488	389.019887784

The last row is an upper bound on the error subtracted in (48), not an additional upper-bound contribution. The tails beyond 512 are included in the first two rows. Their negative logarithms exceed 960 and 947 at rate 1/2, and exceed 514 for both quantities at rate 1/4.

Combining outward endpoints in (48) and (50) gives the four strict comparisons of theorem 4. The checker also establishes that the new complete upper bound is smaller than the old directly summed no-surplus partial count by a factor whose base-two logarithm exceeds 0.0000225617 and 0.0000084146 bits, respectively. Because that partial count lower-bounds the old positive master sum, this verifies a counting improvement independent of how the prior full-precision certificate was rounded.

D Certificate soundness and reproduction

The certificates prove the explicit dense inequalities used above. The upper secants, entropy tangents, and dual bounds are inherited from [8, Appendix H.4]; positive-series logarithm bounds also appear in [6, Lemma 6.7]. We specify the exact domains and acceptance rules for the present statements. The stored proposals are numerical witnesses to finite inequalities, not trusted optimization outputs.

D.1 Binary certificates

The four binary thresholds satisfy $\delta_0 < 1/4$. Thus (63) is the root rectangle

$$[x_0, 4\delta_0] \times [x_0/2, 2\delta_0] \tag{90}$$

intersected with $x \le 2y$; the condition $y \le 1 - x/2$ is automatic there. A certificate partitions this rectangle by exact midpoint splits. On a leaf $[a, b] \times [c, e]$, the data provide $z > 0$ and $p_1, p_2, p_3, p_4 \in (0, 1)$, all rational. Put

$$T_p(A, B) = -B \log p - (A - B) \log(1 - p).$$

Relative-entropy nonnegativity gives $f(A, B) \le T_p(A, B)$. Convexity of $-H$ gives its upper secant

$$S_{a,b}(x) = \frac{b - x}{b - a}[-H(a)] + \frac{x - a}{b - a}[-H(b)].$$

The affine expression

$$\begin{aligned} U(x,y) = m^{-1}\log W_0(z) - x\log z + T_{p_1}(y, x/2) + T_{p_2}(1-y, x/2) \\ + T_{p_3}(\delta_0, y/2) + T_{p_4}(1-\delta_0, y/2) + S_{a,b}(x) + S_{c,e}(y) \end{aligned} \tag{91}$$

bounds Φ_{δ_0} above on the feasible part of that leaf. Its maximum occurs at a feasible member of

$$(a,c),(a,e),(b,c),(b,e),\ (2c,c),(2e,e),(a,a/2),(b,b/2).$$

Thus at most eight rational-logarithm inequalities certify a leaf. The values of z and the p_i need not be stationary points. The checker reconstructs the complete split tree, checks the metadata and spectrum, and verifies every feasible vertex against the margin in Table 2. It also verifies (66). The four accepted trees have 23, 21, 28, and 44 leaves.

The compact spectra used are the following inputs from [8]:

$$\begin{aligned} W_{\mathrm{Golay}}(z) &= 1 + 759z^8 + 2576z^{12} + 759z^{16} + z^{24}, \\ W_{\mathrm{RM}}(z) &= 1 + 620z^8 + 13888z^{12} + 36518z^{16} + 13888z^{20} + 620z^{24} + z^{32}, \\ W_{\mathrm{BCH}}(z) &= 1 + 48z^{12} + 158z^{16} + 48z^{20} + z^{32}. \end{aligned}$$

The QR spectrum is exactly [8, Appendix B], encoded in the artifact. Its construction and distance are imported facts, not conclusions of the entropy checker. Independent small enumerations in the regression suite additionally check the BCH generator and the evaluation construction of $\mathrm{RM}(2,5)$.

D.2 Nonbinary certificates

The only nonconcave terms of (77) are $-H(\alpha)$ and $-H(\eta)$. Feasibility implies that their variables lie in

$$[d/1000, 1] \times [d/2000, \min(1, 2\delta_0)]. \tag{92}$$

The root follows from $m = 16$, $\lambda_0 = 2/125$, $\alpha \geq d\lambda_0/m$, $\eta \geq \alpha/2$, and $\eta \leq 2\delta_0$. Each rational midpoint-split leaf replaces the two negative entropies by upper secants. For positive rational probability vector p^* and rational $0 < \theta < 1$, use the global upper bounds

$$-\sum_i p_i \log p_i \leq -\sum_i p_i \log p_i^*, \qquad f(A,B) \leq -B\log\theta - (A-B)\log(1-\theta).$$

The resulting objective is affine, say $c_* + v_* \cdot z$, over a rational polytope $Az \leq b$, $Ez = e$, $\ell \leq z \leq u$. The variable order is

$$z = (p_0, p_d, \ldots, p_m, \alpha, \beta, \eta, x, y, \zeta, v).$$

Equalities are (75) and $\sum p_i = 1$; inequalities are $p_0 \leq 1 - \lambda_0$, (76), and $x + y - \beta - v \leq 0$. All box bounds are $[0,1]$, except the leaf bounds on (α, η) and the upper bound $\zeta \leq \delta_0$.

For rational multipliers $\mu \geq 0$ and unrestricted rational ξ, set $r_* = v_* - A^{\mathsf{T}}\mu - E^{\mathsf{T}}\xi$. Then throughout the leaf,

$$\mathcal{F} \leq c_* + \mu \cdot b + \xi \cdot e + \sum_i \max\{(r_*)_i \ell_i, (r_*)_i u_i\}. \tag{93}$$

To prove this, subtract the two constraint combinations from the affine objective and maximize the remaining linear functional over the box. The full residual is retained, so exact dual feasibility of a rounded optimizer output is unnecessary.

The checker constructs these constraints from the fixed metadata and reconstructs every leaf from the root. It normalizes p^* exactly, verifies all signs and dimensions, and encloses each logarithm and residual by a rational interval. For a residual interval, both endpoints are multiplied by both box endpoints before taking an upper bound. The upper endpoint of (93) must be at most $-1/1000$ for $\sigma = 8$ or $-1/5000$ for $\sigma = 4$. Every child of every split must be present. The 18-leaf and 21-leaf coverings verify (80); no domain is omitted on the basis of a floating-point infeasibility claim.

D.3 Exact logarithms, data, and the trust boundary

For $1 \leq y \leq 2$, let $z = (y-1)/(y+1) \in [0, 1/3]$. The positive series for $2 \operatorname{artanh} z$ gives, for every integer $J \geq 1$,

$$2\sum_{j=0}^{J-1} \frac{z^{2j+1}}{2j+1} \leq \log y \leq 2\sum_{j=0}^{J-1} \frac{z^{2j+1}}{2j+1} + \frac{2z^{2J+1}}{(2J+1)(1-z^2)}. \tag{94}$$

The upper remainder follows by replacing all omitted denominators by $2J+1$ and summing a geometric series. General positive rational arguments are reduced by reciprocation and powers of two. The routines use $J = 24$ and outward rational-grid rounding at precision 2^{-112} for logarithms; negative coefficients reverse interval endpoints. The finite benchmark sums use the separate 2^{-512} grid described in Section C. No floating-point logarithm is used in an acceptance test.

The companion artifact contains three directories: `binary/`, `nonbinary/`, and `refinement/`. Together they include the six unique rational certificate trees, the exact finite-bound evaluators, regression tests, and recorded outputs. The MDS certificate trees are identical in the latter two directories. The complete leaf data are part of the artifact. Its `MANIFEST.sha256` covers the supplied artifact files, excluding itself and transient verification or build outputs. The companion archive `baa_eprint_v1_artifact.zip` has SHA-256:

```
3996e048f5cf986a14625661209a8014
a14eb2537ce986ae9bd7608c03b199a4
```

The following commands run the mathematical checks from the extracted artifact root:

```
python binary/verify_certificates.py
python binary/check_identities.py
python nonbinary/verify_dense.py
python nonbinary/test_reductions.py --extended
python nonbinary/test_summations.py
python nonbinary/evaluate_bounds.py
python refinement/certify_bounds.py
python refinement/test_refinement.py
```

Only Python's standard library is needed. The binary tests check exact kernels, a complete small ensemble, compact constituent spectra, and consistency of the 116 printed ancillary certificate rows with the JSON data. The nonbinary tests include complete length-three ensembles over $\mathbb{F}_2, \mathbb{F}_3, \mathbb{F}_4, \mathbb{F}_5$ and pointwise checks of the support and summation bounds. The refinement tests check the full-gap formulas, the pointwise witness and upper cover, tail ratios, alternative cutoffs, and malformed certificate rejection.

The certificates establish the stated finite-dimensional inequalities and finite arithmetic comparisons. The written probabilistic reductions connect those comparisons to ensemble failure. Neither the regression suite nor the certificate checker is a proof-assistant verification of those reductions. All numerical assertions in the paper are therefore accompanied by their exact inputs, acceptance inequalities, and explicit trust boundary.