Universal Field-Size Caps and a Two-Sided Sandwich for Mutual Correlated Agreement on Smooth Reed–Solomon Domains

Version 13.2

Przemek Chojecki
ulam.ai

July 19, 2026

Abstract

We study the mutual correlated agreement (MCA), correlated agreement (CA), and list challenges of Arnon–Boneh–Fenzi [ABF26] for Reed–Solomon codes $C = \mathrm{RS}[\mathbb{F}, D, k]$ on smooth evaluation domains. The goal is to locate the threshold $\delta_C^*(\varepsilon)$ at cryptographic targets such as 2^{-128}.

On the unsafe side we prove a field-size-universal cap. For every finite field with $n < |\mathbb{F}| < 2^{256}$ and every smooth multiplicative row of rate $\rho \in \{\frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \frac{1}{16}\}$ with $k \le 2^{40}$,

$$\varepsilon_{\mathrm{mca}}(C, \delta) > \frac{1}{2k}\Big(1 - \frac{n}{|\mathbb{F}|}\Big) \ge 2^{-86} \gg 2^{-128}$$

throughout the printed near-capacity band, with the same bound for $\varepsilon_{\mathrm{ca}}$ on the integer-admissible subinterval. The proof is elementary: quotient-locator pigeonhole floors are composed with a self-contained deep-point conversion. The same mechanism extends to map-smooth domains, Chebyshev line rounds, bivariate circle codes, and rational scale maps relevant to genus-one rows.

At the deployed KoalaBear-sextic and Mersenne-31 circle parameters, identity-prefix floors and a flexible-budget conversion move the exact unsafe agreements to 1116047 and 1116023 for MCA, and to 1116046 and 1116022 for the corresponding list rows. We classify the simple-pole witnesses behind these floors. When the agreement is coprime to the domain order, every threshold witness is aperiodic. An asymptotic family then carries exponentially many aperiodically witnessed bad slopes below the entropy–subfield envelope. This refutes the original unnormalized polynomial aperiodic-band conjecture and forces the safe-side problem to begin only above that envelope, with an explicit reserve.

On the safe side we prove two self-contained reductions. First, for every linear code, MCA has error at most $(\lfloor \delta n \rfloor + 1)/q$ below one third of the minimum distance. Second, below one half of the minimum distance, MCA is bounded by ordinary CA up to an explicit tangent term. Combining this with the classical unique-decoding proximity gap of Ben-Sasson–Carmon–Ishai–Kopparty–Saraf [BCIKS20] gives a conditional half-distance safe frontier; a self-contained half-Johnson CA bound gives an import-free alternative.

The remaining safe-side work is organized as a conditional certificate template rather than a proved three-inequality closure. Its candidate buckets are structured quotient/prefix pullbacks, a base-field-normalized interior split-pencil census, and local primitive shift-pair control. The interior normalization keeps the pigeonhole ceiling before multiplying by the number of size-m supports: the unavoidable discrete floor is distinct from the smooth mean-plus-one upper model. Finite adjacent rows require the *sum* of every paid and residual

numerator to fit inside one row budget. In the asymptotic reserve regime a loss $L(n)$ is absorbable only when $R(n) \to \infty$, $R(n) = o(n)$, and $\log L(n) = o(R(n) \log |\mathbb{B}|)$; polynomial loss is therefore absorbed by $R \gg \log n$, while a generic $e^{o(n)}$ loss is not automatically absorbed. All unconditional deployed-row verdicts carry exact integer certificates; the adjacent safe rows remain explicitly conditional on a complete coverage-and-coalescing certificate.

Keywords: Reed–Solomon codes, proximity gaps, correlated agreement, list decoding, FRI, STARKs.
MSC 2020: 94B35 (primary); 11T71, 68Q25 (secondary).

Contents

1 Introduction

Proximity testing for Reed–Solomon codes underlies the soundness of modern SNARKs, and the sharpest available soundness accounting is phrased through *mutual correlated agreement* (MCA): a random linear combination $f_1 + \gamma f_2$ being δ-close to the code should force f_1, f_2 to be simultaneously explained by codewords on a *common* agreement set, except with small probability $\varepsilon_{\mathrm{mca}}(C, \delta)$ over γ. The survey of Arnon, Boneh and Fenzi [ABF26] crystallizes the state of the art into a *grand MCA challenge*: for $C = \mathrm{RS}[\mathbb{F}, L, k]$ over a smooth domain L, with rate $\rho \in \{1/2, 1/4, 1/8, 1/16\}$, degree $k \le 2^{40}$, and field size $|\mathbb{F}| < 2^{256}$, determine the largest proximity parameter δ^* at which $\varepsilon_{\mathrm{mca}}(C, \delta) \le 2^{-128}$ can hold.

Positive results reach the Johnson bound $1 - \sqrt{\rho}$. On the negative side, [Cho26b] showed $\delta^*_C(2^{-128}) < 1 - \rho - 1/64$ for all prime fields of size up to roughly 2^{150} (per-rate reaches between $2^{150.6}$ and $2^{161.8}$), by exhibiting quotient-locator witnesses whose bad-slope sets are value sets of

restricted symmetric sums over small subgroups; it further proved a *conditional* universal cap over every field of size at most 2^{256} — error 2^{-42} at gap 2^{-11}, assuming $|\mathbb{F}| \geq 2n$ — by composing quotient-core lists with the same conversion isolated and proved directly here (its universal-cap theorem). Above the value-set reach, what remained open was the *error-one* question, identified there with the *per-fiber collision problem*: controlling, at a *fixed* large prime, the collisions of characteristic-zero locator values under reduction. Extension fields, the setting actually deployed ([ABF26, §6.3] works over a sextic extension of KoalaBear), posed an additional obstruction: by the subfield-confinement theorem of [Cho26b], refined in Section 6, every explicit witness of [Cho26a, Cho26b] is $\mathbb{B}$-rational and contributes only $p^{-5} < 2^{-154}$ over the deployed extension.

This paper sharpens the universal cap — gap 2^{-9} (resp. 2^{-10}) in place of 2^{-11}, error 2^{-86} with no condition relating n and $|\mathbb{F}|$, and extension fields handled directly — for every field below 2^{256} at once, by changing what is counted.

Theorem 1.1 (universal unsafe cap, informal; see Theorem 4.1 and Corollary 5.1)**.** *Let $\mathbb{F}$ be any finite field with $|\mathbb{F}| < 2^{256}$, let $D \subseteq \mathbb{F}^\times$ be a smooth multiplicative evaluation domain of order n, and let $C = \mathrm{RS}[\mathbb{F}, D, k]$ have $k = \rho n \leq 2^{40}$ with $\rho \in \{1/2, 1/4, 1/8, 1/16\}$. Assume $2^{10} \mid n$ for $\rho \in \{1/2, 1/4, 1/8\}$, and $2^{11} \mid n$ for $\rho = 1/16$. Then*

$$\varepsilon_{\mathrm{mca}}(C, \delta) > 2^{-86} \gg 2^{-128}$$

for every $\delta \in [1 - \rho - 2^{-9}, 1 - \rho)$ at rates $1/2, 1/4, 1/8$, and for every $\delta \in [1 - \rho - 2^{-10}, 1 - \rho)$ at rate $1/16$. The same lower bound holds for $\varepsilon_{\mathrm{ca}}(C, \delta)$ on the integer-admissible subinterval ending at $1 - \rho - 1/n$. If $|\mathbb{F}| \geq 2n$, the lower bound is $\geq 2^{-42}$. Hence $\delta^*_C(2^{-128}) \leq 1 - \rho - 2^{-9}$ for $\rho \in \{1/2, 1/4, 1/8\}$ and $\delta^*_C(2^{-128}) \leq 1 - \rho - 2^{-10}$ for $\rho = 1/16$.

Combining the unsafe caps, the safe-side theorems, and the certificate refinements gives the following two-sided summary.

Theorem 1.2 (two-sided deployed intervals, informal; see Theorem 10.12, Table 5, Corollary 11.5)**.** *For the two deployed rate-$\frac{1}{2}$ MCA rows the exact active unsafe edges are row-specific. For the KoalaBear-sextic row,*

$$\delta^*_C(2^{-128}) \in \left[\frac{349525}{2^{21}}, \frac{981105}{2^{21}}\right] \qquad \textit{unconditionally,}$$

$$\delta^*_C(2^{-128}) \in \left[\frac{1}{4}, \frac{981105}{2^{21}}\right] \qquad \textit{modulo the single classical import [BCIKS20].}$$

For the Mersenne-31 circle line-round row,

$$\delta^*_C(2^{-100}) \in \left[\frac{349525}{2^{21}}, \frac{981129}{2^{21}}\right] \qquad \textit{unconditionally,}$$

$$\delta^*_C(2^{-100}) \in \left[\frac{1}{4}, \frac{981129}{2^{21}}\right] \qquad \textit{with the same import.}$$

The corresponding list rows have exact unsafe edges $490553/2^{20}$ and $490565/2^{20}$ and the self-contained Johnson-safe edge printed in Table 5. Genus-one and bivariate-circle transports carry analogous two-sided sandwiches with the unsafe edge supplied by their own rational or Chebyshev scale; they do not inherit the KoalaBear numerical endpoint.

The proof composes two ingredients. The first (Lemma 3.1) is a pigeonhole lower bound on list sizes at a single received word: the quotient locators of [Cho26a] attach to each ℓ-subset A of the order-N quotient coset $Q = D^a$ a codeword c_A at distance $\leq 1 - \rho - t/N$ from the word $u_{z_A} = x^{k+ta} + z_A x^{k+(t-1)a}$, where $z_A = -e_1(A)$; pigeonholing the $\binom{N}{\ell}$ subsets over their $\leq |\mathbb{B}|$ slope values produces one word u_z carrying a list of size $\binom{N}{\ell}/|\mathbb{B}|$. Two refinements of [Cho26a, Main Thm. (d)] matter here: we run the construction at *slack two* ($t = 2$), so that the listed codewords have degree $\leq k$ and live in $\mathrm{RS}[\mathbb{F}, D, k+1]$ without any divisibility demand on $k+1$ (which is odd in all challenge instantiations, while a is a 2-power); and we pigeonhole over the *subfield* $\mathbb{B}$ containing D rather than over $\mathbb{F}$, which is what makes the extension-field case go through. The second ingredient is the direct deep-point conversion of Theorem 2.6: in the integer-error range $f \leq n-k-1$, if $\varepsilon_{\mathrm{ca}}(C, \delta)$ is below $\eta(\frac{1}{k} - \frac{n}{k|\mathbb{F}|})$, then *every* list of $\mathrm{RS}[\mathbb{F}, D, k+1]$ at radius δ has size at most $\lceil |\mathbb{F}|\, \varepsilon_{\mathrm{ca}}/(1-\eta) \rceil$. Taking $\eta = \frac{1}{2}$, the heavy fiber of Lemma 3.1 violates this whenever $\binom{N}{\ell} \geq |\mathbb{B}|(|\mathbb{F}|/k + 1)$ — an inequality with hundreds of bits of slack throughout the challenge envelope — and the contrapositive yields $\varepsilon_{\mathrm{ca}}(C, \delta) > \frac{1}{2k}(1 - n/|\mathbb{F}|)$ throughout the integer-admissible range. The MCA cap then persists up to capacity by support-wise MCA monotonicity.

No value-set counting, exponential-sum estimate, or equidistribution hypothesis enters: pigeonhole fibers exist over every field unconditionally, and Theorem 2.6 converts list mass into correlated-agreement error directly. The cost of this route, relative to the error-one results of [Cho26a, Cho26b] below 2^{150}, is an error of order $1/k$ rather than $1 - o(1)$, and a gap 2^{-9} at the first three prize rates (uniformly 2^{-10} across all four), rather than the fixed gap $1/64$; Section 22 quantifies the trade-off and what remains to be settled.

Beyond the universal cap (Corollary 5.1, with a two-tier summary combining it with [Cho26b]), we prove: a deployed-parameter corollary at the KoalaBear-sextic instantiation of [ABF26, §6.3], with $\varepsilon_{\mathrm{mca}} > 2^{-22}$ at gap 2^{-7} and the same $\varepsilon_{\mathrm{ca}}$ lower bound on the integer-admissible subinterval (Corollary 5.4), extended to every interleaved row of [ABF26, Tables 2–3] via an interleaving-transfer lemma (Lemma 3.84 and Corollary 5.5); an independent slacked variant through [BCHKS25, Thm. 1.9] for robustness (Proposition 5.6); and a subfield-confinement lemma (Lemma 6.1) showing that for lines with $\mathbb{B}$-valued words over an extension $\mathbb{F}/\mathbb{B}$, every CA- or MCA-bad slope lies in $\mathbb{B}$. The latter has a structural consequence (Corollary 6.2): at deployed parameters the lines certifying Corollary 5.4 are necessarily *not* $\mathbb{B}$-rational, so the failure proved here is fiber-borne and currently non-constructive; exhibiting explicit certifying lines is posed as the explicit-witness question Question 22.1.

A further axis is that the counting side of the composition is not tied to the multiplicative structure of D. Section 7 proves the fiber lemma for arbitrary polynomial folding maps with complete uniform fibers (Definition 7.1 and Lemma 7.2), removes the divisibility hypothesis $a \mid k$, and instantiates the resulting universal cap (Theorem 7.4) on the deployed isogeny-smooth family: twin-coset x-coordinate domains of the circle group over $p \equiv 3 \pmod 4$ carry exact Chebyshev fibers at every dyadic scale, and over challenge fields containing i the bivariate circle code is diagonally equivalent to a Reed–Solomon code on a torus twin coset. The Mersenne-31/QM31 circle family of [HLP24] thereby acquires the same negative bands as the multiplicative rows — error $> 2^{-23}$ at gap 2^{-7} at the deployed sizes, gaps 2^{-9}–2^{-11} universally (Corollaries 7.12 and 7.13) — answering the circle case of the isogeny-smooth problem with the elliptic (ECFFT) case reduced to a precise lacunarity obstruction (Remark 7.16, the genus-one/circle transport question Question 22.3).

The paper has four main components. First, it gives a field-size-universal unsafe construction:

quotient-fiber and graded-prefix floors produce large lists at one received word, and a direct deep-point conversion turns that list mass into CA and MCA failure near capacity. Second, it supplies the safe side needed to interpret these failures as a threshold statement: a self-contained deep-regime theorem proves MCA soundness below one third of the distance, while a pincer theorem reduces MCA to ordinary CA up to half the distance, with the half-distance endpoint obtained from one isolated imported proximity-gap theorem. Third, it develops exact identity-prefix and simple-pole witnesses, proves aperiodic slope floors, and uses them to refute the original unnormalized band conjecture. Fourth, it packages the deployed consequences as finite certificates and states the remaining normalized safe-side ledger with every additive term exposed.

The final logical state is important. Below the entropy–subfield envelope, the unsafe identity-prefix mechanism itself creates aperiodic mass; no polynomial aperiodic upper bound is possible there. Above the envelope, the paper records a candidate three-bucket safe-side certificate: structured quotient/prefix pullbacks, the base-field-normalized interior split-pencil census, and local primitive shift-pair control. This is a template, not a proof that the buckets are exhaustive or that their support counts coalesce into sampled parameters. Polynomial losses are absorbed only with a reserve large enough to dominate their logarithmic cost, and every finite adjacent row requires one exact summed budget inequality.

Roadmap. Section 2 fixes definitions and proves the direct deep-point conversion. Sections 3 to 5 establish the universal cap and its challenge-level consequences; Sections 6 and 7 treat subfield confinement and map-smooth, circle, and genus-one transports. Sections 8 to 11 develop the staircase, structural ledgers, pincer, explicit witnesses, and finite certificates. Sections 12 and 13 state the active certificate interface and its integer compiler. Sections 14 and 15 prove the identity-scale and flexible-budget deployed floors. Sections 16 and 17 close the dual-weighted and split-kernel algebra, Section 18 proves the aperiodic floors and refutes the original band proposal, and Sections 19 and 20 give the exact collision and subfield-normalized ledgers. Section 21 states the corrected three-bucket closure template and the proper reserve/loss condition; Section 22 summarizes the resulting status. The appendices retain the longer L1/M1 reductions, proof routes, and computational calibration used to reach the final interface. Throughout, companion results of [Cho26a, Cho26b] are cited by their printed theorem names rather than numbers, since the numbering of those manuscripts is not yet frozen.

2 Preliminaries

Codes and distances. Throughout, $\mathbb{B} \subseteq \mathbb{F}$ are finite fields, $q := |\mathbb{F}|$, and $D = \alpha H \subseteq \mathbb{B}^\times$ is a multiplicative coset of a cyclic subgroup H of order n; we write D for the evaluation domain (denoted L in [ABF26]). The subgroup case is $\alpha = 1$, and taking $\mathbb{B} = \mathbb{F}$ is always allowed. For $k \le n$,

$$\mathrm{RS}[\mathbb{F}, D, k] \;:=\; \big\{(p(x))_{x\in D} \;:\; p \in \mathbb{F}[X],\ \deg p < k\big\} \subseteq \mathbb{F}^n, \qquad \rho := k/n.$$

For $u, v \in \mathbb{F}^n$, $\Delta(u, v)$ is the relative Hamming distance, $\Delta(u, C) := \min_{c\in C} \Delta(u, c)$, and the relative minimum distance of $C = \mathrm{RS}[\mathbb{F}, D, k]$ is $\delta_{\min}(C) = 1 - \rho + 1/n$. For $S \subseteq D$, $\Delta_S(u, v)$ and $\Delta_S(u, C)$ denote the corresponding restricted relative distances on S. Lists are denoted

$$\Lambda(C, \delta, u) := \{c \in C : \Delta(u, c) \le \delta\}, \qquad \Lambda(C, \delta) := \max_{u\in\mathbb{F}^n} |\Lambda(C, \delta, u)|.$$

For $s \geq 1$ the s-interleaved code is $C^{\equiv s} := \{(c^{(1)}, \ldots, c^{(s)}) : c^{(i)} \in C\} \subseteq (\mathbb{F}^n)^s$ with the column distance $\Delta_s(F, G) := \frac{1}{n}\left|\{j : \exists i,\ f_j^{(i)} \neq g_j^{(i)}\}\right|$, and $\Delta_s(F, C^{\equiv s}) := \min_{\mathbf{c} \in C^{\equiv s}} \Delta_s(F, \mathbf{c})$.

Definition 2.1 (correlated agreement error; [ABF26, Def. 4.1] up to notation)**.** For $0 \leq \delta_{\mathrm{fld}} \leq \delta_{\mathrm{int}} < 1$,

$$\varepsilon_{\mathrm{ca}}(C, \delta_{\mathrm{fld}}, \delta_{\mathrm{int}}) \ := \ \max_{f_1, f_2 \in \mathbb{F}^n} \Pr_{\gamma \leftarrow \mathbb{F}}\Big[\ \Delta(f_1 + \gamma f_2,\, C) \leq \delta_{\mathrm{fld}} \ \wedge\ \Delta_2\big((f_1, f_2),\, C^{\equiv 2}\big) > \delta_{\mathrm{int}}\ \Big],$$

and the no-proximity-loss variant is $\varepsilon_{\mathrm{ca}}(C, \delta) := \varepsilon_{\mathrm{ca}}(C, \delta, \delta)$. A slope γ realizing the event for a fixed pair (f_1, f_2) is called *CA-bad* for that pair.

Definition 2.2 (mutual correlated agreement error, support-wise; [ABF26, Def. 4.3] up to notation)**.** For $\delta \in (0, 1)$,

$$\varepsilon_{\mathrm{mca}}(C, \delta) := \max_{f_1, f_2 \in \mathbb{F}^n} \Pr_{\gamma \leftarrow \mathbb{F}}\left[\begin{array}{l} \exists S \subseteq D,\ |S| \geq (1-\delta)n,\ \Delta_S(f_1 + \gamma f_2, C) = 0, \\ \text{and } \Delta_S\big((f_1, f_2), C^{\equiv 2}\big) > 0 \end{array}\right].$$

Equivalently, γ is MCA-bad for (f_1, f_2) if some large agreement set S explains the point $f_1 + \gamma f_2$, but the same set S does not simultaneously explain f_1 and f_2 by two codewords of C.

Definition 2.3 (challenge threshold)**.** Following the grand MCA challenge of [ABF26, §1], for $\varepsilon^* \in (0, 1)$ set $\delta_C^*(\varepsilon^*) := \sup\{\delta \in (0, 1-\rho) : \varepsilon_{\mathrm{mca}}(C, \delta) \leq \varepsilon^*\}$, with $\sup \emptyset := 0$.

All displayed threshold intervals use the integer-ball convention: distances are integral, so a radius δ corresponds to the agreement requirement $\lceil (1-\delta)n \rceil$ and to $\lfloor \delta n \rfloor$ permitted errors. When a theorem certifies that every $\delta \in [\delta_0, 1-\rho)$ is unsafe at target ε^*, it follows that $\delta_C^*(\varepsilon^*) \leq \delta_0$ under the supremum convention; when it certifies safety at a single radius $\delta_1 < 1 - \rho$, it follows that $\delta_C^*(\varepsilon^*) \geq \delta_1$.

Fact 2.4 (cf. [ABF26, Fact 4.5])**.** *For every $\delta \in (0, 1)$: $\varepsilon_{\mathrm{ca}}(C, \delta) \leq \varepsilon_{\mathrm{mca}}(C, \delta)$.*

Proof. Fix (f_1, f_2) and let γ be CA-bad: $\Delta(f_1 + \gamma f_2, C) \leq \delta$ and $\Delta_2((f_1, f_2), C^{\equiv 2}) > \delta$. Choose $S \subseteq D$ of size at least $(1-\delta)n$ on which $f_1 + \gamma f_2$ agrees with a codeword of C. Since (f_1, f_2) is not δ-close to $C^{\equiv 2}$, there is no set of size at least $(1-\delta)n$ on which f_1 and f_2 are simultaneously explained by codewords of C; in particular this fails on the chosen S. Thus γ is MCA-bad for the same pair. Taking maxima over pairs gives the claim. □

Lemma 2.5 (support-wise MCA monotonicity)**.** *For every linear code $C \subseteq \mathbb{F}^D$, the function*

$$\delta \longmapsto \varepsilon_{\mathrm{mca}}(C, \delta)$$

is nondecreasing.

Proof. Fix a pair (f_1, f_2) and suppose that γ is MCA-bad at radius δ. Then there is a set $S \subseteq D$ with

$$|S| \geq (1-\delta)n$$

such that $f_1 + \gamma f_2$ is explained by C on S, but (f_1, f_2) is not simultaneously explained by two codewords of C on the same S. If $\delta' \geq \delta$, then

$$|S| \geq (1-\delta)n \geq (1-\delta')n,$$

so the same set S witnesses that γ is MCA-bad at radius δ'. Taking the maximum over pairs gives the claim. □

Conversion theorem. The main composition below uses the deep-point list-to-CA conversion proved here directly. This is the same consequence previously imported from Crites–Stewart, with the integer-radius admissibility condition made explicit; see Remark 2.17. We keep the slacked BCHKS/ABF fallback as a separate imported route.

Theorem 2.6 (deep-point list-to-CA conversion)**.** *Let $C = \mathrm{RS}[\mathbb{F}, D, k]$, let $C^+ := \mathrm{RS}[\mathbb{F}, D, k+1]$, let $q := |\mathbb{F}|$ and $n := |D|$, and assume $q > n$. Let $\delta \in (0,1)$ and set $f_\delta := \lfloor \delta n \rfloor$. Assume*

$$f_\delta \le n - k - 1.$$

For any $\eta \in [0,1)$, if

$$\varepsilon_{\mathrm{ca}}(C,\delta) \le \eta\Big(\frac{1}{k} - \frac{n}{kq}\Big),$$

then

$$\Lambda(C^+,\delta) \le \Big\lceil \frac{q\,\varepsilon_{\mathrm{ca}}(C,\delta)}{1-\eta} \Big\rceil.$$

Proof. Put $\epsilon := \varepsilon_{\mathrm{ca}}(C,\delta)$ and $f = f_\delta$. Let $U : D \to \mathbb{F}$ be any received word and suppose that $P_1, \dots, P_L \in \mathbb{F}[X]_{<k+1}$ are distinct polynomials whose restrictions to D all lie in the Hamming ball of relative radius δ around U. Since distances are integral, each P_i agrees with U on at least

$$a := n - f$$

points of D. The hypothesis $f \le n - k - 1$ gives $a \ge k + 1$, in particular $a > k$.

Let $\Omega := \mathbb{F} \setminus D$, so $|\Omega| = q - n$. For each $\alpha \in \Omega$ define the simple-pole line

$$f_\alpha(x) = \frac{U(x)}{x - \alpha}, \qquad g_\alpha(x) = -\frac{1}{x - \alpha} \qquad (x \in D).$$

If P_i agrees with U on a set S_i of size at least a, then for $z_i(\alpha) := P_i(\alpha)$ and every $x \in S_i$,

$$f_\alpha(x) + z_i(\alpha) g_\alpha(x) = \frac{P_i(x) - P_i(\alpha)}{x - \alpha}.$$

The right-hand side is the restriction of a polynomial of degree $< k$. Thus each distinct value among the $P_i(\alpha)$ gives a slope for which the line point is δ-close to C.

We next check the correlated-agreement far condition. If a polynomial $G \in \mathbb{F}[X]_{<k}$ agreed with g_α on more than k points of D, then

$$(X - \alpha)G(X) + 1$$

would be a polynomial of degree at most k with more than k roots, but its value at $X = \alpha$ is 1, a contradiction. Hence (f_α, g_α) is not δ-close to $C^{\equiv 2}$, because any common explanation on a support of size at least $a > k$ would in particular explain g_α there. Therefore the distinct values of $P_i(\alpha)$ are CA-bad slopes.

It remains to lower-bound the number of distinct values for some α. For $i \ne j$, the polynomial $P_i - P_j$ is nonzero of degree at most k, so it has at most k roots in Ω. Let $M(\alpha)$ be the number of distinct values among $P_1(\alpha), \dots, P_L(\alpha)$. Over all $\alpha \in \Omega$, the total number of colliding unordered pairs (i,j) with $P_i(\alpha) = P_j(\alpha)$ is at most $k\binom{L}{2}$. Equivalently, if the fiber sizes of the map $i \mapsto P_i(\alpha)$ are $m_v(\alpha)$, then

$$\sum_{\alpha \in \Omega} \sum_{v} \binom{m_v(\alpha)}{2} \le k \binom{L}{2}.$$

Choose $\alpha \in \Omega$ for which the inner collision count is at most $k\binom{L}{2}/(q-n)$. If the fiber sizes of the map $i \mapsto P_i(\alpha)$ are $m_1, \dots, m_{M(\alpha)}$, then

$$\sum_r m_r = L, \qquad \sum_r m_r^2 = L + 2\sum_r \binom{m_r}{2} \le L + \frac{kL(L-1)}{q-n}.$$

Cauchy–Schwarz therefore gives

$$L^2 \le M(\alpha) \sum_r m_r^2, \qquad M(\alpha) \ge \frac{L(q-n)}{q-n+k(L-1)} \ge \frac{L(q-n)}{q-n+kL}.$$

Consequently

$$\epsilon \ge \frac{1}{q} \cdot \frac{L(q-n)}{q-n+kL}.$$

Solving this inequality for L gives

$$L \le \frac{\epsilon q(q-n)}{q-n-k\epsilon q},$$

provided $\epsilon < (q-n)/(kq)$. Under the stated hypothesis, $\epsilon \le \eta(q-n)/(kq)$ with $\eta < 1$, so the denominator is at least $(1-\eta)(q-n)$ and hence

$$L \le \frac{q\epsilon}{1-\eta}.$$

Since U and the list were arbitrary, the displayed ceiling bounds the maximum list size of C^+ at radius δ. □

Remark 2.7 (relation to Crites–Stewart). The displayed conversion is the only consequence of Crites–Stewart needed by the universal-cap argument. The proof above is included to make the main cap self-contained. The separate slacked fallback through BCHKS/ABF remains an independent imported route; it is not used in the final deployed certificate grammar and is retained as a robustness check.

Proposition 2.8 (quantitative deep-list floor)**.** *Let $C = \mathrm{RS}[\mathbb{F}, D, k]$, let $C^+ := \mathrm{RS}[\mathbb{F}, D, k+1]$, put $q = |\mathbb{F}|$ and $n = |D|$, and assume $q > n$. Let $\delta_0 \in [0, 1-k/n)$, and suppose that for some received word $U : D \to \mathbb{F}$ there are $L \ge 1$ distinct codewords of C^+ in the closed ball of radius δ_0 around U. Then, for every $\delta \in [\delta_0, 1-k/n)$,*

$$\varepsilon_{\mathrm{ca}}(C, \delta) \ge \frac{L(q-n)}{q(q-n+kL)}.$$

For every such $\delta > 0$,

$$\varepsilon_{\mathrm{mca}}(C, \delta) \ge \frac{L(q-n)}{q(q-n+kL)}$$

as well. More precisely, some received line has at least

$$\left\lceil \frac{L(q-n)}{q-n+kL} \right\rceil$$

CA-bad finite slopes at every such radius; for $\delta > 0$, the same slopes are MCA-bad.

Proof. Fix $\delta \in [\delta_0, 1 - k/n)$ and choose distinct $P_1, \dots, P_L \in \mathbb{F}[X]_{<k+1}$ in the δ_0-ball around U. They are also in the δ-ball. Since $\delta < 1 - k/n$, if $f = \lfloor \delta n \rfloor$ then $f \le n - k - 1$, so every codeword in this δ-ball agrees with U on at least $a = n - f > k$ points.

Repeat the simple-pole construction in the proof of Theorem 2.6. For $\alpha \in \Omega := \mathbb{F} \setminus D$ set

$$f_\alpha(x) = \frac{U(x)}{x - \alpha}, \qquad g_\alpha(x) = -\frac{1}{x - \alpha} \qquad (x \in D).$$

For every distinct value of $P_i(\alpha)$, the word $f_\alpha + P_i(\alpha) g_\alpha$ is δ-close to C. The pair (f_α, g_α) is not δ-close to $C^{\equiv 2}$, because a common explanation on more than k points would in particular explain g_α there, and then $(X - \alpha)G(X) + 1$ would be a polynomial of degree at most k with more than k roots and value 1 at α. Thus the distinct values of the map $i \mapsto P_i(\alpha)$ are CA-bad slopes.

For $i \neq j$, the polynomial $P_i - P_j$ has degree at most k and is nonzero, hence it has at most k roots in Ω. Averaging over $|\Omega| = q - n$ gives an α for which the number of colliding unordered pairs among the L values is at most $k\binom{L}{2}/(q-n)$. If $m_1, \dots, m_M$ are the value multiplicities for this α, then

$$\sum_r m_r = L, \qquad \sum_r m_r^2 \le L + \frac{kL(L-1)}{q-n}.$$

By Cauchy–Schwarz,

$$M \ge \frac{L^2}{L + kL(L-1)/(q-n)} = \frac{L(q-n)}{q - n + k(L-1)} \ge \frac{L(q-n)}{q - n + kL}.$$

Since M is an integer, $M \ge \lceil L(q-n)/(q-n+kL) \rceil$. Dividing by the q possible finite slopes gives the CA lower bound. For $\delta > 0$, the MCA bound follows from Fact 2.4. □

Corollary 2.9 (deep-list trigger and ceiling)**.** *In the setting of Proposition 2.8, define*

$$\mathcal{E}_{q,k}(L) := \frac{L(q-n)}{q(q-n+kL)}, \qquad \mathcal{N}_{q,k}(L) := \left\lceil \frac{L(q-n)}{q-n+kL} \right\rceil.$$

Then $\mathcal{E}_{q,k}(L)$ is increasing in L, and

$$\mathcal{E}_{q,k}(L) > \frac{1}{2k}\left(1 - \frac{n}{q}\right) \quad \Longleftrightarrow \quad L > \frac{q-n}{k}.$$

Moreover

$$\mathcal{E}_{q,k}(L) < \frac{1}{k}\left(1 - \frac{n}{q}\right)$$

for every finite L, and the right-hand side is the limiting value as $L \to \infty$.

Proof. The derivative of $L/(q-n+kL)$ is positive. The trigger inequality is

$$\frac{L}{q-n+kL} > \frac{1}{2k},$$

which is equivalent to $kL > q - n$. The upper bound follows from $q - n + kL > kL$. □

Remark 2.10 (why this is stronger than the contrapositive). Theorem 2.6 is optimized for contradiction: once $L > (q-n)/k$, taking $\eta = 1/2$ already forces the printed $(2k)^{-1}(1-n/q)$ lower bound. Proposition 2.8 keeps the whole interpolation count. It is the form a staircase scanner should use when a quotient fiber is too small to cross the $1/(2k)$ trigger but still contributes a nonzero explicit bad-slope numerator, and also when a very large fiber improves the constant toward $(1/k)(1-n/q)$.

Proposition 2.11 (punctured simple-pole deep-list floor). *Let $C = \mathrm{RS}[\mathbb{F}, D, k]$, let $C^+ := \mathrm{RS}[\mathbb{F}, D, k+1]$, put $q = |\mathbb{F}|$ and $n = |D|$. Let*

$$\Omega_0 \subseteq \mathbb{F} \setminus D, \qquad m := |\Omega_0| > 0.$$

Let $\delta_0 \in [0, 1-k/n)$, and suppose that for some received word $U : D \to \mathbb{F}$ there are $L \geq 1$ distinct codewords of C^+ in the closed ball of radius δ_0 around U. Then, for every $\delta \in [\delta_0, 1-k/n)$,

$$\varepsilon_{\mathrm{ca}}(C, \delta) \geq \frac{Lm}{q(m+kL)}.$$

For every such $\delta > 0$,

$$\varepsilon_{\mathrm{mca}}(C, \delta) \geq \frac{Lm}{q(m+kL)}.$$

More precisely, there is some $\alpha \in \Omega_0$ such that the simple-pole line

$$f_\alpha(x) = \frac{U(x)}{x-\alpha}, \qquad g_\alpha(x) = -\frac{1}{x-\alpha} \qquad (x \in D)$$

has at least

$$\left\lceil \frac{Lm}{m+kL} \right\rceil$$

CA-bad finite slopes at every such radius; for $\delta > 0$, the same slopes are MCA-bad.

Proof. Repeat the proof of Proposition 2.8, but average only over Ω_0 instead of over all of $\mathbb{F} \setminus D$. The simple-pole argument works for every $\alpha \notin D$, so every distinct value among $P_1(\alpha), \ldots, P_L(\alpha)$ gives a CA-bad slope. For $i \neq j$, the nonzero polynomial $P_i - P_j$ has degree at most k, hence has at most k roots inside the smaller set Ω_0. Therefore the total number of colliding unordered pairs over all $\alpha \in \Omega_0$ is at most $k\binom{L}{2}$. Choose $\alpha \in \Omega_0$ for which the collision count is at most $k\binom{L}{2}/m$. If M is the number of distinct values at this α, the same Cauchy–Schwarz calculation gives

$$M \geq \frac{Lm}{m+k(L-1)} \geq \frac{Lm}{m+kL}.$$

Since M is integral, this gives the displayed bad-slope numerator. Dividing by the q finite slopes gives the CA lower bound, and the MCA statement for $\delta > 0$ follows from Fact 2.4. □

Corollary 2.12 (extension-pole deep-list floor). *Assume in addition that $\mathbb{B} \subsetneq \mathbb{F}$, $D \subseteq \mathbb{B}$, and $|D| \geq 2$. Put $b := |\mathbb{B}|$. In the setting of Proposition 2.11, take*

$$\Omega_0 := \mathbb{F} \setminus \mathbb{B}, \qquad m = q - b.$$

Then for every $\delta \in [\delta_0, 1-k/n)$,

$$\varepsilon_{\mathrm{ca}}(C, \delta) \geq \mathcal{E}^{\mathrm{ext}}_{q,b,k}(L) := \frac{L(q-b)}{q(q-b+kL)}.$$

For $\delta > 0$ the same bound holds for $\varepsilon_{\mathrm{mca}}(C, \delta)$. Moreover the witnessing line may be chosen with pole $\alpha \in \mathbb{F} \setminus \mathbb{B}$, and its pole vector $((x-\alpha)^{-1})_{x \in D}$ is not a scalar multiple over $\mathbb{F}$ of any $\mathbb{B}$-valued vector. In particular the witness line is genuinely extension-valued, not merely a $\mathbb{B}$-rational line viewed over $\mathbb{F}$.

Proof. The lower bound is Proposition 2.11 with $\Omega_0 = \mathbb{F} \setminus \mathbb{B}$. It remains to prove genuine extension-valuedness. Suppose that, for some $\lambda \in \mathbb{F}^\times$, all numbers $\lambda/(x-\alpha)$ with $x \in D$ lay in $\mathbb{B}$. Choose two distinct points $x, y \in D$. Then

$$\frac{y-\alpha}{x-\alpha} = \frac{\lambda/(x-\alpha)}{\lambda/(y-\alpha)} \in \mathbb{B}.$$

Writing this ratio as $r \in \mathbb{B}$, we cannot have $r = 1$ because $x \neq y$. Hence

$$\alpha = \frac{rx - y}{r-1} \in \mathbb{B},$$

contradicting $\alpha \notin \mathbb{B}$. Thus the pole vector is not projectively $\mathbb{B}$-valued. In particular the displayed affine parametrization has no $\mathbb{B}$-valued direction vector after rescaling. □

Corollary 2.13 (extension-pole quotient-remainder floor)**.** *Assume $\mathbb{B} \subsetneq \mathbb{F}$, $D \subseteq \mathbb{B}^\times$ is a multiplicative coset of order $n \geq 2$, put $b := |\mathbb{B}|$, and let $C = \mathrm{RS}[\mathbb{F}, D, k]$. Fix a quotient-remainder profile as in Lemma 3.7 with $K = k+1$ and agreement $A_0 > k$. Let*

$$L := H^{k+1}_{c,m,s}$$

be the certified heaviest prefix-fiber size, or replace L by either fallback

$$\left\lceil \frac{M_{c,m,s}}{I^{k+1}_{c,m,s}} \right\rceil, \qquad \left\lceil \frac{M_{c,m,s}}{|\mathbb{B}|^{w_c(s, A_0-k-1)}} \right\rceil.$$

Then, for every

$$\delta \in [1 - A_0/n, 1 - k/n),$$

there is a genuinely extension-valued simple-pole line with at least

$$\mathcal{N}^{\mathrm{ext}}_{q,b,k}(L) := \left\lceil \frac{L(q-b)}{q-b+kL} \right\rceil$$

CA-bad finite slopes. For $\delta > 0$, these slopes are MCA-bad. Equivalently,

$$\varepsilon_{\mathrm{ca}}(C, \delta) \geq \frac{L(q-b)}{q(q-b+kL)}, \qquad \varepsilon_{\mathrm{mca}}(C, \delta) \geq \frac{L(q-b)}{q(q-b+kL)} \quad (\delta > 0).$$

If $L > (q-b)/k$, this gives the clean trigger

$$\varepsilon_{\mathrm{ca}}(C, \delta) > \frac{1}{2k}\left(1 - \frac{b}{q}\right), \qquad \varepsilon_{\mathrm{mca}}(C, \delta) > \frac{1}{2k}\left(1 - \frac{b}{q}\right) \quad (\delta > 0).$$

Proof. Lemma 3.7 gives a $\mathbb{B}$-valued received word with at least L codewords of C^+ in the closed ball of radius $1 - A_0/n$. Apply Corollary 2.12. The trigger inequality is

$$\frac{L(q-b)}{q(q-b+kL)} > \frac{q-b}{2kq},$$

which is equivalent to $kL > q - b$. □

Definition 2.14 (extension-pole certificate)**.** An extension-pole certificate for a list $P_1, \dots, P_L \in \mathbb{F}[X]_{<k+1}$ and a received word U consists of:

(i) a declared proper subfield $\mathbb{B} \subsetneq \mathbb{F}$ containing D and a pole $\alpha \in \mathbb{F} \setminus \mathbb{B}$;

(ii) the agreement supports showing that each P_i lies in the declared ball around U;

(iii) the bucket table for the values $P_i(\alpha)$, with duplicate values coalesced;

(iv) the simple-pole line (f_α, g_α) and the proof that g_α is not projectively $\mathbb{B}$-valued;

(v) the final bad-slope numerator, namely the number of distinct buckets in (iii).

The verifier checks $\alpha \notin \mathbb{B}$, the support agreements, the value buckets, and the simple-pole far condition. For quotient-remainder floors, the list in (ii) may itself be supplied by a prefix-fiber certificate from Definition 3.6.

Remark 2.15 (what the extension-pole construction changes). The subfield-confinement theorem says that $\mathbb{B}$-rational lines are inert over a proper extension. Corollary 2.12 gives the matching existence theorem on the failure side: any large list can be converted using poles restricted to $\mathbb{F} \setminus \mathbb{B}$, so the resulting bad line is genuinely extension-valued. Thus the extension-line bucket in a prize scanner is no longer only an abstract warning; it has an explicit simple-pole source and a stable numerator with $q - b$ replacing $q - n$. What remains to be settled is not existence of extension-valued witnesses, but their sharp classification inside the aperiodic/residue-line ledgers and, for a fully explicit counterexample, printing a concrete pole α with its value-bucket table.

Theorem 2.16 ([BCHKS25, Thm. 1.9], as stated in [ABF26, Thm. 5.2])**.** *Let $C = \mathrm{RS}[\mathbb{F}, D, k]$, let $\rho = k/n$, and let $\delta \in (0, 1-\rho)$ satisfy $\delta + 2/n \le 1 - \rho - 1/n$. If*

$$\varepsilon_{\mathrm{ca}}\Big(C,\ \delta + \tfrac{2}{n},\ 1 - \rho - \tfrac{1}{n}\Big) \ <\ \frac{1}{2n}, \qquad \textit{then} \qquad \Lambda(C, \delta) \ <\ |\mathbb{F}|.$$

Remark 2.17 (on imports and radius conventions). Theorem 2.6 is the only Crites–Stewart-type consequence used by the main cap, and it is now proved above by the simple-pole deep-point argument. The important admissibility condition is the integer-radius condition

$$\lfloor \delta n \rfloor \le n - k - 1.$$

Consequently, the no-loss CA lower bound proved below should be read only in this admissible range. The Proximity-Prize MCA cap is unaffected: we prove the lower bound at the endpoint

$$\delta_N := 1 - \rho - \frac{2}{N}$$

and then use support-wise MCA monotonicity to extend the MCA failure to all larger sub-capacity radii.

Theorem 2.16 remains a separate slacked fallback imported through the BCHKS/ABF route. It is not needed for the main field-size-universal MCA cap.

3 Locator fibers and interleaving transfer

Fix $N \mid n$, set $a := n/N$, and let $Q := D^a := \{x^a : x \in D\} \subseteq \mathbb{B}^\times$, a multiplicative coset of order N. The a-th power map $D \to Q$ is surjective and its fibers $S_b := \{x \in D : x^a = b\}$ have exactly a elements each. Since $a \mid n \mid |\mathbb{B}| - 1$, the characteristic of $\mathbb{F}$ does not divide a, so $X^a - b$ is separable for $b \neq 0$ and its full root set inside D is S_b whenever $b \in Q$. For a set A, $e_j(A)$ denotes the j-th elementary symmetric function of its elements.

Lemma 3.1 (locator fibers are lists)**.** *Let $\mathbb{B} \subseteq \mathbb{F}$, let $D \subseteq \mathbb{B}^\times$ be a multiplicative coset of order n, and let k satisfy $a \mid k$ and $\rho = k/n$. Write $\ell_1 := \rho N + 1$ and $\ell_2 := \rho N + 2$.*

(i) *If $\ell_1 \le N$, there exists $z \in \mathbb{B}$ such that, with $u_z := \big(x^{k+a} + z\,x^k\big)_{x \in D} \in \mathbb{F}^n$,*

$$\big|\Lambda\big(\mathrm{RS}[\mathbb{F}, D, k],\ 1 - \rho - \tfrac{1}{N},\ u_z\big)\big| \ \ge\ \binom{N}{\ell_1} \Big/ |\mathbb{B}|.$$

(ii) *If $\ell_2 \le N$, there exists $z \in \mathbb{B}$ such that, with $u_z := \big(x^{k+2a} + z\,x^{k+a}\big)_{x \in D} \in \mathbb{F}^n$,*

$$\big|\Lambda\big(\mathrm{RS}[\mathbb{F}, D, k+1],\ 1 - \rho - \tfrac{2}{N},\ u_z\big)\big| \ \ge\ \binom{N}{\ell_2} \Big/ |\mathbb{B}|.$$

In both cases the words u_z are $\mathbb{B}$-valued and all listed codewords lie in $\mathrm{RS}[\mathbb{B}, D, k]$ (*resp.* $\mathrm{RS}[\mathbb{B}, D, k+1]$).

Proof. We prove (ii); part (i) is the identical computation one degree rung lower and is [Cho26a, Main Thm. (d)] sharpened by pigeonholing over $\mathbb{B}$ in place of $\mathbb{F}$.

Let $A \subseteq Q$ with $|A| = \ell_2$, and put

$$L_A(X) \ :=\ \prod_{b \in A}\big(X^a - b\big) \ =\ \sum_{j=0}^{\ell_2} (-1)^j e_j(A)\, X^{a(\ell_2 - j)} \ =\ X^{k+2a} - e_1(A)\, X^{k+a} + R_A(X),$$

where $a\ell_2 = k + 2a$, $a(\ell_2 - 1) = k + a$, and R_A collects the terms with $j \ge 2$, so $\deg R_A \le a(\ell_2 - 2) = k$. The polynomial L_A is squarefree with root set $S_A := \bigcup_{b \in A} S_b \subseteq D$ of size $a\ell_2 = k + 2a$.

Set $z_A := -e_1(A) \in \mathbb{B}$ and $c_A := \big(-R_A(x)\big)_{x \in D}$. Since $\deg R_A \le k$, we have $c_A \in \mathrm{RS}[\mathbb{B}, D, k+1] \subseteq \mathrm{RS}[\mathbb{F}, D, k+1]$. For every $x \in S_A$,

$$0 = L_A(x) = x^{k+2a} + z_A\, x^{k+a} + R_A(x), \qquad \text{i.e.} \qquad u_{z_A}(x) = c_A(x),$$

so u_{z_A} and c_A agree on at least $k + 2a$ points of D and $\Delta(u_{z_A}, c_A) \le 1 - \rho - 2/N$.

The map $A \mapsto c_A$ is injective on the ℓ_2-subsets of Q sharing a common slope value: if $z_A = z_{A'}$ and $c_A = c_{A'}$, then R_A and $R_{A'}$ are polynomials of degree $\le k < n$ agreeing on all n points of D, hence equal as polynomials; then $L_A = L_{A'}$, so $S_A = S_{A'}$, and $A = \{x^a : x \in S_A\} = A'$.

Finally, $z_A = -e_1(A)$ takes at most $|\mathbb{B}|$ values as A ranges over the $\binom{N}{\ell_2}$ subsets; some $z \in \mathbb{B}$ is attained by at least $\binom{N}{\ell_2}/|\mathbb{B}|$ of them, and the corresponding codewords c_A are pairwise distinct elements of $\Lambda(\mathrm{RS}[\mathbb{F}, D, k+1],\ 1 - \rho - 2/N,\ u_z)$. □

Remark 3.2 (why slack two). Part (i) applied to $C^+ = \mathrm{RS}[\mathbb{F}, D, k+1]$ directly would require $a \mid k+1$; in every challenge instantiation a is a power of two and $k+1$ is odd, so this fails. At slack two the listed codewords have degree $\le k$, i.e. lie in C^+, while the divisibility demand stays on k. This is exactly the interface needed by Theorem 2.6, whose list conclusion concerns C^+ while its hypothesis concerns C.

Remark 3.3 (the subfield pigeonhole elsewhere). The same one-line strengthening — pigeonholing locator data over the field of definition $\mathbb{B}$ rather than over $\mathbb{F}$ — also upgrades the coefficient-pigeonhole bound of [Cho26b] to the generated field: the prefix map Φ_σ there has image in $\mathbb{B}^\sigma$ whenever $D \subseteq \mathbb{B}$, since the elementary symmetric prefixes of subsets of D are $\mathbb{B}$-valued. This is the form consumed by the "necessary" half of the list-profile conjecture of [Cho26b], which states the entropy floor at the generated field $\tau^*(\rho, q_D)$; Lemma 3.1 supplies the missing subfield step.

Remark 3.4 (lists, not lines). The words u_z of Lemma 3.1 lie on the $\mathbb{B}$-rational line $f + zg$ with $f = (x^{k+ta})_{x\in D}$, $g = (x^{k+(t-1)a})_{x\in D}$, but the lemma is purely a statement about lists at a received word and no line structure is used in Section 4. This matters: by Corollary 6.2 below, over a proper extension $\mathbb{F} \supsetneq \mathbb{B}$ the lines that end up certifying the CA failure cannot be these (or any) $\mathbb{B}$-rational lines.

3.1 Quotient-profile floors with remainder supports

The fiber lemma above uses supports that are unions of complete quotient fibers. For threshold work one also needs agreements that are not multiples of the fiber size. The following locator-prefix pigeonhole gives a finite lower ledger for such remainder supports. It is a lower-floor theorem only; the matching upper ledger over all divisor scales is not settled here; it is the upper-ledger side of the residual band problem.

Let $c \mid n$ and put $Q_c := D^c$, so the map

$$\pi_c : D \to Q_c, \qquad x \mapsto x^c$$

has fibers of size c and $|Q_c| = N = n/c$. For $0 \le s < c$ and an integer $\sigma \ge 0$, define the prefix weight

$$w_c(s,\sigma) := \#\{h : 1 \le h \le \sigma,\ h \bmod c \in \{0,1,\dots,s\}\},$$

where residues are taken in $\{0,1,\dots,c-1\}$. Thus $w_c(0,\sigma) = \lfloor \sigma/c \rfloor$ for full-fiber supports.

Lemma 3.5 (quotient-remainder locator-prefix lower floor). *Let $\mathbb{B} \subseteq \mathbb{F}$, let $D \subseteq \mathbb{B}^\times$ be a multiplicative coset of order n, and let $K < n$. Fix $c \mid n$, write $N = n/c$, and let*

$$A_0 = mc + s, \qquad 0 \le s < c, \qquad 0 \le m \le N, \qquad A_0 \ge K.$$

If $s > 0$, assume $mc + s \le n$. Put $\sigma = A_0 - K$ and

$$M_{c,m,s} := \binom{N}{m}\binom{n-mc}{s}.$$

Then there is a $\mathbb{B}$-valued received word $U : D \to \mathbb{B}$ such that

$$\left|\Lambda(\mathrm{RS}[\mathbb{F}, D, K],\ 1 - A_0/n,\ U)\right| \ge \left\lceil \frac{M_{c,m,s}}{|\mathbb{B}|^{w_c(s,\sigma)}} \right\rceil.$$

The listed codewords may be chosen in $\mathrm{RS}[\mathbb{B}, D, K]$.

Proof. For $M \subseteq Q_c$ with $|M| = m$, write

$$P_M(X) = \prod_{b \in M} (X^c - b).$$

For $R \subseteq D \setminus \pi_c^{-1}(M)$ with $|R| = s$, put

$$E_R(X) = \prod_{r \in R} (X - r), \qquad L_{M,R}(X) = P_M(X) E_R(X).$$

Then $L_{M,R}$ is a squarefree locator of degree A_0 whose roots are exactly $\pi_c^{-1}(M) \cup R$. Write

$$L_{M,R}(X) = X^{A_0} + \sum_{h=1}^{A_0} \lambda_h(M,R) X^{A_0-h}.$$

The coefficients lie in $\mathbb{B}$. Moreover, because P_M has monomials only in degrees congruent to 0 modulo c and E_R has degree s, the coefficient $\lambda_h(M,R)$ is zero unless $h \bmod c \in \{0,1,\dots,s\}$. Consequently the high-degree prefix

$$\big(\lambda_h(M,R) : 1 \le h \le \sigma,\ h \bmod c \in \{0,1,\dots,s\}\big)$$

takes at most $|\mathbb{B}|^{w_c(s,\sigma)}$ values.

Pigeonhole over these prefix values. There is a fixed prefix vector ξ and a subfamily $\mathcal{T}_\xi$ of at least $M_{c,m,s}/|\mathbb{B}|^{w_c(s,\sigma)}$ pairs (M,R) with this prefix. Let

$$U_\xi(X) = X^{A_0} + \sum_{1 \le h \le \sigma} \lambda_h(\xi) X^{A_0-h}$$

be the corresponding high-degree part, restricted to D. For every $(M,R) \in \mathcal{T}_\xi$, the difference

$$R_{M,R}(X) := L_{M,R}(X) - U_\xi(X)$$

has degree $< K$. Hence the codeword $-R_{M,R}$ agrees with U_ξ on all roots of $L_{M,R}$, a set of size A_0.

It remains only to check that the listed codewords are distinct. If two pairs (M,R) and (M',R') in the same prefix class gave the same codeword on D, then the two remainders of degree $< K < n$ would be equal as polynomials. Since the high prefixes are also equal, $L_{M,R} = L_{M',R'}$. Their root sets in D are therefore equal. A full fiber has size c while $|R| = s < c$, so the full fibers contained in this root set are exactly M, and then the residual set is forced as well. Thus $(M,R) = (M',R')$. □

Definition 3.6 (prefix fibers and prefix-image certificates)**.** In the notation of Lemma 3.5, let $\Phi^K_{c,m,s}$ be the map from pairs (M,R) to the high-degree prefix

$$\big(\lambda_h(M,R) : 1 \le h \le A_0 - K,\ h \bmod c \in \{0,1,\dots,s\}\big) \in \mathbb{B}^{w_c(s,A_0-K)}.$$

If $A_0 = K$ this is the unique empty prefix. For a prefix value ξ, put

$$\mu^K_{c,m,s}(\xi) := \big|(\Phi^K_{c,m,s})^{-1}(\xi)\big|.$$

Define the image size and the heaviest prefix fiber by

$$I^K_{c,m,s} := \left|\operatorname{im} \Phi^K_{c,m,s}\right|, \qquad H^K_{c,m,s} := \max_{\xi} \mu^K_{c,m,s}(\xi).$$

Thus

$$H^K_{c,m,s} \geq \left\lceil \frac{M_{c,m,s}}{I^K_{c,m,s}} \right\rceil \geq \left\lceil \frac{M_{c,m,s}}{|\mathbb{B}|^{w_c(s,A_0-K)}} \right\rceil.$$

A prefix-fiber certificate is an orbit decomposition, hash table, or other verifiable bucket table for the finite map $\Phi^K_{c,m,s}$, including the image size and the maximum fiber size. If only $I^K_{c,m,s}$ is certified, the middle lower bound is still valid; if no image certificate is supplied, the ambient-box lower bound is valid.

Lemma 3.7 (heaviest-prefix locator floor)**.** *Under the hypotheses of Lemma 3.5, there is a $\mathbb{B}$-valued received word $U : D \to \mathbb{B}$ such that*

$$\left|\Lambda(\mathrm{RS}[\mathbb{F}, D, K],\, 1 - A_0/n,\, U)\right| \geq H^K_{c,m,s}.$$

Consequently

$$\left|\Lambda(\mathrm{RS}[\mathbb{F}, D, K],\, 1 - A_0/n,\, U)\right| \geq \left\lceil \frac{M_{c,m,s}}{I^K_{c,m,s}} \right\rceil \geq \left\lceil \frac{M_{c,m,s}}{|\mathbb{B}|^{w_c(s,A_0-K)}} \right\rceil$$

when only the coarser certificates are available. The listed codewords may be chosen in $\mathrm{RS}[\mathbb{B}, D, K]$.

Proof. Choose a prefix value ξ with $\mu^K_{c,m,s}(\xi) = H^K_{c,m,s}$ and form the high-degree word

$$U_\xi(X) = X^{A_0} + \sum_{1 \leq h \leq A_0 - K} \lambda_h(\xi) X^{A_0 - h},$$

where $\lambda_h(\xi) = 0$ for the high coefficients outside the allowed residue classes. For every pair (M, R) in this heaviest prefix fiber, the difference $L_{M,R}(X) - U_\xi(X)$ has degree $< K$, so the negative of this remainder is a codeword of $\mathrm{RS}[\mathbb{B}, D, K]$ agreeing with U_ξ on the A_0 roots of $L_{M,R}$. If two pairs in the same prefix fiber gave the same codeword on D, then their degree-$< K < n$ remainders would be equal as polynomials; together with the common prefix this gives $L_{M,R} = L_{M',R'}$. Their root sets in D are equal. A full quotient fiber has size c while the residual set has size $s < c$, so the full fibers contained in the root set determine M, and then the residual set determines R. Thus the codewords are pairwise distinct. The two coarser displayed bounds follow from Definition 3.6. □

For $A \in \{k+1, \ldots, n\}$ and $c \mid n$, write

$$A = m_c(A)c + s_c(A), \qquad 0 \leq s_c(A) < c,$$

and set

$$\begin{aligned}
M_c(A) &:= \binom{n/c}{m_c(A)} \binom{n - cm_c(A)}{s_c(A)}, \\
L^+_{c,\mathrm{heavy}}(A) &:= H^{k+1}_{c,m_c(A),s_c(A)}, \\
L^+_{c,\mathrm{image}}(A) &:= \left\lceil \frac{M_c(A)}{I^{k+1}_{c,m_c(A),s_c(A)}} \right\rceil, \\
L^+_{c,\mathrm{box}}(A) &:= \left\lceil \frac{M_c(A)}{|\mathbb{B}|^{w_c(s_c(A),A-k-1)}} \right\rceil.
\end{aligned}$$

The three levels satisfy

$$L^+_{c,\mathrm{heavy}}(A) \ge L^+_{c,\mathrm{image}}(A) \ge L^+_{c,\mathrm{box}}(A).$$

Here "heavy" uses a full prefix-fiber certificate, "image" uses only the prefix-image size, and "box" is the certificate-free bound already implicit in Lemma 3.5.

Theorem 3.8 (quotient-remainder deep-point lower ledger)**.** *Let $\mathbb{B} \subseteq \mathbb{F}$, let $q = |\mathbb{F}| > n$, let $D \subseteq \mathbb{B}^\times$ be a multiplicative coset of order n, and let $C = \mathrm{RS}[\mathbb{F}, D, k]$. Fix an agreement value $A \in \{k+1, \dots, n\}$ and a divisor $c \mid n$. For each*

$$\star \in \{\mathrm{heavy}, \mathrm{image}, \mathrm{box}\}$$

and every

$$\delta \in [1 - A/n,\ 1 - k/n)$$

one has

$$\varepsilon_{\mathrm{ca}}(C, \delta) \ge \mathcal{E}_{q,k}\big(L^+_{c,\star}(A)\big).$$

If $\delta > 0$, then also

$$\varepsilon_{\mathrm{mca}}(C, \delta) \ge \mathcal{E}_{q,k}\big(L^+_{c,\star}(A)\big).$$

Here $\mathcal{E}_{q,k}$ is as in Corollary 2.9. More precisely, some received line has at least

$$\mathcal{N}_{q,k}\big(L^+_{c,\star}(A)\big)$$

CA-bad finite slopes at that radius, and for $\delta > 0$ these slopes are MCA-bad.

Consequently, for a closed agreement threshold $a \ge k+1$, define

$$B^-_{Q,\star}(a) := \max_{\substack{c \mid n \\ A \in \{a, a+1, \dots, n\}}} \mathcal{N}_{q,k}\big(L^+_{c,\star}(A)\big).$$

Then $B^-_{Q,\star}(a)$ is a lower numerator for $q\,\varepsilon_{\mathrm{ca}}(C, \delta)$ for every $\delta \in [1 - a/n, 1 - k/n)$, and for $q\,\varepsilon_{\mathrm{mca}}(C, \delta)$ on the same interval with $\delta > 0$. The strongest certified ledger uses $\star = \mathrm{heavy}$; the certificate-free ledger uses $\star = \mathrm{box}$.

Proof. Apply Lemma 3.7 with $K = k+1$ and $A_0 = A$. This produces a $C^+ = \mathrm{RS}[\mathbb{F}, D, k+1]$ list of size at least $L^+_{c,\mathrm{heavy}}(A)$ at radius $1 - A/n$, with the image and ambient-box variants giving the corresponding smaller certified list sizes. Then apply Proposition 2.8. If $A \ge a$ and $\delta \ge 1 - a/n$, then $\delta \ge 1 - A/n$, so the same converted floor applies. The maximum over (c, A) is the correct lower-ledger operation: the definitions of $\varepsilon_{\mathrm{ca}}$ and $\varepsilon_{\mathrm{mca}}$ maximize over received lines, so lower witnesses from different divisor scales or exact agreement values need not share one line and must not be summed. □

Corollary 3.9 (quotient-remainder trigger)**.** *In the setting of Theorem 3.8, if $L^+_{c,\star}(A) > (q - n)/k$ for any $\star \in \{\mathrm{heavy}, \mathrm{image}, \mathrm{box}\}$, then, for every $\delta \in [1 - A/n, 1 - k/n)$,*

$$\varepsilon_{\mathrm{ca}}(C, \delta) > \frac{1}{2k}\left(1 - \frac{n}{q}\right),$$

and for every such $\delta > 0$,

$$\varepsilon_{\mathrm{mca}}(C, \delta) > \frac{1}{2k}\left(1 - \frac{n}{q}\right).$$

Proof. This is Corollary 2.9 applied inside Theorem 3.8. □

Corollary 3.10 (quantitative first-grid floor)**.** *Let $C = \mathrm{RS}[\mathbb{F}, D, k]$, let $q = |\mathbb{F}| > n$, and assume only that $D \subseteq \mathbb{F}$ has n distinct points. Then for every $\delta \in [1-(k+1)/n, 1-k/n)$,*

$$\varepsilon_{\mathrm{ca}}(C, \delta) \geq \mathcal{E}_{q,k}\left(\binom{n}{k+1}\right),$$

and for every such $\delta > 0$,

$$\varepsilon_{\mathrm{mca}}(C, \delta) \geq \mathcal{E}_{q,k}\left(\binom{n}{k+1}\right).$$

In particular, if $\binom{n}{k+1} > (q-n)/k$, then the printed $(2k)^{-1}(1-n/q)$ lower bound holds on this entire first grid interval without any multiplicative smoothness assumption.

Proof. Use the ordinary locator identity for all $(k+1)$-subsets of D; this is the $c = 1$ argument behind Lemma 3.7, but it does not use multiplicative smoothness. The prefix is empty, so the list size is at least $\binom{n}{k+1}$. Apply Proposition 2.8 and then Corollary 2.9. □

Definition 3.11 (quotient lower-ledger certificate)**.** A quotient-remainder lower-ledger certificate for a row and a threshold a contains, for each divisor scale c and each exact agreement $A \geq a$ that is used in the maximum:

(i) the decomposition $A = m_c(A)c + s_c(A)$;

(ii) the support count $M_c(A)$ and prefix weight $w_c(s_c(A), A-k-1)$;

(iii) either a full prefix-fiber bucket table certifying $H^{k+1}_{c,m_c(A),s_c(A)}$, or an image-size certificate for $I^{k+1}_{c,m_c(A),s_c(A)}$, or the declaration that the ambient-box denominator is being used;

(iv) the resulting list floor $L^{+}_{c,\star}(A)$;

(v) the converted slope numerator $\mathcal{N}_{q,k}(L^{+}_{c,\star}(A))$ and error floor $\mathcal{E}_{q,k}(L^{+}_{c,\star}(A))$.

The verifier checks the locator-prefix formula, distinctness of codewords inside a prefix fiber, the integer arithmetic, and the final maximum over divisor scales and exact agreement values.

Remark 3.12 (quotient lower ledger versus quotient upper ledger)*.* The quotient-profile lower side is now scanner-ready at every agreement value $A = mc+s$: the generated-field contribution is the actual heaviest prefix fiber $H^{k+1}_{c,m,s}$ when certified, the exact prefix-image denominator $I^{k+1}_{c,m,s}$ when only image data are certified, and the ambient box $|\mathbb{B}|^{w_c(s,A-k-1)}$ when no certificate is supplied. The challenge-field size q enters only through the independent conversion functions $\mathcal{E}_{q,k}$ and $\mathcal{N}_{q,k}$. Across divisor scales the lower-side operation is a maximum, not a union count, because the witness line may change. The remaining upper-ledger problem is sharper and unchanged in logical status: replace the safe support ledger and the exact declared-branch image ledger by an exhaustive quotient-periodic classification, with duplicate slopes coalesced across all divisor scales.

Corollary 3.13 (augmented slack-one quotient floor)**.** *Let $\mathbb{B} \subseteq \mathbb{F}$ and let $D \subseteq \mathbb{B}^{\times}$ be a multiplicative coset of order n. Put $C = \mathrm{RS}[\mathbb{F}, D, k]$ and $C^{+} = \mathrm{RS}[\mathbb{F}, D, k+1]$. If $c \mid \gcd(n, k)$, then there is a $\mathbb{B}$-valued received word U_c such that*

$$\left|\Lambda(C^{+},\, 1-(k+c)/n,\, U_c)\right| \geq \binom{n/c}{k/c+1}.$$

One may take $U_c(x) = x^{k+c}$. In particular, for $c = 1$ this gives the universal first-grid floor

$$\left|\Lambda(C^+,\, 1-(k+1)/n,\, x^{k+1})\right| \geq \binom{n}{k+1},$$

with no quotient denominator. The $c = 1$ statement in fact holds for any set $D \subseteq \mathbb{B}$ of n distinct points, not only for multiplicative cosets.

Proof. For $c > 1$, apply Lemma 3.5 with $K = k+1$, $s = 0$, and $m = k/c+1$. Then $A_0 = k+c$ and $\sigma = A_0 - K = c-1$, so $w_c(0, c-1) = 0$. The high-degree prefix is only the monomial X^{k+c}, giving the displayed word. For $c = 1$, the same argument is just the ordinary locator identity for all $(k+1)$-subsets of D: the high-degree prefix of a monic locator of degree $k+1$ is X^{k+1}, and its remainder has degree at most k. □

Corollary 3.14 (first-grid deep-point cap)**.** *Let $C = \mathrm{RS}[\mathbb{F}, D, k]$, $q = |\mathbb{F}|$, and suppose $q > n$. Let $c \mid \gcd(n, k)$. If $c > 1$, assume $D \subseteq \mathbb{B}^\times$ is a multiplicative coset as above; for $c = 1$, no smoothness is needed. If*

$$\binom{n/c}{k/c+1} \geq q/k+1,$$

then

$$\varepsilon_{\mathrm{ca}}\big(C, 1-(k+c)/n\big) > \frac{1}{2k}\left(1-\frac{n}{q}\right), \qquad \varepsilon_{\mathrm{mca}}\big(C, 1-(k+c)/n\big) > \frac{1}{2k}\left(1-\frac{n}{q}\right).$$

For $c = 1$ this is the first closed grid point below capacity.

Proof. Corollary 3.13 gives a C^+-list of size $L \geq q/k+1$ at agreement $k+c$. The proof of the deep-point conversion Theorem 2.6 applies whenever the agreement is $> k$, so it applies here for every $c \geq 1$. If the displayed CA lower bound failed, then

$$\varepsilon_{\mathrm{ca}}(C, 1-(k+c)/n) \leq \frac{q-n}{2kq}.$$

Theorem 2.6 with $\eta = 1/2$ would give

$$\Lambda(C^+, 1-(k+c)/n) \leq \left\lceil \frac{q-n}{k} \right\rceil < q/k+1,$$

contradicting the list lower bound. Every no-loss CA-bad slope is also support-wise MCA-bad on the same explaining support, giving the MCA inequality. □

Corollary 3.15 (first-grid cap for large challenge-envelope rows)**.** *Let $\rho \in \{\frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \frac{1}{16}\}$, let $C = \mathrm{RS}[\mathbb{F}, D, k]$ with $|D| = n$, $k = \rho n$, $q = |\mathbb{F}| < 2^{256}$, and $q > n$. Assume*

ρ	$\frac{1}{2}$	$\frac{1}{4}$	$\frac{1}{8}$	$\frac{1}{16}$
$k \geq$	127	78	58	47

for the corresponding rate. Then the first closed grid point below capacity is already CA- and MCA-unsafe:

$$\varepsilon_{\mathrm{ca}}\big(C, 1-\rho-1/n\big) > \frac{1}{2k}\left(1-\frac{n}{q}\right), \qquad \varepsilon_{\mathrm{mca}}\big(C, 1-\rho-1/n\big) > \frac{1}{2k}\left(1-\frac{n}{q}\right).$$

If also $k \leq 2^{40}$, then this lower bound is $> 2^{-86}$ and hence is far above the prize target 2^{-128}.

Proof. By Corollary 3.14 with $c = 1$, it is enough to check

$$\binom{n}{k+1} \geq q/k + 1.$$

Since $q < 2^{256}$, the stronger sufficient condition is $k(\binom{n}{k+1} - 1) \geq 2^{256}$. The four printed lower bounds on k are exactly small finite checks of this inequality at the corresponding rate; the left side is increasing thereafter. For the error size, $q \geq n + 1$, $n = k/\rho \leq 16k$, and $k \leq 2^{40}$ give

$$\frac{1}{2k}\left(1 - \frac{n}{q}\right) \geq \frac{1}{2k(n+1)} > 2^{-86}.$$

□

3.2 A conservative quotient-support upper ledger

The remainder lower floor above supplies the unsafe side of the quotient-profile ledger. For a prize threshold one also needs a safe-side accounting rule: if a later argument declares a bad parameter to be quotient-periodic, the declaration must come with a divisor-lattice numerator that is counted with overlaps under control. The next statements give such a conservative upper ledger. They do not prove that every bad line parameter is quotient-periodic; they only say that once a permitted support family has been specified, every ledger term witnessed inside that family is bounded by the size of the family, with a factor d for finite-parameter degree-d power curves.

Let $\mathcal{S}$ be a family of subsets of D, all of size at least a. For a received line $f + zg$, say that a finite slope is *$\mathcal{S}$-witnessed support-wise bad* if it has a support-wise noncontained explanation on some $S \in \mathcal{S}$. Define the analogous $\mathcal{S}$-witnessed no-loss CA slopes by requiring that $f + zg$ be explained on some $S \in \mathcal{S}$ while (f, g) is not globally explained on any support of size at least a.

Lemma 3.16 (one support pays for at most one line parameter)**.** *Let $C = \mathrm{RS}[\mathbb{F}, D, k]$ and let $a \geq k$. Fix a support $S \subseteq D$ with $|S| \geq a$. For any received line $f + zg$, at most one finite slope $z \in \mathbb{F}$ can be support-wise noncontained on the fixed support S. Likewise, at most one finite slope can be no-loss CA-bad while being explained on this fixed support S. Consequently, for every support family $\mathcal{S}$,*

$$\#\{\mathcal{S}\text{-witnessed support-wise MCA-bad finite slopes}\} \leq |\mathcal{S}|,$$

and the same bound holds for no-loss CA slopes witnessed by $\mathcal{S}$.

Proof. Suppose two distinct finite slopes $z_1 \neq z_2$ are both explained on the same support S, by codewords $p_1, p_2 \in C$:

$$f + z_i g = p_i \qquad \text{on } S, \quad i = 1, 2.$$

Then on S,

$$g = \frac{p_1 - p_2}{z_1 - z_2}, \qquad f = p_1 - z_1 g.$$

Both right-hand sides are codewords of C. Thus (f, g) is simultaneously explained on S. This contradicts support-wise noncontainment on S. It also contradicts no-loss CA badness at agreement a, because $|S| \geq a$ gives a global simultaneous explanation support of the forbidden size. Hence each fixed support contributes at most one bad finite line parameter, and summing over supports gives the displayed bound. □

For a finite-parameter power curve

$$W_\gamma = f_0 + \gamma f_1 + \cdots + \gamma^d f_d, \qquad \gamma \in \mathbb{F},$$

say that a parameter is support-wise curve-bad on S if W_γ is explained by a codeword on S, but the coefficient tuple $(f_0, \ldots, f_d)$ is not simultaneously explained by $C^{\equiv(d+1)}$ on S. The no-loss curve-CA variant replaces this support-wise noncontainment by the corresponding global far condition at agreement a.

Lemma 3.17 (one support pays for at most d curve parameters)**.** *Let $C = \mathrm{RS}[\mathbb{F}, D, k]$, let $a \geq k$, and fix a support $S \subseteq D$ with $|S| \geq a$. For a degree-d power curve W_γ, at most d finite parameters $\gamma \in \mathbb{F}$ can be support-wise curve-bad on S. The same bound holds for no-loss curve-CA parameters explained on S. Consequently any support family $\mathcal{S}$ witnesses at most $d|\mathcal{S}|$ bad finite curve parameters.*

Proof. If $d+1$ distinct parameters $\gamma_0, \ldots, \gamma_d$ were explained on the same support S, with codewords $P_j \in C$, then Lagrange interpolation in the parameter variable gives codewords $Q_0, \ldots, Q_d \in C$ such that

$$\sum_{i=0}^{d} \gamma_j^i Q_i = P_j \qquad (0 \leq j \leq d).$$

On every point of S, the same interpolation applied to the values of W_{γ_j} gives $f_i = Q_i$ for every coefficient i. Hence $(f_0, \ldots, f_d)$ is simultaneously explained by $C^{\equiv(d+1)}$ on S. This contradicts support-wise curve-noncontainment on S, and also contradicts no-loss curve-CA badness at agreement a. Thus every support contributes at most d parameters. □

Lemma 3.18 (one support pays for at most one interleaved tuple)**.** *Let $C = \mathrm{RS}[\mathbb{F}, D, k]$, let $a \geq k$, and let $\mu \geq 1$. Fix a μ-row received word $U = (U_1, \ldots, U_\mu)$. For a fixed support $S \subseteq D$ with $|S| \geq a$, there is at most one tuple $(P_1, \ldots, P_\mu) \in C^\mu$ agreeing with U on every row over S. Hence the number of interleaved list tuples whose common agreement support belongs to a support family $\mathcal{S}$ is at most $|\mathcal{S}|$.*

Proof. For each row i, two degree-$< k$ polynomials agreeing with U_i on the fixed support S agree with each other on at least k points, so they are equal. Thus each row has at most one codeword compatible with S, and therefore the whole interleaved tuple is unique if it exists. □

We now instantiate the support-family ledger on quotient-remainder supports. For $c \mid n$ and an integer A with $0 \leq A \leq n$, write

$$A = m_c(A)c + s_c(A), \qquad 0 \leq s_c(A) < c,$$

and define

$$\mathcal{S}_c(A) := \{\pi_c^{-1}(M) \cup R : M \subseteq Q_c, |M| = m_c(A), \\ R \subseteq D \setminus \pi_c^{-1}(M), |R| = s_c(A)\}.$$

For a divisor set $\mathcal{C} \subseteq \{c : c \mid n\}$ put

$$\mathcal{S}_{\mathcal{C}}^{\geq a} := \bigcup_{c \in \mathcal{C}} \bigcup_{A=a}^{n} \mathcal{S}_c(A)$$

and

$$U_{\mathcal{C}}^{\mathrm{supp}}(a) := \left|\mathcal{S}_{\mathcal{C}}^{\geq a}\right|.$$

When the exact union is inconvenient, use the safe sum

$$U_{\mathcal{C}}^{\mathrm{sum}}(a) := \sum_{c\in\mathcal{C}} \sum_{A=a}^{n} \binom{n/c}{\lfloor A/c\rfloor} \binom{n - c\lfloor A/c\rfloor}{A - c\lfloor A/c\rfloor},$$

so that $U_{\mathcal{C}}^{\mathrm{supp}}(a) \leq U_{\mathcal{C}}^{\mathrm{sum}}(a)$.

Proposition 3.19 (divisor-union quotient support ledger)**.** *Let $C = \mathrm{RS}[\mathbb{F}, D, k]$ and let $a \geq k$. Fix a divisor set $\mathcal{C}$ and let $\mathcal{S}_{\mathcal{C}}^{\geq a}$ be the quotient-remainder support family above. Then, for every received line, the number of finite support-wise MCA-bad slopes whose witness support lies in $\mathcal{S}_{\mathcal{C}}^{\geq a}$ is at most*

$$U_{\mathcal{C}}^{\mathrm{supp}}(a) \leq U_{\mathcal{C}}^{\mathrm{sum}}(a).$$

The same bound holds for no-loss CA slopes witnessed by this family. For a finite-parameter degree-d power curve, the corresponding support-wise curve-MCA and no-loss curve-CA numerator is at most

$$d\, U_{\mathcal{C}}^{\mathrm{supp}}(a) \leq d\, U_{\mathcal{C}}^{\mathrm{sum}}(a).$$

For every interleaving arity $\mu \geq 1$, the number of interleaved list tuples whose common agreement support belongs to $\mathcal{S}_{\mathcal{C}}^{\geq a}$ is at most $U_{\mathcal{C}}^{\mathrm{supp}}(a)$.

Proof. Apply Lemmas 3.16, 3.17, and 3.18 to the support family $\mathcal{S}_{\mathcal{C}}^{\geq a}$. The inequality $U_{\mathcal{C}}^{\mathrm{supp}}(a) \leq U_{\mathcal{C}}^{\mathrm{sum}}(a)$ is only the union bound over the divisor and agreement-size indices. □

Remark 3.20 (what this does and does not finish). Proposition 3.19 supplies a certified safe-side quotient ledger at the support-family level, and it is already enough for a staircase scanner to avoid double-counting divisor scales: one may compute the exact union $U_{\mathcal{C}}^{\mathrm{supp}}$ or use the printed safe sum. This is not yet the sharp quotient upper theorem needed for a full prize solution. Near capacity the support count can be exponentially larger than the number of distinct bad slopes, so the next refinement is a slope-collision quotient ledger: classify when many quotient-remainder supports map to the same line or curve parameter, and replace the support count by the corresponding distinct-parameter union count.

Definition 3.21 (block support-profile polynomial)**.** Let D be a multiplicative coset of order n, let $\mathcal{C}$ be a nonempty finite set of divisors of n, and put

$$h := \operatorname{lcm} \mathcal{C}, \qquad N_h := n/h.$$

For each $c \in \mathcal{C}$ let $H_c \leq H_h$ denote the subgroup of the cyclic kernel of order h with $|H_c| = c$. If $T \subseteq H_h$, define

$$\phi_c(T) := \#\{\text{cosets } uH_c \subseteq H_h : \ uH_c \subseteq T\},$$

the number of complete c-subfibers contained in the block subset T. The block support-profile polynomial of $\mathcal{C}$ is

$$G_{\mathcal{C},h}(Y, (Z_c)_{c\in\mathcal{C}}) := \sum_{T\subseteq H_h} Y^{|T|} \prod_{c\in\mathcal{C}} Z_c^{\phi_c(T)}.$$

For a vector $\nu = (\nu_c)_{c\in\mathcal{C}}$ write $Z^\nu := \prod_{c\in\mathcal{C}} Z_c^{\nu_c}$, and define $N_{\mathcal{C},h}(A;\nu)$ by

$$G_{\mathcal{C},h}(Y, Z)^{N_h} = \sum_{A=0}^{n} \sum_{\nu} N_{\mathcal{C},h}(A;\nu) Y^A Z^\nu.$$

Lemma 3.22 (block polynomial counts exact quotient profiles)**.** *For a support $S \subseteq D$ and $c \mid n$, let*

$$F_c(S) := \#\{c\text{-fibers of } D \to D^c \text{ contained in } S\}.$$

Then $N_{\mathcal{C},h}(A;\nu)$ is exactly the number of supports $S \subseteq D$ such that

$$|S| = A, \qquad F_c(S) = \nu_c \quad (c \in \mathcal{C}).$$

Proof. The h-fibers partition D into N_h blocks, and every c-fiber with $c \in \mathcal{C}$ is contained in a unique h-fiber because $c \mid h$. After choosing an identification of each h-fiber with the cyclic kernel H_h, the number of selected points and the number of complete internal c-subfibers are recorded by the monomial

$$Y^{|T|} \prod_{c \in \mathcal{C}} Z_c^{\phi_c(T)}$$

for the block subset T. The choices in distinct h-fibers are independent, and the global statistics $|S|$ and $F_c(S)$ are the sums of the block statistics. Multiplying the block polynomial over the N_h blocks and extracting the coefficient gives the asserted count. □

Lemma 3.23 (membership by full-fiber profile)**.** *Fix A and $c \mid n$, and write*

$$A = m_c(A)c + s_c(A), \qquad 0 \le s_c(A) < c.$$

For a support $S \subseteq D$ of size A, one has

$$S \in \mathcal{S}_c(A) \qquad \Longleftrightarrow \qquad F_c(S) = m_c(A) = \lfloor A/c \rfloor.$$

Proof. If $S \in \mathcal{S}_c(A)$, then by definition it is the union of $m_c(A)$ complete c-fibers and $s_c(A) < c$ residual points outside those fibers; there is no room for any further complete c-fiber. Hence $F_c(S) = m_c(A)$. Conversely, if $F_c(S) = m_c(A)$, take the complete c-fibers contained in S as the set M in the definition of $\mathcal{S}_c(A)$; the remaining $A - m_c(A)c = s_c(A)$ selected points form the residual set. □

Theorem 3.24 (exact divisor-lattice support-union formula)**.** *Let $\mathcal{C}$ be a nonempty finite set of divisors of n, and define $G_{\mathcal{C},h}$ and $N_{\mathcal{C},h}(A;\nu)$ as in Definition 3.21. Then the exact support-family union from Proposition 3.19 is*

$$\boxed{U_{\mathcal{C}}^{\mathrm{supp}}(a) = \sum_{A=a}^{n} \sum_{\nu} N_{\mathcal{C},h}(A;\nu)\, \mathbf{1}[\exists c \in \mathcal{C} : \nu_c = \lfloor A/c \rfloor]\,.}$$

Consequently the quotient-support branch numerators in Proposition 3.19 may be computed by this coefficient formula instead of by the safe sum $U_{\mathcal{C}}^{\mathrm{sum}}(a)$.

Proof. For a fixed exact size A, the union $\bigcup_{c \in \mathcal{C}} \mathcal{S}_c(A)$ consists precisely of those supports S of size A for which $S \in \mathcal{S}_c(A)$ for at least one divisor $c \in \mathcal{C}$. By Lemma 3.23, this is exactly the condition

$$F_c(S) = \lfloor A/c \rfloor$$

for at least one $c \in \mathcal{C}$. By Lemma 3.22, the number of supports with a fixed profile vector $\nu = (F_c(S))_{c \in \mathcal{C}}$ is $N_{\mathcal{C},h}(A;\nu)$. Summing over all exact sizes $A \ge a$ and over the profile vectors satisfying the displayed predicate gives the formula. □

Corollary 3.25 (single-divisor sanity check)**.** *For $\mathcal{C} = \{c\}$ the block polynomial is*

$$G_{\{c\},c}(Y, Z) = (1 + Y)^c + (Z - 1)Y^c,$$

and Theorem 3.24 gives

$$|\mathcal{S}_c(A)| = \binom{n/c}{\lfloor A/c \rfloor} \binom{n - c\lfloor A/c \rfloor}{A - c\lfloor A/c \rfloor},$$

the summand used in the safe ledger.

Proof. Inside one c-fiber, every proper subset contributes only $Y^{|T|}$, while the full subset contributes $Y^c Z$. Thus the block polynomial is $(1 + Y)^c - Y^c + Y^c Z$. Raising to n/c and selecting exactly $m = \lfloor A/c \rfloor$ full blocks and $s = A - mc$ residual points in the remaining blocks gives the displayed count. □

Definition 3.26 (divisor-block support-ledger certificate)**.** A divisor-block support-ledger certificate for a divisor set $\mathcal{C}$ and a threshold a consists of:

(i) the lcm block size $h = \operatorname{lcm} \mathcal{C}$ and the block count $N_h = n/h$;

(ii) the internal coset table of every subgroup $H_c \leq H_h$ for $c \in \mathcal{C}$;

(iii) the block polynomial $G_{\mathcal{C},h}$, either as a full coefficient table or as an orbit-compressed table under translations of H_h;

(iv) the coefficient extraction from $G_{\mathcal{C},h}^{N_h}$ for every exact size $A \geq a$;

(v) the predicate table $\mathbf{1}[\exists c \in \mathcal{C} : \nu_c = \lfloor A/c \rfloor]$ and the final integer $U_{\mathcal{C}}^{\mathrm{supp}}(a)$.

The verifier checks the block enumeration, the subgroup containment relation $c \mid h$, coefficient arithmetic, and the final predicate sum. This is an exact support-union certificate; it does not by itself coalesce distinct line or curve parameters that arise from different supports.

Remark 3.27 (support ledger versus slope ledger)*.* Theorem 3.24 closes the divisor-lattice overlap problem at the support-family level: intersections between quotient descriptions at incomparable divisor scales are now counted by one coefficient extraction in the common h-block. This is still weaker than the desired distinct-slope quotient ledger. To finish that stronger ledger one must apply the exact support-image maps of Propositions 3.30 and 3.32 to the support set counted here and then prove or certify the resulting parameter image size after duplicate coalescing.

3.3 Distinct-parameter quotient image ledger

The support-union ledger of Section 3.2 is deliberately safe but often very loose: many quotient-remainder supports can certify the same line slope or the same curve parameter. The next refinement is to count the image of the actual parameter map. This subsection gives an exact image ledger for the quotient branch. It is still a branch ledger, not an aperiodic upper bound: it counts only witnesses whose agreement support lies in a declared quotient-remainder support family.

We first record the support-to-parameter map in syndrome form. Let $C = \mathrm{RS}[\mathbb{F}, D, k]$, $|D| = n$, and put $r = n - k$. For each $x \in D$ set

$$\lambda_x := \left(\prod_{y \in D,\ y \neq x} (x - y) \right)^{-1}.$$

For a word $Y : D \to \mathbb{F}$, define its syndrome vector

$$\mathrm{Syn}(Y)_m := \sum_{x \in D} \lambda_x x^m Y(x), \qquad 0 \le m < r.$$

Thus $\mathrm{Syn}(Y) = 0$ if and only if $Y \in C$. If $T \subseteq D$ has size $j \le r$, write

$$L_T(X) = \prod_{x \in T} (X - x) = \ell_0 + \ell_1 X + \cdots + \ell_j X^j, \qquad \boldsymbol{\ell}_T = (\ell_0, \ldots, \ell_j)^t,$$

and put $t = r - j$. For $w \in \mathbb{F}^r$ let $H_{t,j}(w)$ be the $t \times (j+1)$ Hankel window

$$H_{t,j}(w)_{a,b} = w_{a+b}, \qquad 0 \le a < t, \quad 0 \le b \le j.$$

When $S = D \setminus T$, the integer t is exactly $|S| - k$.

Lemma 3.28 (support-locator syndrome recurrence)**.** *Let $Y : D \to \mathbb{F}$ and let $T \subseteq D$ with $|T| = j \le r$. Put $S = D \setminus T$ and $t = r - j$. Then Y is explained by C on S if and only if*

$$H_{t,j}(\mathrm{Syn}(Y))\boldsymbol{\ell}_T = 0.$$

Proof. Let $W_T \subseteq \mathbb{F}^r$ be the span of the parity-check columns indexed by T,

$$W_T = \mathrm{span}_{\mathbb{F}}\{\lambda_x(1, x, \ldots, x^{r-1}) : x \in T\}.$$

The word Y is explained by C on S if and only if $Y - c$ is supported on T for some $c \in C$, equivalently $\mathrm{Syn}(Y) \in W_T$. If $w \in W_T$, write $w_m = \sum_{x \in T} a_x x^m$. Then, for $0 \le h < t$,

$$\big(H_{t,j}(w)\boldsymbol{\ell}_T\big)_h = \sum_{b=0}^{j} \ell_b w_{h+b} = \sum_{x \in T} a_x x^h L_T(x) = 0.$$

Conversely, because L_T is monic, the displayed recurrence determines all coordinates $w_j, \ldots, w_{r-1}$ from the first j coordinates, so its solution space has dimension at most j. The j parity-check columns indexed by the distinct points of T have Vandermonde rank j, hence W_T has dimension j and is exactly the recurrence space. □

Lemma 3.29 (exact affine and projective support-image map)**.** *Let $f, g : D \to \mathbb{F}$, put $u = \mathrm{Syn}(f)$ and $v = \mathrm{Syn}(g)$, and fix $T \subseteq D$ with $|T| = j \le r$ and $S = D \setminus T$. Define*

$$A_T := H_{t,j}(u)\boldsymbol{\ell}_T, \qquad B_T := H_{t,j}(v)\boldsymbol{\ell}_T.$$

A finite slope $z \in \mathbb{F}$ is explained on the fixed support S and is support-wise noncontained there if and only if

$$A_T + zB_T = 0, \qquad B_T \ne 0.$$

Consequently the fixed support contributes either no finite bad slope or the unique value

$$\theta_T = -A_{T,m}/B_{T,m}$$

for any coordinate m with $B_{T,m} \ne 0$, provided all nonzero coordinates give the same value.

For the projective line $\alpha f + \beta g$, the fixed support contributes a support-wise noncontained point $[\alpha : \beta] \in \mathbb{P}^1(\mathbb{F})$ if and only if A_T and B_T are linearly dependent and not both zero. In that case the contributed projective parameter is the unique point satisfying

$$\alpha A_T + \beta B_T = 0.$$

Proof. A word Y is explained by C on $S = D \setminus T$ if and only if $Y - c$ is supported on T for some $c \in C$, which is equivalent to $\mathrm{Syn}(Y)$ lying in the span of the parity-check columns indexed by T. The standard locator recurrence says that this is equivalent to

$$H_{t,j}(\mathrm{Syn}(Y))\boldsymbol{\ell}_T = 0.$$

Applying this to $Y = f + zg$ gives $A_T + zB_T = 0$. Under this incidence, simultaneous explanation of f and g on S is equivalent to $B_T = 0$, because $A_T = (A_T + zB_T) - zB_T$. This proves the finite-slope statement. The projective statement is the same calculation with $Y = \alpha f + \beta g$; if $A_T = B_T = 0$ then both coefficient words are already explained on S, so the point is contained rather than bad. □

Let $\mathcal{S}_{\mathcal{C}}^{\geq a}$ be the quotient-remainder support family from Section 3.2, and let

$$\mathcal{T}_{\mathcal{C}}^{\geq a} := \{D \setminus S : S \in \mathcal{S}_{\mathcal{C}}^{\geq a}\}$$

be the corresponding co-support family. For a received line define

$$\Theta_{\mathcal{C}}^{\mathrm{aff}}(f, g; a) := \{\theta_T : T \in \mathcal{T}_{\mathcal{C}}^{\geq a}, \theta_T \text{ is defined by } \textit{Lemma 3.29}\}$$

and define $\Theta_{\mathcal{C}}^{\mathrm{proj}}(f, g; a)$ analogously using the projective point of Lemma 3.29.

Proposition 3.30 (distinct-parameter quotient ledger for lines)**.** *For every received line $f + zg$, the finite affine slopes that are support-wise MCA-bad and have a witness support in $\mathcal{S}_{\mathcal{C}}^{\geq a}$ are exactly the elements of $\Theta_{\mathcal{C}}^{\mathrm{aff}}(f, g; a)$. Hence the exact quotient-branch line numerator is*

$$|\Theta_{\mathcal{C}}^{\mathrm{aff}}(f, g; a)| \leq U_{\mathcal{C}}^{\mathrm{supp}}(a).$$

The corresponding projective-slope numerator is

$$|\Theta_{\mathcal{C}}^{\mathrm{proj}}(f, g; a)| \leq U_{\mathcal{C}}^{\mathrm{supp}}(a).$$

The same image ledgers upper-bound no-loss CA parameters whose explaining support lies in the same quotient family.

Proof. This is only Lemma 3.29 applied to every co-support in the family, followed by taking a union of the resulting parameter values. The CA claim follows because no-loss CA-badness is stronger than having an explained slope on the support and forbids simultaneous explanation on any support of size at least a. □

The curve analogue is just as explicit. For a degree-d power curve

$$W_\gamma = f_0 + \gamma f_1 + \cdots + \gamma^d f_d, \qquad \gamma \in \mathbb{F},$$

put $u_i = \mathrm{Syn}(f_i)$ and, for $T \subseteq D$, define

$$V_{i,T} := H_{t,j}(u_i)\boldsymbol{\ell}_T \in \mathbb{F}^t.$$

Let

$$\Gamma_T^{(d)} := \{\gamma \in \mathbb{F} : V_{0,T} + \gamma V_{1,T} + \cdots + \gamma^d V_{d,T} = 0,\ (V_{0,T}, \ldots, V_{d,T}) \neq 0\}.$$

The nonzero tuple condition is exactly support-wise noncontainment of the coefficient tuple on the fixed support.

Lemma 3.31 (exact support-image map for finite-parameter power curves)**.** *For a fixed support $S = D \setminus T$, the set of support-wise noncontained curve-bad finite parameters is exactly $\Gamma_T^{(d)}$. In particular $|\Gamma_T^{(d)}| \le d$.*

Proof. The same syndrome-locator recurrence applied to W_γ gives the displayed vector equation. If every $V_{i,T}$ is zero, then every coefficient word f_i is already explained on S, so the parameter is contained. Otherwise at least one scalar coordinate of $V_{0,T} + \gamma V_{1,T} + \cdots + \gamma^d V_{d,T}$ is a nonzero polynomial in γ of degree at most d, so it has at most d roots. □

Define the exact quotient curve image

$$\Gamma_{\mathcal{C}}^{(d)}(f_0, \ldots, f_d; a) := \bigcup_{T \in \mathcal{T}_{\mathcal{C}}^{\ge a}} \Gamma_T^{(d)}.$$

Proposition 3.32 (distinct-parameter quotient ledger for curves)**.** *For a finite-parameter degree-d power curve, the quotient-branch support-wise curve-MCA numerator is exactly*

$$|\Gamma_{\mathcal{C}}^{(d)}(f_0, \ldots, f_d; a)|.$$

It satisfies the safe bounds

$$|\Gamma_{\mathcal{C}}^{(d)}(f_0, \ldots, f_d; a)| \le d\, U_{\mathcal{C}}^{\mathrm{supp}}(a) \le d\, U_{\mathcal{C}}^{\mathrm{sum}}(a).$$

The same image ledger upper-bounds no-loss curve-CA parameters whose explaining support lies in the quotient family.

Proof. Apply Lemma 3.31 to each support and take the union of the resulting parameter sets. The displayed inequalities follow from $|\Gamma_T^{(d)}| \le d$ and from the support-union bound of Proposition 3.19. □

Definition 3.33 (fixed-support parameter gcd)**.** Let $C = \mathrm{RS}[\mathbb{F}, D, k]$, let $q = |\mathbb{F}|$, and let

$$W_\Gamma = f_0 + \Gamma f_1 + \cdots + \Gamma^d f_d, \qquad \Gamma \in \mathbb{F},$$

be a finite-parameter degree-d power curve of received words. Fix a co-support $T \subseteq D$, put $S = D \setminus T$, and assume $|S| > k$. With the notation of Lemma 3.31, write

$$P_{T,m}(\Gamma) := \sum_{i=0}^{d} \Gamma^i (V_{i,T})_m, \qquad 0 \le m < t, \qquad t = |S| - k.$$

Define the fixed-support contribution polynomial $g_T^{(d)}(\Gamma) \in \mathbb{F}[\Gamma]$ as follows. If all coordinate polynomials $P_{T,m}$ are identically zero, set $g_T^{(d)} = 1$. Otherwise let $g_T^{(d)}$ be the monic greatest common divisor of the nonzero polynomials among $P_{T,0}, \ldots, P_{T,t-1}$.

Lemma 3.34 (fixed-support gcd equals the noncontained parameter set)**.** *In the setting of Definition 3.33, the finite parameters $\gamma \in \mathbb{F}$ for which W_γ is explained on $S = D \setminus T$ and the coefficient tuple $(f_0, \ldots, f_d)$ is not simultaneously explained on S are exactly the roots in $\mathbb{F}$ of $g_T^{(d)}$. In particular*

$$\deg g_T^{(d)} \le d,$$

with equality possible only before taking common factors across different supports.

Proof. By Lemma 3.31, the fixed support explains W_γ exactly when

$$P_{T,m}(\gamma) = 0 \qquad (0 \le m < t).$$

If every coordinate polynomial is identically zero, then every coefficient vector $V_{i,T}$ is zero, so every coefficient word f_i is already explained on S; this is the contained branch and contributes no noncontained parameter. The convention $g_T^{(d)} = 1$ gives the empty root set, as required.

Otherwise the common zeros of the coordinate polynomials are exactly the roots of their monic gcd after zero coordinates are discarded. Since at least one coordinate polynomial is nonzero of degree at most d, the gcd also has degree at most d. The nonzero tuple condition in Lemma 3.31 is precisely the assertion that the coordinate polynomials are not all zero. □

Definition 3.35 (quotient image lcm polynomial)**.** Let $\mathcal{T}$ be any finite family of co-supports $T \subseteq D$ with $|D \setminus T| > k$. For a degree-d finite-parameter power curve define

$$R_{\mathcal{T}}^{(d)}(\Gamma) := \operatorname{rad} \operatorname{lcm}_{T \in \mathcal{T}} g_T^{(d)}(\Gamma),$$

where factors equal to 1 are ignored and the empty lcm is 1. For the quotient branch at agreement threshold $a > k$ and divisor set $\mathcal{C}$, we use

$$\mathcal{T} = \mathcal{T}_{\mathcal{C}}^{\ge a} := \{D \setminus S : S \in \mathcal{S}_{\mathcal{C}}^{\ge a}\}.$$

The finite-field root-count polynomial is

$$R_{\mathcal{T},F}^{(d)}(\Gamma) := \gcd\big(R_{\mathcal{T}}^{(d)}(\Gamma), \Gamma^q - \Gamma\big).$$

Because $R_{\mathcal{T}}^{(d)}$ is squarefree, the degree of $R_{\mathcal{T},F}^{(d)}$ is the number of $\mathbb{F}$-rational roots of $R_{\mathcal{T}}^{(d)}$.

Theorem 3.36 (exact finite-parameter quotient image ledger)**.** *Let $a > k$, let $\mathcal{C}$ be a finite divisor set, and let $\mathcal{T}_{\mathcal{C}}^{\ge a}$ be the quotient co-support family. For every finite-parameter degree-d power curve, the exact quotient-branch support-wise curve-MCA numerator is*

$$\boxed{\left|\Gamma_{\mathcal{C}}^{(d)}(f_0, \dots, f_d; a)\right| = \deg R_{\mathcal{T}_{\mathcal{C}}^{\ge a},F}^{(d)}.}$$

Consequently

$$\left|\Gamma_{\mathcal{C}}^{(d)}(f_0, \dots, f_d; a)\right| \le \deg R_{\mathcal{T}_{\mathcal{C}}^{\ge a}}^{(d)} \le d\, U_{\mathcal{C}}^{\mathrm{supp}}(a),$$

and the final support count may be evaluated by the exact divisor-block formula of Theorem 3.24. The same lcm ledger upper-bounds no-loss curve-CA parameters whose explaining support lies in the same quotient family.

Proof. By Lemma 3.34, for each fixed co-support T the noncontained parameters witnessed on $D \setminus T$ are exactly the $\mathbb{F}$-roots of $g_T^{(d)}$. Taking the union over $T \in \mathcal{T}_{\mathcal{C}}^{\ge a}$ is therefore the same as taking the $\mathbb{F}$-rational root set of the squarefree lcm $R_{\mathcal{T}_{\mathcal{C}}^{\ge a}}^{(d)}$. Over a finite field, the $\mathbb{F}$-rational roots of a squarefree polynomial are counted by the degree of its gcd with $\Gamma^q - \Gamma$, proving the displayed equality.

The degree bound follows from $\deg g_T^{(d)} \le d$ and from the fact that a squarefree lcm has degree at most the sum of the degrees of the factors before coalescing. The no-loss curve-CA claim follows because a no-loss CA parameter with an explaining support in the quotient family is, on that support, noncontained; otherwise the coefficient tuple would have a simultaneous explanation on a support of size at least a. □

Corollary 3.37 (affine lines as the degree-one lcm ledger)**.** *For an affine received line $f + Zg$, take $d = 1$, $f_0 = f$, and $f_1 = g$. Then $R^{(1)}_{\mathcal{T}_{\mathcal{C}}^{\geq a}}(Z)$ is the squarefree product of the linear factors $Z - z$ for the distinct finite quotient-witnessed support-wise MCA-bad slopes. Hence*

$$\#\{\textit{finite quotient-witnessed MCA-bad slopes}\} = \deg R^{(1)}_{\mathcal{T}_{\mathcal{C}}^{\geq a}}.$$

The same polynomial upper-bounds no-loss CA finite slopes whose explaining support lies in the declared quotient family.

Proof. For $d = 1$, each nonconstant fixed-support gcd has degree at most one. Its root, if present, is exactly the unique finite slope supplied by Lemma 3.29. Taking the squarefree lcm over the declared co-support family therefore multiplies one copy of $Z - z$ for each distinct finite slope and no copies for supports that contribute no finite noncontained slope. □

Definition 3.38 (projective line lcm ledger)**.** For a projective received line $Z_0 f + Z_1 g$ and a co-support T with $|D \setminus T| > k$, form the homogeneous linear coordinate forms

$$L_{T,m}(Z_0, Z_1) := Z_0(A_T)_m + Z_1(B_T)_m, \qquad 0 \leq m < t.$$

If all $L_{T,m}$ are zero, set $h_T^{\mathrm{proj}} = 1$. Otherwise let h_T^{proj} be the gcd of the nonzero linear forms, normalized by a fixed monomial order. Thus h_T^{proj} is either 1 or a homogeneous linear form. For a co-support family $\mathcal{T}$ define

$$R_{\mathcal{T}}^{\mathrm{proj}}(Z_0, Z_1) := \operatorname{rad} \operatorname{lcm}_{T \in \mathcal{T}} h_T^{\mathrm{proj}}.$$

Corollary 3.39 (exact projective quotient image ledger)**.** *For projective slopes, the exact quotient-branch support-wise MCA numerator is*

$$\deg R^{\mathrm{proj}}_{\mathcal{T}_{\mathcal{C}}^{\geq a}}.$$

The same degree upper-bounds no-loss projective-CA parameters whose explaining support lies in the quotient family.

Proof. A nonzero family of homogeneous linear forms on $\mathbb{P}^1$ has a common zero if and only if the forms have a common homogeneous linear factor; this factor cuts out exactly one $\mathbb{F}$-rational projective point. The all-zero case is the contained branch and contributes no bad point, by convention. The squarefree lcm therefore contains one linear factor for each distinct projective parameter and no duplicates. □

Definition 3.40 (lcm quotient-image certificate)**.** An lcm quotient-image certificate for a line or finite-parameter curve consists of:

(i) the divisor set $\mathcal{C}$, the agreement threshold a, and either the exact support-union certificate of Definition 3.26 or an explicit hash/orbit enumeration of the co-support family $\mathcal{T}_{\mathcal{C}}^{\geq a}$;

(ii) for every represented co-support T, the locator vector $\boldsymbol{\ell}_T$ and the syndrome-window coefficient vectors $V_{i,T}$;

(iii) the coordinate polynomials $P_{T,m}(\Gamma)$, the fixed-support gcd $g_T^{(d)}$, and a proof that the all-zero case has been assigned to the contained branch;

(iv) the squarefree lcm $R_{\mathcal{T}}^{(d)}$ and, over the finite field, the root-count factor $R_{\mathcal{T},F}^{(d)} = \gcd(R_{\mathcal{T}}^{(d)}, \Gamma^q - \Gamma)$;

(v) the declared numerator $\deg R_{\mathcal{T},F}^{(d)}$ for finite curves or $\deg R_{\mathcal{T}}^{\mathrm{proj}}$ for projective lines.

The verifier checks the locator split condition, the syndrome recurrence, the gcd and lcm arithmetic, squarefreeness, and the final finite-field root count. This certificate replaces a support count by an exact distinct-parameter count inside the declared quotient branch.

Remark 3.41 (status after the lcm image ledger). The quotient ledgers now separate three tasks that were previously entangled. First, Theorem 3.8 supplies lower floors by heaviest generated-field prefix fibers. Second, Theorem 3.24 counts the declared quotient support union with divisor overlaps handled exactly. Third, Theorem 3.36 and Corollaries 3.37 and 3.39 push that support family through the fixed-locator equations and coalesce duplicate finite, projective, and curve parameters by gcd/lcm arithmetic. The remaining quotient-side obstruction is not duplicate coalescing inside a declared quotient family; it is the structural safe-side theorem that every bad parameter in the intended regime either belongs to this quotient image, to the tangent branch, or to the aperiodic/extension buckets bounded elsewhere.

Definition 3.42 (quotient image certificate)**.** A quotient image certificate for a line or curve ledger consists of:

(i) the divisor set $\mathcal{C}$ and the agreement threshold a;

(ii) a deterministic description, hash, or orbit decomposition of $\mathcal{T}_{\mathcal{C}}^{\geq a}$;

(iii) for every represented co-support T, the locator coefficients $\boldsymbol{\ell}_T$ and the syndrome-window vectors A_T, B_T for a line, or $V_{0,T}, \ldots, V_{d,T}$ for a curve;

(iv) the declared parameter value, parameter bucket, or the declaration that no parameter is contributed;

(v) a bucket table proving the number of distinct parameters after duplicate coalescing.

The verifier checks the locator split condition, the Hankel equations, the noncontainment vector, and the bucket consistency. This certificate is the preferred replacement for a raw quotient-support count whenever the support count is much larger than the actual parameter image.

Remark 3.43 (status of the refinement). The propositions above are exact for the declared quotient branch and are a real refinement over the support ledger: the numerator is an image size, not a support count. They still do not prove that all bad parameters are quotient-periodic. The remaining safe-side theorem for the full prize is the aperiodic Hankel-packing bound: after tangent and quotient image branches are removed, one must bound the parameters coming from co-supports outside the declared quotient families.

3.4 Aperiodic Hankel chart atlas

The quotient image ledger of Section 3.3 counts bad parameters whose witness support lies in a declared quotient-remainder family. The complementary safe-side problem is to control all other supports. This subsection is the contributor-facing formulation of that problem. It does not claim that the needed aperiodic eliminants always exist. It gives a finite chart atlas, with

precise ideals and pivots, so that a verifier can check a supplied eliminant and convert it into a bad-parameter upper bound.

The main point is to avoid one huge saturated incidence ideal. We split the aperiodic branch into:

(i) a regular overdetermined Hankel bucket, where a single maximal minor in the line parameter already gives an eliminant;

(ii) finite affine pivot charts, one for each noncontainment coordinate;

(iii) a separate projective-infinity chart, if projective slopes are being counted;

(iv) finite-parameter curve pivot charts, one for a nonzero coefficient of one coordinate polynomial;

(v) a singular bucket, where all regular minors vanish or the projection to the parameter line remains positive-dimensional after the previous splits.

Only the last bucket is genuinely unresolved. A staircase scanner should keep it as a named obstruction rather than hiding it inside the word "aperiodic."

Fix an exact agreement value $A \geq k$ and put

$$j = j_A := n - A, \qquad t = t_A := A - k, \qquad R = n - k.$$

Let

$$P_D(X) := \prod_{x \in D} (X - x).$$

A monic degree-j co-support locator is written

$$L_\ell(X) = X^j + \ell_{j-1} X^{j-1} + \cdots + \ell_0, \qquad \ell = (\ell_0, \ldots, \ell_{j-1}) \in \mathbb{A}^j.$$

Let $R_{D,j}(\ell; X)$ be the remainder of $P_D(X)$ modulo $L_\ell(X)$, and let $I_{\mathrm{split,j}}$ be the ideal generated by the j coefficients of $R_{D,j}$. Since P_D is squarefree, the F-points of $V(I_{\mathrm{split,j}})$ are exactly the squarefree D-split monic degree-j locators.

For a word $y : D \to F$ write $\mathrm{Syn}(y) \in F^R$ for the parity-check syndrome used above. Let

$$\widetilde{\ell} = (\ell_0, \ldots, \ell_{j-1}, 1)^t.$$

For an affine line $f + zg$, put $u = \mathrm{Syn}(f)$ and $v = \mathrm{Syn}(g)$ and define

$$A(\ell) := H_{t,j}(u)\widetilde{\ell}, \qquad B(\ell) := H_{t,j}(v)\widetilde{\ell} \qquad \in F^t.$$

For a fixed split locator, the support-wise finite-slope condition is exactly

$$A(\ell) + ZB(\ell) = 0, \qquad B(\ell) \neq 0.$$

The inequality is the noncontainment condition.

3.4.1 Regular overdetermined bucket

Assume first that

$$t_A \geq j_A + 1, \qquad \text{equivalently} \qquad 2A \geq n + k + 1.$$

Set

$$M_A(Z) := H_{t_A, j_A}(u) + Z H_{t_A, j_A}(v).$$

This is a $t_A \times (j_A + 1)$ matrix with entries affine-linear in Z.

Definition 3.44 (regular Hankel minor certificate)**.** A regular Hankel minor certificate at exact agreement A is a row set $I_A \subseteq \{0, \ldots, t_A - 1\}$ of size $j_A + 1$ such that

$$\Delta_A(Z) := \det M_A(Z)_{I_A, \{0, \ldots, j_A\}}$$

is not the zero polynomial in $F[Z]$. If no such row set exists, the line is called *Hankel-singular at A.*

Lemma 3.45 (regular exact-agreement eliminant)**.** *Assume $t_A \geq j_A + 1$, and suppose $\Delta_A(Z)$ is a regular Hankel minor certificate. Every finite support-wise MCA-bad slope with a witness support of exact size A is a root of Δ_A. Hence the number of such slopes is at most*

$$\deg \Delta_A \leq j_A + 1 = n - A + 1.$$

The same root containment holds for no-loss CA-bad slopes with exact agreement A.

Proof. Let z be a bad finite slope with witness support S of size A, and put $T = D \setminus S$. Then $|T| = j_A$, and the locator vector $\widetilde{\ell}_T$ is a nonzero kernel vector of $M_A(z)$. Therefore the column rank of $M_A(z)$ is less than $j_A + 1$, so every $(j_A + 1) \times (j_A + 1)$ maximal minor vanishes at z, including $\Delta_A(z)$. The degree bound follows because Δ_A is the determinant of a $(j_A+1) \times (j_A+1)$ matrix whose entries are affine-linear in Z. □

Theorem 3.46 (regular closed-range Hankel packing)**.** *Let $a \geq k$ and suppose that for each exact agreement $A \in \{a, a+1, \ldots, n\}$ with $t_A \geq j_A + 1$ a regular Hankel minor certificate Δ_A is supplied. Then all bad finite slopes whose exact witness size is in this regular bucket lie in the root set of*

$$\Delta_{\geq a}(Z) := \prod_{A=a}^{n} \Delta_A(Z),$$

with factors omitted for exact agreements outside the overdetermined range. Consequently their number is at most

$$\sum_A \deg \Delta_A.$$

If exact root sets are supplied, roots already charged to the tangent or quotient-image ledgers may be removed from this count by duplicate coalescing.

Proof. Apply Lemma 3.45 at the exact size of the witness support and take the union of the resulting root sets. The degree sum is a safe union bound. □

Canonical regular rank-drop gcds. The regular-minor certificate below can be made canonical. Instead of choosing one nonzero maximal minor, take the greatest common divisor of all nonzero maximal minors. This gives exactly the finite-parameter values at which the Hankel pencil can drop rank, before the split-locator and noncontainment tests are imposed. Thus it is an upper ledger for the regular bucket, and often a much smaller one than the degree of a single selected minor.

Fix an exact agreement value $A \geq k$, and keep the notation

$$j = j_A := n - A, \qquad t = t_A := A - k, \qquad r = n - k.$$

The regular overdetermined range is

$$t_A > j_A, \qquad \text{i.e.} \qquad t_A \geq j_A + 1, \qquad \text{equivalently} \qquad 2A \geq n + k + 1.$$

Definition 3.47 (canonical affine rank-drop gcd)**.** Fix an affine received line $f + Zg$, put $u = \mathrm{Syn}(f)$ and $v = \mathrm{Syn}(g)$, and assume $t_A \geq j_A + 1$. Set

$$M_A(Z) := H_{t_A, j_A}(u) + Z H_{t_A, j_A}(v).$$

Let $\mathfrak{M}_A(Z)$ be the finite set of all $(j_A + 1) \times (j_A + 1)$ maximal minors of $M_A(Z)$. If every element of $\mathfrak{M}_A(Z)$ is the zero polynomial, the exact-agreement bucket is called *rank-drop singular*. Otherwise define

$$G_A^{\mathrm{aff}}(Z) := \mathrm{rad}\,\gcd\{\Delta(Z) : \Delta \in \mathfrak{M}_A(Z),\ \Delta \neq 0\},$$

with the gcd normalized to be monic. Over the finite field $\mathbb{F}$ of size q set

$$G_{A,F}^{\mathrm{aff}}(Z) := \gcd\big(G_A^{\mathrm{aff}}(Z), Z^q - Z\big).$$

Thus $G_{A,F}^{\mathrm{aff}}$ is the squarefree polynomial whose roots are precisely those F-rational parameters at which all nonzero maximal minors vanish.

Theorem 3.48 (canonical exact-agreement rank-drop ledger)**.** *Assume $t_A \geq j_A + 1$ and the exact-agreement bucket is not rank-drop singular. Every finite support-wise MCA-bad slope with an exact witness support of size A is a root of G_A^{aff}. Consequently the number of such slopes is at most*

$$\deg G_{A,F}^{\mathrm{aff}} \leq \deg G_A^{\mathrm{aff}} \leq j_A + 1 = n - A + 1.$$

The same bound holds for no-loss CA-bad slopes with exact agreement A.

Proof. Let $z \in F$ be support-wise bad with exact witness support S of size A, and put $T = D \setminus S$. By Lemma 3.28, the locator vector $\boldsymbol{\ell}_T$ is a nonzero kernel vector of $M_A(z)$. Therefore $\mathrm{rank}\, M_A(z) < j_A + 1$, so every maximal minor of $M_A(Z)$ vanishes at $Z = z$. In particular every nonzero maximal minor vanishes at z, and hence their gcd G_A^{aff} vanishes at z.

The finite-field root count is exactly $\deg G_{A,F}^{\mathrm{aff}}$ because $G_{A,F}^{\mathrm{aff}}$ is squarefree and divides $Z^q - Z$. Each maximal minor has degree at most $j_A + 1$, since it is the determinant of a matrix whose entries are affine-linear in Z; the gcd cannot have larger degree. The no-loss CA claim is identical, because no-loss CA-badness with exact agreement A also supplies an explaining support of size A and hence the same locator-kernel condition. □

Remark 3.49 (why this improves the one-minor certificate). Lemma 3.45 used one certified nonzero maximal minor and therefore counted the roots of a single determinant. The polynomial G_A^{aff} uses all maximal minors and counts their common finite-field root set. It is never worse than any one-minor certificate and can be much smaller; it is also canonical, so two scanners that use the same syndrome convention print the same regular-bucket polynomial.

Definition 3.50 (closed-ball rank-drop lcm)**.** For an agreement threshold $a > k$, let

$$\mathcal{A}_{\mathrm{reg}}(a) := \{A \in \{a, a+1, \dots, n\} : t_A \geq j_A + 1 \text{ and the bucket } A \text{ is not rank-drop singular}\}.$$

Define

$$G_{\geq a}^{\mathrm{aff}}(Z) := \operatorname{rad} \operatorname{lcm}_{A \in \mathcal{A}_{\mathrm{reg}}(a)} G_A^{\mathrm{aff}}(Z), \qquad G_{\geq a,F}^{\mathrm{aff}}(Z) := \gcd\big(G_{\geq a}^{\mathrm{aff}}(Z), Z^q - Z\big),$$

with the empty lcm equal to 1. Exact agreement values $A \in \{a, \dots, n\} \setminus \mathcal{A}_{\mathrm{reg}}(a)$ are recorded as residual buckets: they are either underdetermined ($t_A < j_A + 1$) or rank-drop singular.

Corollary 3.51 (canonical regular closed-ball Hankel packing)**.** *All finite support-wise MCA-bad slopes whose exact witness size lies in $\mathcal{A}_{\mathrm{reg}}(a)$ are roots of $G_{\geq a}^{\mathrm{aff}}$. Hence their number is at most*

$$\deg G_{\geq a,F}^{\mathrm{aff}} \leq \sum_{A \in \mathcal{A}_{\mathrm{reg}}(a)} (n - A + 1).$$

The same bound holds for no-loss CA-bad slopes in the same regular buckets.

Proof. Apply Theorem 3.48 at the exact agreement size of the witness and take the union over A. The squarefree lcm coalesces duplicate roots across exact-agreement buckets; the final gcd with $Z^q - Z$ counts only finite-field roots. The displayed degree sum is the safe bound obtained before coalescing. □

Definition 3.52 (paid-root removal for the regular branch)**.** Suppose a squarefree polynomial $P_{\mathrm{paid,F}}(Z)$ is supplied whose roots are exactly the finite-field parameters already charged to tangent/common-line and quotient-image ledgers. Define the regular residual polynomial

$$G_{\mathrm{reg}\setminus\mathrm{paid},F}^{\mathrm{aff}}(Z) := \frac{G_{\geq a,F}^{\mathrm{aff}}(Z)}{\gcd(G_{\geq a,F}^{\mathrm{aff}}(Z), P_{\mathrm{paid,F}}(Z))}.$$

Its degree is the number of regular rank-drop parameters not already paid by the previous ledgers.

Corollary 3.53 (quotient/tangent-removed regular aperiodic numerator)**.** *If $P_{\mathrm{paid,F}}$ is exact in the sense of Definition 3.52, then the number of finite slopes in the regular overdetermined branch that remain after removing the paid quotient/tangent parameters is at most*

$$\deg G_{\mathrm{reg}\setminus\mathrm{paid},F}^{\mathrm{aff}}.$$

This is a safe aperiodic numerator for the regular buckets. The underdetermined and rank-drop singular exact-agreement buckets must still be handled by pivot charts or by a separately named residual obstruction.

Proof. By Corollary 3.51, every regular-bucket bad slope is a root of $G^{\mathrm{aff}}_{\ge a,F}$. Removing the roots of the exact paid polynomial deletes precisely those already charged parameters from this finite root set. The remaining root set is counted by the displayed quotient because all polynomials are squarefree divisors of $Z^q - Z$. □

Definition 3.54 (finite-parameter curve rank-drop gcd)**.** For a finite-parameter degree-d curve

$$W_\Gamma = f_0 + \Gamma f_1 + \cdots + \Gamma^d f_d,$$

put $u_i = \mathrm{Syn}(f_i)$ and, at exact agreement A, define

$$M_A^{(d)}(\Gamma) := \sum_{i=0}^{d} \Gamma^i H_{t_A, j_A}(u_i).$$

When $t_A \ge j_A + 1$ and not all maximal minors of $M_A^{(d)}$ vanish identically, set

$$G_A^{(d)}(\Gamma) := \operatorname{rad} \gcd\{\Delta(\Gamma) : \Delta \text{ is a nonzero maximal minor of } M_A^{(d)}\},$$

and

$$G_{A,F}^{(d)}(\Gamma) := \gcd\big(G_A^{(d)}(\Gamma), \Gamma^q - \Gamma\big).$$

Theorem 3.55 (canonical curve rank-drop ledger)**.** *In a non-singular overdetermined exact-agreement bucket, every finite curve-MCA-bad parameter is a root of $G_A^{(d)}$. Consequently the number of such parameters is at most*

$$\deg G_{A,F}^{(d)} \le \deg G_A^{(d)} \le d(j_A + 1).$$

The same bound holds for no-loss curve-CA parameters in the same bucket. The closed-ball form is obtained by taking the squarefree lcm of the polynomials $G_A^{(d)}$ over the regular exact-agreement buckets and then taking the gcd with $\Gamma^q - \Gamma$.

Proof. A curve-bad parameter with exact witness support S gives, for $T = D \setminus S$, a nonzero locator vector in the kernel of $M_A^{(d)}(\gamma)$. Hence every maximal minor vanishes at γ, and the same gcd argument as in Theorem 3.48 applies. Each matrix entry has degree at most d in Γ, so each maximal minor has degree at most $d(j_A + 1)$. □

Definition 3.56 (regular rank-drop certificate)**.** A regular rank-drop certificate for a row, a received affine line or finite curve, and an agreement threshold a consists of:

(i) the syndrome vectors of the coefficient words and the exact-agreement range $A \in \{a, \ldots, n\}$;

(ii) for every overdetermined exact agreement A, either a declaration that all maximal minors vanish identically, or the canonical gcd polynomial G_A^{aff}, respectively $G_A^{(d)}$, together with ideal-membership or raw minor arithmetic proving that it is the gcd of all nonzero maximal minors;

(iii) the squarefree lcm over all non-singular regular buckets and the finite-root factor obtained by gcd with $Z^q - Z$ or $\Gamma^q - \Gamma$;

(iv) if paid branches are removed, the exact paid-root polynomial and the finite root-set quotient of Definition 3.52;

(v) a residual-bucket table listing all underdetermined and rank-drop singular exact agreements that are not covered by the regular gcd ledger.

The verifier checks the syndrome convention, the Hankel matrices, the maximal minor gcds, squarefree lcms, finite-field root factors, paid-root removal, and the residual-bucket table.

Remark 3.57 (status after the rank-drop ledger). The regular overdetermined part of the aperiodic Hankel problem is now canonical: one no longer has to choose a minor or accept a degree sum when the common root set of all maximal minors is smaller. This still does not solve the near-capacity underdetermined buckets, which form the hard aperiodic residual class. Those buckets remain assigned to the affine pivot charts, projective-infinity charts, curve coefficient pivots, or named singular residual obstructions.

3.4.2 Finite affine pivot charts

The regular minor test is intentionally coarse and only works directly in the overdetermined range. The next atlas is the one contributors should use to construct actual aperiodic eliminants.

Definition 3.58 (aperiodic locator chart)**.** An aperiodic locator chart at exact agreement A is a pair

$$X = (E_X, \Delta_X),$$

where E_X is a finite list of polynomial equations in the locator coefficients ℓ, and $\Delta_X(\ell)$ is a product of declared nonzero tests. The chart represents the constructible set

$$V(I_{\mathrm{split,j}} + \langle E_X \rangle) \setminus V(\Delta_X).$$

The equations and inequations must include, or reference, the removal of already paid tangent/common-code-line and quotient-image branches. A family of charts $\mathfrak{X}_A$ is a valid aperiodic chart cover if the union of their constructible sets contains every remaining split locator of exact co-support size j_A not assigned to those earlier ledgers.

For a coordinate $h \in \{0, \dots, t-1\}$ define the affine pivot consistency ideal

$$\begin{aligned} J^{\mathrm{aff}}_{X,h} := I_{\mathrm{split,j}} + \langle E_X \rangle + \langle A_m(\ell)B_h(\ell) - A_h(\ell)B_m(\ell) : 0 \le m < t \rangle \\ + \langle ZB_h(\ell) + A_h(\ell) \rangle \subseteq F[\ell_0, \dots, \ell_{j-1}, Z], \end{aligned}$$

and saturate it by the pivot denominator:

$$\widehat{J}^{\mathrm{aff}}_{X,h} := J^{\mathrm{aff}}_{X,h} : (\Delta_X B_h)^{\infty}.$$

The equations $A_mB_h - A_hB_m = 0$ force A and B to be collinear on the pivot chart, and $ZB_h + A_h = 0$ records the unique finite slope.

Theorem 3.59 (pivoted affine-line eliminant certificate)**.** *Let X be an aperiodic locator chart and let h be a noncontainment pivot. If there is a nonzero polynomial*

$$Q_{X,h}(Z) \in \widehat{J}^{\mathrm{aff}}_{X,h} \cap F[Z]$$

of degree $D_{X,h}$, then the number of finite support-wise MCA-bad slopes in that chart with $B_h \neq 0$ is at most $\min\{|F|, D_{X,h}\}$. The same bound holds for no-loss CA-bad slopes in that pivot chart.

Proof. A bad slope in the chart has a split locator with $\Delta_X \neq 0$ and, on this pivot chart, $B_h \neq 0$. The Hankel incidence equation gives $A + ZB = 0$, hence the point (ℓ, Z) satisfies all generators of $J^{\mathrm{aff}}_{X,h}$ and survives the saturation. Therefore $Q_{X,h}(Z) = 0$. A nonzero degree-$D_{X,h}$ polynomial over F has at most $D_{X,h}$ roots. □

Remark 3.60 (why the pivot split is useful). The saturation by the whole vector B in the unsplit incidence ideal is replaced by the explicit open cover $B_h \neq 0$. A scanner can try all h, produce small univariate eliminants on successful pivots, and report the residual locus $B_0 = \cdots = B_{t-1} = 0$ as contained for support-wise MCA. For no-loss CA, that residual locus also contributes no bad slope at the same support, because if $A = B = 0$ then both f and g are explained on the witness support.

3.4.3 Projective-slope split

For projective parameters $[Z_0 : Z_1]$ the incidence equation is

$$Z_0 A(\ell) + Z_1 B(\ell) = 0.$$

The finite patch $Z_0 \neq 0$ is exactly the affine pivot atlas above with $Z = Z_1/Z_0$. The extra point at infinity is $[0 : 1]$.

Definition 3.61 (projective infinity chart)**.** For a chart X, the projective-infinity chart is the constructible locus

$$I_{\mathrm{split,j}} + \langle E_X \rangle + \langle B_0, \ldots, B_{t-1} \rangle, \qquad \Delta_X \neq 0, \quad A \neq 0.$$

If it is nonempty, it contributes the single projective parameter $[0 : 1]$. If it is empty, the infinity point contributes nothing on X.

Corollary 3.62 (projective eliminant certificate)**.** *For a chart cover $\mathfrak{X}_A$, the projective-slope numerator is bounded by the union of the finite-patch root sets from Theorem 3.59 together with at most one extra point $[0 : 1]$, provided the infinity charts are certified either empty or nonempty. Equivalently, a homogeneous nonzero eliminant*

$$Q_X(Z_0, Z_1)$$

of degree D_X bounds the number of projective bad parameters in that chart by D_X.

Proof. The finite patch is the affine theorem. The complement of that patch is the single point $[0 : 1]$, which is governed by Definition 3.61. For the homogeneous formulation, every bad projective parameter is a zero of the homogeneous eliminant, and a nonzero homogeneous polynomial of degree D_X on $\mathbb{P}^1(F)$ has at most D_X zeros. □

3.4.4 Finite-parameter curve charts

For a degree-d finite power curve

$$W_\Gamma = f_0 + \Gamma f_1 + \cdots + \Gamma^d f_d,$$

put

$$V_i(\ell) := H_{t,j}(\mathrm{Syn}(f_i))\widetilde{\ell} \in F^t, \qquad 0 \leq i \leq d.$$

The curve incidence equation is

$$V_0(\ell) + \Gamma V_1(\ell) + \cdots + \Gamma^d V_d(\ell) = 0.$$

Curve-MCA noncontainment means that not all coefficient vectors $V_i(\ell)$ are zero. Thus the noncontained branch is covered by coefficient pivots (i,h) with $(V_i)_h \neq 0$.

For a chart X and a coefficient pivot (i,h), define

$$J^{\mathrm{curve,d}}_{X,i,h} := I_{\mathrm{split,j}} + \langle E_X \rangle + \left\langle \sum_{r=0}^{d} \Gamma^r (V_r(\ell))_m : 0 \le m < t \right\rangle \subseteq F[\ell_0, \ldots, \ell_{j-1}, \Gamma],$$

and set

$$\widehat{J}^{\mathrm{curve,d}}_{X,i,h} := J^{\mathrm{curve,d}}_{X,i,h} : (\Delta_X (V_i)_h)^\infty.$$

Theorem 3.63 (pivoted curve eliminant certificate)**.** *If a chart X and coefficient pivot (i,h) have a nonzero eliminant*

$$Q_{X,i,h}(\Gamma) \in \widehat{J}^{\mathrm{curve,d}}_{X,i,h} \cap F[\Gamma]$$

of degree $D_{X,i,h}$, then the number of finite degree-d curve-MCA bad parameters in that chart and pivot is at most $\min\{|F|, D_{X,i,h}\}$. The same bound applies to no-loss curve-CA parameters in the same pivot chart.

Proof. Every bad curve parameter on the pivot chart gives a split locator satisfying all curve-incidence equations and the nonzero pivot condition $(V_i)_h \neq 0$. It therefore survives the saturation and is a root of the supplied eliminant. □

Residual rank-pivot charts. The gcd ledger above removes the regular overdetermined branch. What remains is not a formless exceptional set: every remaining bad parameter is still a point of a finite collection of determinantal rank charts. The point of the next statements is to make this residual bucket verifier-facing. A certificate may now say exactly which rank chart it uses, which noncontainment coordinate is not settled here, and which univariate or homogeneous eliminant cuts out the parameter projection. If no such eliminant is supplied, the chart is a named residual obstruction rather than an unaccounted term.

Fix an exact agreement value $A \ge k$ and keep

$$j = j_A := n - A, \qquad t = t_A := A - k, \qquad r = n - k.$$

For a projective received line $Z_0 f + Z_1 g$, put

$$U := H_{t,j}(\mathrm{Syn}(f)), \qquad V := H_{t,j}(\mathrm{Syn}(g)), \qquad M^{\mathrm{proj}}_A(Z_0, Z_1) := Z_0 U + Z_1 V.$$

For a locator vector $\widetilde{\ell}$ write

$$A(\ell) := U\widetilde{\ell}, \qquad B(\ell) := V\widetilde{\ell}.$$

Thus the projective support equation is

$$Z_0 A(\ell) + Z_1 B(\ell) = 0,$$

and projective noncontainment is $(A(\ell), B(\ell)) \neq (0,0)$.

Let $X = (E_X, \Delta_X)$ be a locator chart in the sense of Definition 3.58. For row and column sets $I \subseteq \{0, \dots, t-1\}$ and $J \subseteq \{0, \dots, j\}$ with $|I| = |J| = s$, write

$$\Delta_{I,J}(Z_0, Z_1) := \det M_A^{\mathrm{proj}}(Z_0, Z_1)_{I,J},$$

with the convention that the unique 0×0 minor is 1. Let $\mathcal{M}_{s+1}(M_A^{\mathrm{proj}})$ denote the set of all $(s+1) \times (s+1)$ minors, empty when $s = \min\{t, j+1\}$.

For a noncontainment pivot

$$\pi \in \{A_h : 0 \le h < t\} \cup \{B_h : 0 \le h < t\},$$

let $\Pi_\pi(\ell)$ denote the corresponding scalar coordinate. Define the saturated projective rank-pivot ideal

$$\widehat{J}^{\mathrm{proj}}_{X,I,J,\pi} := \big(I_{\mathrm{split},j} + \langle E_X \rangle + \langle M_A^{\mathrm{proj}}(Z_0, Z_1)\widetilde{\ell}\rangle + \langle \Delta : \Delta \in \mathcal{M}_{s+1}(M_A^{\mathrm{proj}})\rangle\big) : \big(\Delta_X \Delta_{I,J} \Pi_\pi\big)^\infty \subseteq F[\ell_0, \dots, \ell_{j-1}, Z_0, Z_1].$$

The open condition $\Delta_{I,J} \neq 0$ and the vanishing of all larger minors put M_A^{proj} on the exact-rank-s stratum on this chart. The saturation by Π_π chooses one of the finitely many open noncontainment coordinates.

Theorem 3.64 (residual rank-pivot cover)**.** *Fix an exact agreement value A and a chart cover $\mathfrak{X}_A$ of the remaining locator branch after tangent/common-line and quotient-image removals. Every projective support-wise MCA-bad parameter with exact witness size A in that remaining branch lies on one of the rank-pivot charts*

$$\widehat{J}^{\mathrm{proj}}_{X,I,J,\pi}$$

for some $X \in \mathfrak{X}_A$, some rank s, some nonzero $s \times s$ minor $\Delta_{I,J}$ at the parameter, and some nonzero noncontainment coordinate Π_π. The same cover contains every no-loss projective CA-bad parameter in that branch.

Proof. Let $[z_0 : z_1]$ be bad with witness support S, co-support T, and locator vector $\widetilde{\ell}_T$. Since the chart family covers the remaining branch, $\widetilde{\ell}_T$ lies in some $X \in \mathfrak{X}_A$ with $\Delta_X \neq 0$. The support equation gives $M_A^{\mathrm{proj}}(z_0, z_1)\widetilde{\ell}_T = 0$. Let $s = \operatorname{rank} M_A^{\mathrm{proj}}(z_0, z_1)$. Choose a nonzero $s \times s$ minor at this parameter; then every $(s+1)\times(s+1)$ minor vanishes. Finally, support-wise noncontainment says that $(A(\ell_T), B(\ell_T)) \neq (0, 0)$, so some coordinate pivot Π_π is nonzero. Thus the point survives the displayed saturation and belongs to the indicated rank-pivot chart. The no-loss CA statement uses the same explaining support and same noncontainment conclusion. □

Theorem 3.65 (projective residual eliminant certificate)**.** *If a rank-pivot chart has a nonzero homogeneous eliminant*

$$Q_{X,I,J,\pi}(Z_0, Z_1) \in \widehat{J}^{\mathrm{proj}}_{X,I,J,\pi} \cap F[Z_0, Z_1]$$

of degree $D_{X,I,J,\pi}$, then the number of projective line parameters in that chart is at most $D_{X,I,J,\pi}$. More exactly, over a finite field F of size q, the number of finite-patch roots is

$$\deg \gcd\big(Q_{X,I,J,\pi}(1, Z), Z^q - Z\big),$$

and the point at infinity contributes precisely when $Q_{X,I,J,\pi}(0, 1) = 0$. Thus the exact finite-field count is

$$\deg \gcd\big(Q_{X,I,J,\pi}(1, Z), Z^q - Z\big) + \mathbf{1}_{Q_{X,I,J,\pi}(0,1)=0}.$$

Proof. Every parameter represented by the chart gives a point of the saturated incidence variety, hence is a zero of every element of the elimination ideal, including $Q_{X,I,J,\pi}$. A nonzero homogeneous polynomial of degree D on $\mathbb{P}^1(F)$ has at most D projective zeros. On the affine patch $Z_0 = 1$ the finite parameters are exactly the roots of $Q(1, Z)$ in F, which are counted by the gcd with $Z^q - Z$ after squarefree reduction. The complement of the finite patch is the single point $[0 : 1]$, giving the displayed indicator. □

Corollary 3.66 (finite and infinity patches)**.** *The finite affine rank-pivot certificate is obtained from Theorem 3.65 by setting $Z_0 = 1$ and using only noncontainment pivots $B_h \neq 0$. The infinity chart is obtained by setting $Z_0 = 0$, $Z_1 = 1$, imposing*

$$B(\ell) = 0, \qquad A_h(\ell) \neq 0$$

for some h, and applying the same saturated rank-pivot test to the constant matrix V. Consequently the projective sampler has no hidden extra case: every finite or infinite bad point is represented by one of these two patch types.

Proof. On the finite patch $[1 : Z]$, the support equation is $A(\ell) + ZB(\ell) = 0$. If $B(\ell) = 0$, then also $A(\ell) = 0$, so the fixed support is contained; hence every finite noncontained point has some $B_h \neq 0$. At infinity $[0 : 1]$, the support equation is $B(\ell) = 0$ and noncontainment is exactly $A(\ell) \neq 0$. These alternatives exhaust $\mathbb{P}^1(F)$. □

Curve coefficient-pivot rank charts. For a finite-parameter degree-d power curve

$$W_\Gamma = f_0 + \Gamma f_1 + \cdots + \Gamma^d f_d,$$

put $U_i := H_{t,j}(\mathrm{Syn}(f_i))$ and

$$M_A^{(d)}(\Gamma) := \sum_{i=0}^{d} \Gamma^i U_i, \qquad V_i(\ell) := U_i \widetilde{\ell}.$$

For row and column sets I, J of size s, write $\Delta_{I,J}^{(d)}(\Gamma)$ for the corresponding $s \times s$ minor of $M_A^{(d)}(\Gamma)$, and let $\mathcal{M}_{s+1}(M_A^{(d)})$ be the set of all larger minors. For a coefficient noncontainment pivot (i, h) define

$$\begin{aligned} \widehat{J}^{\mathrm{curve,d}}_{X,I,J,i,h} := \big(&I_{\mathrm{split},j} + \langle E_X \rangle + \langle M_A^{(d)}(\Gamma)\widetilde{\ell} \rangle \\ &+ \langle \Delta : \Delta \in \mathcal{M}_{s+1}(M_A^{(d)}) \rangle\big) : \big(\Delta_X \Delta_{I,J}^{(d)}(V_i)_h\big)^\infty. \end{aligned}$$

Theorem 3.67 (curve residual rank-pivot eliminant)**.** *The curve-bad parameters in a residual exact-agreement branch are covered by the charts $\widehat{J}^{\mathrm{curve,d}}_{X,I,J,i,h}$. If such a chart has a nonzero eliminant*

$$Q_{X,I,J,i,h}(\Gamma) \in \widehat{J}^{\mathrm{curve,d}}_{X,I,J,i,h} \cap F[\Gamma]$$

of degree $D_{X,I,J,i,h}$, then that chart contributes at most

$$\deg \gcd\big(Q_{X,I,J,i,h}(\Gamma), \Gamma^q - \Gamma\big) \leq D_{X,I,J,i,h}$$

finite parameters over F. The same bound applies to no-loss curve-CA parameters in the same chart.

Proof. A curve-bad parameter with locator $\widetilde{\ell}_T$ satisfies $M_A^{(d)}(\gamma)\widetilde{\ell}_T = 0$. Taking s to be the rank of $M_A^{(d)}(\gamma)$ gives a nonzero $s \times s$ minor and vanishing of all larger minors. Noncontainment of the coefficient tuple means that some coordinate $(V_i(\ell_T))_h$ is nonzero. Thus the point lies on one of the saturated rank-pivot charts and any univariate eliminant for that chart must vanish at γ. The finite-field root count is the gcd with $\Gamma^q - \Gamma$. □

Extension-line Weil restriction. The previous charts are intrinsic over the line field. For extension-field samplers one also needs a generated-field ledger: a low-dimensional set over the base field $\mathbb{B}$ can be negligible after division by $|F|$, even when it is large as a subset of $\mathbb{B}$. The following is the certificate form of that accounting.

Let $F/\mathbb{B}$ have degree m, choose a $\mathbb{B}$-basis $\omega_0, \dots, \omega_{m-1}$ of F, and write a finite affine line parameter as

$$Z = Z_0\omega_0 + \cdots + Z_{m-1}\omega_{m-1}, \qquad Z_i \in \mathbb{B}.$$

Assume $D \subseteq \mathbb{B}$ and write every locator coefficient in the same base field coordinates when necessary. Replacing each F-valued equation in a rank-pivot chart by its m coordinate equations over $\mathbb{B}$ gives a Weil-restricted saturated ideal

$$\widehat{J}_Y^{\mathrm{WR}} \subseteq \mathbb{B}[\ell\text{-coordinates}, Z_0, \dots, Z_{m-1}]$$

for a chart Y; noncontainment is split by one nonzero $\mathbb{B}$-coordinate of the chosen F-valued pivot and included in the saturation. Let π_Z denote projection to the parameter-coordinate space $\mathbb{A}^m_{\mathbb{B}}$.

Theorem 3.68 (extension-line dimension–degree ledger)**.** *Suppose that for a Weil-restricted residual chart Y a certificate gives a closed set*

$$Y_Z \supseteq \pi_Z\big(V(\widehat{J}_Y^{\mathrm{WR}})\big) \subseteq \mathbb{A}^m_{\mathbb{B}}$$

of degree at most Δ_Y and dimension at most e_Y. Then the number of extension-field affine parameters contributed by this chart is at most

$$\Delta_Y |\mathbb{B}|^{e_Y}.$$

If Y_Z is zero-dimensional and radical, this bound may be replaced by the exact number of its $\mathbb{B}$-rational points. After normalization by the line field, this chart contributes at most

$$\frac{\Delta_Y |\mathbb{B}|^{e_Y}}{|F|} = \Delta_Y |\mathbb{B}|^{e_Y - m}$$

to the MCA or CA error numerator.

Proof. The chosen basis identifies F with $\mathbb{B}^m$, so affine line parameters over F are exactly the $\mathbb{B}$-rational points of $\mathbb{A}^m_{\mathbb{B}}$. Every bad parameter represented by the chart lies in the projection of the Weil-restricted incidence variety and hence in $Y_Z(\mathbb{B})$. The standard affine degree point-count bound gives $|Y_Z(\mathbb{B})| \le \Delta_Y |\mathbb{B}|^{e_Y}$. The normalized statement divides by $|F| = |\mathbb{B}|^m$. □

Corollary 3.69 (base-rational line inertness as a chart certificate)**.** *If f and g are $\mathbb{B}$-valued, then every finite affine bad slope on every rank-pivot chart lies in $\mathbb{B}$. Equivalently, in the above coordinates all extension-valued charts with $(Z_1, \dots, Z_{m-1}) \neq 0$ are empty.*

Proof. For a fixed support and locator over $\mathbb{B}$, the vectors $A(\ell)$ and $B(\ell)$ are $\mathbb{B}$-valued. On a finite noncontained chart some coordinate B_h is nonzero, and the slope is forced by $Z = -A_h/B_h \in \mathbb{B}$. Hence no genuinely extension-valued finite parameter can survive the chart equations. □

Residual aperiodic certificate theorem. We can now state the safe-side theorem in the form a scanner should print. Let $P_{\mathrm{paid,F}}$ be the exact squarefree finite-root polynomial already charged to tangent/common-line and quotient-image branches, and let $G^{\mathrm{aff}}_{\mathrm{reg}\setminus\mathrm{paid},F}$ be the regular residual polynomial of Definition 3.52. For the remaining exact agreement levels, let $\mathfrak{R}_A$ be a finite list of rank-pivot residual charts, each equipped with either a finite/projective/curve eliminant or a Weil-restricted extension projection certificate.

Theorem 3.70 (scanner-checkable residual aperiodic ledger)**.** *Assume that the chart families $\mathfrak{R}_A$, for $A \geq a$, cover every remaining locator after tangent/common-line and quotient-image removals. Suppose each residual chart Y is certified by one of the following data:*

(i) *an affine eliminant $Q_Y(Z)$, contributing* $\deg \gcd(Q_Y, Z^q - Z)$*;*

(ii) *a homogeneous projective eliminant $Q_Y(Z_0, Z_1)$, contributing* $\deg \gcd(Q_Y(1, Z), Z^q - Z) + \mathbf{1}_{Q_Y(0,1)=0}$*;*

(iii) *a finite curve eliminant $Q_Y(\Gamma)$, contributing* $\deg \gcd(Q_Y, \Gamma^q - \Gamma)$*;*

(iv) *an extension Weil-restricted projection certificate (Δ_Y, e_Y), contributing $\Delta_Y |\mathbb{B}|^{e_Y}$ parameters.*

Then the finite affine aperiodic numerator after quotient/tangent removals is bounded by

$$\deg G^{\mathrm{aff}}_{\mathrm{reg}\setminus\mathrm{paid},F} + \sum_{A \geq a} \sum_{Y \in \mathfrak{R}_A} B_Y,$$

where B_Y is the corresponding contribution above. The projective and curve variants are obtained by using the projective or curve contributions for all charts. If exact root sets or exact zero-dimensional projections are printed, the displayed safe sum may be replaced by their union after duplicate coalescing.

Proof. The regular nonsingular overdetermined branch is covered by Corollary 3.53. Every remaining bad parameter is covered by the rank-pivot cover theorem, or by its curve analogue, after applying the declared chart cover. Each chart contribution is an upper bound by Theorems 3.65, 3.67 and 3.68. Summing the chart bounds is a safe union bound; replacing the sum by an exact union is valid when the certificate supplies the corresponding finite root or projection sets. □

Definition 3.71 (singular residual certificate)**.** A singular or underdetermined residual aperiodic certificate consists of:

(i) the paid tangent/common-line and quotient-image root ledgers already removed;

(ii) the regular rank-drop gcd/lcm ledger and its paid-root removal;

(iii) for each remaining exact agreement A, a chart cover of the residual locator set after the paid branches are removed;

(iv) for every chart, its rank s, pivot minor $\Delta_{I,J}$, noncontainment pivot, and either an affine, projective, or curve eliminant, or a Weil-restricted projection degree/dimension certificate;

(v) a residual table for every chart on which no eliminant or projection bound is supplied, labelled as quotient, tangent/common-line, base-field exceptional, genuinely extension-valued, or new singular obstruction.

The verifier checks chart coverage, rank-stratum equations, saturations by the pivot and non-containment coordinates, eliminant membership, finite-field root counts, extension projection bounds, and duplicate coalescing when claimed.

Remark 3.72 (status after the singular-bucket charts). This does not prove that the hard aperiodic residual buckets are small. It does remove one ambiguity: underdetermined and rank-drop singular buckets now have a precise finite certificate language, including finite slopes, projective infinity, finite-parameter curves, and genuinely extension-valued parameters. The next mathematical step is to construct low-degree eliminants, or prove structural classification, for these rank-pivot charts in the near-capacity regimes where $t_A = A - k$ is small.

Kernel-compressed rank-pivot eliminants. The rank-pivot charts above are already finite certificate objects, but the variables are still the full locator coordinates. The next reduction removes the linear kernel equations on a pivot chart by Cramer's rule. It is the algebraic form a scanner should use before attempting Gröbner elimination or resultants: all remaining equations are split-locator equations, chart-removal equations, rank-stratum equations, and noncontainment pivots in the free kernel coordinates and the line parameter.

We write the statement first on the finite affine patch. Fix an exact agreement value A, put $j = n - A$ and $t = A - k$, and set

$$M_A(Z) := U + ZV, \qquad U := H_{t,j}(\mathrm{Syn}(f)), \quad V := H_{t,j}(\mathrm{Syn}(g)).$$

Let $I \subseteq \{0, \dots, t-1\}$ and $J \subseteq \{0, \dots, j\}$ have the same cardinality s, and put

$$\Delta_{I,J}(Z) := \det M_A(Z)_{I,J}.$$

Let $C_J := \{0, \dots, j\} \setminus J$ and introduce free kernel coordinates $Y = (Y_c)_{c \in C_J}$. On the open set $\Delta_{I,J} \neq 0$, define the polynomial vector

$$\Lambda^{\#}_{I,J}(Z, Y) \in F[Z, Y]^{j+1}$$

by

$$\big(\Lambda^{\#}_{I,J}\big)_{C_J} := \Delta_{I,J} Y, \qquad \big(\Lambda^{\#}_{I,J}\big)_{J} := -\operatorname{adj}\big(M_A(Z)_{I,J}\big) M_A(Z)_{I,C_J} Y.$$

Thus

$$M_A(Z)_{I,*} \Lambda^{\#}_{I,J}(Z, Y) = 0.$$

If, in addition, all $(s+1) \times (s+1)$ minors of $M_A(Z)$ vanish, then every row of $M_A(Z)$ lies in the span of the pivot rows, and hence

$$M_A(Z) \Lambda^{\#}_{I,J}(Z, Y) = 0.$$

To encode splitting without dividing by the leading coefficient, use homogeneous locator coordinates. For $\lambda = (\lambda_0, \dots, \lambda_j)$ set

$$L_\lambda(X) := \sum_{b=0}^{j} \lambda_b X^b.$$

Let $Q(X) = \sum_{r=0}^{n-j} q_r X^r$ be an auxiliary quotient polynomial and let $I_{\mathrm{hsplit},j}$ be the homogeneous split incidence

$$P_D(X) = L_\lambda(X)Q(X), \qquad \lambda_j \neq 0,$$

viewed as the coefficient equations saturated by λ_j. Equivalently, its points with $\lambda_j \neq 0$ are precisely the nonzero scalar multiples of the monic locators $\prod_{x\in T}(X-x)$ with $T \subseteq D$ and $|T| = j$. For a monic locator chart $X = (E_X, \Delta_X)$, write $E_X^{\hom}$ and $\Delta_X^{\hom}$ for the usual homogenizations in the coordinates $\lambda_0, \dots, \lambda_j$.

For a finite noncontainment pivot $B_h \neq 0$, set

$$\Pi_h^{\#}(Z, Y) := \big(V\Lambda_{I,J}^{\#}(Z, Y)\big)_h.$$

The *kernel-compressed affine rank-pivot ideal* is

$$\begin{aligned}\mathcal{K}^{\mathrm{aff}}_{X,I,J,h} := \big(&\langle \mathrm{coeff}_X(P_D - L_{\Lambda_{I,J}^{\#}} Q)\rangle + \langle E_X^{\hom}(\Lambda_{I,J}^{\#})\rangle \\ &+ \langle \Delta : \Delta \in \mathcal{M}_{s+1}(M_A)\rangle\big) : \big(\Delta_{I,J}\, \Delta_X^{\hom}(\Lambda_{I,J}^{\#})\, (\Lambda_{I,J}^{\#})_j\, \Pi_h^{\#}\big)^\infty\end{aligned}$$

in the ring $F[Z, Y, (q_r)]$. Here $L_{\Lambda_{I,J}^{\#}}$ means L_λ after substituting $\lambda = \Lambda_{I,J}^{\#}(Z, Y)$.

Theorem 3.73 (kernel-compressed affine eliminant certificate)**.** *Every finite affine bad slope represented by the rank-pivot chart (X, I, J, h) is in the projection of $V(\mathcal{K}^{\mathrm{aff}}_{X,I,J,h})$ to the Z-line. Consequently, if*

$$Q_{X,I,J,h}(Z) \in \mathcal{K}^{\mathrm{aff}}_{X,I,J,h} \cap F[Z]$$

is nonzero, then the chart contributes at most

$$\deg \gcd(Q_{X,I,J,h}(Z), Z^q - Z)$$

finite slopes over F_q.

Proof. Let z be a bad slope in the chart, and let λ be a nonzero scalar multiple of the corresponding locator vector. Since the pivot minor is nonzero, put $Y_c = \lambda_c/\Delta_{I,J}(z)$ for $c \in C_J$. Cramer's rule and the pivot row equations give

$$\Lambda_{I,J}^{\#}(z, Y) = \lambda.$$

The locator is split, lies in the chart X, has nonzero leading coefficient, and has the chosen nonzero noncontainment coordinate. The larger minors vanish on the exact-rank stratum. Thus the point survives the displayed saturation and lies in the compressed incidence. Every element of the elimination ideal must vanish on its Z-coordinate, and the finite-field count is obtained by intersecting with $Z^q - Z$. □

Remark 3.74 (why the compression is useful). The original chart has j locator variables, one parameter variable, and the split equations. The compressed chart has only $j + 1 - s$ free kernel coordinates before imposing the leading-coordinate and chart opens. In the overdetermined rank-drop case this is often a large saving. In the genuinely underdetermined near-capacity case the saving may be small, but the construction still separates three causes of difficulty: many split locators, a dominant parameter projection, or a high-degree finite eliminant.

The projective and curve variants are identical after replacing the parameter ring. For projective lines, replace $M_A(Z)$ by $M_A^{\mathrm{proj}}(Z_0, Z_1)$, use homogeneous pivot minors, and take homogeneous eliminants in $F[Z_0, Z_1]$. On the infinity patch $[0 : 1]$, this is just the constant matrix V with equations $B(\ell) = 0$ and an open pivot $A_h \neq 0$. For a degree-d finite power curve, replace $M_A(Z)$ by

$$M_A^{(d)}(\Gamma) = \sum_{i=0}^{d} \Gamma^i H_{t,j}(\mathrm{Syn}(f_i))$$

and use coefficient pivots $(V_i)_h \neq 0$.

Corollary 3.75 (kernel-compressed projective and curve eliminants)**.** *The projective rank-pivot chart obtained from the compressed construction has the same parameter projection as the uncompressed rank-pivot chart on the declared opens. A nonzero homogeneous eliminant $Q(Z_0, Z_1)$ contributes at most*

$$\deg \gcd(Q(1, Z), Z^q - Z) + \mathbf{1}_{Q(0,1)=0}$$

projective parameters over F_q. For a finite degree-d curve chart, a nonzero eliminant $Q(\Gamma)$ contributes at most

$$\deg \gcd(Q(\Gamma), \Gamma^q - \Gamma)$$

parameters.

Proof. The same Cramer parametrization applies to the projective pencil and to the curve pencil. The root-count statements are the finite-field root counts for a homogeneous polynomial on $\mathbb{P}^1(F_q)$ and a univariate polynomial on $\mathbb{A}^1(F_q)$. □

Exact residual support-image lcm ledger. The compressed eliminant is a constructive way to find a small parameter polynomial. There is also an exact, definition-level residual image ledger that works for any finite co-support family, quotient or not. This is the clean object against which scanner-produced eliminants should be audited.

Let $\mathcal{T}$ be any family of co-supports $T \subseteq D$ of size $j = n - A$. For a line $f + Zg$, put

$$A_T := H_{t,j}(\mathrm{Syn}(f))\boldsymbol{\ell}_T, \qquad B_T := H_{t,j}(\mathrm{Syn}(g))\boldsymbol{\ell}_T.$$

Define the finite affine fixed-support polynomial

$$g_T^{\mathrm{aff}}(Z) := \begin{cases} \mathrm{monic}\,\gcd\{(A_T)_h + Z(B_T)_h : 0 \le h < t\}, & (A_T, B_T) \neq (0, 0), \\ 1, & (A_T, B_T) = (0, 0), \end{cases}$$

with the convention that the gcd of constants with no common root is 1. Thus g_T^{aff} has degree either 0 or 1. The contained case $(A_T, B_T) = (0, 0)$ is assigned the polynomial 1, because it contributes no support-wise bad slope. Set

$$R_{\mathcal{T}}^{\mathrm{aff}}(Z) := \mathrm{rad}\,\mathrm{lcm}_{T \in \mathcal{T}}\, g_T^{\mathrm{aff}}(Z).$$

For projective lines, define $g_T^{\mathrm{proj}}(Z_0, Z_1)$ as the monic homogeneous gcd of the linear forms

$$Z_0(A_T)_h + Z_1(B_T)_h, \qquad 0 \le h < t,$$

again setting it equal to 1 in the contained case, and put

$$R_{\mathcal{T}}^{\mathrm{proj}} := \operatorname{rad} \operatorname{lcm}_{T\in\mathcal{T}} g_T^{\mathrm{proj}}.$$

For a finite degree-d curve, use the vector polynomials

$$\sum_{i=0}^{d} \Gamma^i (V_{i,T})_h$$

and define $R_{\mathcal{T}}^{(d)}(\Gamma)$ by the same radical lcm rule, with the all-zero coefficient tuple treated as contained.

Theorem 3.76 (exact arbitrary-residual image ledger)**.** *Let $\mathcal{T}$ be any finite co-support family. The exact finite affine line parameters witnessed by $\mathcal{T}$ are the roots in F of $R_{\mathcal{T}}^{\mathrm{aff}}$. Hence their number is*

$$\deg \gcd(R_{\mathcal{T}}^{\mathrm{aff}}(Z), Z^q - Z).$$

The exact projective parameters witnessed by $\mathcal{T}$ are the projective zeros of $R_{\mathcal{T}}^{\mathrm{proj}}$, counted over F_q by

$$\deg \gcd(R_{\mathcal{T}}^{\mathrm{proj}}(1, Z), Z^q - Z) + \mathbf{1}_{R_{\mathcal{T}}^{\mathrm{proj}}(0,1)=0}.$$

For finite degree-d curves, the exact finite parameter count is

$$\deg \gcd(R_{\mathcal{T}}^{(d)}(\Gamma), \Gamma^q - \Gamma).$$

The same ledgers upper-bound no-loss CA parameters witnessed by the same co-support family.

Proof. For a fixed co-support T, the syndrome-locator recurrence says that the line point is explained on $D \setminus T$ precisely when $A_T + ZB_T = 0$. If $A_T = B_T = 0$, then both coefficient words are explained on the same support and the fixed support is contained, so it is correctly charged as 1. Otherwise the common roots of the scalar coordinates are exactly the finite bad slopes contributed by T, and this common-root set is represented by g_T^{aff}. Taking the radical lcm over $T \in \mathcal{T}$ coalesces duplicate slopes. The projective and curve statements are identical, using the homogeneous or degree-d coordinate polynomials. No-loss CA badness is stronger than existence of an explaining support with noncontainment, so the same image set is an upper ledger. □

Corollary 3.77 (closed-ball residual exact image after paid-branch removal)**.** *Fix a closed-ball agreement threshold a. For each exact agreement $A \in \{a, a+1, \dots, n\}$, let $\mathcal{T}_{\mathrm{all},A}$ be the family of co-supports of size $n - A$, let $\mathcal{T}_{\mathrm{paid},A} \subseteq \mathcal{T}_{\mathrm{all},A}$ be a certified paid support family, and put*

$$\mathcal{T}_{\mathrm{res},A} := \mathcal{T}_{\mathrm{all},A} \setminus \mathcal{T}_{\mathrm{paid},A}.$$

Let $R_{\mathrm{res},A}^{\mathrm{aff}}$ be the polynomial of Theorem 3.76 for the fixed-size family $\mathcal{T}_{\mathrm{res},A}$. Define the closed-ball residual polynomial

$$R_{\mathrm{res},\geq a}^{\mathrm{aff}}(Z) := \operatorname{rad} \operatorname{lcm}_{A=a}^{n} R_{\mathrm{res},A}^{\mathrm{aff}}(Z), \qquad R_{\mathrm{res},\geq a,F}^{\mathrm{aff}}(Z) := \gcd\big(R_{\mathrm{res},\geq a}^{\mathrm{aff}}(Z), Z^q - Z\big).$$

Then the exact finite affine residual numerator after paid-support removal is

$$\deg R_{\mathrm{res},\geq a,F}^{\mathrm{aff}}.$$

The projective and finite-curve residual numerators are obtained by the same exact-agreement radical-lcm construction, using $R^{\mathrm{proj}}_{\mathrm{res},A}$ and $R^{(d)}_{\mathrm{res},A}$ respectively, followed by the finite-field root count of Theorem 3.76. If a scanner also prints an exact paid root polynomial $P_{\mathrm{paid,F}}$ for the paid branch, the total distinct finite affine numerator is

$$\deg \operatorname{rad} \operatorname{lcm}\big(P_{\mathrm{paid,F}}, R^{\mathrm{aff}}_{\mathrm{res},\geq a,F}\big),$$

not the sum of the paid and residual degrees.

Proof. Apply Theorem 3.76 separately at each exact agreement size A, where the co-support size $j = n - A$ and the Hankel window height $t = A - k$ are fixed. The radical lcm over A is exactly union and duplicate coalescing across exact agreement levels. Intersecting with $Z^q - Z$ counts only F-rational finite slopes. The final displayed lcm performs the same duplicate coalescing between paid roots and residual roots. □

Elimination trichotomy for residual charts. The exact image lcm exists because split locators are a finite set. A compressed rank-pivot certificate is useful when it finds a much smaller polynomial without enumerating all residual supports. The following trichotomy is the scanner contract.

Proposition 3.78 (compressed-chart elimination trichotomy)**.** *For a compressed finite affine residual chart over the algebraic closure, exactly one of the following occurs:*

(i) *the saturated compressed ideal is the unit ideal, and the chart is empty;*

(ii) *the saturated compressed ideal is proper and the elimination ideal in $F[Z]$ contains a nonzero polynomial, and its squarefree generator gives a finite parameter envelope containing the exact support-image lcm of that chart;*

(iii) *the saturated compressed ideal is proper and the elimination ideal in $F[Z]$ is zero, equivalently the projection of the relaxed chart to the Z-line is Zariski dense.*

If the full homogeneous split equations are present, case (iii) *can occur only through a positive-dimensional relaxation or through a contained component not removed by the noncontainment saturation; an exact split, noncontained chart has finite parameter image.*

Proof. The first three alternatives are the standard elimination theorem for the projection of an affine variety to a line. If the full split equations are present and the leading coordinate is nonzero, the locator coordinate belongs to the finite set of scalar multiples of the finitely many degree-j locators supported on D. After projectivizing the locator scale, each fixed support contributes at most one finite noncontained slope by Lemma 3.16. Hence the exact noncontained split image is a finite subset of the parameter line. A dense projection can therefore only come from a relaxation of the split condition or from an unremoved contained component on which every parameter is explained but no noncontainment pivot should survive. □

Remark 3.79 (compressed-residual status). The singular bucket has now been reduced to two concrete certificate tasks. The first is computational-algebraic: use the kernel-compressed ideals to find small univariate or homogeneous eliminants. The second is enumerative but exact: when compression is not yet sharp enough, print the residual support-image lcm for a finite residual family and coalesce it with the paid quotient/tangent roots. Neither task proves the final aperiodic packing theorem by itself, but both replace the phrase "singular residual bucket" by checkable objects with exact root-count semantics.

3.4.5 Fallback and final certificate

The eliminant certificates above are strongest when the projection to the parameter line is finite. A failed eliminant search should not be discarded; it should be split further or recorded as a singular bucket. The following fallback is safe but usually too weak for a final 2^{-128} threshold.

Proposition 3.80 (dimension-degree fallback)**.** *Let $Y \subseteq \mathbb{A}_F^N$ be an affine variety of degree at most Δ and Krull dimension at most e. Then*

$$|Y(F)| \le \Delta |F|^e.$$

Consequently, if a saturated Hankel incidence chart for line or curve parameters has degree Δ_X and dimension e_X, then the number of bad parameters in that chart is at most $\Delta_X |F|^{e_X}$.

Proof. This is the standard affine degree point-count bound, proved by induction on dimension using generic affine hyperplane sections. The parameter image is no larger than the incidence set. □

For a threshold statement, use finitely many charts and agreement levels. Let $\mathfrak{X}_A$ be a chart cover of the declared aperiodic branch at exact agreement A. If each finite affine pivot chart has a univariate eliminant, set

$$B_{\mathrm{ap}}^{\mathrm{aff,elim}}(a) := \sum_{A=a}^{n} \sum_{X \in \mathfrak{X}_A} \sum_{h=0}^{t_A-1} \deg Q_{X,h}.$$

For projective slopes add the certified infinity contribution, and for degree-d curves replace the inner sum by the coefficient-pivot sum $\sum_{i=0}^{d} \sum_h \deg Q_{X,i,h}$. Exact root-union tables may replace these safe sums whenever duplicate parameters across charts are certified.

Definition 3.81 (aperiodic Hankel-packing certificate)**.** An aperiodic Hankel-packing certificate for a row, a line or curve family, and an agreement threshold a consists of:

(i) the tangent/common-code-line and quotient-image ledgers that have already been removed;

(ii) for each exact agreement $A \in \{a, \dots, n\}$, either a regular minor certificate Δ_A or a declaration that the exact-agreement bucket is Hankel-singular;

(iii) for every singular or near-capacity bucket, a finite constructible locator chart cover $\mathfrak{X}_A$ as in Definition 3.58;

(iv) for affine lines, a list of pivot eliminants $Q_{X,h}(Z)$, or a failed-pivot record explaining which residual locus remains unsplit;

(v) for projective slopes, the finite-patch eliminants plus the infinity-chart empty/nonempty certificate;

(vi) for curves, coefficient-pivot eliminants $Q_{X,i,h}(\Gamma)$;

(vii) a root-union table or a safe degree-sum table producing the declared aperiodic numerator $B_{\mathrm{ap}}(a)$;

(viii) if any chart uses the dimension-degree fallback, its degree, dimension, and an explicit statement that the resulting bound is only a fallback, not a sharp threshold certificate.

The verifier checks locator divisibility by P_D, chart coverage of the declared aperiodic complement, nonzero pivot tests, ideal membership of eliminants after saturation, and final root-count arithmetic.

Remark 3.82 (contributor workflow). For each row and exact agreement A, contributors should try the following in order. First test the overdetermined regular minors. If a nonzero minor is found, record its root set and remove roots already paid by tangent or quotient ledgers. If all maximal minors vanish identically, move the instance to the singular bucket. On the singular or near-capacity bucket, build locator charts, then split by affine pivots $B_h \neq 0$; for projective samplers also test the infinity chart $B = 0$, $A \neq 0$; for curves split by coefficient pivots $(V_i)_h \neq 0$. A chart is complete only when it has either a nonzero eliminant, a certified empty locus, or a named residual obstruction. Residual obstructions should be labelled as quotient-periodic, tangent/common-code-line, subfield/extension exceptional, or candidate new obstruction.

Remark 3.83 (status). Theorems 3.46, 3.59 and 3.63 and Corollary 3.62 are proved certificate theorems. They do not prove that all rows admit small aperiodic eliminants. They make the remaining MCA safe-side task concrete: after tangent and quotient branches are removed, construct small eliminants for the pivot charts, or identify the singular chart that prevents such an eliminant.

Lemma 3.84 (interleaving transfer)**.** *For every linear code $C \subseteq \mathbb{F}^n$, every $s \geq 1$ and every $\delta \in (0,1)$:*

$$\varepsilon_{\mathrm{ca}}\big(C^{\equiv s}, \delta\big) \;\geq\; \varepsilon_{\mathrm{ca}}(C, \delta), \qquad \varepsilon_{\mathrm{mca}}\big(C^{\equiv s}, \delta\big) \;\geq\; \varepsilon_{\mathrm{mca}}(C, \delta).$$

Proof. For $h \in \mathbb{F}^n$ write $h^{\Delta} := (h, \ldots, h) \in (\mathbb{F}^n)^s$. For any codeword tuple $\mathbf{c} = (c^{(1)}, \ldots, c^{(s)}) \in C^{\equiv s}$, the columns where h^{Δ} and $\mathbf{c}$ agree are $\bigcap_i \{j : c_j^{(i)} = h_j\} \subseteq \{j : c_j^{(1)} = h_j\}$, so $\Delta_s(h^{\Delta}, C^{\equiv s}) \geq \Delta(h, C)$; taking $c^{(1)} = \cdots = c^{(s)}$ gives equality: $\Delta_s(h^{\Delta}, C^{\equiv s}) = \Delta(h, C)$.

Fix (f_1, f_2) attaining $\varepsilon_{\mathrm{ca}}(C, \delta)$ (resp. $\varepsilon_{\mathrm{mca}}$) and consider the pair $(f_1^{\Delta}, f_2^{\Delta})$ for $C^{\equiv s}$, with the slope γ acting row-wise, so $f_1^{\Delta} + \gamma f_2^{\Delta} = (f_1 + \gamma f_2)^{\Delta}$. By the displayed equality, $\Delta_s\big((f_1 + \gamma f_2)^{\Delta}, C^{\equiv s}\big) \leq \delta$ iff $\Delta(f_1 + \gamma f_2, C) \leq \delta$. For the correlated condition, the pair $(f_1^{\Delta}, f_2^{\Delta})$ viewed as a $2s$-row interleaved word has rows f_1 (s times) and f_2 (s times); by the same column argument its distance to $(C^{\equiv s})^{\equiv 2} = C^{\equiv 2s}$ equals $\Delta_2((f_1, f_2), C^{\equiv 2})$. For the MCA condition, a common set S explains f_1^{Δ} and f_2^{Δ} iff it explains f_1 and f_2. Hence γ is CA-bad (resp. MCA-bad) for $(f_1^{\Delta}, f_2^{\Delta})$ iff it is so for (f_1, f_2), and the maxima over pairs can only grow. □

4 The main theorem

Theorem 4.1 (universal cap)**.** *Let $\mathbb{B} \subseteq \mathbb{F}$ be finite fields, let $q = |\mathbb{F}|$, let $D \subseteq \mathbb{B}^{\times}$ be a multiplicative coset of order n, and let k have $\rho = k/n \in (0,1)$; set $C := \mathrm{RS}[\mathbb{F}, D, k]$. Let $N \mid n$ satisfy $a := n/N \mid k$ and $(1-\rho)N \geq 3$. If*

$$\binom{N}{\rho N + 2} \geq |\mathbb{B}|\Big(\frac{q}{k} + 1\Big), \tag{1}$$

then, writing

$$\delta_N := 1 - \rho - \frac{2}{N},$$

we have

$$\varepsilon_{\mathrm{ca}}(C, \delta) > \frac{1}{2k}\Big(1 - \frac{n}{q}\Big)$$

for every integer-admissible radius

$$\delta_N \le \delta \le 1 - \rho - \frac{1}{n}.$$

In particular,

$$\varepsilon_{\mathrm{mca}}(C, \delta) > \frac{1}{2k}\Big(1 - \frac{n}{q}\Big)$$

for every

$$\delta_N \le \delta < 1 - \rho.$$

Consequently,

$$\delta_C^*(\varepsilon^*) \le 1 - \rho - \frac{2}{N}$$

for every

$$\varepsilon^* \le \frac{1}{2k}\Big(1 - \frac{n}{q}\Big).$$

Proof. Note first that $\rho N = k/a \in \mathbb{Z}$, $\ell_2 = \rho N + 2 \le N - 1$ by $(1-\rho)N \ge 3$, and $\delta_N = 1 - \rho - 2/N > 0$ for the same reason. Also $k + 1 \le n$.

Fix

$$\delta \in [\delta_N,\, 1 - \rho - \tfrac{1}{n}].$$

Then

$$\lfloor \delta n \rfloor \le n - k - 1,$$

so Theorem 2.6 is applicable.

Suppose toward a contradiction that

$$\varepsilon_{\mathrm{ca}}(C, \delta) \le \frac{1}{2k}\Big(1 - \frac{n}{q}\Big).$$

Apply Theorem 2.6 with $\eta = \frac{1}{2}$. Its hypothesis holds, so

$$\Lambda(C^+, \delta) \le \Big\lceil 2q \cdot \varepsilon_{\mathrm{ca}}(C, \delta) \Big\rceil \le \Big\lceil \frac{q-n}{k} \Big\rceil < \frac{q-n}{k} + 1 \le \frac{q}{k} + 1.$$

On the other hand, by Lemma 3.1(ii) and monotonicity of lists in the radius, since $\delta \ge \delta_N$,

$$\Lambda(C^+, \delta) \ge \big|\Lambda(C^+, \delta_N, u_z)\big| \ge \binom{N}{\ell_2} \Big/ |\mathbb{B}| \ge \frac{q}{k} + 1$$

by (1). This is a contradiction. Hence

$$\varepsilon_{\mathrm{ca}}(C, \delta) > \frac{1}{2k}\Big(1 - \frac{n}{q}\Big)$$

for every $\delta \in [\delta_N, 1 - \rho - 1/n]$.

In particular, the above holds at $\delta = \delta_N$. By Fact 2.4,

$$\varepsilon_{\mathrm{mca}}(C, \delta_N) \ge \varepsilon_{\mathrm{ca}}(C, \delta_N) > \frac{1}{2k}\Big(1 - \frac{n}{q}\Big).$$

By support-wise MCA monotonicity, Lemma 2.5, the same lower bound for $\varepsilon_{\mathrm{mca}}(C,\delta)$ holds for every

$$\delta \in [\delta_N, 1-\rho).$$

Thus no sub-capacity radius

$$\delta \geq 1-\rho-\frac{2}{N}$$

is ε^*-admissible whenever

$$\varepsilon^* \leq \frac{1}{2k}\Big(1-\frac{n}{q}\Big),$$

which proves the claimed bound on $\delta_C^*(\varepsilon^*)$. □

Remark 4.2 (constants). Choosing $\eta = 1-\theta$ in Theorem 2.6 improves the conclusion to $\varepsilon_{\mathrm{ca}} > \frac{1-\theta}{k}(1-\frac{n}{q})$ at the price of strengthening (1) to $\binom{N}{\rho N+2} \geq |\mathbb{B}|\big(\frac{q}{\theta k}+1\big)$; given the hundreds of bits of slack in (1) throughout the challenge envelope (Table 1), the error constant is effectively $(1-o(1))/k$. We fix $\eta = \frac{1}{2}$ for cleanliness.

Remark 4.3 (scope). No smoothness is assumed beyond the existence of a single divisor $N \mid n$ with $a = n/N \mid k$ and (1). The theorem applies verbatim to 2-adic, mixed-radix, or any other FFT-friendly multiplicative domain, and — via the subfield pigeonhole — to any extension field $\mathbb{F} \supseteq \mathbb{B}$ with $D \subseteq \mathbb{B}$.

The deep-point conversion is an integer-radius theorem with the condition $\lfloor \delta n \rfloor \leq n-k-1$. Therefore the no-loss CA lower bound is stated for

$$\delta \leq 1-\rho-\frac{1}{n}.$$

The MCA threshold cap is stronger in the range needed for the Proximity Prize: once the support-wise MCA lower bound is proved at

$$\delta_N = 1-\rho-\frac{2}{N},$$

monotonicity of support-wise MCA extends it to every larger $\delta < 1-\rho$.

5 Consequences for the grand MCA challenge

Corollary 5.1 (universal field-size cap for the challenge envelope)**.** *Let $\rho \in \{\frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \frac{1}{16}\}$ and set*

$$N_\rho := 1024 \quad (\rho \in \{\tfrac{1}{2}, \tfrac{1}{4}, \tfrac{1}{8}\}), \qquad N_{1/16} := 2048.$$

Let $\mathbb{F}$ be any finite field with $q < 2^{256}$, let $\mathbb{B} \subseteq \mathbb{F}$ be any subfield, let $D \subseteq \mathbb{B}^\times$ be a multiplicative coset of order n with $N_\rho \mid n$, and let $k = \rho n \leq 2^{40}$. Then $C = \mathrm{RS}[\mathbb{F}, D, k]$ satisfies

$$\varepsilon_{\mathrm{mca}}(C,\delta) > \frac{1}{2k}\Big(1-\frac{n}{q}\Big) \geq 2^{-86} \gg 2^{-128}$$

for every

$$\delta \in \Big[1-\rho-\frac{2}{N_\rho}, 1-\rho\Big).$$

Moreover, the same lower bound holds for $\varepsilon_{\mathrm{ca}}(C,\delta)$ throughout the integer-admissible subinterval

$$\delta \in \Big[1-\rho-\frac{2}{N_\rho},\, 1-\rho-\frac{1}{n}\Big].$$

If $q \geq 2n$, the lower bound improves to $\geq 2^{-42}$. In particular,

$$\delta^*_C(2^{-128}) \;\leq\; 1-\rho-2^{-9} \quad (\rho \in \{\tfrac{1}{2},\tfrac{1}{4},\tfrac{1}{8}\}), \qquad \delta^*_C(2^{-128}) \;\leq\; 1-\rho-2^{-10} \quad (\rho = \tfrac{1}{16}).$$

Proof. We verify the hypotheses of Theorem 4.1 with $N = N_\rho$. Divisibility: $a = n/N_\rho$ divides k because $k/a = \rho N_\rho \in \mathbb{Z}$ (equal to $512, 256, 128, 128$ respectively); $(1-\rho)N_\rho \geq 128 \geq 3$. For (1), the entropy bound $\binom{N}{\ell} \geq 2^{N\mathrm{H}_2(\ell/N)}/(N+1)$ together with the conservative evaluations of H_2 recorded in Table 1 gives $\log_2 \binom{N_\rho}{\rho N_\rho+2} \geq 552$ in all four cases, while

$$|\mathbb{B}|\Big(\frac{q}{k}+1\Big) \;\leq\; q\Big(\frac{q}{k}+1\Big) \;\leq\; q^2+q \;<\; 2^{512}+2^{256} \;<\; 2^{513} \;\leq\; 2^{552},$$

so (1) holds with at least 39 bits to spare (and over 500 bits at rate $\frac{1}{2}$). Since $D \subseteq \mathbb{F}^\times$, we have $n \leq q-1$; for fixed n, the function $1-n/q$ is minimized at $q = n+1$, hence $1-n/q \geq 1/(n+1)$. Also $n = k/\rho \leq 16k \leq 2^{44}$. Therefore

$$\frac{1}{2k}\Big(1-\frac{n}{q}\Big) \geq \frac{1}{2k(n+1)} > 2^{-86}.$$

If $q \geq 2n$ then $1-n/q \geq 1/2$, giving the sharper $1/(4k) \geq 2^{-42}$. The gap is $2/N_\rho = 2^{-9}$ for $\rho \in \{1/2, 1/4, 1/8\}$ and $2/N_\rho = 2^{-10}$ for $\rho = 1/16$. Theorem 4.1 gives the MCA lower bound throughout the stated sub-capacity interval, and gives the CA lower bound on the integer-admissible subinterval ending at $1-\rho-1/n$. □

ρ	N_ρ	$\ell_2 = \rho N_\rho + 2$	$\mathrm{H}_2(\ell_2/N_\rho) \geq$	$\log_2 \binom{N_\rho}{\ell_2} \geq$	needed	gap $2/N_\rho$
1/2	1024	514	0.99998	1013	513	2^{-9}
1/4	1024	258	0.8143	823	513	2^{-9}
1/8	1024	130	0.5490	552	513	2^{-9}
1/16	2048	130	0.34105	687	513	2^{-10}

Table 1: Numerics for Corollary 5.1: $\log_2 \binom{N}{\ell} \geq N\mathrm{H}_2(\ell/N) - \log_2(N+1)$, rounded down. The "needed" column is the universal upper bound $\log_2\big(|\mathbb{B}|(q/k+1)\big) < 513$ valid throughout the envelope $q < 2^{256}$, $k \geq 1$.

Remark 5.2 (hypothesis $N_\rho \mid n$). For the 2-adic domains of the challenge this holds whenever $n \geq 2^{10}$ (resp. 2^{11}), i.e. for all but very small instances. For mixed-radix smooth n, replace N_ρ by any divisor $N \mid n$ of comparable size with $\rho N \in \mathbb{Z}$; the entropy count in Table 1 is insensitive to the exact value of N at fixed $\log_2 N$. The divisibility-free fiber lemma removes the integrality requirement $\rho N \in \mathbb{Z}$ altogether: Lemma 7.2, Corollary 7.6 caps every mixed-radix row possessing *any* divisor in the printed dyadic windows.

Remark 5.3 (coarse grand-challenge caps and the remaining frontier). Combining Corollary 5.1 with the grand-challenge cap of [Cho26b, Prop. "Grand-challenge cap"] gives the following coarse uniform negative caps for smooth multiplicative domains:

$$\delta_C^*(2^{-128}) \;\le\; \begin{cases} 1-\rho-\frac{1}{64}, & \mathbb{F} \text{ prime},\ q \lesssim 2^{150}, \\ 1-\rho-2^{-9}, & \rho \in \{\frac{1}{2},\frac{1}{4},\frac{1}{8}\},\ \text{every } \mathbb{F},\ q < 2^{256}, \\ 1-\rho-2^{-10}, & \rho = \frac{1}{16},\ \text{every } \mathbb{F},\ q < 2^{256}. \end{cases}$$

These inequalities are not the final frontier studied in the remainder of this paper. They are first-tier caps: robust, uniform, and useful for placing the grand challenge beyond the Johnson region, but far from the capacity edge at the deployed parameters. The rest of the paper has a different purpose. It refines these coarse caps into the identity-scale and flexible-budget frontier certificates, proves the structural ledgers that rule out the already-paid tangent, quotient, extension-line, and split top-chart strata, and isolates the remaining safe-side inputs as quotient/prefix flatness, the base-field-normalized split-pencil census, and primitive shift-pair control.

Thus the quoted caps should be read as the broad negative band available before the frontier analysis, not as a stopping point. Together with the positive results collected in [ABF26, Table 1], they give the initial sandwich; the later frontier package states the sharpened active sandwich. Circle rows — line rounds and bivariate circle codes — have analogous coarse caps; see Corollary 7.12 and Remark 7.15.

Corollary 5.4 (deployed parameters: KoalaBear sextic)**.** *Let $\mathbb{B} = \mathbb{F}_p$ with $p = 2^{31} - 2^{24} + 1$, let $\mathbb{F} = \mathbb{F}_{p^6}$ (so $q = p^6 \approx 2^{185.93}$), let $D \subseteq \mathbb{B}^\times$ be the subgroup of order $n = 2^{21}$, and let $k = 2^{20}$, $\rho = \frac{1}{2}$, as in [ABF26, §6.3]. Then*

$$\varepsilon_{\mathrm{mca}}\big(\mathrm{RS}[\mathbb{F}_{p^6}, D, 2^{20}],\, \delta\big) > \big(1 - 2^{-164}\big)\, 2^{-21} > 2^{-22}$$

for every

$$\delta \in \Big[\frac{1}{2} - 2^{-7},\, \frac{1}{2}\Big).$$

Moreover, the same lower bound holds for $\varepsilon_{\mathrm{ca}}$ throughout the integer-admissible subinterval

$$\delta \in \Big[\frac{1}{2} - 2^{-7},\, \frac{1}{2} - 2^{-21}\Big).$$

Proof. Apply Theorem 4.1 with $N = 256$, $a = 2^{13} \mid k$, $(1-\rho)N = 128 \ge 3$, gap $2/N = 2^{-7}$. Hypothesis (1): $\log_2 \binom{256}{130} \ge 256 \cdot \mathrm{H}_2(130/256) - \log_2 257 \ge 256 \cdot 0.99982 - 8.01 > 247$, while $|\mathbb{B}|(q/k+1) < 2^{31}\,(2^{166}+1) < 2^{198}$. The error bound is $\frac{1}{2k}(1 - n/q)$ with $n/q = 2^{21}/p^6 < 2^{-164}$. Theorem 4.1 gives the stated MCA bound on $[\frac{1}{2} - 2^{-7}, \frac{1}{2})$ and the stated CA bound on $[\frac{1}{2} - 2^{-7}, \frac{1}{2} - 2^{-21})$. □

Corollary 5.5 (all interleaved rows of [ABF26, Tables 2–3])**.** *In the setting of Corollary 5.4, for each $s = 2^j$ with $0 \le j \le 12$ let $C_s := \mathrm{RS}[\mathbb{F}_{p^6}, D_s, k_s]$ with $|D_s| = n_s = 2^{21}/s$ and $k_s = 2^{20}/s$ (rate $\frac{1}{2}$, $s\, n_s = 2^{21}$ as deployed). Then for every $\delta \in [\frac{1}{2} - 2^{-7}, \frac{1}{2})$,*

$$\varepsilon_{\mathrm{mca}}\big(C_s^{\equiv s}, \delta\big) > \frac{s}{2^{21}}\big(1 - 2^{-164}\big) \ge 2^{-21}\big(1 - 2^{-164}\big).$$

Moreover, the same lower bound holds for $\varepsilon_{\mathrm{ca}}(C_s^{\equiv s}, \delta)$ on the integer-admissible subinterval

$$\delta \in \Big[\frac{1}{2} - 2^{-7},\ \frac{1}{2} - \frac{1}{n_s}\Big).$$

Every row of the deployed parameter family therefore fails the 2^{-128} target at gap 2^{-7} below capacity.

Proof. For each $s \le 2^{12}$: $256 \mid n_s$ (as $n_s \ge 2^9$), $a_s = n_s/256 = 2^{13}/s$ divides $k_s = 2^{20}/s$, and $(1-\rho)\cdot 256 = 128 \ge 3$. Hypothesis (1) for C_s: $|\mathbb{B}|(q/k_s + 1) < 2^{31}(2^{166}\, s + 1) < 2^{198}\, s \le 2^{210} \le 2^{247} \le \binom{256}{130}$. Theorem 4.1 gives the MCA lower bound for C_s throughout $[\frac{1}{2} - 2^{-7}, \frac{1}{2})$, and the CA lower bound throughout $[\frac{1}{2} - 2^{-7}, \frac{1}{2} - 1/n_s]$. Lemma 3.84 transfers the corresponding lower bounds to the s-interleaved code. □

Proposition 5.6 (independent slacked variant via Theorem 2.16)**.** *In the setting of Theorem 4.1, assume instead $(1-\rho)N \ge 2$, $n \ge 3N$, and*

$$\binom{N}{\rho N + 1} \ \ge\ |\mathbb{B}| \cdot q.$$

Then

$$\varepsilon_{\mathrm{ca}}\Big(C,\ 1 - \rho - \frac{1}{N} + \frac{2}{n},\ 1 - \rho - \frac{1}{n}\Big) \ \ge\ \frac{1}{2n}.$$

In particular, with $N = N_\rho$ as in Corollary 5.1 and $n \ge 3N_\rho$, every challenge-envelope instantiation satisfies this slacked CA lower bound, and $\frac{1}{2n} \ge 2^{-45} \gg 2^{-128}$.

Proof. Set $\delta_0 := 1 - \rho - \frac{1}{N}$. By $(1-\rho)N \ge 2$, $\delta_0 \in (0, 1-\rho)$. By $n \ge 3N$, we also have $\delta_0 + 2/n \le 1 - \rho - 1/n$, so the parameters in Theorem 2.16 are admissible. By Lemma 3.1(i), some u_z carries at least $\binom{N}{\rho N+1}/|\mathbb{B}| \ge q$ codewords of C within radius δ_0, so $\Lambda(C, \delta_0) \ge q$. The contrapositive of Theorem 2.16 gives

$$\varepsilon_{\mathrm{ca}}\Big(C, \delta_0 + \frac{2}{n}, 1 - \rho - \frac{1}{n}\Big) \ge \frac{1}{2n}.$$

For the numerical claim, the worst case of $\log_2 \binom{N_\rho}{\rho N_\rho + 1}$ over the four rates is at $\rho = 1/8$, where the entropy bound used in Table 1 certifies $\log_2 \binom{1024}{129} \ge 1024\, \mathrm{H}_2(129/1024) - \log_2 1025 \ge 549$ (the exact value is ≈ 554.7); hence $\binom{N_\rho}{\rho N_\rho + 1} \ge 2^{549} > 2^{512} > |\mathbb{B}|q$ throughout $q < 2^{256}$. Finally $n = k/\rho \le 16k \le 2^{44}$, so $1/(2n) \ge 2^{-45}$. □

Remark 5.7. Proposition 5.6 is logically independent of Theorem 2.6 and yields a weaker error ($\frac{1}{2n}$ versus $\frac{1}{2k}$) in the *maximally slacked* form of correlated agreement — the internal radius sits at capacity minus $\frac{1}{n}$ — which is the weakest failure mode one can certify. That even this fails by 83 orders of magnitude (in bits) above the 2^{-128} target gives a separate fallback route through the imported BCHKS/ABF theorem.

6 Subfield confinement and the shape of certifying lines

The deployed instantiations run the protocol over an extension $\mathbb{F}/\mathbb{B}$ while keeping the domain in the base: $D \subseteq \mathbb{B}$. The subfield-confinement theorem of [Cho26b, Thm. "Subfield confinement"] pins down where bad slopes can live for lines whose words are base-rational: all support-wise

MCA-bad slopes lie in $\mathbb{B}$. The following lemma refines this by tracking two radii, so that it also covers the proximity-loss form of correlated agreement (Definition 2.1); part (b) is the statement of [Cho26b], reproved here for completeness in the present notation.

Lemma 6.1 (subfield confinement; refines [Cho26b])**.** *Let $\mathbb{B} \subseteq \mathbb{F}$ be finite fields, $D \subseteq \mathbb{B}^\times$, $C = \mathrm{RS}[\mathbb{F}, D, k]$, and let $f, g \in \mathbb{B}^n \subseteq \mathbb{F}^n$ be $\mathbb{B}$-valued words. Then:*

(a) *for any $0 \le \delta_{\mathrm{fld}} \le \delta_{\mathrm{int}} < 1$, every CA-bad slope for (f, g) at radii $(\delta_{\mathrm{fld}}, \delta_{\mathrm{int}})$ lies in $\mathbb{B}$;*

(b) *for any $\delta \in (0, 1)$, every MCA-bad slope for (f, g) at radius δ lies in $\mathbb{B}$.*

Consequently the bad-slope set of any $\mathbb{B}$-valued pair has density at most $|\mathbb{B}|/q$ in $\mathbb{F}$. The same holds for pairs $(\lambda f, \lambda g)$ with $\lambda \in \mathbb{F}^\times$ and f, g $\mathbb{B}$-valued, by $\mathbb{F}$-linearity of C.

Proof. Fix a $\mathbb{B}$-basis $1 = \beta_0, \beta_1, \ldots, \beta_{m-1}$ of $\mathbb{F}$. Every word $w \in \mathbb{F}^n$ decomposes uniquely as $w = \sum_i \beta_i w_i$ with $w_i \in \mathbb{B}^n$, and every codeword $c \in \mathrm{RS}[\mathbb{F}, D, k]$ decomposes as $c = \sum_i \beta_i c_i$ with $c_i \in \mathrm{RS}[\mathbb{B}, D, k]$: write the coefficient vector of the underlying polynomial in the basis and evaluate at $D \subseteq \mathbb{B}$.

Let $z \in \mathbb{F} \setminus \mathbb{B}$, say $z = \sum_i z_i \beta_i$ with $z_{i^*} \neq 0$ for some $i^* \ge 1$. We first record the support-wise consequence. Suppose that $c \in C$ agrees with $f + zg$ on some set $S \subseteq D$. Since $f_j, g_j \in \mathbb{B}$, the basis components of $(f + zg)_j$ are $f_j + z_0 g_j$ (component 0) and $z_i g_j$ (components $i \ge 1$). Agreement on S holds componentwise, so on S:

$$f + z_0 g = c_0, \qquad z_{i^*}\, g = c_{i^*}.$$

Hence g agrees on S with $z_{i^*}^{-1} c_{i^*} \in C$, and then $f = (f + z_0 g) - z_0 g$ agrees on S with $c_0 - z_0 z_{i^*}^{-1} c_{i^*} \in C$. Thus every support on which the line point $f + zg$ is explained by C also explains the pair (f, g) on that same support.

For CA, if $\Delta(f + zg, C) \le \delta_{\mathrm{fld}}$, choose such an S with $|S| \ge (1 - \delta_{\mathrm{fld}})n$; the preceding paragraph gives $\Delta_2((f, g), C^{=2}) \le \delta_{\mathrm{fld}} \le \delta_{\mathrm{int}}$, so z is not CA-bad. For MCA, the preceding paragraph rules out every possible bad support S in the support-wise definition. Hence z is not MCA-bad either. □

Corollary 6.2 (certifying lines over extensions are genuinely $\mathbb{F}$-valued)**.** *At the parameters of Corollary 5.4, $|\mathbb{B}|/q = p^{-5} < 2^{-154}$. Hence any pair (f, g) whose CA- or MCA-bad-slope density at some radius*

$$\delta \in \left[\frac{1}{2} - 2^{-7}, \frac{1}{2}\right)$$

exceeds 2^{-154} contains a word that is not $\mathbb{B}$-valued, even after removing a common scalar factor. In particular, any pair attaining the MCA lower bound $> 2^{-22}$ of Corollary 5.4 is genuinely $\mathbb{F}$-valued. On the integer-admissible subinterval

$$\delta \in \left[\frac{1}{2} - 2^{-7}, \frac{1}{2} - 2^{-21}\right),$$

the same statement applies to pairs attaining the CA lower bound. All explicit failure witnesses currently known on these domains (the locator lines of [Cho26a, Cho26b] and the fiber words of Lemma 3.1) are $\mathbb{B}$-rational and are therefore inert here: the failure certified by Corollary 5.4 is fiber-borne and, at present, non-constructive.

Proof. Immediate from Lemma 6.1 and Corollary 5.4: a $\mathbb{B}$-valued pair, up to common scalar, has CA- and MCA-bad-slope density at most $|\mathbb{B}|/q = p^{-5} < 2^{-154} < 2^{-22}$. Thus any pair witnessing the deployed lower bound cannot be of this form. □

Remark 6.3. Lemma 6.1 and Corollary 6.2 sharpen the division of labor between the two grand challenges of [ABF26] over deployed extensions: the *list* side does not deflate ($\mathrm{RS}[\mathbb{B}, D, k] \subseteq \mathrm{RS}[\mathbb{F}, D, k]$, so all base-field list lower bounds, including Lemma 3.1, persist verbatim), while every sub-2^{-128} *MCA* attack by explicit base-rational witnesses is killed by confinement — and yet MCA still fails at 2^{-22}. The failure can also be made constructive once a large list is supplied: Corollaries 2.12 and 2.13 allow one to choose a simple pole in $\mathbb{F} \setminus \mathbb{B}$ and print its value-bucket table. The deflation and decoupling corollaries of [Cho26b] record the same division at the no-loss radius, using the same corrected rounding $p^{-5} = 2^{-154.94\ldots} < 2^{-154}$ as in Corollary 6.2. Exhibiting deployed-size bucket tables for such lines is the explicit-witness question Question 22.1.

7 Beyond multiplicative smoothness: map-smooth domains, Chebyshev fibers, and the circle-code cap

The fiber lemma Lemma 3.1 uses exactly two features of the power map $x \mapsto x^a$ on a multiplicative coset: its fibers inside the domain are complete and of uniform size a, and the locators $\prod_{b \in A}(X^a - b)$ are *lacunary* — expanded in powers of X^a, their two top terms occupy degrees $a\ell$ and $a(\ell-1)$, leaving everything from degree $a(\ell-2)$ downward to the listed codeword. This section isolates exactly these two features. Lemma 7.2 proves the fiber lemma for an arbitrary polynomial folding map φ whose fibers inside the domain are complete and uniform, and simultaneously removes the divisibility hypothesis $a \mid k$ by rounding the locator length; Theorem 7.4 is the corresponding universal cap, containing Theorem 4.1 as the special case $\varphi = X^a$ with $a \mid k$. We then verify both features for the deployed isogeny-smooth family, the circle-group domains of Circle STARKs over $p = 2^{31} - 1$ [HLP24]: twin-coset x-coordinate domains carry exact Chebyshev fibers at every dyadic scale (Lemma 7.8), and over any challenge field containing i the bivariate circle code is diagonally equivalent to a Reed–Solomon code on a torus twin coset, an equivalence preserving lists and both correlated-agreement errors exactly (Lemmas 7.10 and 7.11). The resulting universal caps (Corollary 7.12) and deployed instantiations (Corollary 7.13) settle the circle-group case of the isogeny-smooth problem posed below; what remains of that problem is isolated, with a precise obstruction, in Remark 7.16 and restated as the genus-one/circle transport question Question 22.3.

7.1 Map-smooth domains and a divisibility-free fiber lemma

Definition 7.1 (map-smooth domain)**.** Let $\mathbb{B} \subseteq \mathbb{F}$ be finite fields, let $\varphi \in \mathbb{B}[X]$ have degree $a \geq 1$, and let $D \subseteq \mathbb{B}$ be a finite set of size n. The set D is (φ, a)-*smooth* if every nonempty fiber $\varphi^{-1}(w) \cap D$, $w \in \varphi(D)$, has exactly a elements. In that case $Q := \varphi(D) \subseteq \mathbb{B}$ has size $N := n/a$, and D is the disjoint union of the N fibers $S_b := \{x \in D : \varphi(x) = b\}$, $b \in Q$.

Two remarks on scope. First, a multiplicative coset $D \subseteq \mathbb{B}^\times$ of order n is (X^a, a)-smooth for every $a \mid n$: the fibers are the cosets of the order-a subgroup contained in D's underlying group. Second, nothing is assumed about the roots of $\varphi - w$ outside D: the definition constrains fibers inside D only, so separability or squarefreeness of $\varphi - w$ plays no role below.

Lemma 7.2 (map-smooth fiber lemma, divisibility-free)**.** *Let $\mathbb{B} \subseteq \mathbb{F}$ and let $D \subseteq \mathbb{B}$ be a (φ, a)-smooth domain of size n, with $Q = \varphi(D)$ and $N = n/a$ as above. Let $1 \le k < n$ and $\rho := k/n$. Put*

$$\ell_1 := \Big\lfloor \frac{k-1}{a} \Big\rfloor + 2, \qquad \ell_2 := \Big\lfloor \frac{k}{a} \Big\rfloor + 2, \qquad A_1 := a\ell_1, \qquad A_2 := a\ell_2,$$

so that $k + a \le A_1 \le k + 2a - 1$ and $k + a + 1 \le A_2 \le k + 2a$, with $A_1 = k + a$ and $A_2 = k + 2a$ when $a \mid k$.

(i) *If $\ell_1 \le N - 1$, there exists $z \in \mathbb{B}$ such that, with $u_z := \big(\varphi(x)^{\ell_1} + z\,\varphi(x)^{\ell_1 - 1}\big)_{x \in D} \in \mathbb{B}^n$,*

$$\big|\Lambda\big(\mathrm{RS}[\mathbb{F}, D, k],\, 1 - \tfrac{A_1}{n},\, u_z\big)\big| \;\ge\; \binom{N}{\ell_1} \Big/ |\mathbb{B}|.$$

(ii) *If $\ell_2 \le N - 1$, there exists $z \in \mathbb{B}$ such that, with $u_z := \big(\varphi(x)^{\ell_2} + z\,\varphi(x)^{\ell_2 - 1}\big)_{x \in D} \in \mathbb{B}^n$,*

$$\big|\Lambda\big(\mathrm{RS}[\mathbb{F}, D, k+1],\, 1 - \tfrac{A_2}{n},\, u_z\big)\big| \;\ge\; \binom{N}{\ell_2} \Big/ |\mathbb{B}|.$$

In both cases u_z is $\mathbb{B}$-valued, and the listed codewords lie in $\mathrm{RS}[\mathbb{B}, D, a(\ell_1 - 2) + 1] \subseteq \mathrm{RS}[\mathbb{B}, D, k]$ (resp. $\mathrm{RS}[\mathbb{B}, D, a(\ell_2 - 2) + 1] \subseteq \mathrm{RS}[\mathbb{B}, D, k + 1]$).

Proof. First the bracketing of A_1, A_2: writing $k = a\lfloor k/a \rfloor + s$ with $0 \le s \le a - 1$ gives $A_2 = k - s + 2a \in [k + a + 1, k + 2a]$; writing $k - 1 = a\lfloor (k-1)/a \rfloor + s_1$ with $0 \le s_1 \le a - 1$ gives $A_1 = k - 1 - s_1 + 2a \in [k + a, k + 2a - 1]$; and $a \mid k$ forces $s = 0$ and $s_1 = a - 1$.

We prove (ii); part (i) is the identical computation one degree rung lower, with $k+1$ replaced by k throughout. Let $A \subseteq Q$ with $|A| = \ell_2$, and put

$$P_A(Y) := \prod_{b \in A} (Y - b) = Y^{\ell_2} - e_1(A)\, Y^{\ell_2 - 1} + R_A(Y), \qquad R_A(Y) := \sum_{j=2}^{\ell_2} (-1)^j e_j(A)\, Y^{\ell_2 - j},$$

so $\deg_Y R_A \le \ell_2 - 2$, and set $L_A(X) := P_A(\varphi(X)) \in \mathbb{B}[X]$. The polynomial L_A vanishes on

$$S_A := \varphi^{-1}(A) \cap D = \bigsqcup_{b \in A} S_b, \qquad |S_A| = a\ell_2 = A_2,$$

since the fibers S_b are complete, pairwise disjoint, and of size exactly a.

Set $z_A := -e_1(A) \in \mathbb{B}$ and $c_A := \big(-R_A(\varphi(x))\big)_{x \in D}$. Since $\deg_X R_A(\varphi(X)) \le a(\ell_2 - 2) = A_2 - 2a \le k$, we have $c_A \in \mathrm{RS}[\mathbb{B}, D, a(\ell_2 - 2) + 1] \subseteq \mathrm{RS}[\mathbb{B}, D, k + 1] \subseteq \mathrm{RS}[\mathbb{F}, D, k + 1]$. For every $x \in S_A$,

$$0 = L_A(x) = \varphi(x)^{\ell_2} + z_A\, \varphi(x)^{\ell_2 - 1} + R_A(\varphi(x)), \qquad \text{i.e.} \qquad u_{z_A}(x) = c_A(x),$$

so u_{z_A} and c_A agree on at least A_2 points of D and $\Delta(u_{z_A}, c_A) \le 1 - A_2/n$.

The assignment $A \mapsto c_A$ is injective among subsets sharing a common value of z_A. Indeed, suppose $z_A = z_{A'}$ and $c_A = c_{A'}$. Then $R_A(\varphi(X))$ and $R_{A'}(\varphi(X))$ are polynomials of degree $\le k < n$ agreeing at all n points of D, hence equal as polynomials. Composition with the nonconstant φ is injective on $\mathbb{F}[Y]$: if $R := R_A - R_{A'} \ne 0$ has degree $d \ge 1$, then $R(\varphi(X))$ has degree $ad \ge 1$ with leading coefficient $\mathrm{lc}(R)\,\mathrm{lc}(\varphi)^d \ne 0$, and if $d = 0$ then $R(\varphi(X)) = R$ is a

nonzero constant; either way $R(\varphi(X)) \neq 0$. Hence $R_A = R_{A'}$, so $P_A = Y^{\ell_2} + z_A Y^{\ell_2-1} + R_A(Y) = P_{A'}$, and the root multisets give $A = A'$.

Finally, $z_A = -e_1(A)$ lies in $\mathbb{B}$ because $A \subseteq Q \subseteq \mathbb{B}$, so it takes at most $|\mathbb{B}|$ values as A ranges over the $\binom{N}{\ell_2}$ subsets of Q of size ℓ_2 (here $\ell_2 \leq N-1 \leq N$ is used to have this many subsets). Some value $z \in \mathbb{B}$ is attained by at least $\binom{N}{\ell_2}/|\mathbb{B}|$ subsets, and the corresponding codewords c_A are pairwise distinct elements of $\Lambda(\mathrm{RS}[\mathbb{F}, D, k+1], 1 - A_2/n, u_z)$. □

Remark 7.3 (the divisibility hypothesis is removed). With $\varphi = X^a$, D a multiplicative coset, and $a \mid k$, Lemma 7.2 is Lemma 3.1 verbatim: $A_1 = k + a$, $A_2 = k + 2a$, and u_z has the printed monomial shape. The rounding $\ell_2 = \lfloor k/a \rfloor + 2$ removes the hypothesis $a \mid k$ at the cost of a certified agreement $A_2 \in [k + a + 1, k + 2a]$ rather than exactly $k + 2a$; the certified radius therefore satisfies

$$1 - \rho - \frac{2}{N} \;\leq\; 1 - \frac{A_2}{n} \;\leq\; 1 - \rho - \frac{1}{N} - \frac{1}{n},$$

so the listed ball always sits at least one quotient step $\approx 1/N$ below capacity. This matters precisely where $a \mid k$ genuinely fails: odd-dimension circle codes (Corollary 7.12) and mixed-radix multiplicative rows without an integral ρN (Corollary 7.6). Remark 3.2 already used slack two to move the divisibility demand from $k + 1$ onto k; the present rounding removes it altogether.

Theorem 7.4 (universal cap for map-smooth domains)**.** *Let $\mathbb{B} \subseteq \mathbb{F}$ be finite fields, $q := |\mathbb{F}|$, and let $D \subseteq \mathbb{B}$ be a (φ, a)-smooth domain of size $n < q$. Let $1 \leq k < n$, $\rho := k/n$, $C := \mathrm{RS}[\mathbb{F}, D, k]$, and let $N = n/a$, $\ell_2 = \lfloor k/a \rfloor + 2$, $A_2 = a\ell_2$ be as in Lemma 7.2. Assume $\ell_2 \leq N - 1$ and*

$$\binom{N}{\ell_2} \;\geq\; |\mathbb{B}|\Big(\frac{q}{k} + 1\Big). \tag{2}$$

Then

$$\varepsilon_{\mathrm{ca}}(C, \delta) \;>\; \frac{1}{2k}\Big(1 - \frac{n}{q}\Big) \qquad \text{for every } \; \delta \in \Big[1 - \frac{A_2}{n},\; 1 - \rho - \frac{1}{n}\Big],$$

and

$$\varepsilon_{\mathrm{mca}}(C, \delta) \;>\; \frac{1}{2k}\Big(1 - \frac{n}{q}\Big) \qquad \text{for every } \; \delta \in \Big[1 - \frac{A_2}{n},\; 1 - \rho\Big).$$

Consequently

$$\delta^*_C(\varepsilon^*) \;\leq\; 1 - \frac{A_2}{n} \;\leq\; 1 - \rho - \frac{a+1}{n} \qquad \text{for every } \; \varepsilon^* \leq \frac{1}{2k}\Big(1 - \frac{n}{q}\Big).$$

Proof. Since $\ell_2 \leq N - 1$ we have $A_2 \leq n - a < n$, and since $A_2 \geq k + a + 1 \geq k + 2$ both stated intervals are nonempty, with $1 - A_2/n \leq 1 - \rho - (a+1)/n \leq 1 - \rho - 2/n$.

Fix $\delta \in [1 - A_2/n,\, 1 - \rho - 1/n]$; then $\lfloor \delta n \rfloor \leq n - k - 1$, so Theorem 2.6 is applicable to C at radius δ. Suppose toward a contradiction that $\varepsilon_{\mathrm{ca}}(C, \delta) \leq \frac{1}{2k}(1 - \frac{n}{q})$. Applying Theorem 2.6 with $\eta = \frac{1}{2}$ gives, for $C^+ := \mathrm{RS}[\mathbb{F}, D, k+1]$,

$$\Lambda(C^+, \delta) \;\leq\; \big\lceil 2q\,\varepsilon_{\mathrm{ca}}(C, \delta) \big\rceil \;\leq\; \Big\lceil \frac{q-n}{k} \Big\rceil \;<\; \frac{q-n}{k} + 1 \;\leq\; \frac{q}{k} + 1.$$

On the other hand, by Lemma 7.2(ii) and monotonicity of lists in the radius (using $\delta \geq 1 - A_2/n$),

$$\Lambda(C^+, \delta) \;\geq\; \big|\Lambda\big(C^+, 1 - \tfrac{A_2}{n}, u_z\big)\big| \;\geq\; \binom{N}{\ell_2}\Big/|\mathbb{B}| \;\geq\; \frac{q}{k} + 1$$

by (2), a contradiction. Hence the $\varepsilon_{\mathrm{ca}}$ lower bound holds throughout the stated integer-admissible interval. In particular it holds at $\delta_0 := 1 - A_2/n$, where Fact 2.4 gives the same lower bound for $\varepsilon_{\mathrm{mca}}(C, \delta_0)$, and support-wise MCA monotonicity (Lemma 2.5) extends it to every $\delta \in [\delta_0, 1-\rho)$. The bound on δ_C^* follows exactly as in Theorem 4.1, using $A_2 \geq k+a+1$. $\square$

Remark 7.5 (relation to Theorem 4.1). With $\varphi = X^a$, D a multiplicative coset, and $a \mid k$, Theorem 7.4 is Theorem 4.1 verbatim: $A_2 = k + 2a$, the left endpoint is $\delta_N = 1 - \rho - 2/N$, hypothesis (2) is (1), and $\ell_2 \leq N - 1$ is the printed condition $(1-\rho)N \geq 3$. The generalization is strict in two independent directions — the folding map and the removal of $a \mid k$ — at no cost in constants. As in Remark 4.2, the constant $\frac{1}{2k}$ can be pushed to $\frac{1-o(1)}{k}$ given slack in (2).

Corollary 7.6 (mixed-radix multiplicative rows, divisibility-free)**.** *Let $\rho \in \{\frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \frac{1}{16}\}$, let $\mathbb{F}$ be any finite field with $q := |\mathbb{F}| < 2^{256}$, let $\mathbb{B} \subseteq \mathbb{F}$ be any subfield, let $D \subseteq \mathbb{B}^\times$ be a multiplicative coset of order n, and let $k = \rho n \leq 2^{40}$, $C = \mathrm{RS}[\mathbb{F}, D, k]$. Suppose n has* some *divisor N with*

$$2^{10} \leq N \leq 2^{11} \quad \big(\rho \in \{\tfrac{1}{2}, \tfrac{1}{4}, \tfrac{1}{8}\}\big), \qquad \textit{resp.} \qquad 2^{11} \leq N \leq 2^{12} \quad \big(\rho = \tfrac{1}{16}\big),$$

with no divisibility condition relating N and k (in particular ρN need not be an integer). Then

$$\varepsilon_{\mathrm{mca}}(C, \delta) \;>\; \frac{1}{2k}\Big(1 - \frac{n}{q}\Big) \;\geq\; 2^{-86} \;\gg\; 2^{-128} \qquad \textit{for every} \;\; \delta \in \Big[1 - \rho - \frac{1}{N},\; 1 - \rho\Big),$$

and the same lower bound holds for $\varepsilon_{\mathrm{ca}}(C, \delta)$ on the integer-admissible subinterval $[1 - \frac{A_2}{n}, 1 - \rho - \frac{1}{n}]$, with A_2 as in Lemma 7.2 for $a = n/N$. In particular

$$\delta_C^*(2^{-128}) \;\leq\; 1 - \rho - 2^{-11} \quad \big(\rho \in \{\tfrac{1}{2}, \tfrac{1}{4}, \tfrac{1}{8}\}\big), \qquad \delta_C^*(2^{-128}) \;\leq\; 1 - \rho - 2^{-12} \quad \big(\rho = \tfrac{1}{16}\big).$$

Proof. Apply Theorem 7.4 with $\varphi = X^{n/N}$ and $a = n/N$; the coset D is (X^a, a)-smooth, and $n < q$ as in Corollary 5.1. Since $\lfloor k/a \rfloor \in (k/a - 1, k/a]$, we have $\ell_2/N \in (\rho + \frac{1}{N}, \rho + \frac{2}{N}]$, so $\ell_2 \leq \rho N + 2 \leq N/2 + 2 \leq N - 1$. For (2): the interval $(\rho + \frac{1}{N}, \rho + \frac{2}{N}]$ is contained in $[\frac{1}{8}, 0.502]$ for the first three rates and in $[\frac{1}{16}, 0.0645]$ for $\rho = \frac{1}{16}$; on each interval the concave function H_2 attains its minimum at an endpoint, and $\mathrm{H}_2(0.502) \geq 0.999 \geq \mathrm{H}_2(\frac{1}{8})$ while $\mathrm{H}_2(0.0645) \geq \mathrm{H}_2(\frac{1}{16})$, so $\mathrm{H}_2(\ell_2/N) \geq \mathrm{H}_2(\frac{1}{8}) \geq 0.5435$, resp. $\mathrm{H}_2(\ell_2/N) \geq \mathrm{H}_2(\frac{1}{16}) \geq 0.3372$. Hence, using the worst N in each window,

$$\log_2 \binom{N}{\ell_2} \;\geq\; N\,\mathrm{H}_2(\ell_2/N) - \log_2(N+1) \;\geq\; \begin{cases} 2^{10} \cdot 0.5435 - \log_2(2^{11}+1) \;\geq\; 545, \\ 2^{11} \cdot 0.3372 - \log_2(2^{12}+1) \;\geq\; 678, \end{cases}$$

while $|\mathbb{B}|(q/k+1) \leq q(q/k+1) \leq q^2 + q < 2^{513}$, exactly as in Corollary 5.1. Theorem 7.4 gives the lower bounds on $[1 - \frac{A_2}{n}, 1 - \rho - \frac{1}{n}]$ and $[1 - \frac{A_2}{n}, 1 - \rho)$; since $A_2 \geq k + a + 1$, the MCA interval contains $[1 - \rho - \frac{1}{N}, 1 - \rho)$, and $\frac{1}{N} \geq 2^{-11}$, resp. 2^{-12}. The error numerics $\frac{1}{2k}(1 - \frac{n}{q}) \geq \frac{1}{2k(n+1)} > 2^{-86}$ (using $n \leq 16k \leq 2^{44}$) are those of Corollary 5.1. $\square$

For FFT-friendly mixed-radix orders $n = 2^\alpha 3^\beta 5^\gamma \cdots$ the divisor set is dense on the logarithmic scale, so the window hypothesis is mild; and taking the smallest available N in the window recovers the smallest gap.

7.2 Twin cosets on the circle, Chebyshev fibers, and torus uniformization

Throughout this subsection p is a prime with $p \equiv 3 \pmod 4$, so that X^2+1 is irreducible over $\mathbb{F}_p$ and $\mathbb{F}_{p^2} = \mathbb{F}_p(i)$ with $i^2 = -1$. Let

$$\mathbb{U} \;:=\; \{u \in \mathbb{F}_{p^2}^{\times} :\; u^{p+1} = 1\}$$

be the norm-one torus, a cyclic group of order $p+1$; for the deployed Mersenne prime $p = 2^{31}-1$ one has $p+1 = 2^{31}$. The circle $\mathcal{C}(\mathbb{F}_p) := \{(x,y) \in \mathbb{F}_p^2 : x^2+y^2 = 1\}$ is in bijection with $\mathbb{U}$ via

$$\iota(x,y) \;:=\; x + iy, \qquad \iota^{-1}(u) \;=\; \Big(\frac{u+u^{-1}}{2},\; \frac{u-u^{-1}}{2i}\Big),$$

using $u^p = u^{-1}$ for $u \in \mathbb{U}$ and $i^p = -i$. Write $\chi(u) := \frac{u+u^{-1}}{2} \in \mathbb{F}_p$ for the x-coordinate map; $\chi(u) = \chi(v)$ holds if and only if $v \in \{u, u^{-1}\}$. The Chebyshev polynomials of the first kind, normalized by $T_0 = 1$, $T_1 = X$, $T_{a+1} = 2XT_a - T_{a-1}$, have integer coefficients and degree a, reduce modulo p to polynomials in $\mathbb{F}_p[X]$, and satisfy the semiconjugacy

$$T_a(\chi(u)) = \chi(u^a) \quad (u \in \mathbb{U},\; a \ge 0), \qquad\qquad T_{ab} = T_a \circ T_b.$$

Following the twin-coset geometry of circle FFTs [HLP24], for a subgroup $H \le \mathbb{U}$ of order M and $g \in \mathbb{U}$ with $g^2 \notin H$ we call

$$\mathcal{D} \;:=\; gH \;\cup\; g^{-1}H \;\subseteq\; \mathbb{U}$$

a *twin coset*. It has size $2M$, is closed under inversion, and contains no self-inverse element: if $u \in gH$ satisfied $u = u^{-1}$, then $u \in gH \cap (gH)^{-1} = gH \cap g^{-1}H$, forcing $g^2 \in H$; likewise for $u \in g^{-1}H$. Consequently χ is exactly two-to-one on $\mathcal{D}$, and the x-coordinate domain $D := \chi(\mathcal{D}) \subseteq \mathbb{F}_p$ has size exactly M. For $a \mid M$ write $H^{(a)} := \{h^a : h \in H\}$ for the image subgroup, of order M/a, and $\mathcal{D}^{(a)} := \{u^a : u \in \mathcal{D}\}$.

Lemma 7.7 (power fibers on torus twin cosets)**.** *Let $\mathcal{D} = gH \cup g^{-1}H$ be a twin coset of size $2M$ and let $a \mid M$.*

(a) *If $g^{2a} \notin H^{(a)}$, then $\mathcal{D}$ is (X^a, a)-smooth (over any field containing $\mathbb{F}_{p^2}$), and $\mathcal{D}^{(a)} = g^aH^{(a)} \sqcup g^{-a}H^{(a)}$ is again a twin coset, of size $2M/a$.*

(b) *If $g^{2a} \in H^{(a)}$, then $\mathcal{D}$ is $(X^a, 2a)$-smooth, with $\mathcal{D}^{(a)} = g^aH^{(a)}$ a single coset of size M/a.*

Moreover, $g^{2a} \notin H^{(a)}$ holds for every $a \mid M$ simultaneously if and only if $g^{2M} \ne 1$, i.e. $\mathrm{ord}(g) \nmid 2M$.

Proof. Since $a \mid M \mid p+1$ and $\mathbb{U}$ is cyclic of order $p+1$, the kernel K_a of $u \mapsto u^a$ on $\mathbb{U}$ has order a and, being the unique subgroup of $\mathbb{U}$ of that order, is contained in H. Hence the a-th power map sends the coset gH onto $g^aH^{(a)}$ with every fiber a coset of K_a, i.e. exactly a-to-one, and likewise sends $g^{-1}H$ onto $g^{-a}H^{(a)}$. The two image cosets of the subgroup $H^{(a)}$ are disjoint or equal according to whether $g^{2a} \notin H^{(a)}$ or $g^{2a} \in H^{(a)}$. In the disjoint case, every $v \in \mathcal{D}^{(a)}$ lies in exactly one of the two image cosets and has exactly a preimages in $\mathcal{D}$; the twin-coset condition for $\mathcal{D}^{(a)}$, namely $(g^a)^2 \notin H^{(a)}$, is the hypothesis itself. In the coincident case, every $v \in g^aH^{(a)}$ has exactly a preimages in gH and exactly a in $g^{-1}H$, hence exactly $2a$ in $\mathcal{D}$. For the last claim: at $a = M$ we have $H^{(M)} = \{1\}$, so the scale-M condition reads $g^{2M} \ne 1$; conversely, if $g^{2a} = h^a$ for some $a \mid M$ and $h \in H$, then $g^{2M} = (g^{2a})^{M/a} = h^M = 1$. □

Lemma 7.8 (exact Chebyshev fibers on x-coordinate twin-coset domains)**.** *Let $\mathcal{D} = gH \cup g^{-1}H$ be a twin coset of size $2M$, let $D := \chi(\mathcal{D}) \subseteq \mathbb{F}_p$, $|D| = M$, and let $a \mid M$ satisfy $g^{2a} \notin H^{(a)}$. Then D is (T_a, a)-smooth over $\mathbb{F}_p$, and*

$$T_a(D) \;=\; \chi\big(\mathcal{D}^{(a)}\big)$$

is the x-coordinate domain of the twin coset $\mathcal{D}^{(a)}$, of size M/a. In particular, if $\mathrm{ord}(g) \nmid 2M$, this holds at every scale $a \mid M$ simultaneously, and for $M = 2^m$ the tower

$$D \;\xrightarrow{T_2}\; T_2(D) \;\xrightarrow{T_2}\; T_4(D) \;\xrightarrow{T_2}\; \cdots$$

consists of x-coordinate twin-coset domains of sizes $M, M/2, M/4, \ldots$ — the line-round tower of circle FRI [HLP24].

Proof. The semiconjugacy gives $T_a(D) = T_a(\chi(\mathcal{D})) = \chi(\mathcal{D}^{(a)})$; by Lemma 7.7(a), $\mathcal{D}^{(a)}$ is a twin coset, so χ is exactly two-to-one on it and $|T_a(D)| = |\mathcal{D}^{(a)}|/2 = M/a$.

Fix $w \in T_a(D)$ and write $w = \chi(v)$ with $v \in \mathcal{D}^{(a)} = g^a H^{(a)} \sqcup g^{-a} H^{(a)}$. Since $\big(g^a H^{(a)}\big)^{-1} = g^{-a} H^{(a)}$, the set $\mathcal{D}^{(a)}$ is closed under inversion, and after replacing v by v^{-1} if necessary (which does not change w) we may assume $v \in g^a H^{(a)}$. Then $v \neq v^{-1}$: otherwise v would lie in both of the disjoint cosets $g^a H^{(a)}$ and $g^{-a} H^{(a)}$. For $x = \chi(u) \in D$ with $u \in \mathcal{D}$,

$$T_a(x) = w \iff \chi(u^a) = \chi(v) \iff u^a \in \{v, v^{-1}\}.$$

By the fiber structure in the proof of Lemma 7.7, the equation $u^a = v$ has exactly a solutions $u \in gH$ and none in $g^{-1}H$ (whose a-th powers fill the disjoint coset $g^{-a}H^{(a)}$), and the equation $u^a = v^{-1}$ has exactly a solutions in $g^{-1}H$ — the inverses of the former — and none in gH. Hence

$$E_w \;:=\; \{u \in \mathcal{D} :\; u^a \in \{v, v^{-1}\}\}$$

has exactly $2a$ elements, is closed under inversion, and, being a subset of $\mathcal{D}$, contains no self-inverse element. Since $\chi(u) = \chi(u')$ forces $u' \in \{u, u^{-1}\}$, the image $\chi(E_w)$ — which is precisely the fiber $T_a^{-1}(w) \cap D$ — has exactly $|E_w|/2 = a$ elements. As $w \in T_a(D)$ was arbitrary, D is (T_a, a)-smooth with $|T_a(D)| = M/a = |D|/a$, as required. The tower statement follows from the all-scales equivalence in Lemma 7.7 together with $T_{2^{j+1}} = T_2 \circ T_{2^j}$. □

Remark 7.9 (standard-position cosets are twin cosets with $\mathrm{ord}(g) \nmid 2M$). A standard-position coset of circle FFT [HLP24] is a coset gH_0, $H_0 := \langle g^2 \rangle$, of size $2M$ with $\mathrm{ord}(g) = 4M$ (so $|H_0| = 2M$). Let $K \leq H_0$ be the index-two subgroup, of order M. Then $g^{-1}K = g\,g^{-2}K$, and g^{-2} generates H_0, hence $g^{-2} \notin K$, so

$$gH_0 \;=\; gK \;\sqcup\; g^{-1}K$$

is exactly the twin coset attached to (K, g): the condition $g^2 \notin K$ holds because g^2 generates $H_0 \supsetneq K$, and $\mathrm{ord}(g) = 4M \nmid 2M$. Thus every standard-position domain of [HLP24], at every folding level, satisfies the hypotheses of Lemmas 7.7 and 7.8 at every dyadic scale simultaneously.

Lemma 7.10 (diagonal invariance of lists and correlated agreement)**.** *Let $C \subseteq \mathbb{F}^n$ be a linear code and $t \in (\mathbb{F}^\times)^n$; write $t \circ v$ for the coordinatewise product and $t \circ C := \{t \circ c : c \in C\}$. Then for every received word U, every radius δ, and every pair of radii,*

$$\Lambda(t \circ C, \delta, t \circ U) = t \circ \Lambda(C, \delta, U), \qquad \varepsilon_{\mathrm{ca}}(t \circ C, \cdot, \cdot) = \varepsilon_{\mathrm{ca}}(C, \cdot, \cdot), \qquad \varepsilon_{\mathrm{mca}}(t \circ C, \cdot) = \varepsilon_{\mathrm{mca}}(C, \cdot).$$

Proof. The map $v \mapsto t \circ v$ is an $\mathbb{F}$-linear bijection of $\mathbb{F}^n$ that preserves coordinatewise agreement: $(t \circ v)_j = (t \circ w)_j \iff v_j = w_j$. Hence it preserves Δ and every restricted distance Δ_S, maps C onto $t \circ C$, acts row-wise on pairs, and commutes with the line structure: $t \circ (f_1 + \gamma f_2) = (t \circ f_1) + \gamma\,(t \circ f_2)$. It therefore carries every witness of the defining events of Λ, $\varepsilon_{\mathrm{ca}}$, and $\varepsilon_{\mathrm{mca}}$ for C bijectively onto a witness for $t \circ C$ with the same slope γ and the same agreement supports, and conversely via t^{-1}. □

Lemma 7.11 (torus uniformization of circle codes)**.** *Let $p \equiv 3 \pmod 4$, let $\mathbb{F} \supseteq \mathbb{F}_{p^2}$ be a finite field (so $i \in \mathbb{F}$), let $E \subseteq \mathcal{C}(\mathbb{F}_p)$ be a set of circle points, and let $E' := \iota(E) \subseteq \mathbb{U} \subseteq \mathbb{F}_{p^2}^{\times}$. For $w \geq 0$ define the $\mathbb{F}$-coefficient circle code*

$$\mathcal{C}_w(\mathbb{F}, E) \;:=\; \big\{(f(P))_{P \in E} \;:\; f \in \mathbb{F}[x,y]/(x^2+y^2-1),\ \deg f \leq w\big\}.$$

Then, under the coordinate bijection $\iota : E \to E'$ and with the diagonal vector $t := (u^{-w})_{u \in E'}$,

$$\mathcal{C}_w(\mathbb{F}, E) \;=\; t \circ \mathrm{RS}[\mathbb{F}, E', 2w+1].$$

Consequently, by Lemma 7.10, the circle code and the Reed–Solomon code $\mathrm{RS}[\mathbb{F}, E', 2w+1]$ of odd dimension $2w+1$ on the torus domain E' have identical list sizes at every radius and identical $\varepsilon_{\mathrm{ca}}$ and $\varepsilon_{\mathrm{mca}}$ at every (pair of) radii.

Proof. Since $i \in \mathbb{F}$ and $x + iy$ is invertible in the quotient ring (with inverse $x - iy$, as $(x + iy)(x - iy) = x^2 + y^2 = 1$), the substitution $u := x + iy$ defines a ring isomorphism

$$\mathbb{F}[x,y]/(x^2+y^2-1) \;\xrightarrow{\sim}\; \mathbb{F}[u,u^{-1}], \qquad x \mapsto \frac{u+u^{-1}}{2}, \quad y \mapsto \frac{u-u^{-1}}{2i},$$

with inverse $u \mapsto x + iy$; it matches evaluation at $P \in E$ on the left with evaluation at $\iota(P) \in E'$ on the right. Every class of degree $\leq w$ has the canonical form $f_0(x) + y\,f_1(x)$ with $\deg f_0 \leq w$ and $\deg f_1 \leq w-1$ (reduce $y^2 = 1 - x^2$), and $\{h_0(x) + y\,h_1(x)\}$ is a free-module canonical form of the quotient ring over $\mathbb{F}[x]$, so this space has dimension exactly $2w+1$. Under the isomorphism,

$$x^j = 2^{-j} u^{-j} (u^2+1)^j \in u^{-j}\,\mathbb{F}[u]_{\leq 2j}, \qquad y\,x^j \in u^{-(j+1)}\,\mathbb{F}[u]_{\leq 2j+2},$$

so the degree-$\leq w$ subspace maps into $u^{-w}\mathbb{F}[u]_{\leq 2w}$; both spaces have dimension $2w+1$ and the map is injective, hence the image is exactly $u^{-w}\mathbb{F}[u]_{\leq 2w}$. Evaluating on E' gives

$$\mathcal{C}_w(\mathbb{F}, E) = \big\{(u^{-w}h(u))_{u \in E'} \;:\; h \in \mathbb{F}[u],\ \deg h \leq 2w\big\} = t \circ \mathrm{RS}[\mathbb{F}, E', 2w+1]. \qquad \square$$

Note the parity: the uniformized dimension $k_c = 2w+1$ is always odd, while the useful folding scales a on a dyadic torus twin coset are even, so $a \nmid k_c$ *always*. The divisibility-free form of Lemma 7.2 is therefore not a luxury here: the original Lemma 3.1 can never be applied to a uniformized circle code directly.

7.3 Universal caps for circle codes and circle-FRI line rounds

Corollary 7.12 (universal circle-row cap)**.** *Let $p \equiv 3 \pmod 4$ be prime, let $\mathcal{D} = gH \cup g^{-1}H \subseteq \mathbb{U}$ be a twin coset of size $2M$ with $\mathrm{ord}(g) \nmid 2M$ (e.g. any standard-position coset of [HLP24], by Remark 7.9), and let $\mathbb{F}$ be any finite field with $q := |\mathbb{F}| < 2^{256}$.*

(a) (Line rounds.) *Assume* $\mathbb{F} \supseteq \mathbb{F}_p$. *Let* $D := \chi(\mathcal{D})$, $n := M$, *and* $C := \mathrm{RS}[\mathbb{F}, D, k]$ *with* $k \le 2^{40}$ *and* $\rho := k/n \in [\frac{1}{16}, \frac{1}{2}]$. *If* $\rho \ge \frac{1}{8}$, *assume* $2^{10} \mid M$ *and set* $N := 2^{10}$; *if* $\rho < \frac{1}{8}$, *assume* $2^{11} \mid M$ *and set* $N := 2^{11}$. *Then*

$$\varepsilon_{\mathrm{mca}}(C,\delta) \;>\; \frac{1}{2k}\Big(1-\frac{n}{q}\Big) \;\ge\; 2^{-86} \;\gg\; 2^{-128} \qquad \text{for every } \; \delta \in \Big[1-\rho-\frac{1}{N},\; 1-\rho\Big),$$

and the same lower bound holds for $\varepsilon_{\mathrm{ca}}(C,\delta)$ *on the integer-admissible subinterval* $[1-\frac{A_2}{n},\, 1-\rho-\frac{1}{n}]$, *with* A_2 *as in Lemma 7.2 for* $a = n/N$. *If moreover* $\frac{n}{N} \mid k$ *— as in every 2-power instantiation at the four challenge rates — the MCA band extends to* $[1-\rho-\frac{2}{N},\, 1-\rho)$. *In particular*

$$\delta^*_C(2^{-128}) \le 1-\rho-2^{-10} \;\; (\rho \ge \tfrac{1}{8}), \qquad \delta^*_C(2^{-128}) \le 1-\rho-2^{-11} \;\; (\rho < \tfrac{1}{8}),$$

improving to $1-\rho-2^{-9}$ (*resp.* $1-\rho-2^{-10}$) *under* $\frac{n}{N} \mid k$.

(b) (Circle codes.) *Assume* $\mathbb{F} \supseteq \mathbb{F}_{p^2}$. *Let* $C := \mathcal{C}_w(\mathbb{F}, \iota^{-1}(\mathcal{D}))$ *with* $k_c := 2w+1 \le 2^{41}$, $n_c := 2M$, *and* $\rho_c := k_c/n_c \in [\frac{1}{16}, \frac{1}{2}]$. *If* $\rho_c \ge \frac{1}{8}$, *assume* $2^9 \mid M$ *and set* $N := 2^{10}$; *if* $\rho_c < \frac{1}{8}$, *assume* $2^{10} \mid M$ *and set* $N := 2^{11}$. *Then*

$$\varepsilon_{\mathrm{mca}}(C,\delta) \;>\; \frac{1}{2k_c}\Big(1-\frac{n_c}{q}\Big) \;\ge\; 2^{-88} \;\gg\; 2^{-128} \qquad \text{for every } \; \delta \in \Big[1-\rho_c-\frac{1}{N},\; 1-\rho_c\Big),$$

with the corresponding $\varepsilon_{\mathrm{ca}}$ *bound on* $[1-\frac{A_2}{n_c},\, 1-\rho_c-\frac{1}{n_c}]$ (*here* $a := n_c/N \nmid k_c$ *always, since* k_c *is odd*). *In particular* $\delta^*_C(2^{-128}) \le 1-\rho_c-2^{-10}$ *for* $\rho_c \ge \frac{1}{8}$, *and* $\delta^*_C(2^{-128}) \le 1-\rho_c-2^{-11}$ *for* $\rho_c < \frac{1}{8}$.

Proof. (a) Put $a := n/N = M/N$, so $a \mid M$, and $g^{2a} \notin H^{(a)}$ by the all-scales equivalence of Lemma 7.7; Lemma 7.8 then makes D a (T_a, a)-smooth domain over $\mathbb{B} := \mathbb{F}_p \subseteq \mathbb{F}$. We check the hypotheses of Theorem 7.4. First, $n < q$: since $M \mid p+1$ and $\mathrm{ord}(g) \mid p+1$, the case $M = p+1$ would give $\mathrm{ord}(g) \mid p+1 \mid 2M$, contradicting $\mathrm{ord}(g) \nmid 2M$; hence $M \le \frac{p+1}{2}$ and $n \le \frac{p+1}{2} < p \le q$. Second, $\ell_2 \le \rho N + 2 \le N/2 + 2 \le N-1$. Third, for (2): $\ell_2/N \in (\rho + \frac{1}{N}, \rho + \frac{2}{N}]$, which is contained in $[\frac{1}{8}, 0.502]$ when $\rho \ge \frac{1}{8}$ and in $[\frac{1}{16}, 0.126]$ when $\rho < \frac{1}{8}$; on each interval the concave H_2 is minimized at an endpoint, and both right endpoints have larger H_2-value than the left, so $\mathrm{H}_2(\ell_2/N) \ge \mathrm{H}_2(\frac{1}{8}) \ge 0.5435$, resp. $\mathrm{H}_2(\ell_2/N) \ge \mathrm{H}_2(\frac{1}{16}) \ge 0.3372$. Hence

$$\log_2 \binom{N}{\ell_2} \;\ge\; N\,\mathrm{H}_2(\ell_2/N) - \log_2(N+1) \;\ge\; 546, \;\; \text{resp. } \; 679,$$

while $|\mathbb{B}|(q/k+1) \le q(q/k+1) \le q^2 + q < 2^{513}$. The error numerics are those of Corollary 5.1: $\rho \ge \frac{1}{16}$ gives $n \le 16k \le 2^{44}$, so $\frac{1}{2k}(1-\frac{n}{q}) \ge \frac{1}{2k(n+1)} > 2^{-86}$. Theorem 7.4 now yields the $\varepsilon_{\mathrm{ca}}$ and $\varepsilon_{\mathrm{mca}}$ lower bounds on $[1-\frac{A_2}{n}, 1-\rho-\frac{1}{n}]$ and $[1-\frac{A_2}{n}, 1-\rho)$ respectively; since $A_2 \ge k+a+1$, the MCA interval contains $[1-\rho-\frac{1}{N}, 1-\rho)$, and when $\frac{n}{N} \mid k$ one has $A_2 = k+2a$, so it equals $[1-\rho-\frac{2}{N}, 1-\rho)$.

(b) By Lemma 7.11, all list sizes, $\varepsilon_{\mathrm{ca}}$, and $\varepsilon_{\mathrm{mca}}$ of C coincide with those of $C' := \mathrm{RS}[\mathbb{F}, \mathcal{D}, k_c]$ on the torus twin coset $\mathcal{D} \subseteq \mathbb{F}_{p^2}$. Put $a := n_c/N = 2M/N$; then $a \mid M$ (indeed $M/a = N/2 \in \mathbb{Z}$), and $g^{2a} \notin H^{(a)}$ as before, so $\mathcal{D}$ is (X^a, a)-smooth over $\mathbb{B} := \mathbb{F}_{p^2} \subseteq \mathbb{F}$ by Lemma 7.7(a). Also $n_c = 2M < p+1 \le p^2 \le q$: the case $2M = p+1$ would again give $\mathrm{ord}(g) \mid p+1 = 2M$. The verification of (2) is identical to (a), with the same two entropy cases and $|\mathbb{B}| = p^2 \le q$; the error numerics use $n_c \le 16k_c \le 2^{45}$, so $\frac{1}{2k_c}(1-\frac{n_c}{q}) \ge \frac{1}{2k_c(n_c+1)} > 2^{-88}$. Theorem 7.4 applied to C', transported back through Lemmas 7.10 and 7.11, gives the stated bands; $A_2 \ge k_c + a + 1$ again yields the containment of $[1-\rho_c-\frac{1}{N}, 1-\rho_c)$. □

Corollary 7.13 (deployed circle parameters: Mersenne-31 with QM31)**.** *Let $p = 2^{31} - 1$, let $\mathbb{F} = \mathbb{F}_{p^4}$ be the deployed QM31 extension (so $q = p^4 > 2^{123.9}$ and $\mathbb{F} \supseteq \mathbb{F}_{p^2}$), and let $\mathcal{D}$ be a standard-position twin coset with $M = 2^{21}$ (so $\mathrm{ord}(g) = 2^{23} \nmid 2^{22} = 2M$).*

(a) (Line round, rate $\frac{1}{2}$.) *With $D = \chi(\mathcal{D})$, $|D| = 2^{21}$, and $C = \mathrm{RS}[\mathbb{F}_{p^4}, D, 2^{20}]$,*

$$\varepsilon_{\mathrm{mca}}(C, \delta) \;>\; \left(1 - 2^{-102}\right) 2^{-21} \;>\; 2^{-22} \qquad \text{for every } \; \delta \in \Big[\frac{1}{2} - 2^{-7},\; \frac{1}{2}\Big),$$

and the same lower bound holds for $\varepsilon_{\mathrm{ca}}(C, \delta)$ for every $\delta \in [\frac{1}{2} - 2^{-7}, \frac{1}{2} - 2^{-21}]$.

(b) (Circle code.) *With $C = \mathcal{C}_w(\mathbb{F}_{p^4}, \iota^{-1}(\mathcal{D}))$, $w = 2^{20}$, $k_c = 2^{21} + 1$, $n_c = 2^{22}$, $\rho_c = \frac{1}{2} + 2^{-22}$,*

$$\varepsilon_{\mathrm{mca}}(C, \delta) \;>\; 2^{-23} \qquad \text{for every } \; \delta \in \Big[\frac{1}{2} - 2^{-7},\; 1 - \rho_c\Big),$$

and the same lower bound holds for $\varepsilon_{\mathrm{ca}}(C, \delta)$ for every $\delta \in [\frac{1}{2} - 2^{-7}, \frac{1}{2} - 2^{-21}]$.

Proof. Both parts use $N = 256$ and, as in Corollary 5.4, $\log_2 \binom{256}{130} \geq 247$.

(a) Here $a = n/N = 2^{13}$ divides $k = 2^{20}$, so $\ell_2 = \rho N + 2 = 130$, $A_2 = k + 2a = 2^{20} + 2^{14}$, and $1 - A_2/n = \frac{1}{2} - 2^{-7}$. The standard-position hypothesis gives exact Chebyshev fibers at scale a (Remark 7.9 and Lemma 7.8), with $\mathbb{B} = \mathbb{F}_p$. Hypothesis (2): $|\mathbb{B}|(q/k + 1) = p\,(p^4/2^{20} + 1) < 2^{31} \cdot 2^{105} = 2^{136} \leq 2^{247} \leq \binom{256}{130}$. The error bound is $\frac{1}{2k}(1 - n/q)$ with $n/q = 2^{21}/p^4 < 2^{-102}$. Theorem 7.4 gives the MCA band $[\frac{1}{2} - 2^{-7}, \frac{1}{2})$ and the $\varepsilon_{\mathrm{ca}}$ band $[\frac{1}{2} - 2^{-7}, \frac{1}{2} - 2^{-21}]$, since $1 - \rho - \frac{1}{n} = \frac{1}{2} - 2^{-21}$.

(b) Here $a = n_c/N = 2^{14} \mid M = 2^{21}$, $\ell_2 = \lfloor (2^{21} + 1)/2^{14} \rfloor + 2 = 2^7 + 2 = 130$, $A_2 = 2^{14} \cdot 130 = 2^{21} + 2^{15}$, and $1 - A_2/n_c = \frac{1}{2} - 2^{-7}$. Lemma 7.11 transports everything to $C' = \mathrm{RS}[\mathbb{F}_{p^4}, \mathcal{D}, 2^{21} + 1]$ on the torus twin coset, which is (X^a, a)-smooth over $\mathbb{B} = \mathbb{F}_{p^2}$ by Lemma 7.7 and Remark 7.9. Hypothesis (2): $|\mathbb{B}|(q/k_c + 1) = p^2\left(p^4/(2^{21} + 1) + 1\right) < 2^{62} \cdot 2^{104} = 2^{166} \leq 2^{247}$. The error bound: $n_c/q = 2^{22}/p^4 < 2^{-101}$, so

$$\frac{1}{2k_c}\Big(1 - \frac{n_c}{q}\Big) \;>\; \frac{1 - 2^{-101}}{2^{22} + 2} \;>\; 2^{-22}\big(1 - 2^{-21}\big)\big(1 - 2^{-101}\big) \;>\; 2^{-23}.$$

Theorem 7.4 gives the MCA band $[\frac{1}{2} - 2^{-7}, 1 - \rho_c)$ and the $\varepsilon_{\mathrm{ca}}$ band $[\frac{1}{2} - 2^{-7}, 1 - \rho_c - \frac{1}{n_c}] = [\frac{1}{2} - 2^{-7}, \frac{1}{2} - 2^{-21}]$. □

Corollary 7.14 (circle certifying lines are genuinely extension-valued)**.** *The subfield-confinement lemma Lemma 6.1 holds verbatim for an arbitrary evaluation set $D \subseteq \mathbb{B}$: its proof uses only that the evaluation points lie in $\mathbb{B}$, not that they are invertible. Consequently, at the parameters of Corollary 7.13:*

(a) *any pair of words on $D = \chi(\mathcal{D})$ that is $\mathbb{F}_p$-valued up to a common scalar factor has MCA-bad-slope density at most $p/q = p^{-3} < 2^{-92}$ at every radius; in particular, every pair attaining the bound $> 2^{-22}$ of Corollary 7.13(a) is genuinely $\mathbb{F}_{p^4}$-valued;*

(b) *the diagonal vector $t = (u^{-w})_u$ of Lemma 7.11 is $\mathbb{F}_{p^2}$-valued, so any pair of circle words whose ι-transports are $\mathbb{F}_{p^2}$-valued up to a common scalar factor has MCA-bad-slope density at most $p^2/q = p^{-2} < 2^{-61}$ at every radius; in particular, every circle pair attaining the bound $> 2^{-23}$ of Corollary 7.13(b) contains a word whose ι-transport is not $\mathbb{F}_{p^2}$-valued, even after removing a common scalar.*

As in Remark 2.15, the extension-pole floors Corollaries 2.12 and 2.13 supply the matching constructive mechanism: the fiber received words of Lemma 7.2 on circle domains are $\mathbb{B}$-valued with $\mathbb{B}$-valued lists, and any such list converts through a simple pole $\alpha \in \mathbb{F} \setminus \mathbb{B}$ into genuinely extension-valued bad lines, with denominator $q - b$ ($b = |\mathbb{B}| \in \{p, p^2\}$) in place of $q - n$.

Proof. Inspecting the proof of Lemma 6.1: it decomposes words and codewords over a $\mathbb{B}$-basis of $\mathbb{F}$ and evaluates polynomials at the points of $D \subseteq \mathbb{B}$; no step uses $D \subseteq \mathbb{B}^\times$, so the two-radius statement holds for arbitrary $D \subseteq \mathbb{B}$, in particular for the Chebyshev domain $\chi(\mathcal{D}) \subseteq \mathbb{F}_p$ and the torus domain $\mathcal{D} \subseteq \mathbb{F}_{p^2}$. (a) is then Lemma 6.1 with $\mathbb{B} = \mathbb{F}_p$ inside $q = p^4$: the bad-slope density of a $\mathbb{B}$-confined pair is at most $|\mathbb{B}|/q = p^{-3} < 2^{-92} < 2^{-22}$. (b) A circle pair whose ι-transports are $\mathbb{F}_{p^2}$-valued up to a common scalar corresponds, after multiplying by the $\mathbb{F}_{p^2}$-valued diagonal t, to an $\mathbb{F}_{p^2}$-confined pair for $\mathrm{RS}[\mathbb{F}_{p^4}, \mathcal{D}, k_c]$; Lemma 7.10 preserves the bad-slope set, and Lemma 6.1 with $\mathbb{B} = \mathbb{F}_{p^2}$ bounds its density by $p^2/q = p^{-2} < 2^{-61} < 2^{-23}$. The final claim is Corollary 2.12 applied with $\mathbb{B} \in \{\mathbb{F}_p, \mathbb{F}_{p^2}\}$ to the $\mathbb{B}$-valued received words and lists of Lemma 7.2. □

Remark 7.15 (interleaving, slacked fallback, and the two-tier picture for circle rows). The interleaving-transfer lemma Lemma 3.84 is stated for arbitrary linear codes, so every bound of Corollaries 7.12 and 7.13 transfers verbatim to s-interleaved circle rows, exactly as in Corollary 5.5. The slacked BCHKS/ABF fallback also carries over: Lemma 7.2(i) supplies a list of size $\geq \binom{N}{\ell_1}/|\mathbb{B}|$ in $\mathrm{RS}[\mathbb{F}, D, k]$ at radius $1 - \frac{A_1}{n}$, and whenever $A_1 \geq k + 3$ (automatic in the 2-power line rounds with $a \geq 4$, where $A_1 = k + a$) and $\binom{N}{\ell_1} \geq |\mathbb{B}|\, q$, the contrapositive of Theorem 2.16 gives $\varepsilon_{\mathrm{ca}}\big(C,\, 1 - \frac{A_1}{n} + \frac{2}{n},\, 1 - \rho - \frac{1}{n}\big) \geq \frac{1}{2n}$, as in Proposition 5.6. Finally, the two-tier summary of Remark 5.3 acquires a circle column: the deployed Mersenne-31 circle family carries the same kind of near-capacity negative band as the multiplicative rows — gap 2^{-7} at the deployed sizes, 2^{-9}–2^{-11} universally — over every challenge field $q < 2^{256}$ containing the relevant base field.

7.4 The lacunarity obstruction for genus-one (ECFFT) domains

Remark 7.16 (what blocks the elliptic case). The proof of Lemma 7.2 works because φ is a *polynomial*: in the expansion $\prod_{b \in A}(\varphi - b) = \sum_j (-1)^j e_j(A)\, \varphi^{\ell - j}$, consecutive terms drop by a full $a = \deg \varphi$ in degree, so the received word occupies the top two degree slots and everything else is a codeword sitting $2a$ degrees down — gap $\approx 2/N$. Consider instead a rational folding map $\psi = f/g$ of degree a, with $e := \deg g < \deg f = a$ and g nonvanishing on D, whose fibers inside D are complete and of uniform size a. The natural locators are $\Lambda_A := \prod_{b \in A}(f - b\, g)$, of degree $a\ell$, vanishing on the ψ-fibers over A; their expansion $\sum_j (-1)^j e_j(A)\, f^{\ell - j} g^j$ has j-th term of degree $a\ell - j(a - e)$, so the slack-two window has width $2(a - e)$ in degree and the certified gap is $2(a - e)/n$. Polynomial maps have $e = 0$ and recover gap $2/N$. The FFTree domains of ECFFT [BCKL21] fold by x-coordinate maps of 2-isogenies of elliptic curves, which are rational of degree type $(a, e) = (2, 1)$; every composition of such maps is again the x-map of an isogeny, of type $(a, a - 1)$. For these, $a - e = 1$ at every scale: the window collapses to width 2, and the construction certifies only the first staircase step — a large list at agreement $k + 2$, hence an $\varepsilon_{\mathrm{ca}}$ floor at gap $2/n$ — a nonvacuous but microscopic band, far from the 2^{-9}–2^{-11} bands above. The circle escapes this because its folding tower becomes polynomial after the x-projection: T_2 fixes ∞ and is totally ramified there. The elliptic x-maps also fix ∞ but are *not* totally ramified there ($\psi^{-1}(\infty) = \{\infty, x(T)\}$ for a 2-isogeny with kernel generator T), and the locator expansion sees exactly this defect as $e = a - 1$. Obtaining a macroscopic universal cap for ECFFT rows

therefore requires either a lacunary basis of isogeny locators, a different received-word family, or a matching safe-side theorem showing the collapse is essential. The second exit is realized by the rational graded floor of Proposition 11.1 (Corollary 11.2); the residual band form is recorded as the genus-one/circle transport question Question 22.3.

8 Threshold formulation and certificate interface

The Proximity Prize asks for a threshold, not merely for a near-capacity obstruction. We therefore use the integer staircase below to state both sides of the result: lower certificates mark agreements where MCA or list decoding is unsafe, while upper certificates mark agreements where the corresponding error is below the target. The paper does not determine the staircase to one step in the remaining band, but it reduces the unknown part to a single aperiodic correlated-agreement problem and gives exact finite certificates for all printed edges.

8.1 Established inputs

For the grand MCA challenge at target error $\varepsilon^* = 2^{-128}$, the unconditional output of Corollary 5.1 is the following negative certificate: for every challenge-envelope row satisfying the printed divisor hypothesis,

$$\delta \geq 1 - \rho - 2^{-9} \quad (\rho \in \{1/2, 1/4, 1/8\})$$

and

$$\delta \geq 1 - \rho - 2^{-10} \quad (\rho = 1/16)$$

are impossible MCA radii. Equivalently, any full determination of $\delta^*_C(2^{-128})$ for these smooth multiplicative rows must search below these gaps. This conclusion uses only the self-contained deep-point conversion, the quotient-fiber pigeonhole, and support-wise MCA monotonicity; it no longer rests on importing the Crites–Stewart conversion.

The theorem also distinguishes two different grades of information. First, the *threshold grade* is already enough for the prize inequality: the certified error is $\gg 2^{-128}$. Second, the *failure-profile grade* is still weak: the present proof gives error of order $1/k$, not error $1 - o(1)$. Closing the error-one profile at fixed gaps such as $1/64$ is a different, stronger problem. It would improve our understanding of the failure landscape, but is not needed to obtain the stated upper bound on $\delta^*_C(2^{-128})$.

The negative certificate also covers a wider row family. The same conclusion — an unsafe near-capacity band at error $\gg 2^{-128}$ — now holds for mixed-radix multiplicative rows without the integrality condition $\rho N \in \mathbb{Z}$ (Corollary 7.6, gaps $2^{-11}/2^{-12}$), for the x-coordinate line rounds of circle FRI over any $p \equiv 3 \pmod 4$, and for the bivariate circle codes over any challenge field containing i (Corollary 7.12, gaps 2^{-9}–2^{-11}), including the deployed Mersenne-31/QM31 family at gap 2^{-7} (Corollary 7.13). These additions use the same three ingredients; the only new inputs are the exact Chebyshev-fiber geometry of twin cosets and the torus uniformization, both elementary.

The treatment in Section 9 adds the safe side. Theorem 9.6 certifies $\varepsilon_{\mathrm{mca}}(C, \delta) \leq (\lfloor \delta n \rfloor + 1)/q$ whenever $3\lfloor \delta n \rfloor \leq n - k$, for every row at once, and Theorem 9.15 certifies the interleaved-list challenge up to the Johnson radius uniformly in the arity; Table 4 records the per-obligation outcome, and Theorem 9.27 assembles the resulting two-sided sandwich.

8.2 Integer-staircase formulation

Let

$$a(\delta) := \lceil (1-\delta)n \rceil, \qquad r(\delta) := n - a(\delta)$$

be the agreement and integer-radius forms of the same radius. For a fixed finite-slope line field F_{line}, write $q_{\mathrm{line}} = |F_{\mathrm{line}}|$ and let

$$B_{\mathrm{mca}}(a)$$

be the exact maximum number of MCA-bad line parameters at agreement at least a, for the line family used in the challenge. Then

$$\varepsilon_{\mathrm{mca}}(C, 1 - a/n) = B_{\mathrm{mca}}(a)/q_{\mathrm{line}}$$

under the closed finite-slope convention. For projective-slope or curve samplers, the same staircase formulation applies after replacing q_{line} and B_{mca} by the corresponding sampler denominator and numerator. Thus the MCA prize problem is an integer staircase problem: find the first agreement a for which

$$B_{\mathrm{mca}}(a) \le 2^{-128} q_{\mathrm{line}}.$$

The present paper supplies lower bounds on $B_{\mathrm{mca}}(a)$ near capacity by routing quotient-fiber list mass through deep points.

Proposition 8.1 (one-step threshold certificate)**.** *Fix a row C and a denominator q_{line}. Suppose that for some agreement a_0 one has a lower certificate*

$$B_{\mathrm{mca}}(a_0) > 2^{-128} q_{\mathrm{line}}$$

and an upper theorem

$$B_{\mathrm{mca}}(a_0 + 1) \le 2^{-128} q_{\mathrm{line}}.$$

Then, under the closed integer-radius convention, radius $1 - a_0/n$ is unsafe and radius $1 - (a_0 + 1)/n$ is safe. As a real-radius supremum, the threshold is $1 - a_0/n$ and is not attained if the closed ball at that endpoint is counted.

Proof. The function $B_{\mathrm{mca}}(a)$ is nonincreasing in a, since raising the required agreement only removes witnesses. At the closed radius $1 - a/n$, the MCA error is $B_{\mathrm{mca}}(a)/q_{\mathrm{line}}$. The two displayed inequalities are therefore exactly the unsafe/safe transition at consecutive integer agreement levels. □

8.3 Conditional MCA one-step criterion

The following theorem is a deliberately explicit proof blueprint. Each hypothesis is a concrete mathematical task, not a black-box "large field" assumption.

Theorem 8.2 (conditional MCA threshold theorem)**.** *Fix a smooth multiplicative row $C = \mathrm{RS}[F, D, k]$, a line field F_{line}, and a target $\varepsilon^* = 2^{-128}$. Suppose the following four inputs are proved for this row, uniformly in the received line:*

(i) ***Quotient-profile lower and upper ledger.*** *For every divisor $c \mid n$ and every agreement $a = mc + s$ with $0 \le s < c$, the quotient-periodic bad slopes, including remainder supports consisting of m full fibers plus s residual points, are counted with overlaps across divisors accounted for. Write the resulting divisor-lattice upper numerator as $B_Q^{\mathrm{all}}(a)$; individual divisor contributions $B_Q(a; c)$ may also be printed, but the upper ledger must count their union or a certified safe sum rather than only their maximum.*

(ii) ***Aperiodic residue-line packing.*** *After removing tangent and quotient-periodic locators, the Hankel/residue-line incidence variety has at most $B_{\mathrm{ap}}(a)$ bad line parameters.*

(iii) ***Extension-line accounting.*** *If F_{line} is a proper extension of the field generated by D, either all genuinely extension-valued line witnesses are included in $B_{\mathrm{ap}}(a)$, or a separate extension term $B_{\mathrm{ext}}(a)$ is proved.*

(iv) ***No missing line family.*** *The line family in the theorem is the same finite-slope or projective-slope family used by the challenge or protocol under consideration.*

Assume moreover that the resulting upper numerator

$$B^+(a) := B_{\tan}(a) + B_Q^{\mathrm{all}}(a) + B_{\mathrm{ap}}(a) + B_{\mathrm{ext}}(a)$$

and a corresponding lower numerator $B^-(a)$ satisfy

$$B^-(a_0) > \varepsilon^* q_{\mathrm{line}}, \qquad B^+(a_0+1) \le \varepsilon^* q_{\mathrm{line}}$$

for some a_0. Then $\delta_C^*(\varepsilon^*)$ is determined at the one-step accuracy of Proposition 8.1.

Proof. The four hypotheses classify every possible MCA-bad line parameter into one of four buckets: tangent, quotient-periodic, aperiodic, or genuinely extension-line. The quotient bucket is counted by the divisor-lattice union ledger B_Q^{all}, so overlaps or simultaneous quotient descriptions cannot invalidate the upper bound. The displayed $B^+(a)$ is therefore an upper bound for $B_{\mathrm{mca}}(a)$, while $B^-(a)$ is a certified lower bound. The conclusion is exactly Proposition 8.1. □

The present paper supplies part of the lower numerator $B^-(a)$ in the near-capacity band: quotient-fiber list mass forces MCA mass above the printed universal caps. It does not prove the aperiodic upper input (ii), and therefore does not by itself determine the exact staircase.

8.4 Conditional interleaved-list one-step criterion

The grand list-decoding challenge asks for the largest δ such that an interleaved list, divided by the field used for the relevant challenge or list ledger, is below 2^{-128}. In the simplest same-field form this is

$$|\Lambda(C^{\equiv m}, \delta)| \le 2^{-128}|F|.$$

The quotient-fiber construction in this paper already gives lower floors for ordinary lists, and ordinary lists embed into interleaved lists by placing the same received word in one row and zero words in the others. On the MCA side, Section 3.4 supplies a certified eliminant route for aperiodic Hankel packing, but it does not prove that the needed eliminants always exist. The missing positive list direction is an arbitrary-word locator-fiber theorem.

Theorem 8.3 (conditional interleaved-list threshold theorem)**.** *Fix an interleaving arity m. Suppose that for every received word U and agreement $a = k + \sigma$ one has a base-code list bound of the form*

$$|\Lambda(C, 1 - a/n, U)| \le L_Q(a, U) + L_{\mathrm{ap}}(a),$$

where L_Q is an explicit quotient-profile list contribution and L_{ap} is a uniform aperiodic locator-fiber bound. Suppose also that the same bound is available at the higher agreement levels that occur in the codegree recursion, starting with $2a - k$. Then the interleaved-list threshold is obtained by applying the codegree recursion row by row and comparing the resulting numerator to $2^{-128}|F|$.

Proof. For two rows, split the second-row list according to whether its agreement set has size below or above $2a - k$. Below $2a - k$, two first-row codewords would agree on more than k positions and hence are equal; above $2a - k$, use the base-code list bound at agreement $2a - k$. Iterating gives the stated m-row recursion. This is the same codegree mechanism used in the experimental ledger; the only missing input is the arbitrary-word locator-fiber bound. □

8.5 Residual ingredients for one-step threshold certificates

The one-step criteria above consume the following residual ingredients, which we record with their current status. The lower-floor half of the periodic-support ledger is supplied in scanner-ready form by Lemma 3.7 and Theorem 3.8: every quotient-plus-remainder agreement $a = mc+s$ has a generated-field prefix-fiber floor, an image-size fallback, an ambient-box fallback, and a converted CA/MCA lower numerator. The support-family upper half is exact at the divisor-lattice level by Theorem 3.24: overlaps between quotient descriptions at different divisor scales are counted by a common-block coefficient extraction rather than by a raw safe sum. The distinct-parameter branch image is exact inside every declared quotient family by the lcm ledger of Theorem 3.36 and Corollaries 3.37 and 3.39; duplicate finite slopes, projective slopes, and finite curve parameters are coalesced before the numerator is printed. The residual-band analysis is split into four certified pieces. The regular overdetermined nonsingular branch is canonical by Theorems 3.48 and 3.55 and Corollary 3.51. The underdetermined and rank-drop singular branches are covered by the rank-pivot certificate theorem Theorems 3.64 and 3.70. The new compressed-residual layer adds the kernel-compressed eliminant construction of Theorem 3.73 and Corollary 3.75 and the exact arbitrary-residual support-image lcm ledger of Theorem 3.76 and Corollary 3.77. Thus a residual chart must now be discharged by a small compressed eliminant, by an exact residual image polynomial, by a Weil-restricted extension projection bound, or by a named structural obstruction. Extension-line accounting is handled on the safe side by Theorem 3.68: after choosing a B-basis of F, each extension chart prints a parameter-projection degree and dimension over B, or an exact zero-dimensional point count. The curve-sampler ledger is upgraded for finite-parameter curves and projective lines in parallel: quotient branches have exact gcd/lcm image ledgers, regular curve rank-drop branches have canonical maximal-minor gcd ledgers, and residual curve branches have both coefficient-pivot rank charts and kernel-compressed curve eliminants. What remains to be settled is no longer a bookkeeping ambiguity: prove that the compressed residual eliminants are small enough in the near-capacity rows, or classify any large residual image as quotient, tangent/common-line, base-field exceptional, genuinely extension-valued, or a new obstruction.

T1. Remainder quotient-profile theorem. Generalize the fiber lemma from full quotient-fiber unions to all agreements $a = mc+s$. The target is an exact lower and upper quotient ledger over the divisor lattice of D. The present quotient-ledger package supplies the scanner-ready lower half via heaviest prefix fibers and the support-union upper half via the exact divisor-block coefficient formula; inside a declared quotient family, the gcd/lcm image ledger also coalesces duplicate finite, projective, and curve parameters. What remains is the structural exhaustion theorem saying that the declared quotient family captures all quotient-periodic bad parameters in the intended regime, with the subfield denominator $|B|$ printed separately from the line denominator q.

T2. Aperiodic Hankel-packing theorem. Use the syndrome/Hankel normal form of [Cho26b] to prove that, after quotient-periodic and tangent locators are removed, the

number of remaining bad line parameters is at most $B_{\mathrm{ap}}(a)$. The regular overdetermined nonsingular branch is now canonical by Theorem 3.48 and Corollary 3.51: a scanner takes the gcd of all nonzero maximal Hankel minors and then an lcm over exact agreement levels. The underdetermined and rank-drop singular buckets now have the rank-pivot cover of Theorem 3.64, projective-infinity and curve coefficient pivots, and the kernel-compressed eliminant ideals of Theorem 3.73 and Corollary 3.75. The remaining proof problem is constructive and structural: produce small eliminants for these compressed residual charts, use the exact residual image lcm of Theorem 3.76 when finite enumeration is feasible, or classify every large residual component as quotient, tangent/common-line, base-field exceptional, genuinely extension-valued, or a new obstruction.

T3. Extension-line lift or counterexample. For $B \subset F$ and $D \subset B$, either prove that genuinely F-valued lines are covered by the aperiodic ledger with a stable numerator, or exhibit explicit extension-valued counterexamples. Corollary 6.2 shows that deployed failures certified by this paper must be genuinely extension-valued; Corollaries 2.12 and 2.13 provide an explicit lower-side mechanism from any large list by taking the pole in $F \setminus B$. On the safe side, Theorem 3.68 gives the certificate format: Weil-restrict each residual extension chart to B, print a parameter-projection degree/dimension pair or an exact zero-dimensional point count, and divide by the line field size. The remaining structural problem is to prove that all genuinely extension-valued residual charts have small projected dimension/degree or to exhibit the chart that violates this expectation.

T4. Arbitrary-word locator-fiber theorem. Prove the base-code list bound needed by Theorem 8.3, first for monomial-prefix received words and quotient-free dimensions, then for arbitrary received words with quotient cores budgeted.

T5. Curve and projective-line ledgers. Extend the affine-line Hankel classification to projective slopes and to finite-parameter power curves $f_0 + \gamma f_1 + \cdots + \gamma^d f_d$. The projective and curve ledger package gives exact gcd/lcm ledgers for projective and finite-curve quotient branches, canonical rank-drop gcds for regular overdetermined finite-curve buckets, projective-infinity residual charts, coefficient-pivot curve rank charts, and kernel-compressed projective/curve eliminant certificates. The remaining work is to prove smallness or structural classification for the residual projective and curve charts, including exceptional vanishing-vector strata and genuinely extension-valued components.

T6. Certificate scanner. Implement a finite staircase scanner that, for each challenge row, prints tangent, quotient, generated-field entropy, aperiodic, extension, curve, and interleaved-list numerators at every agreement a. The desired output is an interval

$$a_{\mathrm{unsafe}} < a_C^* \leq a_{\mathrm{safe}}$$

that shrinks to one step when Theorem 8.2 or Theorem 8.3 applies.

T7. Formal finite core. Formalize the self-contained deep-point conversion, the quotient-fiber pigeonhole, support-wise MCA monotonicity, and the integer staircase lemma. These are small enough for proof-assistant checking and directly strengthen a formal certificate package.

T8. Circle-row completion and genus-one floors. Extend the quotient-support and quotient-image upper ledgers (item 1) and the aperiodic Hankel classification (item 2)

from power locators on multiplicative divisor lattices to Chebyshev locator towers on x-coordinate twin-coset domains, where the divisor lattice of D is replaced by the lattice of dyadic Chebyshev scales; and settle the genus-one lower-floor question of the genus-one/circle transport question Question 22.3, for which the lacunarity collapse of Remark 7.16 blocks the present route beyond the first staircase step.

8.6 Scope of claims

We state precisely what this paper establishes and what it does not. Established: a universal unsafe band at capacity for the grand MCA challenge, over smooth multiplicative rows in the stated field-size envelope (Corollaries 5.1 and 7.6) and, by the later sections, over Chebyshev, circle, and genus-one rows as well (Corollaries 11.2, 11.5 and 7.12); an unconditional safe region below one third of the distance (Theorem 9.6) with a list-side safe region up to the Johnson radius, uniform in the interleaving arity (Theorem 9.15); a safe region up to one half of the distance modulo a single classical import (Theorem 10.1 and Corollary 10.7); and the resulting row-specific sandwich, which determines each deployed MCA threshold within a factor smaller than two modulo that import, and within a factor of three unconditionally (Theorems 9.27 and 10.12, Table 5). Not established, and not claimed: the exact threshold, a positive theorem inside the remaining frontier band, protocol soundness, or a curve-MCA theorem. What separates the sandwich from a one-step determination is a complete normalized upper certificate of the form templated in Section 21; the original broader proposal is refuted in Section 18.

9 Structural ledgers, residual bands, and certificate machinery

This section collects the technical infrastructure used by the two-sided threshold statements, organized by mathematical function. Stabilizer and quotient-descent lemmas isolate the periodic supports and make their numerators field-size independent. The deep-regime theorem controls MCA directly below one third of the distance and removes the singular Hankel branches there. The restriction-of-scalars lift shows that extension-valued lines are not a separate class of witnesses. The Johnson and recursion statements control the list side outside the band. The curve theorem covers projective and finite-parameter samplers in the deep regime. Finally, the scanner and finite-certificate grammar package all deployed verdicts as exact integer comparisons. The remaining mathematical component is the normalized residual ledger of Remark 9.12 and Section 21; the sub-envelope range is already unsafe by Section 18.

9.1 Stabilizer exhaustion and exact quotient descent

Throughout this subsection, *multiplicative case* means: $D = \alpha H$ is a multiplicative coset of a cyclic group H of order n, acting freely on D by multiplication; for $c \mid n$, $H_c \leq H$ is the unique subgroup of order c, equal to the full group μ_c of c-th roots of unity in $\mathbb{B}$, and the fibers of $\pi_c : x \mapsto x^c$ on D are exactly the H_c-orbits. *Chebyshev case* means: $D = \chi(\mathcal{D})$ for a twin coset $\mathcal{D} = gH \cup g^{-1}H$ of size $2M$ with $\mathrm{ord}(g) \nmid 2M$ as in Section 7.2; for $c \mid M$, $K_c \leq \mathbb{U}$ is the order-c kernel of the c-th power map, and for $S \subseteq D$ we write $\widetilde{S} := \chi^{-1}(S) \cap \mathcal{D}$ for the (automatically inversion-closed) lift. By Lemma 7.8, the T_c-fibers in D are exactly the χ-images of the $\langle K_c,\, u \mapsto u^{-1}\rangle$-orbits on $\mathcal{D}$. In both cases the *scale lattice* is the divisor lattice of n (resp. of M).

Definition 9.1 (periodicity scale of a support)**.** For $S \subseteq D$, the periodicity scale $c(S)$ is $|\{h \in H : hS = S\}|$ in the multiplicative case and $|\{\zeta \in K_M : \zeta\widetilde{S} = \widetilde{S}\}|$ in the Chebyshev case. A support is *aperiodic* if $c(S) = 1$.

Theorem 9.2 (stabilizer partition and structural exhaustion)**.** *In either case:*

(i) *The stabilizer in the definition is a subgroup of a cyclic group, so $c(S)$ divides n (resp. M), and for every scale c,*

$$S \text{ is a union of complete scale-}c \text{ fibers} \iff c \mid c(S).$$

In particular $c(S)$ is the unique maximal such scale, and the supports of each size are partitioned by the value of $c(S)$, with $c(S) = 1$ the aperiodic class.

(ii) *Fix a received line and an agreement threshold $a \geq k$, and call a bad parameter* periodically witnessed *if at least one of its witness supports S (in the sense of Definitions 2.1 and 2.2) has $c(S) \geq 2$. Then every bad parameter is either periodically witnessed — and hence counted, with multiplicity one per support, by the family of supports of size $\geq a$ that are unions of complete fibers at some scale ≥ 2 — or all of its witness supports are aperiodic, in which case it belongs to the aperiodic bucket by definition. The two classes are exhaustive and disjoint.*

Proof. (i) The set $\{h \in H : hS = S\}$ is a subgroup of the cyclic group H (resp. $\{\zeta \in K_M : \zeta\widetilde{S} = \widetilde{S}\} \leq K_M$), hence equals H_{c_0} (resp. K_{c_0}) for a unique divisor c_0, which is $c(S)$. In the multiplicative case, the π_c-fibers in D are the H_c-orbits, and a set is a union of complete H_c-orbits if and only if it is H_c-invariant, i.e. $H_c \leq H_{c_0}$, i.e. $c \mid c_0$. In the Chebyshev case, S is a union of complete T_c-fibers if and only if $\widetilde{S}$ is a union of complete $\langle K_c, \mathrm{inv}\rangle$-orbits; since $\widetilde{S}$ is inversion-closed automatically, this holds if and only if $\widetilde{S}$ is K_c-invariant, i.e. $c \mid c_0$. The partition statement is immediate.

(ii) Exhaustiveness and disjointness are definitional given (i). The counting clause is Lemma 3.16 applied to the support family $\{S : |S| \geq a,\ c(S) \geq 2\}$: each fixed support pays for at most one finite line parameter, so the number of periodically witnessed bad parameters is at most the size of that family. □

Corollary 9.3 (exact size of the periodic support family; q-free numerator)**.** *Let P be the set of primes dividing n (resp. M). The number of supports of exact size A with $c(S) \geq 2$ is*

$$\Pi(A) \;:=\; \sum_{\emptyset \neq T \subseteq P} (-1)^{|T|+1}\, \mathbf{1}\big[\mathrm{lcm}(T) \mid A\big] \binom{n/\,\mathrm{lcm}(T)}{A/\,\mathrm{lcm}(T)},$$

and consequently, for every received line and every agreement threshold $a \geq k$,

$$\#\{\text{periodically witnessed bad finite parameters}\} \;\leq\; \sum_{A=a}^{n} \Pi(A),$$

with an extra factor d for degree-d power-curve samplers (Lemma 3.17). The numerator is independent of q and of $|\mathbb{B}|$; the subfield order enters only the lower floors, as the periodic-support ledger demands.

Proof. $c(S) \geq 2$ if and only if $H_p S = S$ for some $p \in P$, i.e. S lies in the union over $p \in P$ of the families of complete-p-fiber unions. For $T \subseteq P$, invariance under every H_p, $p \in T$, is invariance under the subgroup they generate, which is $H_{\mathrm{lcm}(T)}$; the number of unions of complete $\mathrm{lcm}(T)$-fibers of total size A is $\binom{n/\mathrm{lcm}(T)}{A/\mathrm{lcm}(T)}$ when $\mathrm{lcm}(T) \mid A$ and 0 otherwise. Inclusion–exclusion gives $\Pi(A)$; the parameter bound is Theorem 9.2(ii) (resp. Lemma 3.17). The Chebyshev case is identical with $K_p \leq K_M$. □

This complements Theorem 3.24: that formula counts a *declared* remainder-profile family exactly; Corollary 9.3 counts the *intrinsic* periodic family exactly, and Theorem 9.2 proves that nothing quotient-periodic escapes it. What the periodic-support ledger still lacked was the passage from support-level to image-level counting. The next theorem supplies it exactly on the class of received data for which the quotient family was designed.

Theorem 9.4 (exact quotient descent for fiber-constant data)**.** *Work in the multiplicative or Chebyshev case, write φ_c for the scale-c map (X^c, resp. T_c) and $Q_c := \varphi_c(D)$, $N_c := n/c$. Let c be a scale, let $p \in \mathbb{F}[X]_{<k}$, let $S \subseteq D$ be a union of complete c-fibers with $|S| \geq k$, and suppose p agrees on S with a fiber-constant word $W \circ \varphi_c$. Then there is a unique $g \in \mathbb{F}[Y]$ with*

$$p = g \circ \varphi_c, \qquad \deg g < \lceil k/c \rceil,$$

and g agrees with W on $\varphi_c(S)$, a set of $|S|/c$ points of Q_c. Consequently, for a received line $f + zg_0$ with f and g_0 both constant on c-fibers, and any agreement threshold $a \geq k$: a finite slope z is MCA-bad with some witness support that is a union of complete c-fibers if and only if z is MCA-bad for the quotient line $(f_c, g_{0,c})$ on $\mathrm{RS}[\mathbb{F}, Q_c, \lceil k/c \rceil]$ at the corresponding agreement $\geq \lceil a/c \rceil$, with witness $\varphi_c(S)$; and the same holds verbatim for lists: $p \mapsto g$ is a bijection between scale-c-witnessed list elements and quotient list elements.

Proof. Multiplicative case. For $h \in H_c = \mu_c$ and $x \in S$ we have $hx \in S$ and $\varphi_c(hx) = \varphi_c(x)$, so $p(hx) = W(\varphi_c(hx)) = W(\varphi_c(x)) = p(x)$. Thus $p(hX) - p(X)$, of degree $< k$, has at least $|S| \geq k$ roots, hence vanishes identically: $p(hX) = p(X)$ for every $h \in \mu_c$. Comparing coefficients, $a_i h^i = a_i$ for all $h \in \mu_c$, so $a_i = 0$ whenever $c \nmid i$ (take h a generator of μ_c), i.e. $p = g(X^c)$ with $c \deg g = \deg p < k$, so $\deg g < \lceil k/c \rceil$; uniqueness is clear. For $x \in S$, $g(\varphi_c(x)) = p(x) = W(\varphi_c(x))$, and $\varphi_c(S)$ has $|S|/c$ elements since S is a union of complete c-fibers.

Chebyshev case. Work over the compositum $\mathbb{F}' := \mathbb{F} \cdot \mathbb{F}_{p^2}$ and put $P(u) := p(\chi(u)) \in \mathbb{F}'[u, u^{-1}]$, a Laurent polynomial with exponents in $[-(k-1), k-1]$. For $\zeta \in K_c = \mu_c$ and $u \in \widetilde{S}$ we have $\zeta u \in \widetilde{S}$ and $T_c(\chi(\zeta u)) = \chi(\zeta^c u^c) = T_c(\chi(u))$, so $P(\zeta u) = P(u)$ on $\widetilde{S}$; since $|\widetilde{S}| = 2|S| \geq 2k$ exceeds the degree $2(k-1)$ of $u^{k-1}\big(P(\zeta u) - P(u)\big)$, the identity $P(\zeta u) = P(u)$ holds in $\mathbb{F}'[u, u^{-1}]$. Hence only exponents divisible by c occur, and P is also inversion-symmetric because χ is, so P lies in the symmetric Laurent ring $\mathbb{F}'[u^c + u^{-c}]$: $P(u) = G(u^c + u^{-c})$ for a unique G. Since $u^c + u^{-c} = 2\,T_c(\chi(u))$, setting $g(Y) := G(2Y)$ gives $p(\chi(u)) = g(T_c(\chi(u)))$ identically, and injectivity of the ring map $r(x) \mapsto r((u+u^{-1})/2)$ yields $p = g \circ T_c$ in $\mathbb{F}'[X]$; as g is the unique solution of an $\mathbb{F}$-linear system with data in $\mathbb{F}$, in fact $g \in \mathbb{F}[Y]$, and $c \deg g = \deg p < k$ as before. Agreement on $\varphi_c(S) = T_c(S)$ follows as above, with $|T_c(S)| = |S|/c$ by Lemma 7.8.

Lines and lists. If z has an MCA witness S that is a union of complete c-fibers, the explaining codeword for the point $f + zg_0$ (a fiber-constant word) descends by the above; conversely a quotient explanation composes with φ_c. For noncontainment, if the pair (f, g_0) were explained on S by (p_1, p_2), both p_i descend (each agrees on S with a fiber-constant word), so the quotient pair

would be explained on $\varphi_c(S)$; and conversely quotient explanations compose. Hence badness with a scale-c witness corresponds exactly. For lists, $p \mapsto g$ is well defined and injective (uniqueness of g), and $g \mapsto g \circ \varphi_c$ is its inverse from the quotient list, since the full agreement set of $g \circ \varphi_c$ with $W \circ \varphi_c$ is the full φ_c-preimage of the agreement set of g with W. □

Corollary 9.5 (structural exhaustion is exact)**.** *With the intrinsic Definition 9.1, the structural exhaustion is a theorem: every bad parameter is quotient-periodic (periodically witnessed, hence counted exactly at the support level by Corollary 9.3, and exactly at the image level by Theorem 9.4 whenever the received line is fiber-constant) or belongs to the aperiodic bucket. The subfield order $|\mathbb{B}|$ appears only in lower floors; all upper numerators above are q-free. The entire residual content of the periodic-support ledger is thereby the aperiodic bucket, treated next.*

9.2 The deep regime, singular pencils, and the residual band

The following theorem is the safe-side counterpart of the paper's unsafe caps. It is self-contained, uses nothing beyond interpolation and double counting, and holds for arbitrary linear codes — in particular for every multiplicative row, every Chebyshev line-round row, every bivariate circle code over any field (including fields not containing i), and every ECFFT row.

Theorem 9.6 (deep-regime mutual correlated agreement, unconditional)**.** *Let $C \subseteq \mathbb{F}^n$ be any $\mathbb{F}$-linear code of minimum Hamming weight w, let $q = |\mathbb{F}|$, let $\delta \in (0,1)$, and put $r := \lfloor \delta n \rfloor$. Assume*

$$3r \;\le\; w - 1.$$

Then for every pair $(f_1, f_2) \in \mathbb{F}^n \times \mathbb{F}^n$, at least one of the following alternatives holds:

(a) (Tangent pair.) *There are codewords $c_1, c_2 \in C$ and a set $T' \subseteq \{1, \dots, n\}$ with $|T'| \le r$ such that $f_i = c_i$ off T' for $i = 1, 2$; in this case (f_1, f_2) has no CA-bad slope at radii $(\delta, \delta_{\mathrm{int}})$ for any $\delta_{\mathrm{int}} \ge \delta$, and at most $|T'| \le r$ MCA-bad slopes at radius δ, namely a subset of $\{-e_{1,j}/e_{2,j} : j \in T',\ e_{2,j} \ne 0\}$ where $e_i := f_i - c_i$.*

(b) (Generic pair.) $\#\{\gamma \in \mathbb{F} :\ \Delta(f_1 + \gamma f_2, C) \le \delta\} \;\le\; r + 1$.

Consequently

$$\varepsilon_{\mathrm{ca}}(C, \delta) \;\le\; \frac{r+1}{q}, \qquad \varepsilon_{\mathrm{mca}}(C, \delta) \;\le\; \frac{r+1}{q}.$$

For $C = \mathrm{RS}[\mathbb{F}, D, k]$ on any evaluation set D of size n, the hypothesis reads $3\lfloor \delta n \rfloor \le n - k$.

Proof. Write $S_{\mathrm{cl}} := \{\gamma : \Delta(f_1 + \gamma f_2, C) \le \delta\}$; note $\Delta \le \delta$ means disagreement on at most r coordinates, by integrality. If $|S_{\mathrm{cl}}| \le 1$, alternative (b) holds ($r + 1 \ge 1$). Otherwise fix $\gamma_1 \ne \gamma_2$ in S_{cl} with codewords p_1, p_2 and error sets E_1, E_2 ($|E_i| \le r$), and set

$$c_2 := \frac{p_1 - p_2}{\gamma_1 - \gamma_2}, \qquad c_1 := p_1 - \gamma_1 c_2, \qquad e_i := f_i - c_i.$$

Off $T := E_1 \cup E_2$ ($|T| \le 2r$) we can solve the two linear equations $f_1 + \gamma_i f_2 = p_i$ coordinatewise, giving $f_2 = c_2$ and $f_1 = c_1$ there; hence e_1, e_2 are supported in T. Let $T' := \{j : (e_{1,j}, e_{2,j}) \ne (0,0)\} \subseteq T$.

Step 1: every close slope rides the recovered line. Let $\gamma \in S_{\mathrm{cl}}$ with codeword p and error set E, $|E| \le r$. Off $T \cup E$ we have $p = f_1 + \gamma f_2 = c_1 + \gamma c_2$, so the codeword $p - (c_1 + \gamma c_2)$ has weight at most $|T| + |E| \le 3r \le w - 1 < w$ and is therefore zero. Hence

$$f_1 + \gamma f_2 - (c_1 + \gamma c_2) \;=\; e_1 + \gamma e_2 \quad \text{has weight at most } r \qquad \text{for every } \gamma \in S_{\mathrm{cl}}. \tag{$*$}$$

Step 2: dichotomy. Suppose $|T'| \ge r + 1$. For $j \in T'$, the affine function $\gamma \mapsto e_{1,j} + \gamma e_{2,j}$ is not identically zero, so it vanishes for at most one γ. By $(*)$, each $\gamma \in S_{\mathrm{cl}}$ needs at least $|T'| - r \ge 1$ vanishing coordinates inside T'. Double counting vanishings,

$$|S_{\mathrm{cl}}|\,\big(|T'| - r\big) \;\le\; |T'|, \qquad \text{hence} \qquad |S_{\mathrm{cl}}| \;\le\; \frac{|T'|}{|T'| - r} \;\le\; r + 1,$$

which is alternative (b). Otherwise $|T'| \le r$ and (f_1, f_2) is a tangent pair as in (a) with this T'.

Step 3: bad slopes of tangent pairs. Let (f_1, f_2) be tangent with data (c_1, c_2, T'), $|T'| \le r$. First, $\Delta_2((f_1, f_2), C^{=2}) \le |T'|/n \le r/n \le \delta \le \delta_{\mathrm{int}}$, so no slope is CA-bad at radii $(\delta, \delta_{\mathrm{int}})$. For MCA, note $2r \le w - 1$ (from $3r \le w - 1$ and $r \ge 0$), so two codewords agreeing on at least $n - 2r$ coordinates are equal. Let γ be MCA-bad with witness S, $|S| \ge n - r$. On $S \setminus T'$, which has at least $n - 2r$ elements, $f_1 + \gamma f_2 = c_1 + \gamma c_2$; hence the codeword explaining the point on S equals $c_1 + \gamma c_2$, and thus $e_1 + \gamma e_2$ vanishes on $S \cap T'$. Moreover, any pair explanation (p_1', p_2') on S satisfies $p_i' = c_i$ on the $\ge n - 2r$ points of $S \setminus T'$, forcing $p_i' = c_i$; so the pair fails to be explained on S exactly when $S \cap T' \ne \emptyset$ (every $j \in T'$ has $(e_{1,j}, e_{2,j}) \ne 0$). Therefore γ is MCA-bad if and only if $e_{1,j} + \gamma e_{2,j} = 0$ for some $j \in T'$: necessity was just shown; for sufficiency take $S := (\{1, \dots, n\} \setminus T') \cup \{j\}$, of size $\ge n - r$, on which $c_1 + \gamma c_2$ explains the point while the pair is not explained. Each $j \in T'$ contributes at most one slope, giving at most $|T'| \le r$ MCA-bad slopes.

Finally, generic pairs have at most $r + 1$ close slopes, hence at most $r + 1$ CA- or MCA-bad slopes (an MCA witness of size $\ge n - r$ implies $\Delta \le r/n \le \delta$); tangent pairs have 0 CA-bad and at most r MCA-bad slopes. Taking the maximum over pairs gives the two displayed error bounds. For $\mathrm{RS}[\mathbb{F}, D, k]$, $w = n - k + 1$, so $3r \le w - 1$ is $3r \le n - k$. □

Remark 9.7 (two radii; scope). The CA statement holds in the two-radius form: for any $\delta_{\mathrm{int}} \ge \delta_{\mathrm{fld}} = \delta$ with $3\lfloor \delta n \rfloor \le w - 1$, $\varepsilon_{\mathrm{ca}}(C, \delta, \delta_{\mathrm{int}}) \le (r + 1)/q$, since the tangent alternative kills the far condition at every $\delta_{\mathrm{int}} \ge \delta$. The constant $\frac{1}{3}$ can be improved toward $\frac{1}{2}$ of the distance by the methods of [BCHKS25]; we deliberately keep the three-line-per-step, formalization-ready proof — it feeds the formal finite core — and record the import separately. Note also that Theorem 9.6 applies to the bivariate circle code over $\mathbb{F}_p$ itself (minimum weight $n_c - k_c + 1$, by extending scalars to $\mathbb{F}_{p^2}$ and using Lemma 7.11), so the deep regime of the genus-one/circle transport question Question 22.3(b) is closed as well.

We next resolve a named obstruction of the residual-band analysis — the singular Hankel buckets — in the regime where Theorem 9.6 applies. First, the structural normal form of the identically singular branch.

Lemma 9.8 (identically singular pencils are locator families)**.** *In the setting of Section 3.4, fix an exact agreement A with $t_A \ge j_A + 1$ and set $M(Z) := H_{t_A, j_A}(u) + Z H_{t_A, j_A}(v)$. The following are equivalent:*

(i) *every $(j_A + 1) \times (j_A + 1)$ minor of $M(Z)$ vanishes identically in Z;*

(ii) *there is a nonzero polynomial vector* $\widetilde{v}(Z) = \sum_{i=0}^{d} Z^i v_i \in \mathbb{F}[Z]^{j_A+1}$ *with* $M(Z)\widetilde{v}(Z) \equiv 0$, *equivalently a chain*

$$H_{t_A,j_A}(u)\, v_0 = 0, \qquad H_{t_A,j_A}(u)\, v_i = -H_{t_A,j_A}(v)\, v_{i-1} \ (1 \le i \le d), \qquad H_{t_A,j_A}(v)\, v_d = 0.$$

In that case, writing $\lambda(Z)$ *for the last coordinate of* $\widetilde{v}(Z)$, *every* $Z = \gamma$ *with* $\lambda(\gamma) \neq 0$ *admits a monic degree-*j_A *syndrome annihilator for* $u + \gamma v$, *i.e. every such slope passes the rank test at exact agreement* A; *the at most* $\deg \lambda$ *remaining slopes pass it after column truncation. The rank test is necessary for badness at agreement* A *but not sufficient (the annihilator must in addition be D-split), so membership in the singular bucket never certifies badness by itself.*

Proof. (i) says $\operatorname{rank}_{\mathbb{F}(Z)} M(Z) \le j_A$, which holds if and only if the kernel of $M(Z)$ over the rational function field is nonzero; clearing denominators and dividing by common factors produces a nonzero polynomial kernel vector, and expanding $M(Z)\widetilde{v}(Z) \equiv 0$ in powers of Z gives the chain. The converse direction is immediate. Specializing $Z = \gamma$ with $\lambda(\gamma) \neq 0$ and normalizing the last coordinate to 1 produces the vector $\widetilde{\ell}$ of the chart atlas with $A(\ell) + \gamma B(\ell) = 0$, i.e. a monic annihilator; $\lambda \neq 0$ has at most $\deg \lambda$ roots. The final sentence restates the convention of Section 3.4. □

Corollary 9.9 (the singular buckets are resolved in the deep regime)**.** *Let* $C = \mathrm{RS}[\mathbb{F}, D, k]$ *and let* A *be an exact agreement with* $3(n - A) \le n - k$. *Then for every received line — in particular for every line in the underdetermined or identically singular buckets of Section 3.4 — the number of finite slopes that are bad at agreement threshold* A *is at most* $n - A + 1$ *if the pair is not* $\frac{n-A}{n}$*-tangent, and at most* $n - A$ *MCA-bad (zero CA-bad) slopes if it is. In particular, in this regime the singular bucket contributes at most* $n - A + 1$ *to* $B_{\mathrm{ap}}(A)$ *per line, and no line contributes* $\Theta(q)$*: the "totally bad line" scenario that the singular bucket was guarding against does not occur below one third of the distance.*

Proof. Immediate from Theorem 9.6 with $r = n - A$ and $\delta = r/n$: badness at agreement threshold A implies distance at most r/n. □

Remark 9.10 (calibration: quotient floors do not exhaust unsafe mass). The quotient-fiber witnesses of Lemmas 3.1 and 7.2 are periodic: their agreement supports are unions of complete nontrivial fibers and are charged to the quotient ledger by Theorems 9.2 and 9.4. The identity scale $c = 1$ is different. The prefix witnesses of Lemma 14.1 need not be periodic, and Theorem 18.4 and Corollary 18.6 later show that at odd agreement every witness in the relevant pole-line fiber is aperiodic and that the resulting slope count can be exponential. Thus the periodic quotient ledger remains valid, but it cannot support a polynomial aperiodic bound below the entropy–subfield envelope.

Remark 9.11 (Refuted unnormalized band proposal). An earlier draft proposed the following statement. For a multiplicative or Chebyshev row $C = \mathrm{RS}[\mathbb{F}, D, k]$, one asks for a polynomial P, independent of q, such that for every agreement

$$k + \frac{2n}{N} < a < n - \left\lfloor \frac{n-k}{3} \right\rfloor,$$

the number of MCA-bad finite slopes of any received line whose witness supports are all aperiodic is at most $P(n)$. This proposal is retained only as a historical comparison. It is false below the entropy–subfield envelope by Corollary 18.6; no theorem or conditional closure in version 13.2 assumes it.

Remark 9.12 (Normalized aperiodic band formulation). Let $g^*(\rho, \beta)$ be the frontier function of Definition 14.4, with $\beta = \log_2 |\mathbb{B}|$. The operative aperiodic upper problem begins only at

$$a \geq k + g^*(\rho, \beta)n + R(n).$$

The sub-envelope range is assigned to the unsafe identity-prefix floor. In an asymptotic theorem the reserve is required to satisfy $R(n) \to \infty$ and $R(n) = o(n)$, and a multiplicative loss $L(n)$ is permitted only when

$$\log L(n) = o\big(R(n) \log |\mathbb{B}|\big).$$

Thus a polynomial loss is absorbed by $R \gg \log n$, while a generic $e^{o(n)}$ loss is not automatically absorbed. At a deployed adjacent value $a_0 + 1$, no asymptotic absorption is allowed: every multiplier, additive error, and bucket-summing cost must be included in one exact integer upper ledger.

9.3 The exact extension-line lift

Theorem 9.13 (restriction of scalars for lines)**.** *Let $\mathbb{B} \subseteq \mathbb{F}$ with $e := [\mathbb{F} : \mathbb{B}]$ and $D \subseteq \mathbb{B}$, and fix a $\mathbb{B}$-basis $1 = \beta_0, \beta_1, \ldots, \beta_{e-1}$ of $\mathbb{F}$. Decompose words as $w = \sum_i \beta_i w_i$ with $w_i \in \mathbb{B}^n$, and let $\mu_{ij}(\gamma) \in \mathbb{B}$ be the multiplication tensor, $\gamma\beta_i = \sum_j \mu_{ij}(\gamma)\beta_j$, which is $\mathbb{B}$-linear and injective in γ. Then:*

(i) *A word c lies in $\mathrm{RS}[\mathbb{F}, D, k]$ if and only if every component c_j lies in $\mathrm{RS}[\mathbb{B}, D, k]$, and for all words w,*

$$\Delta\big(w, \mathrm{RS}[\mathbb{F}, D, k]\big) \;=\; \Delta_e\big((w_0, \ldots, w_{e-1}),\, \mathrm{RS}[\mathbb{B}, D, k]^{\equiv e}\big).$$

More generally, a set $S \subseteq D$ explains w over $\mathbb{F}$ if and only if S column-explains the component tuple over $\mathbb{B}$.

(ii) *For a pair (f_1, f_2) over $\mathbb{F}$ and $\gamma \in \mathbb{F}$, the components of $f_1 + \gamma f_2$ are*

$$(f_1 + \gamma f_2)_j \;=\; f_{1,j} + \sum_i \mu_{ij}(\gamma)\, f_{2,i},$$

an e-dimensional $\mathbb{B}$-linear (in γ) family of e-interleaved words over $\mathbb{B}$.

(iii) *For every radius (and pair of radii), γ is CA-bad (resp. MCA-bad) for (f_1, f_2) over $\mathbb{F}$ if and only if the parameter γ of the sampler in* (ii) *is bad in the corresponding interleaved sense over $\mathbb{B}$, with pair noncontainment read on the $2e$-tuple $(f_{1,0}, \ldots, f_{1,e-1}, f_{2,0}, \ldots, f_{2,e-1})$ against $\mathrm{RS}[\mathbb{B}, D, k]^{\equiv 2e}$, on the same supports. For $\gamma \in \mathbb{B}$ the sampler restricts to the diagonal e-interleaved line $(f_{1,j} + \gamma f_{2,j})_j$.*

Consequently the extension term $B_{\mathrm{ext}}(a)$ of Theorem 8.2 is not an independent quantity: it equals, definitionally, a base-field interleaved-sampler numerator. In particular no "extension-valued counterexample" can exist outside the base ledgers: every genuinely $\mathbb{F}$-valued line is charged, exactly, to a base bucket.

Proof. (i) is the componentwise decomposition already used in the proof of Lemma 6.1 (valid for any $D \subseteq \mathbb{B}$, cf. Corollary 7.14): the coefficients of a polynomial over $\mathbb{F}$ decompose in the basis, evaluation at $D \subseteq \mathbb{B}$ commutes with the decomposition, and pointwise agreement over $\mathbb{F}$

is simultaneous componentwise agreement, i.e. column agreement. (ii) is the definition of μ. (iii) The point condition, the pair condition, and every support-restricted condition in Definitions 2.1 and 2.2 are, by (i), equalities of components on sets of coordinates, and these translate verbatim; injectivity of $\gamma \mapsto \mu(\gamma)$ makes the parameterizations match one-to-one. The diagonal claim is $\mu_{ij}(\gamma) = \gamma\delta_{ij}$ for $\gamma \in \mathbb{B}$. □

Corollary 9.14 (status of the extension-line lift)**.** *The extension-line question is discharged on the "lift" side of its dichotomy: extension lines are covered — exactly, not merely up to a stable numerator — by base-field ledgers, namely the interleaved-sampler ledgers of the list and curve sides via Theorem 9.13. In the deep regime nothing at all remains: Theorem 9.6 applies over $\mathbb{F}$ directly, uniformly in the field of definition of the pair. On the lower side, Corollaries 2.12 and 2.13 already produce genuinely extension-valued bad lines from any supplied list, and Lemma 6.1 (with Corollaries 6.2 and 7.14) shows the certified deployed failures must be of this kind. What the extension-line lift leaves behind is the base-field interleaved instance of the normalized residual ledger Remark 9.12 and Section 21; extension-valued lines do not create a fourth bucket.*

9.4 Uniform interleaved-list bounds, exact recursion, and calibration

Theorem 9.15 (elementary Johnson list bound, uniform in the interleaving arity)**.** *Let $C = \mathrm{RS}[\mathbb{F}, D, k]$ on any evaluation set D of size n, let $m \geq 1$, and let a be an agreement threshold with $a^2 > (k-1)n$. Then for every received m-tuple U,*

$$\left|\Lambda\big(C^{\equiv m},\, 1 - \tfrac{a}{n},\, U\big)\right| \;\leq\; \frac{n\,(a-k+1)}{a^2-(k-1)n},$$

where the list is taken in the column distance Δ_m. The same bound holds for any code family in which two distinct codewords column-agree on at most $k-1$ coordinates; in particular it holds for interleaved bivariate circle codes with k replaced by k_c.

Proof. Let $\mathbf{c}^{(1)}, \ldots, \mathbf{c}^{(L)}$ be distinct listed codewords with column-agreement sets $S_1, \ldots, S_L \subseteq D$, $|S_i| \geq a$. Two distinct interleaved codewords column-agree exactly where all their rows agree; some pair of rows consists of distinct polynomials of degree $< k$, so $|S_i \cap S_j| \leq k-1$ for $i \neq j$. Let $\deg(x) := \#\{i : x \in S_i\}$, so $\sum_x \deg(x) \geq La$ and

$$\sum_{x \in D} \binom{\deg(x)}{2} = \sum_{i<j} |S_i \cap S_j| \leq \binom{L}{2}(k-1).$$

By convexity of $t \mapsto t(t-1)$, $\sum_x \deg(x)(\deg(x)-1) \geq n \cdot \frac{La}{n}\left(\frac{La}{n} - 1\right)$ (the right side is nonpositive when $La \leq n$, and the average-value bound applies otherwise). Hence $La(La-n) \leq L(L-1)(k-1)n$, i.e. $a(La-n) \leq (L-1)(k-1)n$, i.e.

$$L\big(a^2-(k-1)n\big) \;\leq\; an-(k-1)n \;=\; n(a-k+1),$$

which is the claim when $a^2 > (k-1)n$. □

Theorem 9.16 (exact divisor-lattice list recursion for fiber-constant words)**.** *Multiplicative case; let $c_0 \mid n$, let $U = W \circ \pi_{c_0}$ be constant on c_0-fibers, and let $a \geq k$. For $p \in \Lambda(C, 1-a/n, U)$ let $S_p := \{x \in D : p(x) = U(x)\}$ be the full agreement set and put $c_p := \gcd\big(c(S_p), c_0\big)$. For $c \mid c_0$*

write $W_c(y) := W\big(y^{c_0/c}\big)$ *on* Q_c*, so* $U = W_c \circ \pi_c$*. Then, for each* $c \mid c_0$*, the descent map of Theorem 9.4 is a bijection from* $\{p \in \Lambda(C, 1 - a/n, U) : \ c_p = c\}$ *onto*

$$\Big\{ g \in \Lambda\big(\mathrm{RS}[\mathbb{F}, Q_c, \lceil k/c \rceil],\, 1 - \lceil a/c \rceil / N_c,\, W_c\big) \ : \ \gcd\big(c(\bar{S}_g),\, c_0/c\big) = 1 \Big\},$$

where $\bar{S}_g$ *is the full agreement set of* g *with* W_c *and its scale is computed in the quotient lattice. In particular*

$$\big|\Lambda(C, 1 - \tfrac{a}{n}, U)\big| = \sum_{c \mid c_0} \#\{\cdots\},$$

and if moreover $a^2 > kn$*, then every bucket obeys the level-*c *Johnson bound and*

$$\big|\Lambda(C, 1 - \tfrac{a}{n}, U)\big| \ \le \ \sum_{c \mid c_0} \frac{n\,(a - k + 2c - 1)}{a^2 - kn} \ \le \ \tau(c_0)\, \frac{n\,(a - k + 2c_0 - 1)}{a^2 - kn},$$

τ *the number-of-divisors function.*

Proof. Fix p with $c_p = c$. Then $c \mid c(S_p)$, so S_p is a union of complete c-fibers by Theorem 9.2(i), and $U = W_c \circ \pi_c$ is constant on c-fibers; Theorem 9.4 (with $S = S_p$, $|S_p| \ge a \ge k$) gives $p = g \circ \pi_c$ with g in the displayed quotient list, and $S_p = \pi_c^{-1}(\bar{S}_g)$ because $p(x) = U(x) \iff g(x^c) = W_c(x^c)$. The power map is equivariant, $\pi_c(hx) = h^c \pi_c(x)$, and Q_c is a coset of $H^{(c)} := \{h^c : h \in H\}$ with kernel H_c; hence $hS_p = S_p \iff h^c \bar{S}_g = \bar{S}_g$ and $c(S_p) = c \cdot c(\bar{S}_g)$. Since $c \mid c_0$, $\gcd(c \cdot c(\bar{S}_g), c_0) = c \cdot \gcd(c(\bar{S}_g), c_0/c)$, so $c_p = c$ is exactly the displayed coprimality. Conversely, for g in the displayed set, $p := g \circ \pi_c$ has full agreement set $\pi_c^{-1}(\bar{S}_g)$ of size $c|\bar{S}_g| \ge c\lceil a/c \rceil \ge a$, and the same computation returns $c_p = c$; the two maps are mutually inverse. Summing the buckets gives the exact formula. For the Johnson clause: at level c the code has dimension $\lceil k/c \rceil$ and length $N_c = n/c$, and

$$\Big\lceil \frac{a}{c} \Big\rceil^2 - \Big(\Big\lceil \frac{k}{c} \Big\rceil - 1 \Big) \frac{n}{c} \ \ge \ \frac{a^2}{c^2} - \frac{kn}{c^2} \ = \ \frac{a^2 - kn}{c^2} \ > \ 0,$$

using $\lceil k/c \rceil - 1 < k/c$, so Theorem 9.15 applies at every level; with $\lceil a/c \rceil - \lceil k/c \rceil + 1 \le (a - k + 2c - 1)/c$ the level-c bound is at most $n(a - k + 2c - 1)/(a^2 - kn)$. □

Proposition 9.17 (calibration: the quotient list term is necessarily word-dependent and exponential)**.** *Let the hypotheses of Lemma 7.2*(ii) *hold and let* $U = u_z$ *be the pigeonholed received word at agreement* A_2*. Every listed codeword* $c_A = -R_A(\varphi)$ *lies in the scale-*a *quotient image* $\{g \circ \varphi : \deg g \le \ell_2 - 2\}$*. Hence in any bound of the shape demanded by the interleaved-list ledger,*

$$|\Lambda(C^+, 1 - \tfrac{A_2}{n}, U)| \ \le \ L_Q(A_2, U) + L_{\mathrm{ap}}(A_2),$$

in which L_Q *upper-bounds the quotient-image sublist and* L_{ap} *the rest, one necessarily has*

$$L_Q(A_2, u_z) \ \ge \ \binom{N}{\ell_2} \Big/ |\mathbb{B}|.$$

No uniform subexponential L_Q *exists in the fiber-unsafe sub-band. Moreover, the identity-scale floors of Theorem 18.4 and Corollary 18.6 rule out a uniform polynomial* L_{ap} *below the entropy–subfield envelope. The residual upper ledger is therefore posed only in the normalized range of Remark 9.12.*

Proof. Immediate from Lemma 7.2: the $\geq \binom{N}{\ell_2}/|\mathbb{B}|$ listed codewords are of the displayed composed form, hence are counted by L_Q. □

Corollary 9.18 (the grand interleaved-list staircase outside the band)**.** *Fix any row satisfying the hypotheses of Corollary 5.1, Corollary 7.6, or Corollary 7.12, and any interleaving arity $m \geq 1$.*

(i) (Unsafe near capacity.) *For every agreement $a \leq k + n/N$ (radius $\geq 1 - \rho - 1/N$),*

$$\left|\Lambda\big(C^{\equiv m},\, 1 - \tfrac{a}{n}\big)\right| \;\geq\; \binom{N}{\ell_1}\Big/|\mathbb{B}| \;>\; 2^{-128}\, q,$$

by Lemma 7.2(i) and the entropy counts of the corresponding corollary, together with the embedding $c \mapsto (c, 0, \ldots, 0)$ of base lists into interleaved lists.

(ii) (Safe in the Johnson regime.) *For every agreement a with $a^2 > (k-1)n$ and $n(a - k + 1)/(a^2 - (k-1)n) \leq 2^{-128}q$,*

$$\left|\Lambda\big(C^{\equiv m},\, 1 - \tfrac{a}{n}\big)\right| \;\leq\; 2^{-128}\, q,$$

by Theorem 9.15, uniformly in m.

Thus the interleaved-list staircase is determined outside the same Johnson-to-capacity band as the MCA staircase; inside the band, Proposition 9.17 shows that any resolution of the interleaved-list ledger must carry an exponential, word-dependent quotient term, and Theorem 9.16 shows that for fiber-constant words this term is exactly a divisor-lattice recursion.

Proof. (i) At agreement $A_1 \leq k + n/N$ the base-field list floor is $\binom{N}{\ell_1}/|\mathbb{B}|$ by Lemma 7.2(i); list sizes are nondecreasing as the agreement threshold decreases, and the zero-padding embedding preserves column agreement, so the interleaved list is at least as large. The entropy counts of the cited corollaries give $\binom{N}{\ell_1} \geq 2^{513} \geq |\mathbb{B}| \cdot 2^{-128}q \cdot 2^0$ throughout the challenge envelope, since $|\mathbb{B}|\, q\, 2^{-128} \leq q^2\, 2^{-128} < 2^{384}$. (ii) is Theorem 9.15. □

9.5 Curves and projective slopes in the deep regime

Theorem 9.19 (deep-regime theorem for degree-d power-curve samplers)**.** *Let $C \subseteq \mathbb{F}^n$ be $\mathbb{F}$-linear with minimum weight w, let $d \geq 1$, let $W_\gamma = f_0 + \gamma f_1 + \cdots + \gamma^d f_d$ be a power-curve sampler, let $\delta \in (0,1)$, $r := \lfloor \delta n \rfloor$, and assume*

$$(d+2)\, r \;\leq\; w - 1.$$

Then either there are codewords $c_0, \ldots, c_d \in C$ and a set T' with $|T'| \leq r$ such that $f_i = c_i$ off T' for every i (curve-tangent tuple; at most $d\,|T'| \leq dr$ MCA-bad and no CA-bad parameters), or

$$\#\{\gamma \in \mathbb{F} :\; \Delta(W_\gamma, C) \leq \delta\} \;\leq\; d\,(r+1).$$

In particular the curve analogues of $\varepsilon_{\mathrm{ca}}$ and $\varepsilon_{\mathrm{mca}}$ (with noncontainment read on the coefficient tuple against $C^{\equiv(d+1)}$, as in Lemma 3.17) are at most $d(r+1)/q$ at radius δ. Adjoining the projective parameter at infinity (the leading word f_d) adds at most one to every count. For $C = \mathrm{RS}[\mathbb{F}, D, k]$ the hypothesis reads $(d+2)\lfloor \delta n \rfloor \leq n - k$.

Proof. Let S_{cl} be the displayed set. If $|S_{\mathrm{cl}}| \le d$ the second alternative holds. Otherwise pick distinct $\gamma_0, \dots, \gamma_d \in S_{\mathrm{cl}}$ with codewords p_j and error sets E_j ($|E_j| \le r$), and let $c_i := \sum_j \lambda_{ij} p_j \in C$, where (λ_{ij}) is the inverse of the Vandermonde matrix (γ_j^i); off $T := \bigcup_j E_j$ ($|T| \le (d+1)r$) the same inversion applied to values gives $f_i = c_i$. Put $e_i := f_i - c_i$, supported in T, and $T' := \{x : (e_{0,x}, \dots, e_{d,x}) \neq 0\}$. For any $\gamma \in S_{\mathrm{cl}}$ with codeword p and error set E: off $T \cup E$, $p = W_\gamma = \sum_i \gamma^i c_i$, so the codeword $p - \sum_i \gamma^i c_i$ has weight at most $(d+2)r < w$, hence is zero, and

$$W_\gamma - \sum_i \gamma^i c_i = \sum_i \gamma^i e_i \quad \text{has weight at most } r.$$

For $x \in T'$ the nonzero polynomial $\gamma \mapsto \sum_i e_{i,x}\gamma^i$ has at most d roots. If $|T'| \ge r+1$, each $\gamma \in S_{\mathrm{cl}}$ needs at least $|T'| - r$ vanishing coordinates in T', so $|S_{\mathrm{cl}}|(|T'| - r) \le d|T'|$ and $|S_{\mathrm{cl}}| \le d|T'|/(|T'| - r) \le d(r+1)$. Otherwise $|T'| \le r$ and the tuple is curve-tangent. In the tangent case, $2r \le w - 1$ makes two codewords agreeing on $\ge n - 2r$ coordinates equal; arguing exactly as in Step 3 of Theorem 9.6 (with $n-(d+2)r \ge n-w+1$ forcing both the point codeword and any tuple explanation to coincide with the c_i off T'), a parameter is MCA-bad if and only if $\sum_i e_{i,x}\gamma^i = 0$ for some $x \in T'$, giving at most $d|T'|$ parameters, and the coefficient tuple is δ-close to $C^{\equiv(d+1)}$ so no parameter is CA-bad. MCA-bad parameters are close, so the counts cover both notions; the parameter at infinity is a single additional value. For Reed–Solomon, $w = n - k + 1$. □

Together with Theorem 9.13, Theorem 9.19 also covers the multi-parameter samplers into which extension lines decompose, since those are handled over $\mathbb{F}$ directly by Theorem 9.6; what the curve-sampler ledger still lacks is only the band-regime aperiodic pivot atlas, i.e. the curve instance of the broad band formulation Remark 9.12, together with the exceptional vanishing-vector strata already flagged in Section 3.4.

9.6 The certificate scanner, its soundness, and deployed output

Proposition 9.20 (small-field degeneracy)**.** *Let $C \subsetneq \mathbb{F}^n$ be any linear code and $q = |\mathbb{F}|$. For every $\delta \in (0,1)$, $\varepsilon_{\mathrm{mca}}(C, \delta) \ge 1/q$. Consequently $\delta^*_C(\varepsilon^*) = 0$ whenever $\varepsilon^* < 1/q$; in particular every row over a field of size $q < 2^{128}$ — including the deployed QM31 circle rows, $q = p^4 < 2^{124}$ — has a degenerate grand-challenge threshold $\delta^*_C(2^{-128}) = 0$, and the meaningful staircase question there is posed at a larger target such as $\varepsilon^* = 2^{-100}$. At equality $1/q = \varepsilon^*$ the construction below gives only a floor meeting the target exactly, not strict unsafety; all degeneracy statements therefore require $1/q > \varepsilon^*$.*

Proof. Pick $f_2 \notin C$, any $c_0 \in C$ and $\gamma_0 \in \mathbb{F}$, and set $f_1 := c_0 - \gamma_0 f_2$. With $S = D$: the point $f_1 + \gamma_0 f_2 = c_0$ is explained on S, while the pair is explained on $S = D$ only if both words lie in C, which fails. Hence γ_0 is MCA-bad at every radius, and $\varepsilon_{\mathrm{mca}} \ge 1/q$. □

Definition 9.21 (staircase scanner)**.** The scanner takes a row header (D-type, $n, k, |\mathbb{B}|, q, q_{\mathrm{line}} = q$), a target ε^*, and the applicable divisor data, and for each agreement $a \in \{k+1, \dots, n\}$ evaluates the following verdicts, each tagged by its certifying statement:

(V1) **Cap-unsafe.** If $a \le A_2$ for an instantiated cap (Theorems 4.1 and 7.4 via Corollaries 5.1, 7.6, 7.12 and 7.13) and $\frac{1}{2k}(1 - \frac{n}{q}) > \varepsilon^*$: mark a MCA-unsafe.

(V2) **Floor-unsafe.** If a remainder floor applies at a (Theorem 3.8) with converted numerator exceeding $\varepsilon^* q$: mark a MCA-unsafe; if a fiber list floor at a exceeds $\varepsilon^* q$ (Lemmas 3.1 and 7.2 with the entropy count): mark a list-unsafe.

(V3) **Deep-safe.** If $3(n-a) \le n-k$ and $n-a+1 \le \varepsilon^* q$: mark a MCA-safe (Theorem 9.6); with $(d+2)(n-a) \le n-k$ and $d(n-a+1) \le \varepsilon^* q$: curve-$d$-safe (Theorem 9.19).

(V4) **Johnson-list-safe.** If $a^2 > (k-1)n$ and $n(a-k+1)/(a^2-(k-1)n) \le \varepsilon^* q$: mark a list-safe for every interleaving arity (Theorem 9.15).

The output per challenge type is the pair $a_{\text{unsafe}} := \max\{a \text{ marked unsafe}\}$ and $a_{\text{safe}} := \min\{a \text{ marked safe}\}$, together with the open interval between them. If $\varepsilon^* < 1/q$ the scanner instead reports the degeneracy of Proposition 9.20.

Theorem 9.22 (scanner soundness)**.** *Every verdict printed by the scanner of Definition 9.21 is correct: agreements marked unsafe satisfy $\varepsilon_{\text{mca}}(C, 1-a/n) > \varepsilon^*$ (resp. list size $> \varepsilon^* q$), agreements marked safe satisfy $\varepsilon_{\text{mca}}(C, 1-a/n) \le \varepsilon^*$ (resp. list size $\le \varepsilon^* q$, for every m), and consequently $a_{\text{unsafe}} < a_C^* \le a_{\text{safe}}$ in the staircase formulation of Proposition 8.1.*

Proof. Each verdict cites a theorem of this paper whose hypotheses are verified numerically by the scanner; monotonicity across agreements is Lemma 2.5 for MCA and radius-monotonicity of lists. The bracketing of a_C^* is Proposition 8.1. □

Row / challenge	Radius range δ	Verdict	Certified numerator	Source
KoalaBear sextic: $n = 2^{21}$, $k = 2^{20}$, $q = p^6 > 2^{185.9}$, $\varepsilon^* = 2^{-128}$, $\varepsilon^* q \approx 2^{57.9}$				
MCA	$[\frac{1}{2} - 2^{-7}, \frac{1}{2})$	unsafe	$> 2^{-22}\, q \approx 2^{164}$	Corollary 5.4
MCA	$\delta \le 349525/2^{21} \approx 0.16666$	safe	$\le 349526 < 2^{-167} q$	Theorem 9.6
MCA	(0.16667, 0.49218)	open	—	the broad band formulation Remark 9.12
List (any m)	$\delta \ge \frac{1}{2} - 2^{-8}$	unsafe	$\ge \binom{256}{129}/p \ge 2^{216}$	Lemma 3.1(i)
List (any m)	$\delta \le 0.29289$ ($a \ge 1482910$)	safe	$\le 2^{20}$	Theorem 9.15
List (any m)	(0.29290, 0.49609)	open	—	the broad band formulation Remark 9.12
M31 circle line round: $n = 2^{21}$, $k = 2^{20}$, $q = p^4 < 2^{124}$; at 2^{-128}: $\delta^* = 0$ (Proposition 9.20); below: $\varepsilon^* = 2^{-100}$				
MCA	$[\frac{1}{2} - 2^{-7}, \frac{1}{2})$	unsafe	$> 2^{-22}\, q \approx 2^{102}$	Corollary 7.13
MCA	$\delta \le 0.16666$	safe	$\le 349526 < 2^{-105} q$	Theorem 9.6
List (any m)	$\delta \ge \frac{1}{2} - 2^{-8}$	unsafe	$\ge 2^{216}$	Lemma 7.2(i)
List (any m)	$\delta \le 0.29289$	safe	$\le 2^{20} \le 2^{-100} q$	Theorem 9.15

Table 2: Baseline scanner output (Definition 9.21) for the two deployed rows, before the pincer and the optimized graded-prefix floors; it is retained to display the self-contained deep-safe and Johnson-list-safe certificates, and the final deployed threshold bands are the narrowed bands of Table 5. The open MCA band for the circle row at $\varepsilon^* = 2^{-100}$ is the same interval (0.16667, 0.49218) as above.

The table realizes the interval output demanded by the staircase interface: for the KoalaBear MCA staircase, $a_{\text{unsafe}} = 1{,}064{,}960$ ($\delta = \frac{1}{2} - 2^{-7}$) and $a_{\text{safe}} = 1{,}747{,}627$ ($\delta \approx 0.16666$); the interval shrinks to one step exactly when the broad band formulation Remark 9.12 is settled, via Theorem 8.2 whose other inputs are now discharged (Remark 9.28).

9.7 The formal finite core

Proposition 9.23 (finite certificates and decidability)**.** *Fix a row (C, q, q_{line}) and an agreement a. The numerator $B_{\text{mca}}(a)$ of Section 8 is a well-defined integer, computable by finite enumeration over pairs, slopes, supports, and codewords; hence the one-step certificate of Proposition 8.1 is a finite object. A verifiable prize certificate consists of: the row header; the integer a_0; a* lower witness *— a pair (f_1, f_2), a set of more than $\varepsilon^* q_{\text{line}}$ slopes, and for each slope a codeword and a support, checkable by evaluation — or a theorem tag from this paper with its numerically verified hypotheses; and an* upper witness *— a theorem tag with numerically verified hypotheses certifying $B_{\text{mca}}(a_0 + 1) \le \varepsilon^* q_{\text{line}}$. Verification of a certificate is polynomial in its length.*

Proof. All quantifiers in Definitions 2.1 and 2.2 range over the finite sets $\mathbb{F}^n$, $\mathbb{F}$, 2^D, and C; the maximum defining $B_{\text{mca}}(a)$ is therefore attained and computable. Checking a lower witness is evaluation and counting; checking a theorem tag is checking finitely many integer inequalities. □

Core statement	Proof ingredients
Theorem 2.6 (deep-point conversion)	polynomial division, pigeonhole over $\mathbb{F} \setminus D$
Lemma 3.1, Lemma 7.2 (fiber floors)	elementary symmetric functions, subfield pigeonhole
Fact 2.4, Lemma 2.5	set manipulation
Proposition 8.1 (staircase)	monotone integer staircase
Theorem 9.6 (deep MCA)	linear interpolation, double counting of vanishings
Theorem 9.15 (Johnson lists)	convexity double counting of incidences
Theorem 9.2, Theorem 9.4	cyclic group actions, root counting
Theorem 9.13 (restriction of scalars)	basis decomposition of $\mathbb{F}/\mathbb{B}$
Definition 9.21, Theorem 9.22	finite case analysis over the statements above

Table 3: The formal finite core: each statement quantifies over explicitly bounded finite sets and uses only the listed elementary ingredients; the dependency order is top to bottom. Formalizing this table in a proof assistant is the outstanding transcription task.

The list deliberately includes the new safe-side theorems: a certificate package that contains both an unsafe certificate (Corollary 5.1 and its relatives) and a machine-checked safe region (Theorems 9.6 and 9.15) documents the staircase from both sides in the sense of Proposition 8.1.

9.8 Tower transfer and the genus-one first step

The tower-transfer ledger half is already discharged: Theorems 9.2 and 9.4 and Corollary 9.3 were stated and proved for the Chebyshev case alongside the multiplicative case, with the divisor lattice of D replaced by the dyadic scale lattice of the twin coset, exactly as the tower transfer requires; and the broad band formulation Remark 9.12 is likewise stated for both cases. It remains to treat the genus-one floor. We first extend Definition 7.1 to rational maps.

Definition 9.24 (rationally smooth domain)**.** Let $\mathbb{B} \subseteq \mathbb{F}$, let $\psi = f/g$ with $f, g \in \mathbb{B}[X]$ coprime, $a := \deg f > e := \deg g \geq 0$, and let $D \subseteq \mathbb{B}$ with g nonvanishing on D. The set D is (ψ, a)-*smooth* if every nonempty fiber $\psi^{-1}(w) \cap D$, $w \in \psi(D)$, has exactly a elements; then $N := |D|/a = |\psi(D)|$. The FFTree domains of ECFFT [BCKL21] are $(\psi, 2)$-smooth for the x-coordinate map ψ of each 2-isogeny in the chain, of degree type $(a, e) = (2, 1)$.

Proposition 9.25 (rational-map fiber floor; unified (a, e) form)**.** *Let $D \subseteq \mathbb{B}$ be (ψ, a)-smooth of size n, with $\psi = f/g$ of degree type (a, e), and let $2 \leq \ell \leq N - 1$. Put*

$$k' \;:=\; a\ell - 2(a-e), \qquad u_z \;:=\; \big(f(x)^\ell + z\, f(x)^{\ell-1} g(x)\big)_{x \in D} \in \mathbb{B}^n.$$

Then there exists $z \in \mathbb{B}$ with

$$\big|\Lambda\big(\mathrm{RS}[\mathbb{F}, D, k'+1],\, 1 - \tfrac{a\ell}{n},\, u_z\big)\big| \;\geq\; \binom{N}{\ell} \Big/ |\mathbb{B}|,$$

i.e. a list at agreement $a\ell = k' + 2(a-e)$. For $e = 0$ this is Lemma 7.2(ii) *(window $2a$, gap $\approx 2/N$); for $e = a - 1$ the window is 2 and the certified agreement is $k' + 2$ — the first staircase step, in accordance with the collapse of Remark 7.16.*

Proof. For $A \subseteq \psi(D)$ with $|A| = \ell$ put $\Lambda_A := \prod_{b \in A}(f - bg)$, of degree $a\ell$ with leading coefficient $\mathrm{lc}(f)^\ell \neq 0$, vanishing at every $x \in D$ with $\psi(x) \in A$ — exactly $a\ell$ points by smoothness (and $g(x) \neq 0$). Expanding,

$$\Lambda_A = \sum_{j=0}^{\ell} (-1)^j e_j(A)\, f^{\ell-j} g^j,$$

with j-th term of degree $a\ell - j(a-e)$. Set $z_A := -e_1(A) \in \mathbb{B}$ and

$$c_A := \Big(-\sum_{j \geq 2} (-1)^j e_j(A)\, f^{\ell-j}(x)\, g^j(x)\Big)_{x \in D},$$

a polynomial of degree at most $a\ell - 2(a-e) = k'$, so $c_A \in \mathrm{RS}[\mathbb{B}, D, k'+1]$; on the $a\ell$ fiber points, $u_{z_A} = c_A$. For injectivity at fixed z: if $c_A = c_{A'}$ as words then as polynomials (deg $\leq k' \leq a(N-1) < n$), so $\sum_{j \geq 2} (-1)^j \big(e_j(A) - e_j(A')\big) f^{\ell-j} g^j = 0$; dividing by g^ℓ in $\mathbb{F}(X)$ gives a polynomial identity in ψ, and ψ is nonconstant, hence transcendental over $\mathbb{F}$ inside $\mathbb{F}(X)$, so all coefficients vanish: $e_j(A) = e_j(A')$ for $j \geq 2$, and e_1 agrees via z; then $A = A'$ from the root sets of $\prod_b (Y - b)$. Pigeonholing z_A over $\mathbb{B}$ finishes as in Lemma 7.2. □

Corollary 9.26 (every ECFFT row has an unsafe first staircase step)**.** *Let D be an ECFFT domain that is $(\psi, 2)$-smooth (Definition 9.24) over $\mathbb{B}$, let $C = \mathrm{RS}[\mathbb{F}, D, k]$ with $2 \mid k$, $\rho = k/n \in [\frac{1}{16}, \frac{1}{2}]$, $n < q = |\mathbb{F}| < 2^{256}$, and $n \geq 2^{12}$. Then, with $\ell = (k+2)/2 \leq N - 1$,*

$$\binom{n/2}{(k+2)/2} \;\geq\; |\mathbb{B}| \Big(\frac{q}{k} + 1\Big),$$

and consequently

$$\varepsilon_{\mathrm{ca}}\Big(C,\, 1 - \rho - \tfrac{2}{n}\Big) \;>\; \frac{1}{2k}\Big(1 - \frac{n}{q}\Big), \qquad \varepsilon_{\mathrm{mca}}(C, \delta) \;>\; \frac{1}{2k}\Big(1 - \frac{n}{q}\Big) \;\text{ for every } \delta \in \Big[1 - \rho - \tfrac{2}{n},\, 1 - \rho\Big),$$

so $\delta_C^*(\varepsilon^*) \leq 1 - \rho - 2/n$ for every $\varepsilon^* \leq \frac{1}{2k}(1 - \frac{n}{q})$; and, provided $k < 2^{128}$, the list challenge is likewise unsafe at that step, since $\binom{n/2}{(k+2)/2}/|\mathbb{B}| \geq q/k > 2^{-128} q$. For odd k the same argument with $\ell = \lfloor k/2 \rfloor + 1$ gives the step at $1 - \rho - 1/n$.

Proof. Apply Proposition 9.25 with $(a, e) = (2, 1)$: $k' = 2\ell - 2 = k$, so the list lives in $\mathrm{RS}[\mathbb{F}, D, k+1] = C^+$ at radius $\delta_0 := 1 - (k+2)/n$, and $\lfloor \delta_0 n \rfloor = n - k - 2 \le n - k - 1$ is admissible for Theorem 2.6. The entropy count: $\ell/(n/2) = (k+2)/n \in (\rho, \rho + 2^{-11}] \subseteq [\frac{1}{16}, 0.502]$, so $\mathrm{H}_2(\ell/(n/2)) \ge \mathrm{H}_2(\frac{1}{16}) \ge 0.3372$ by the endpoint argument of Corollary 7.6, and

$$\log_2 \binom{n/2}{\ell} \;\ge\; \frac{n}{2}\,\mathrm{H}_2\Big(\frac{\ell}{n/2}\Big) - \log_2\Big(\frac{n}{2}+1\Big) \;\ge\; 0.3372 \cdot \frac{n}{2} - \log_2\Big(\frac{n}{2}+1\Big) \;\ge\; 690 - 12 \;\ge\; 513,$$

since the middle expression is increasing in n and equals $690.6\ldots - 11.0\ldots$ at $n = 2^{12}$, while $|\mathbb{B}|(q/k+1) \le q^2 + q < 2^{513}$. The contradiction with Theorem 2.6 at $\eta = \frac{1}{2}$, then Fact 2.4 and Lemma 2.5, run exactly as in Theorem 7.4. The list-challenge clause follows because $\binom{n/2}{\ell}/|\mathbb{B}| \ge q/k + 1 > 2^{-128} q$ whenever $k < 2^{128}$, which covers every challenge-envelope row. □

9.9 The two-sided picture and what remains

Theorem 9.27 (two-sided sandwich for the grand MCA challenge)**.** *Let $C = \mathrm{RS}[\mathbb{F}, D, k]$ be a row for which one of the capacity caps of this paper applies at target $\varepsilon^* = 2^{-128}$ — Corollary 5.1, Corollary 7.6, Corollary 7.12 or Corollary 7.13, or Corollary 9.26 — with printed gap G (respectively 2^{-9}/2^{-10}, 2^{-11}/2^{-12}, 2^{-9}–2^{-11} or 2^{-7}, and $2/n$). Put*

$$r^* \;:=\; \min\Big(\Big\lfloor \frac{n-k}{3} \Big\rfloor,\; \lfloor 2^{-128} q \rfloor - 1\Big).$$

If $q > 2^{128}$ then

$$\frac{r^*}{n} \;\le\; \delta^*_C(2^{-128}) \;\le\; 1 - \rho - G,$$

so, whenever $\lfloor 2^{-128} q \rfloor - 1 \ge \lfloor (n-k)/3 \rfloor$ — e.g. whenever $q \ge 2^{128} n$, as at every deployed multiplicative row — the grand MCA threshold of such a row is determined up to a multiplicative factor of three, up to the microscopic gap G; for q nearer 2^{128} the lower edge degrades to r^*/n and no factor-of-three claim is made. If $q < 2^{128}$ then $\delta^*_C(2^{-128}) = 0$ (Proposition 9.20); at the boundary $q = 2^{128}$ the pair of Proposition 9.20 gives $\varepsilon_{\mathrm{mca}} \ge 2^{-128}$ at every radius, which meets the threshold condition with equality and does not force degeneracy.

Proof. Lower bound: $3r^* \le n - k$ and $r^* + 1 \le \lfloor 2^{-128} q \rfloor \le 2^{-128} q$, so Theorem 9.6 gives $\varepsilon_{\mathrm{mca}}(C, r^*/n) \le (r^*+1)/q \le 2^{-128}$, and $r^*/n < (n-k)/n = 1 - \rho$, so $\delta^* \ge r^*/n$ by Definition 2.3 (the case $r^* = 0$ being trivial). Upper bound: the applicable cap certifies $\varepsilon_{\mathrm{mca}}(C, \delta) > 2^{-128}$ for every $\delta \in [1 - \rho - G,\, 1 - \rho)$, so no such δ belongs to the supremum set, and by Lemma 2.5 neither does any larger sub-capacity radius; hence $\delta^* \le 1 - \rho - G$. The degenerate clause is Proposition 9.20 with $\varepsilon^* = 2^{-128} < 1/q$; the boundary remark is the same construction, whose error floor $1/q$ equals the target exactly at $q = 2^{128}$. □

For $\rho = \frac{1}{2}$ and $q \ge 2^{128} n$ — e.g. the KoalaBear-sextic row, $q > 2^{185}$ — the sandwich reads $\delta^*_C(2^{-128}) \in [0.1666\ldots, \frac{1}{2} - G]$: the first two-sided determination of a grand-challenge MCA threshold, exact to within a factor of three.

Remark 9.28 (dependency status of the one-step criteria). The one-step MCA criterion of Theorem 8.2 has the following dependency status. The periodic-support input is supplied by Corollary 9.3 at the support level and by Theorem 9.4 at the image level for fiber-constant lines. The deep-regime input is supplied by Theorem 9.6 and Corollary 9.9; outside that range it is exactly

Component	Outcome	Where
Periodic-support ledger	Stabilizer exhaustion is exact at the support level; for fiber-constant lines it descends to an exact quotient-image ledger; upper numerators are independent of q and $\|\mathbb{B}\|$.	Section 9.1
Deep MCA core	MCA is safe below one third of the distance; singular and rank-drop Hankel branches are removed in that range; the remaining band is the broad band formulation Remark 9.12, with rank-pivot and compressed-eliminant certificates.	Section 9.2 and Theorem 9.6
Extension-line lift	Extension-valued lines are handled exactly by restriction of scalars; there is no independent extension bucket.	Section 9.3
Interleaved-list side	Johnson safety is uniform in the interleaving arity; fiber-constant recursion is exact; the quotient term is necessarily exponential in the band.	Section 9.4
Curve/projective samplers	The deep regime is controlled for degree-d curves and projective parameters; residual curve charts are part of the same aperiodic band problem.	Section 9.5
Scanner	The finite staircase scanner is specified, proved sound, and run on the deployed rows.	Section 9.6
Certificate core	The finite certificate format, decidability statement, and dependency chart are supplied; machine formalization is separate implementation work.	Section 9.7
Tower transfer	The ledger half is transferred to the Chebyshev tower; the genus-one first-step cap is proved, later strengthened to a macroscopic cap by the rational graded floors.	Sections 9.8 and 11.1

Table 4: Structural and certificate ingredients used in the two-sided threshold statements.

the broad band formulation Remark 9.12, represented by the residual rank-pivot and kernel-compressed certificate language of Theorems 3.64, 3.73 and 3.76. The extension-line input is supplied by Theorem 9.13. The sampler input is definitional once the protocol's line sampler is fixed. For the one-step list criterion of Theorem 8.3, the required base-code bound is unconditional throughout the Johnson regime and uniform in m (Theorem 9.15, making the codegree recursion unnecessary there), is exactly recursive for fiber-constant words (Theorem 9.16), and is provably exponential in the band (Proposition 9.17), so any resolution must take that shape.

After these ingredients are assembled, the remaining mathematical gap has a compact form. The central residue is the broad band formulation Remark 9.12, the aperiodic band conjecture, now represented by the finite rank-pivot, compressed-eliminant, and exact-image certificate language of Theorems 3.64, 3.73 and 3.76; this is the analytic kernel separating the two-sided interval of Theorem 9.27 from a one-step threshold determination. The other residues are narrower: the failure-profile upgrade from error $1/k$ to error $1-o(1)$ (the error-one question Question 22.2) and explicit deployed-size certifying pairs (the explicit-witness question Question 22.1); the genus-one band, formerly a separate residue, is settled at macroscopic scale in Section 11.1. The next section attacks the main band directly, from both sides at once.

10 Safe-side pincer and scale-optimized unsafe floors

The residue identified at the end of Section 9 is a single open band per row. This section attacks the band from both sides and narrows it substantially, with three new results. On the safe side, we prove that *mutual* correlated agreement coincides with plain correlated agreement, up to an explicit tangent term, all the way up to half the minimum distance (Theorem 10.1); this eliminates the support-mismatch layer of the problem on that range and, combined with the imported unique-decoding proximity gap of Ben-Sasson–Carmon–Ishai–Kopparty–Saraf [BCIKS20], moves the certified safe frontier of every deployed row from one third to one half of the distance. We also prove a self-contained correlated-agreement bound at half the Johnson radius (Theorem 10.4), which improves the *unconditional* safe frontier for every row of rate below $\frac{1}{4}$. On the unsafe side, we extend the graded locator-prefix pigeonhole of Lemma 3.5 to every map-smooth scale — in particular to the Chebyshev scales of the circle tower — and optimize it over the scale (Proposition 10.8); at deployed sizes this widens the certified unsafe band by a factor of four with a hand-checkable certificate, and by a factor of about sixteen with an exactly verified one, saturating an intrinsic entropy–subfield envelope that we identify in closed form (Remark 10.11). Theorem 10.12 and Table 5 assemble the narrowed sandwich.

10.1 Mutual equals correlated agreement up to half the distance

The deep-regime theorem (Theorem 9.6) controls $\varepsilon_{\mathrm{mca}}$ directly, but only below one third of the distance: its alternatives identify every close point-codeword with a value of a single interpolated line, and that identification costs a factor of three in the radius. The next theorem decouples the two layers of the problem. It shows that below *half* the distance, the support-mismatch phenomenon that separates mutual from plain correlated agreement is carried entirely by the tangent slopes of Theorem 9.6(a), of which there are at most $\lfloor \delta n \rfloor$; everything else is already visible to $\varepsilon_{\mathrm{ca}}$.

Theorem 10.1 (mutual from correlated agreement, up to half the distance)**.** *Let $C \subseteq \mathbb{F}^n$ be any $\mathbb{F}$-linear code of minimum Hamming weight w, let $q = |\mathbb{F}|$, let $\delta \in (0,1)$, and put $r := \lfloor \delta n \rfloor$. Assume*

$$2r \;\le\; w - 1.$$

Then for every pair $(f_1, f_2) \in \mathbb{F}^n \times \mathbb{F}^n$:

(a) *if $\Delta_2\big((f_1,f_2), C^{\equiv 2}\big) \le \delta$, the pair has* no *CA-bad slope at radius δ, and at most r MCA-bad slopes at radius δ: writing (p_1,p_2) for a pair of codewords column-agreeing with (f_1,f_2) outside a set E with $|E| \le r$ and $e_i := f_i - p_i$, the MCA-bad slopes form a subset of $\{-e_{1,j}/e_{2,j} : \; j \in E,\; e_{2,j} \neq 0\}$;*

(b) *if $\Delta_2\big((f_1,f_2), C^{\equiv 2}\big) > \delta$, every MCA-bad slope at radius δ is CA-bad at radii (δ,δ).*

Consequently

$$\varepsilon_{\mathrm{mca}}(C,\delta) \;\le\; \max\Big(\varepsilon_{\mathrm{ca}}(C,\delta),\; \frac{r}{q}\Big).$$

For $C = \mathrm{RS}[\mathbb{F}, D, k]$ on any evaluation set D of size n, the hypothesis reads $2\lfloor \delta n \rfloor \le n - k$: mutual and plain correlated agreement coincide, up to the explicit tangent term r/q, on the entire unique-decoding range of the interleaved code $C^{\equiv 2}$.

Proof. Part (b) is immediate from the definitions: an MCA-bad slope γ has a witness set S with $|S| \geq \lceil (1-\delta)n \rceil = n - r$ on which $f_1 + \gamma f_2$ agrees with a codeword, so $\Delta(f_1 + \gamma f_2, C) \leq r/n \leq \delta$; if additionally the pair is δ-far in column distance, γ is CA-bad at radii (δ, δ) by Definition 2.1.

For part (a), fix codewords (p_1, p_2) column-agreeing with (f_1, f_2) outside the column error set $E := \{j : (e_{1,j}, e_{2,j}) \neq (0,0)\}$, $e_i := f_i - p_i$, $|E| \leq r$. The pair is δ-close, so it has no CA-bad slope. Let γ be MCA-bad with witness (S, c): $|S| \geq n - r$, $c \in C$, $c = f_1 + \gamma f_2$ on S, and no pair of codewords of C explains (f_1, f_2) on S.

First, $c = p_1 + \gamma p_2$. Indeed, $z := c - (p_1 + \gamma p_2) \in C$ vanishes on $S \setminus E$: there $f_i = p_i$, so $f_1 + \gamma f_2 = p_1 + \gamma p_2$, while $c = f_1 + \gamma f_2$ on all of S. Since $|S \setminus E| \geq n - 2r$, the codeword z has weight at most $2r \leq w - 1$, hence $z = 0$.

Next, $S \not\subseteq D \setminus E$: otherwise (p_1, p_2) would explain the pair on S. So there is $j \in S \cap E$. On S we have $0 = (f_1 + \gamma f_2) - c = (e_1 + \gamma e_2)$, so at this j,

$$e_{1,j} + \gamma\, e_{2,j} = 0, \qquad (e_{1,j}, e_{2,j}) \neq (0,0).$$

If $e_{2,j} = 0$ the left side is the nonzero constant $e_{1,j}$, which cannot vanish; hence $e_{2,j} \neq 0$ and $\gamma = -e_{1,j}/e_{2,j}$. The MCA-bad slopes of the pair therefore form a subset of the at most $|E| \leq r$ listed values. Taking the maximum over pairs gives the displayed bound. □

Remark 10.2 (scope and sharpness). Three comments. (i) The hypothesis $2r \leq w - 1$ is exactly the unique-decoding radius of the interleaved code $C^{\equiv 2}$ in column weight, and the proof uses it twice, both times through the statement "a codeword of weight $\leq 2r$ vanishes"; beyond half the distance a δ-close pair can admit several column explanations with incomparable agreement sets, and witness codewords need not ride any explanation line, which is precisely where the band begins. (ii) The theorem is an unconditional *implication*, with no list bound and no field-size hypothesis; instantiated with the self-contained $\varepsilon_{\mathrm{ca}}$ bound of Theorem 9.6 it returns the deep theorem (with the marginally better error $\max((r+1)/q, r/q)$ replaced by $(r+1)/q$), and its value is that it accepts *any* stronger correlated-agreement input on the range $(w-1)/3 < 2r \leq w - 1$, where no self-contained input is available. (iii) The same proof gives the two-radius statement: for $\delta_{\mathrm{int}} \geq \delta$, every MCA-bad slope of a δ_{int}-far pair is CA-bad at radii $(\delta, \delta_{\mathrm{int}})$, while a δ_{int}-close pair that is also δ-close contributes at most r slopes; only pairs in the annulus between the two radii require the no-loss form.

Corollary 10.3 (the band conjecture loses its support-mismatch layer below half the distance)**.** *Let $C = \mathrm{RS}[\mathbb{F}, D, k]$ and let δ satisfy $2\lfloor \delta n \rfloor \leq n - k$. Then any correlated-agreement bound $\varepsilon_{\mathrm{ca}}(C, \delta) \leq \varepsilon$ implies $\varepsilon_{\mathrm{mca}}(C, \delta) \leq \max(\varepsilon, \lfloor \delta n \rfloor / q)$. In particular, on the agreement subrange $a \geq \lceil (n + k)/2 \rceil$ of the broad band formulation Remark 9.12, the aperiodic band conjecture reduces to its correlated-agreement form: column-close pairs contribute at most $\lfloor \delta n \rfloor$ explicit tangent slopes, and what must be bounded polynomially is only, for each column-far pair, the number of slopes γ for which $f_1 + \gamma f_2$ has an aperiodically witnessed agreement set of size at least $\lceil (1 - \delta)n \rceil$.*

Proof. Immediate from Theorem 10.1: case (a) pairs contribute at most $r = \lfloor \delta n \rfloor$ slopes, each explicitly a tangent ratio; case (b) identifies every remaining MCA-bad slope, together with its witness support, as a CA-bad slope with the same witness. □

10.2 A self-contained correlated-agreement bound at half the Johnson radius

Theorem 10.1 converts the safe side of the band problem into a demand for correlated-agreement inputs. The strongest known inputs are the proximity-gap theorems of [BCIKS20] (unique-

decoding range with error n/q; Johnson range with error $\mathrm{poly}(n)/q$), whose proofs run through a Berlekamp–Welch decoder over the rational function field $\mathbb{F}(Z)$ and an induction over curves, and the recent strengthening [BCHKS25]. Those we import. But the certificate interface of Section 8 also prices *self-contained* inputs, since the formal core (Table 3) is built from elementary counting; the following theorem is the elementary member of the family. Its only ingredients are the pair-interleaved Johnson bound of Theorem 9.15 and a rider double count, and it reaches half the Johnson radius with an explicit list-size error.

Theorem 10.4 (elementary correlated agreement at half the Johnson radius)**.** *Let* $C = \mathrm{RS}[\mathbb{F}, D, k]$ *on any evaluation set* D *of size* n*, let* $\delta \in (0,1)$*, put* $r := \lfloor \delta n \rfloor$*, and assume*

$$n - 2r \;>\; 0 \qquad \text{and} \qquad (n-2r)^2 \;>\; (k-1)\,n.$$

Then, with

$$L_2 \;:=\; \frac{n\,\big(n-2r-k+1\big)}{(n-2r)^2-(k-1)n}$$

the pair-interleaved Johnson bound of Theorem 9.15 at agreement $n-2r$*,*

$$\varepsilon_{\mathrm{ca}}(C,\delta) \;\le\; \frac{1+(r+1)\,L_2}{q},$$

and the same bound holds for $\varepsilon_{\mathrm{ca}}(C,\delta,\delta_{\mathrm{int}})$ *for every* $\delta_{\mathrm{int}} \ge \delta$*.*

Proof. Fix a pair (f_1, f_2) with $\Delta_2((f_1,f_2), C^{\equiv 2}) > \delta_{\mathrm{int}} \ge \delta$; only such pairs have CA-bad slopes. Let $B \subseteq \mathbb{F}$ be the set of CA-bad slopes; for each $\gamma \in B$ fix a codeword $c_\gamma \in C$ and an agreement set S_γ with $|S_\gamma| \ge n - r$ and $c_\gamma = f_1 + \gamma f_2$ on S_γ. If $|B| \le 1$ we are done, so fix $\gamma_1 \in B$ and consider any $\gamma \in B \setminus \{\gamma_1\}$. Define

$$P_2 := \frac{c_{\gamma_1} - c_\gamma}{\gamma_1 - \gamma} \in C, \qquad P_1 := c_{\gamma_1} - \gamma_1 P_2 \in C,$$

so that, identically as codewords,

$$P_1 + \gamma_1 P_2 = c_{\gamma_1}, \qquad P_1 + \gamma P_2 = c_\gamma. \tag{3}$$

On $S_{\gamma_1} \cap S_\gamma$, which has size at least $n-2r$, the two coordinatewise linear equations $f_1 + \gamma_1 f_2 = c_{\gamma_1}$ and $f_1 + \gamma f_2 = c_\gamma$ have the unique solution $f_2 = P_2$, $f_1 = P_1$; hence the codeword pair (P_1, P_2) column-agrees with (f_1, f_2) on at least $n - 2r$ coordinates, i.e.

$$(P_1,P_2) \;\in\; \mathcal{L} \;:=\; \Lambda\Big(C^{\equiv 2},\, 1 - \tfrac{n-2r}{n},\, (f_1,f_2)\Big), \qquad |\mathcal{L}| \le L_2$$

by Theorem 9.15 with $m = 2$ and $a = n - 2r$. Say that γ *rides* a member $(P_1, P_2) \in \mathcal{L}$ if $c_\gamma = P_1 + \gamma P_2$ as codewords; by (3), every $\gamma \in B \setminus \{\gamma_1\}$ rides some member of $\mathcal{L}$.

Fix $(P_1, P_2) \in \mathcal{L}$ and let $E := \{j : (f_{1,j}, f_{2,j}) \ne (P_{1,j}, P_{2,j})\}$ be its column error set. Since the pair is δ-far, $|E| \ge r + 1$. If γ rides (P_1, P_2), then $(f_1 - P_1) + \gamma(f_2 - P_2) = (f_1 + \gamma f_2) - c_\gamma$ vanishes on S_γ, hence on at least $|E| - r \ge 1$ coordinates of E; and at each $j \in E$ the map $\gamma \mapsto (f_{1,j} - P_{1,j}) + \gamma(f_{2,j} - P_{2,j})$ is a nonzero affine function, vanishing for at most one slope. Double counting vanishing incidences between riders of (P_1, P_2) and coordinates of E,

$$\#\{\text{riders of } (P_1,P_2)\} \cdot \big(|E| - r\big) \;\le\; |E|, \qquad \text{so} \qquad \#\{\text{riders}\} \;\le\; \frac{|E|}{|E|-r} \;\le\; r+1.$$

Summing over $\mathcal{L}$ and adding back γ_1: $|B| \le 1 + (r+1)L_2$. Taking the maximum over pairs completes the proof. □

Remark 10.5 (provenance and comparison). The argument is elementary and in the folklore spirit "a list bound implies a proximity gap with a list-size loss"; we do not claim novelty of the mechanism, only of the exact constants in this normalization and of its place in the formal core: its only inputs are Theorem 9.15 and counting, so it enters the dependency chart of Table 3 without enlarging the trusted base. The theorems of [BCIKS20, BCHKS25] are strictly stronger where they apply — they reach the full Johnson radius $1-\sqrt{\rho}$ with error $\mathrm{poly}(n)/q$ — at the price of a substantially deeper proof. Note also the built-in factor of two: the recovered pair (P_1, P_2) lives at column radius $2r$, so the list bound is invoked at agreement $n-2r$, which is why the theorem stops at half the Johnson radius rather than at the Johnson radius.

Corollary 10.6 (self-contained safe frontier at half the Johnson radius)**.** *Let $C = \mathrm{RS}[\mathbb{F}, D, k]$ with $n \geq k+2$, and let δ, r, L_2 be as in Theorem 10.4, so $n-2r>0$ and $(n-2r)^2 > (k-1)n$. Then*

$$\varepsilon_{\mathrm{mca}}(C,\delta) \;\leq\; \frac{1+(r+1)L_2}{q}\,.$$

For the challenge rows this improves the self-contained safe frontier exactly when $\rho < \frac{1}{4}$: the frontier moves from $(1-\rho)/3$ to $\approx (1-\sqrt{\rho})/2$, i.e. from 0.2917 to 0.3232 at $\rho = \frac{1}{8}$ and from 0.3125 to 0.3750 at $\rho = \frac{1}{16}$, while at $\rho \in \{\frac{1}{2}, \frac{1}{4}\}$ the deep frontier $(1-\rho)/3$ remains the better self-contained certificate. At the deployed KoalaBear row the half-Johnson certificate reads: $r = 307121$, $\delta = r/n \approx 0.14645$, $L_2 \leq 1001282$, error at most $(1+307122\cdot 1001282)/q < 2^{-147}$ — dominated by Theorem 9.6 at this rate, but recorded as an independent certificate of the same verdict.

Proof. The hypothesis gives $n-2r>0$ and $(n-2r)^2 > (k-1)n$, hence $n-2r > \sqrt{(k-1)n}$; since $n \geq k+2$ we also have $\sqrt{(k-1)n} \geq k$, so $n-2r > k$, i.e. $2r \leq n-k-1$: the radius lies below half the distance. Theorem 10.1 therefore applies, and with the input of Theorem 10.4,

$$\varepsilon_{\mathrm{mca}}(C,\delta) \leq \max\Big(\tfrac{1+(r+1)L_2}{q},\ \tfrac{r}{q}\Big) = \tfrac{1+(r+1)L_2}{q}.$$

The frontier comparison is the elementary inequality $(1-\sqrt{\rho})/2 > (1-\rho)/3 \iff \sqrt{\rho} < \frac{1}{2}$; the KoalaBear numbers are $n = 2^{21}$, $k = 2^{20}$, $(n-2r)^2 - (k-1)n = 909700$ at $r = 307121$, and $L_2 = \lfloor n(n-2r-k+1)/909700 \rfloor \leq 1001282$, with $(1+307122\cdot 1001282) < 2^{38.2}$ and $q > 2^{185.9}$. □

Corollary 10.7 (conditional safe frontier at half the distance)**.** *Let $C = \mathrm{RS}[\mathbb{F}, D, k]$, $\rho = k/n$. Import the unique-decoding proximity gap of [BCIKS20, Thm. 4.1]: in the normalization of Definition 2.1, $\varepsilon_{\mathrm{ca}}(C,\delta) \leq n/q$ for every $\delta \leq (1-\rho)/2$. (Only $\varepsilon_{\mathrm{ca}}(C,\delta) \leq \mathrm{poly}(n)/q$ is used below, so the conclusion is robust to the exact constant.) Then for every δ with $2\lfloor \delta n \rfloor \leq n-k$,*

$$\varepsilon_{\mathrm{mca}}(C,\delta) \;\leq\; \max\Big(\frac{n}{q},\ \frac{\lfloor \delta n \rfloor}{q}\Big) \;=\; \frac{n}{q}\,.$$

Deployed instantiations: for the KoalaBear sextic row, $\varepsilon_{\mathrm{mca}}(C,\delta) \leq n/q < 2^{-164} < 2^{-128}$ for every $\delta \leq \frac{1}{4}$ (the endpoint $r = 2^{19}$, $2r = n-k$, is included); for the circle line-round row over $\mathbb{F}_{p^4}$, $p = 2^{31}-1$, $\varepsilon_{\mathrm{mca}} \leq n/q < 2^{-102} < 2^{-100}$ for every $\delta \leq \frac{1}{4}$. In both cases the certified safe frontier moves from one third to one half of the distance.

Proof. Combine the import with Theorem 10.1; $\lfloor \delta n \rfloor \leq n$ makes the maximum n/q. For the numbers: $n/q = 2^{21}/p^6 < 2^{-164.9}$ with $p = 2^{31} - 2^{24} + 1$, and $n/q = 2^{21}/(2^{31}-1)^4 < 2^{-102.9}$.

The circle case uses the diagonal equivalence of Lemma 7.11, which preserves both correlated-agreement errors exactly, so the import applies verbatim to the equivalent Reed–Solomon code on the torus twin coset. □

10.3 Widening the unsafe band: graded prefix floors on every scale

The deployed caps of Corollaries 5.4 and 7.13 certify the unsafe band $[\frac{1}{2} - 2^{-7}, \frac{1}{2})$ using only the top *two* slots of the locator expansion at the single top scale ($N = 256$). The quotient-remainder ledger already prices deeper prefixes at the multiplicative level (Lemma 3.5); two observations remain to be made. First, the graded pigeonhole runs verbatim on *every* map-smooth scale — the locator expansion $\prod_{b\in M}(\varphi - b) = \sum_j (-1)^j e_j(M)\varphi^{m-j}$ is graded by $\deg \varphi$ for any polynomial φ, in particular for the Chebyshev scales of the circle tower — which extends the tower ledger from the first staircase step to every depth. Second, the floor can be *optimized over the scale*: small scales trade prefix cost $|\mathbb{B}|$ per slot against fiber entropy $\binom{n/c}{m}$, and the optimum widens the certified unsafe band by an order of magnitude at deployed sizes. The route has an intrinsic ceiling, which we identify exactly in Remark 10.11.

Proposition 10.8 (graded prefix floor for map-smooth scales)**.** *Let $\mathbb{B} \subseteq \mathbb{F}$ be a subfield, let $D \subseteq \mathbb{B}$ with $|D| = n$, and let $\varphi \in \mathbb{B}[X]$ of degree $c \geq 2$ whose fibers within D are complete and of uniform size c (the map-smooth hypothesis of Lemma 7.2); put $N := n/c$ and $Q := \varphi(D) \subseteq \mathbb{B}$, so $|Q| = N$. Let $1 \leq K \leq n$ and let m satisfy $\lceil K/c \rceil \leq m \leq N$; put*

$$w \;:=\; m - \lceil K/c \rceil \;\geq\; 0.$$

Then there is a $\mathbb{B}$-valued received word $U : D \to \mathbb{B}$ with

$$\big|\Lambda\big(\mathrm{RS}[\mathbb{F}, D, K],\, 1 - \tfrac{mc}{n},\, U\big)\big| \;\geq\; \left\lceil \frac{\binom{N}{m}}{|\mathbb{B}|^{\,w}} \right\rceil .$$

For $\varphi = X^c$ on a multiplicative coset this is Lemma 3.5 with $s = 0$ (note $w = \lfloor (mc - K)/c \rfloor = m - \lceil K/c \rceil$); the content here is that the same graded pigeonhole runs on every polynomial scale, in particular on the Chebyshev scales T_c of the circle tower (Lemma 7.8), whose locators are not lacunary in the monomial basis but are perfectly graded in powers of T_c.

Proof. For $M \subseteq Q$ with $|M| = m$, the locator $\Lambda_M := \prod_{b\in M}(\varphi - b) \in \mathbb{B}[X]$ vanishes exactly on the union of the φ-fibers over M inside D, a set of size mc, and expands as $\Lambda_M = \sum_{j=0}^{m} (-1)^j e_j(M)\, \varphi^{m-j}$ with $e_j(M) \in \mathbb{B}$ the elementary symmetric functions of $M \subseteq \mathbb{B}$. For $z = (z_1, \ldots, z_w) \in \mathbb{B}^w$ define the $\mathbb{B}$-valued word

$$U_z \;:=\; \Big(\varphi^m + \textstyle\sum_{j=1}^{w} z_j\, \varphi^{m-j}\Big)\Big|_D .$$

Assign to each M the vector $z(M) := \big((-1)^j e_j(M)\big)_{1\leq j\leq w} \in \mathbb{B}^w$; by pigeonhole some z^* is assigned by at least $\binom{N}{m}/|\mathbb{B}|^w$ of the $\binom{N}{m}$ sets. For each such M put

$$c_M \;:=\; U_{z^*} - \Lambda_M\big|_D \;=\; -\sum_{j=w+1}^{m} (-1)^j e_j(M)\, \varphi^{m-j}\Big|_D ,$$

a polynomial of degree at most $c(m - w - 1) = c\big(\lceil K/c \rceil - 1\big) \leq K - 1$, hence a codeword of $\mathrm{RS}[\mathbb{F}, D, K]$; it agrees with U_{z^*} on the mc zeros of Λ_M in D, so it lies in the displayed list.

The map $M \mapsto c_M$ is injective: if $c_M = c_{M'}$ then $\sum_{j>w}(-1)^j\big(e_j(M) - e_j(M')\big)\varphi^{m-j} = 0$ as a polynomial of degree $< n$ vanishing on D, hence identically; the powers φ^{m-j} have pairwise distinct degrees $c(m-j)$ and are therefore linearly independent over $\mathbb{F}$, so $e_j(M) = e_j(M')$ for $j > w$, while for $j \le w$ equality holds by the choice of z^*; thus $\prod_{b\in M}(Y-b) = \prod_{b\in M'}(Y-b)$ and $M = M'$. □

Corollary 10.9 (widened unsafe band at deployed sizes)**.** (i) (KoalaBear sextic, MCA.) *With $\mathbb{B}$, $\mathbb{F}$, D, $n = 2^{21}$, $k = 2^{20}$ as in Corollary 5.4,*

$$\varepsilon_{\mathrm{mca}}\big(\mathrm{RS}[\mathbb{F}_{p^6}, D, 2^{20}],\, \delta\big) \;>\; 2^{-22} \qquad \textit{for every} \qquad \delta \in \Big[\tfrac{15}{32},\, \tfrac{1}{2}\Big),$$

a certified gap of 2^{-5}, four times the gap 2^{-7} of Corollary 5.4, by a hand-checkable certificate at scale $c = 16$.

(ii) (KoalaBear sextic, grand list challenge.) *For every $\delta \in [\frac{15}{32}, 1-\rho)$ there is a $\mathbb{B}$-valued word U with $|\Lambda(C, \delta, U)| > 2^{-128}q$.*

(iii) (Circle line-round row.) *With $p' = 2^{31}-1$, $\mathbb{F} = \mathbb{F}_{p'^4}$, and the twin-coset x-domain of size $n = 2^{21}$, $k = 2^{20}$ as in Corollary 7.13, the same two statements hold verbatim on $[\frac{15}{32}, \frac{1}{2})$ resp. $[\frac{15}{32}, 1-\rho)$, at the Chebyshev scale $\varphi = T_{16}$, against the target 2^{-100}.*

Proof. All four claims instantiate Proposition 10.8 at scale $c = 16$, $N = n/c = 2^{17}$, $m = 69632$, so $mc = k + 2^{16}$ and $m/N = \frac{17}{32}$, and convert the floor as in Corollary 5.4. The common entropy estimate is

$$\log_2 \binom{N}{m} \;\ge\; N\,\mathrm{H}_2\big(\tfrac{17}{32}\big) - \log_2(N+1) \;\ge\; 131072 \cdot 0.99718 - 17 \;>\; 130{,}685.$$

(i) Take $K = k+1$, so $w = m - \lceil (k+1)/16 \rceil = 69632 - 65537 = 4095$. The threshold for the deep-point conversion (Theorem 2.6 with $\eta = \frac{1}{2}$, admissible since $n - mc \le n - k - 1$) is a floor exceeding $q/k + 1$:

$$w \log_2 p + \log_2(q/k+1) \;<\; 4095 \cdot 30.989 + 166 \;<\; 127{,}116 \;<\; 130{,}685,$$

with over 3500 bits of room. So $\varepsilon_{\mathrm{ca}}(C, 1 - mc/n) > \frac{1}{2k}(1 - n/q) > 2^{-22}$, and Fact 2.4 with Lemma 2.5 extend the same lower bound to $\varepsilon_{\mathrm{mca}}(C, \delta)$ for every $\delta \in [1 - mc/n, \frac{1}{2}) = [\frac{15}{32}, \frac{1}{2})$. (ii) Take $K = k$, so $w = 4096$; the list threshold is $\binom{N}{m}/|\mathbb{B}|^w > 2^{-128}q$, i.e. $w \log_2 p + \log_2 q - 128 < 4096 \cdot 30.989 + 58 < 127{,}000 < 130{,}685$; lists only grow with the radius, giving the full range. (iii) T_{16} is a polynomial of degree 16 over $\mathbb{B} = \mathbb{F}_{p'}$ with complete uniform fibers on the twin-coset x-domain at every dyadic scale (Lemma 7.8), so Proposition 10.8 applies with the same N, m, w; the thresholds are weaker ($\log_2(q/k+1) < 104$ and $\log_2 q - 100 < 24$), and the conversion runs through the diagonal equivalence of Lemma 7.11 exactly as in Corollary 7.13. □

Remark 10.10 (exactly verified frontier). The hand-checkable rows above deliberately stop at the round gap 2^{-5}. Optimizing (c, m) under exact integer verification of $\binom{N}{m} > p^w(q/k+1)$ resp. $\binom{N}{m} \cdot 2^{128} > p^w q$ — a finite certificate in the sense of Definition 9.21 and Proposition 9.23 — widens the band further, to the following edges (all verified by exact arithmetic; $\Delta := mc - k$):

row / challenge	scale c	m	w	Δ	certified unsafe range
KoalaBear, MCA	16	69748	4211	67392	$\delta \in [\frac{15331}{32768}, \frac{1}{2}) = [0.46786\ldots, \frac{1}{2})$
KoalaBear, list	32	34874	2106	67392	$\delta \in [\frac{15331}{32768}, 1-\rho)$
circle $p'{=}2^{31}{-}1$, MCA	32	34873	2104	67360	$\delta \in [\frac{30663}{65536}, \frac{1}{2}) = [0.46787\ldots, \frac{1}{2})$
circle $p'{=}2^{31}{-}1$, list	32	34874	2106	67392	$\delta \in [\frac{15331}{32768}, 1-\rho)$

The certified gap is thus $\approx 0.03214 \approx 2^{-4.96}$ throughout, versus 2^{-7} before. These scale-16/32 rows are retained as intermediate general-scale certificates; on the deployed identity-scale rows their endpoints are superseded by Propositions 14.2 and 15.2.

Remark 10.11 (the entropy–subfield frontier). The optimized route has a closed-form ceiling. At scale c the floor beats the conversion threshold iff

$$\log_2 \binom{n/c}{m} \;\gtrsim\; w \log_2 |\mathbb{B}| + \log_2(q/k),$$

and writing $mc =: (\rho+g)n$ — so $m/N = \rho+g$ and $w = gN + O(1)$ — and dividing by N, this reads

$$\mathrm{H}_2(\rho+g) \;\geq\; g\, \log_2 |\mathbb{B}| + o(1)$$

for challenge-envelope field sizes ($\log_2 q \leq 256$, so the additive terms are $O(1)$ against $N \to \infty$). Define

$$g^*(\rho,\beta) \;:=\; \sup\big\{\, g \in (0, 1-\rho] :\; \mathrm{H}_2(\rho+g) \geq \beta g \,\big\}, \qquad \beta := \log_2 |\mathbb{B}|.$$

Every gap $g < g^*$ is certified at large enough deployed sizes, and *no* gap beyond g^* is certifiable by graded prefix floors at any scale: the frontier is intrinsic to the route, not to our optimization. For $\beta = 31$ (both deployed base primes): $g^*(\frac{1}{2}) \approx 0.03216$, $g^*(\frac{1}{4}) \approx 0.02748$, $g^*(\frac{1}{8}) \approx 0.01920$, $g^*(\frac{1}{16}) \approx 0.01238$. The deployed KoalaBear row achieves 0.03214 at $n = 2^{21}$: the envelope is saturated to three digits already at deployed sizes. Two structural consequences: the left edge of the band in the broad band formulation Remark 9.12 moves from $k+2n/N$ to $k+g^*n-o(n)$; and the frontier depends on the base field only through β, so larger base fields *narrow* the certifiable unsafe band — a quantitative form of the subfield-confinement phenomenon of Section 6.

10.4 The narrowed sandwich

Theorem 10.12 (narrowed row-specific two-sided sandwich)**.** *Let $\delta^*_C(\varepsilon^*)$ be the challenge threshold of Definition 2.3.*

(i) (KoalaBear sextic MCA row, $\varepsilon^* = 2^{-128}$.) *Self-contained,*

$$\delta^*_C(2^{-128}) \in \left[\frac{349525}{2^{21}}, \frac{981105}{2^{21}}\right] \approx [0.1666665, 0.4678273].$$

With the imported unique-decoding proximity gap of Corollary 10.7,

$$\delta^*_C(2^{-128}) \in \left[\frac{1}{4}, \frac{981105}{2^{21}}\right],$$

a band of width 0.217828.

row	challenge	exact unsafe edge	safe, self-contained	safe, with import
KoalaBear sextic	MCA, 2^{-128}	$[981105/2^{21}, 1/2)$	$\le 349525/2^{21}$	$\le 1/4$
KoalaBear sextic	list, 2^{-128}	$\ge 490553/2^{20}$	≤ 0.29289	—
Mersenne-31 circle	MCA, 2^{-100}	$[981129/2^{21}, 1/2)$	$\le 349525/2^{21}$	$\le 1/4$
Mersenne-31 circle	list, 2^{-100}	$\ge 490565/2^{20}$	≤ 0.29289	—

Table 5: Active deployed scanner output. The exact unsafe edges are the identity-prefix list certificates of Proposition 14.2 and the flexible-budget MCA certificates of Proposition 15.2; they supersede the earlier scale-16/32 endpoints retained in Remark 10.10. The self-contained MCA edge is Theorem 9.6, the optional imported edge is Corollary 10.7, and the list-safe edge is the self-contained Johnson certificate of Theorem 9.15. The corresponding open intervals are printed in Theorem 10.12.

(ii) (Mersenne-31 circle line-round MCA row, $\varepsilon^* = 2^{-100}$.) *Self-contained,*

$$\delta_C^*(2^{-100}) \in \left[\frac{349525}{2^{21}}, \frac{981129}{2^{21}}\right] \approx [0.1666665, 0.4678388].$$

With the same import,

$$\delta_C^*(2^{-100}) \in \left[\frac{1}{4}, \frac{981129}{2^{21}}\right],$$

a band of width 0.217839.

(iii) (General envelope rows.) *For any row with $q \ge 2^{128}(n+1)$ and a map-smooth scale lattice, the safe edge is $\max(\lfloor (n-k)/3 \rfloor, r_{\mathrm{hJ}})/n$ self-contained, where r_{hJ} is the largest radius passing Corollary 10.6, and $\lfloor (n-k)/2 \rfloor/n$ with the import. The unsafe edge is $1-\rho-g$ for every exactly certified g, with envelope $g^*(\rho, \log_2 |\mathbb{B}|)$ of Remark 10.11; the numerical edge is row-specific.*

Proof. The lower edges are supplied by Theorem 9.6 at $r = 349525 = \lfloor (n-k)/3 \rfloor$ and by Corollary 10.7 at $r = 2^{19}$ for the imported clauses; their errors are below the stated row targets as recorded in Table 2. The exact upper edges are the two MCA entries of Proposition 15.2. Part (iii) collects the corresponding general safe and unsafe statements without identifying numerical endpoints across different domain families. □

The reduced kernel, restated once. After the pincer, the open MCA band of each deployed row is $(\frac{1}{4},\ \text{edge})$ with edge ≈ 0.4679, modulo one import whose own proof is independent of everything here; self-contained, it is $(0.1667,\ \text{edge})$, and at rates below $\frac{1}{4}$ the left endpoint improves to half the Johnson radius. Below half the distance the mutual and plain correlated-agreement questions are now provably the same question up to an explicit tangent term (Theorem 10.1), so the normalized residual ledger Remark 9.12 is needed only above half the distance, and there the certified unsafe mass presses down to the entropy–subfield envelope g^* (Remark 10.11), which the deployed rows already saturate. Closing the band therefore requires one of exactly two kinds of new mathematics: a floor construction whose received words carry more than a graded locator prefix — anything beating g^* — or a correlated-agreement theorem crossing half the distance toward capacity, to which Theorem 10.1 would immediately add the mutual layer. This is, for the first time in the program, a two-sided quantitative target with both envelopes made explicit.

11 Further refinements and finite certificate closure

This section records three groups of consequences that are not part of the main unsafe/safe sandwich but complete the account. First, graded rational floors extend the unsafe construction to genus-one scale maps, so the ECFFT obstruction of Remark 7.16 is a one-slot obstruction rather than a macroscopic barrier (Section 11.1), and a stereographic model treats circle codes over challenge fields that do not contain i, completing the circle-code transfer (Section 11.2). Second, explicit head words, simple-pole certifying pairs, and the optimized failure profile clarify what is constructive in the failure mechanism and what remains beyond the present method (Sections 11.3 and 11.4). Third, Section 11.5 closes the finite certificate grammar over every verdict used in the narrowed sandwich, prints the deployed certificates as integer tuples, bounds verification complexity, and isolates the single imported input to one optional clause.

11.1 Graded floors on rational scales: the genus-one band is macroscopic

The graded prefix floor of Proposition 10.8 pays one subfield factor per locator slot and buys c degrees of agreement per slot on a polynomial scale of degree c. On a rational scale of degree type (a, e) the slot spacing is $a - e$ rather than a; nothing else changes.

Proposition 11.1 (graded prefix floor for rational scales)**.** *Let $\mathbb{B} \subseteq \mathbb{F}$ be a subfield and let $D \subseteq \mathbb{B}$ with $|D| = n$ be (ψ, a)-smooth in the sense of Definition 9.24, with $\psi = f/g$ of degree type (a, e), $a > e \geq 0$; put $N := n/a$ and $Q := \psi(D) \subseteq \mathbb{B}$. Let $1 \leq K \leq n$ and let m satisfy*

$$am \;\geq\; K, \qquad em \;\leq\; K, \qquad m \;\leq\; N, \qquad \text{and put} \qquad w \;:=\; \left\lfloor \frac{am - K}{a - e} \right\rfloor \;\leq\; m.$$

Then there is a $\mathbb{B}$-valued received word $U : D \to \mathbb{B}$ with

$$\left|\Lambda\big(\mathrm{RS}[\mathbb{F}, D, K],\, 1 - \tfrac{am}{n},\, U\big)\right| \;\geq\; \left\lceil \frac{\binom{N}{m}}{|\mathbb{B}|^{\,w}} \right\rceil.$$

For $e = 0$ this is Proposition 10.8 (then $w = m - \lceil K/a \rceil$); for $w \leq 1$ it recovers the slack-two window of Proposition 9.25.

Proof. For $M \subseteq Q$ with $|M| = m$ put $\Lambda_M := \prod_{b \in M}(f - bg)$. Each factor has degree exactly a with leading coefficient $\mathrm{lc}(f)$, so $\deg \Lambda_M = am$ with leading coefficient $\mathrm{lc}(f)^m \neq 0$; and $\Lambda_M \in \mathbb{B}[X]$ vanishes at $x \in D$ if and only if $\psi(x) \in M$, i.e. at exactly am points of D, by smoothness and $g \neq 0$ on D. Expanding,

$$\Lambda_M \;=\; \sum_{j=0}^{m} (-1)^j e_j(M)\, f^{\,m-j} g^{\,j},$$

with j-th term of degree $am - j(a - e)$, strictly decreasing in j. For $z = (z_1, \ldots, z_w) \in \mathbb{B}^w$ define the $\mathbb{B}$-valued word

$$U_z \;:=\; \Big(f(x)^m + \sum_{j=1}^{w} z_j\, f(x)^{\,m-j} g(x)^{\,j}\Big)_{x \in D},$$

which makes sense because $w \leq m$: indeed $am - K \leq m(a - e)$ is the hypothesis $em \leq K$. By pigeonhole over the prefix map $M \mapsto z(M) := \big((-1)^j e_j(M)\big)_{j \leq w} \in \mathbb{B}^w$ there is z^* attained by

at least $\lceil \binom{N}{m}/|\mathbb{B}|^w \rceil$ sets M. For each such M set

$$c_M \;:=\; U_{z^*} - \Lambda_M|_D \;=\; \Big(-\sum_{j>w}(-1)^j e_j(M)\, f^{\,m-j} g^j\Big)_{x\in D},$$

of degree at most $am-(w+1)(a-e) \le K-1$ by the choice of w, so $c_M \in \mathrm{RS}[\mathbb{B},D,K] \subseteq \mathrm{RS}[\mathbb{F},D,K]$; and U_{z^*} agrees with c_M on the am fiber points of M, i.e. within relative radius $1-am/n$. Injectivity of $M \mapsto c_M$ on the pigeonhole class: equality of words of degree $\le K-1 \le n-1$ forces equality of polynomials, hence $\sum_{j>w}(-1)^j\big(e_j(M)-e_j(M')\big)f^{\,m-j}g^j = 0$; dividing by g^m in $\mathbb{F}(X)$ gives a polynomial identity in ψ, and ψ is nonconstant, hence transcendental over $\mathbb{F}$ inside $\mathbb{F}(X)$, so all coefficients vanish, exactly as in the proof of Proposition 9.25. Thus $e_j(M)=e_j(M')$ for $j>w$; for $j\le w$ they agree through z^*; hence $\prod_{b\in M}(Y-b)=\prod_{b\in M'}(Y-b)$ and $M=M'$. The c_M are therefore distinct listed codewords for the single word U_{z^*}. □

Corollary 11.2 (macroscopic universal cap for genus-one rows)**.** *Let D be $(\psi,2)$-smooth over $\mathbb{B}$ with ψ of degree type $(2,1)$ (Definition 9.24; in particular every FFTree domain of ECFFT [BCKL21]), let $C=\mathrm{RS}[\mathbb{F},D,k]$ with $\rho := k/n \in [\frac{1}{16},\frac{1}{2}]$, $n\ge 2^{14}$, $n<q=|\mathbb{F}|<2^{256}$, and $\beta := \log_2|\mathbb{B}| \le 64$. Put $m := \lceil (k+1+2^{-9}n)/2 \rceil$ and $N := n/2$. Then*

$$\binom{N}{m} \;>\; |\mathbb{B}|^{\,2m-k-1}\Big(\frac{q}{k}+1\Big) \qquad \text{and} \qquad \binom{N}{m} \;>\; |\mathbb{B}|^{\,2m-k}\,2^{-128}q,$$

and consequently:

(i) *(MCA/CA cap) $\varepsilon_{\mathrm{ca}}\big(C,\delta\big) > \frac{1}{2k}\big(1-\frac{n}{q}\big)$ for every integer-admissible $\delta \in [1-\frac{2m}{n},\, 1-\rho-\frac{1}{n}]$, and $\varepsilon_{\mathrm{mca}}(C,\delta)$ obeys the same lower bound for every $\delta \in [1-\rho-2^{-9},\, 1-\rho)$;*

(ii) *(list cap) there is a $\mathbb{B}$-valued word U with $|\Lambda(C, 1-\frac{2m}{n}, U)| > 2^{-128}q$, so the grand list challenge is unsafe at target 2^{-128} on the same band.*

The gap 2^{-9} coincides with the universal multiplicative gap of Corollary 5.1: genus-one rows are no longer behind the polynomial-scale rows.

Proof. Write $w_1 := 2m-k-1$ and $w_0 := 2m-k$ for the two prefix depths ($K=k+1$ and $K=k$). From $k+1+2^{-9}n \le 2m \le k+2^{-9}n+3$: first, $w_0 \le 2^{-9}n+3 = N/256+3$, so $\beta w_0 \le N/4+192$; second, $m/N \in [\rho+2^{-9},\, \rho+2^{-9}+2^{-12}]$ using $3/n \le 3\cdot 2^{-14} < 2^{-12}$. On the interval $[\frac{1}{16}+2^{-9},\, \frac{1}{2}+2^{-9}+2^{-12}]$ the concave H_2 is minimized at an endpoint, and $\mathrm{H}_2(\frac{1}{16}+2^{-9}) \ge 0.344$, $\mathrm{H}_2(\frac{1}{2}+2^{-9}+2^{-12}) \ge 0.999$, so $\mathrm{H}_2(m/N) \ge 0.344$. Hence

$$\log_2\binom{N}{m} \;\ge\; N\,\mathrm{H}_2(m/N) - \log_2(N+1) \;\ge\; 0.344\,N - \log_2(N+1),$$

while both right-hand sides are at most $2^{\beta w_0+256} \le 2^{N/4+448}$ crudely (using $q/k+1 \le q < 2^{256}$ and $2^{-128}q < 2^{128}$). Since $0.094\,N \ge 0.094\cdot 2^{13} = 770 > 448+\log_2(N+1)$ for all $N \ge 2^{13}$, both displayed inequalities hold. The hypotheses of Proposition 11.1 hold for both $K \in \{k, k+1\}$: $2m \ge k+1 \ge K$; $em = m \le k \le K$ because $2^{-9}n+3 \le k$ for $\rho \ge \frac{1}{16}$, $n \ge 2^{14}$; and $m \le N$ because $2m \le k+2^{-9}n+3 \le n$. Part (ii) is the floor at $K=k$. For (i), the floor at $K=k+1$ exceeds $q/k+1 \ge \lceil 2q\,\varepsilon_{\mathrm{ca}} \rceil$ whenever $\varepsilon_{\mathrm{ca}}(C,\delta) \le \frac{1}{2}(\frac{1}{k}-\frac{n}{kq})$, contradicting Theorem 2.6 at every integer-admissible $\delta \ge \delta_0 := 1-2m/n$ (admissible since $n-2m \le n-k-1$); the assembly through Fact 2.4 at δ_0 and Lemma 2.5 up to $1-\rho$ is verbatim the proof of Theorem 7.4. The band containment is $1-2m/n \le 1-\rho-2^{-9}$, i.e. $2m \ge k+2^{-9}n$. □

Remark 11.3 (the trichotomy of Remark 7.16, resolved). Remark 7.16 offered three exits from the lacunarity collapse: a lacunary basis of isogeny locators, a different received-word family, or a matching safe-side theorem showing the collapse essential. The second exit is the one that exists, and the collapse was not essential: the one-slot analysis priced the window width $2(a-e)$ as the *reach* of the route, when it only sets the *exchange rate* — $a-e$ degrees of agreement per subfield factor, against an entropy budget of H_2 bits per fiber. Optimizing the number of slots gives the closed-form frontier

$$\mathrm{H}_2(\rho + g^*) \;=\; g^* \cdot \beta \cdot \frac{a}{a-e},$$

so the degree type enters only through the ramification-defect ratio $a/(a-e)$: polynomial scales ($e=0$) have ratio 1 and recover Remark 10.11; genus-one x-maps have $e=a-1$, ratio a, optimized at the first fold $a=2$, i.e. an effective subfield cost of 2β. Numerically $g^*(\frac{1}{2}, 62) \approx 0.0161 \approx 2^{-5.96}$ at $\beta = 31$ and $g^*(\frac{1}{2}, 128) \approx 0.0078 \approx 2^{-7}$ at $\beta = 64$: macroscopic in every case, against the $2/n$ of the one-slot route. The safe side needs no new work: Theorems 9.6, 10.1 and 10.4 and Corollary 10.7 are domain-agnostic, so every $(\psi, 2)$-smooth row carries the full narrowed sandwich of Theorem 10.12's shape — self-contained safe frontier at $(1-\rho)/3$, safe at $(1-\rho)/2$ modulo the single import, unsafe from gap 2^{-9} (and, row by row, from the optimized frontier $g^*(\rho, 2\beta)$). This answers the genus-one/circle transport question Question 22.3(a) as posed.

11.2 The stereographic model: circle codes over every challenge field

The torus uniformization of Lemma 7.11 needs $i \in \mathbb{F}$. The circle is a conic with the rational point $(1,0)$, so it also has a uniformization defined over $\mathbb{F}_p$ itself — the stereographic projection — and this one transports the code theory to *every* challenge field. The price is that the resulting Reed–Solomon domain is no longer a multiplicative coset; but it is rationally smooth of the same $(2,1)$ type as the genus-one rows, so Section 11.1 covers it.

Lemma 11.4 (stereographic uniformization of circle codes; no i required)**.** *Let $p \equiv 3 \pmod 4$, let $\mathcal{D} = gH \cup g^{-1}H \subseteq \mathbb{U}$ be a twin coset of size $2M$ with $\mathrm{ord}(g) \nmid 2M$ and $4 \mid M$, let $E := \iota^{-1}(\mathcal{D}) \subseteq \mathcal{C}(\mathbb{F}_p)$, and let $\mathbb{F} \supseteq \mathbb{F}_p$ be* any *finite field. Define $s(x,y) := y/(1+x)$ and $T := s(E)$.*

(i) *The four torsion points $(\pm 1, 0), (0, \pm 1)$ lie in $\iota^{-1}(H)$ and hence not in E; s is injective on E with inverse $s \mapsto \left(\frac{1-s^2}{1+s^2}, \frac{2s}{1+s^2}\right)$; and $T \subseteq \mathbb{F}_p \setminus \{0, \pm 1\}$ with $|T| = 2M$. (Over $\mathbb{F}_p$ one has $1+s^2 \neq 0$ throughout, since $p \equiv 3 \pmod 4$.)*

(ii) *With the diagonal vector $d := \left((1+s(P)^2)^{-w}\right)_{P \in E}$ and the coordinate bijection $s|_E : E \to T$,*

$$\mathcal{C}_w(\mathbb{F}, E) \;=\; d \circ \mathrm{RS}[\mathbb{F}, T, 2w+1].$$

Consequently, by Lemma 7.10, the bivariate circle code on E and $\mathrm{RS}[\mathbb{F}, T, 2w+1]$ have identical list sizes at every radius and identical $\varepsilon_{\mathrm{ca}}$ and $\varepsilon_{\mathrm{mca}}$ at every (pair of) radii — with no hypothesis on $i \in \mathbb{F}$.

(iii) *T is $(\psi, 2)$-smooth over $\mathbb{B} := \mathbb{F}_p$ in the sense of Definition 9.24 for the tangent-halving map $\psi(s) := \frac{s^2-1}{2s}$, of degree type $(2,1)$: the ψ-fibers within T are the pairs $\{s(P), s(P\tau)\} = \{s, -1/s\}$, $\tau := (-1, 0)$, complete because $\tau \in \iota^{-1}(H)$ and of uniform size 2 because $s^2 \neq -1$ in $\mathbb{F}_p$; and $N = M$, $Q = \psi(T) \subseteq \mathbb{F}_p$.*

Proof. (i) H is the unique subgroup of order M of the cyclic group $\mathbb{U}$ of order $p+1$; since $4 \mid M$ it contains all elements of order dividing 4, whose ι^{-1}-images are exactly the four torsion points. From $\mathrm{ord}(g) \nmid 2M$ we get $g, g^{-1} \notin H$, so $gH \cap H = g^{-1}H \cap H = \emptyset$ and $\mathcal{D} \cap H = \emptyset$; in particular $(-1,0) \notin E$, so s is defined on E, and $(1,0), (0,\pm 1) \notin E$, so $0, \pm 1 \notin T$. For the inverse formula: from $x^2+y^2=1$, $1+s^2 = \frac{(1+x)^2+y^2}{(1+x)^2} = \frac{2}{1+x}$, whence $\frac{1-s^2}{1+s^2} = (1+x)-1 = x$ and $\frac{2s}{1+s^2} = s(1+x) = y$; and $1+s^2 \neq 0$ for $s \in \mathbb{F}_p$ because -1 is a nonsquare when $p \equiv 3 \pmod 4$.

(ii) Every class of degree $\leq w$ has the canonical form $f_0(x) + y\, f_1(x)$ with $\deg f_0 \leq w$, $\deg f_1 \leq w-1$, a space of dimension $2w+1$ (proof of Lemma 7.11). Substituting the inverse of (i) and clearing denominators,

$$(1+s^2)^w \big(f_0(x) + y\, f_1(x)\big)\Big|_{x=\frac{1-s^2}{1+s^2},\, y=\frac{2s}{1+s^2}}$$

is a polynomial in s of degree at most $2w$: the monomial x^j contributes $(1-s^2)^j(1+s^2)^{w-j}$ and $y\,x^j$ contributes $2s\,(1-s^2)^j(1+s^2)^{w-1-j}$. These $2w+1$ polynomials form a basis of $\mathbb{F}[s]_{\leq 2w}$: the first family is even, the second odd, so it suffices to check each family separately; and in the variable $v := s^2$ the first family is $\{(1-v)^j(1+v)^{w-j}\}_{j \leq w}$, the image of the monomial basis $\{Y^j Z^{w-j}\}$ under the invertible substitution $(Y,Z) = (1-v, 1+v)$ (p odd), and similarly for the odd family. Hence $F \mapsto (1+s^2)^w F$ is a linear isomorphism from the circle space onto $\mathbb{F}[s]_{\leq 2w}$, and evaluating at $P \in E$ gives $F(P) = (1+s(P)^2)^{-w}\, f(s(P))$ with $f \in \mathbb{F}[s]_{\leq 2w}$ ranging over all of $\mathrm{RS}[\mathbb{F}, T, 2w+1]$: the displayed code identity. The vector d is $\mathbb{F}_p$-valued and everywhere nonzero, so Lemma 7.10 applies.

(iii) The group law is $(x_1,y_1)\cdot(x_2,y_2) = (x_1x_2 - y_1y_2,\, x_1y_2 + x_2y_1)$, so $P\tau = (-x,-y)$ and

$$s(P)\, s(P\tau) \;=\; \frac{y}{1+x} \cdot \frac{-y}{1-x} \;=\; \frac{-y^2}{1-x^2} \;=\; -1,$$

using $x^2+y^2=1$ and $y \neq 0$ on E. Since $\tau \in \iota^{-1}(H)$ ($2 \mid M$) and $\mathcal{D}H = \mathcal{D}$ coset-wise, $E\tau = E$, so the involution $s \mapsto -1/s$ preserves T. The fibers of ψ: $\psi(s) = b$ if and only if $s^2 - 2bs - 1 = 0$, whose two roots multiply to -1; hence every fiber is $\{s, -1/s\}$, of size exactly 2 since $s^2 = -1$ is impossible in $\mathbb{F}_p$, and complete within T by the closure just shown. Finally $f = X^2 - 1$ and $g = 2X$ are coprime of degrees $2 > 1$, and g is nonvanishing on T because $0 \notin T$. □

Corollary 11.5 (circle codes over every challenge field; the genus-one/circle transport question Question 22.3(b))**.** *Let E be a standard-position circle domain as in Lemma 11.4 with $M = 2^{21}$ (the deployed shape: $n_c = 2M = 2^{22}$, $k_c = 2^{21}+1$), and let $\mathbb{F} \supseteq \mathbb{F}_p$, $p = 2^{31}-1$, be a challenge field with $n_c < q := |\mathbb{F}| < 2^{256}$.*

(i) *If $q < 2^{100}$ — in particular $\mathbb{F} = \mathbb{F}_p$ and $\mathbb{F} = \mathbb{F}_{p^3}$, hence every field without i among the extensions $\mathbb{F}_{p^r}$, $r \leq 4$, that deployed verifiers sample from, since $i \in \mathbb{F}_{p^r}$ if and only if $2 \mid r$ — the row is degenerately unsafe at every deployed target by Proposition 9.20, and the genus-one/circle transport question Question 22.3(b) is out of scope there as posed.*

(ii) *If $i \in \mathbb{F}$, Corollaries 7.12 and 7.13 apply verbatim through Lemma 7.11.*

(iii) *If $i \notin \mathbb{F}$ and $q \geq 2^{100}$ (the first such field is $\mathbb{F}_{p^5}$), the stereographic model of Lemma 11.4 transports the row to $\mathrm{RS}[\mathbb{F}, T, k_c]$ on the $(\psi, 2)$-smooth tangent domain T, and Proposition 11.1 with*

$$m = 2^{20} + 2^{15} + 1, \qquad w = 2m - (k_c+1) = 2^{16}, \qquad \binom{2^{21}}{m} \;>\; p^{\,2^{16}} \cdot 2^{256}$$

(*verified exactly in integer arithmetic; the margin exceeds* 2^{63000}) *gives, for every* $q \in (n_c, 2^{256})$ *at once, an MCA and list cap on*

$$\delta \in \big[1-\rho_c-2^{-6},\ 1-\rho_c\big), \qquad \textit{error} > \tfrac{1}{2k_c}\big(1-\tfrac{n_c}{q}\big) > 2^{-23},$$

by the assembly of Corollary 11.2 (here $\beta = \log_2 p$*, effective cost* 2β*, asymptotic frontier* $g^* \approx 2^{-5.96}$*). The safe side is generic: Theorem 9.6 applies to the circle code directly (as already noted after its proof), and Theorems 10.1 and 10.4 and Corollary 10.7 apply to the diagonal model. Every bivariate circle row over every challenge field therefore carries a two-sided sandwich, and the deep-point conversion Theorem 2.6 does extend to the bivariate setting — through the stereographic model rather than a new proof.*

Proof. (i) is Proposition 9.20 with $q = p < 2^{31}$ resp. $q = p^3 < 2^{93}$: $\varepsilon_{\mathrm{mca}} \geq 1/q > 2^{-100}$ at every radius, so $\delta_C^*(\varepsilon^*) = 0$ for every deployed target $\varepsilon^* \leq 2^{-100}$; the parity criterion is $\mathrm{ord}(i) = 4 \mid p^r - 1$ if and only if r is even, as $p \equiv 3 \pmod 4$. (ii) is quoted. (iii): the transport is Lemma 11.4(ii), the smoothness is (iii), and the hypotheses of Proposition 11.1 hold with $K \in \{k_c, k_c+1\}$: $2m = k_c+1+2^{16} \geq K$, $em = m \leq k_c \leq K$, $m \leq N = 2^{21}$. The single displayed integer inequality dominates both thresholds for every $q < 2^{256}$, since $q/k_c + 1 \leq q < 2^{256}$ and $2^{-100}q < 2^{156}$; the edge satisfies $1 - 2m/n_c = 1 - \rho_c - 2^{-6} - 2^{-21} + 2^{-22} \leq 1 - \rho_c - 2^{-6}$, and $\lfloor \delta_0 n_c \rfloor = n_c - 2m \leq n_c - k_c - 1$, so Theorem 2.6, Fact 2.4, and Lemma 2.5 assemble exactly as in Corollary 11.2; the error numerics use $q \geq p^5$, so $\frac{1}{2k_c}\big(1-\frac{n_c}{q}\big) > 2^{-22}(1-2^{-21})(1-2^{-133}) > 2^{-23}$. □

11.3 Explicit witness words and explicit certifying pairs

Every floor so far selects its received word by pigeonhole over a prefix. At the first staircase step the pigeonhole can be removed entirely: if the support is assembled from *antipodal pairs*, its first elementary symmetric function vanishes identically, and the received word collapses to a pure power.

Theorem 11.6 (explicit head-and-pairs floor)**.** *Let* D *be* (φ, c)*-smooth over* $\mathbb{B}$ (*Definition 7.1*) *with* $N := n/c$ *even, and suppose* $Q := \varphi(D)$ *satisfies* $-Q = Q$ *and* $0 \notin Q$*, so that* Q *is partitioned into* $N/2$ *antipodal classes* $\{y, -y\}$*. Let* $K \leq n$ *and let* m *satisfy* $2 \leq m \leq N-2$ *and* $cm \leq K - 1 + 2c$*.*

(i) *If* m *is even, the* pure power word $u := \varphi^m|_D$ *satisfies*

$$\big|\Lambda\big(\mathrm{RS}[\mathbb{F}, D, K],\ 1-\tfrac{cm}{n},\ u\big)\big| \ \geq\ \binom{N/2}{m/2},$$

the listed codewords being $c_M := u - \Lambda_M|_D$ *for* M *a union of* $m/2$ *antipodal classes.*

(ii) *If* m *is odd, then for every* $t \in Q$ *the* one-head word $u := \big(\varphi^m - t\,\varphi^{m-1}\big)|_D$ *satisfies*

$$\big|\Lambda\big(\mathrm{RS}[\mathbb{F}, D, K],\ 1-\tfrac{cm}{n},\ u\big)\big| \ \geq\ \binom{N/2-1}{(m-1)/2},$$

with $M = \{t\} \cup ((m-1)/2$ *antipodal classes distinct from* $\{t, -t\})$*.*

In both cases the word, the family, and every listed codeword are explicit closed forms, with no pigeonhole and no subfield division. Both deployed towers qualify: on a multiplicative coset, Q is a coset of the even-order subgroup μ_N, so $-Q = Q$ and $0 \notin Q$; on a Chebyshev domain $D = \chi(\mathcal{D})$ with $2c \mid M$, the antipodal classes of $Q = T_c(D)$ are exactly the T_2-fibers inside Q, complete of size 2 by Lemma 7.8 at scales c and $2c$ (two elements with equal T_2-value satisfy $y' = -y$, and $0 \notin Q$ since a T_2-fiber $\{0\}$ would have size 1).

Proof. In case (i), M is a union of antipodal pairs, so $e_1(M) = 0$ and $\Lambda_M = \varphi^m + \sum_{j\geq 2}(-1)^j e_j(M)\varphi^{m-j}$, whence $c_M := u - \Lambda_M|_D = -\sum_{j\geq 2}(-1)^j e_j(M)\varphi^{m-j}|_D$ has degree at most $c(m-2) \leq K-1$ and agrees with u on the cm fiber points of M. In case (ii), $e_1(M) = t + \sum(\text{pair sums}) = t$, so the slot-1 coefficient of Λ_M is $-t\varphi^{m-1}$ for every M in the family, and $c_M := u - \Lambda_M|_D$ again consists of the slots $j \geq 2$ only, of degree at most $c(m-2) \leq K-1$; note that $e_2(M)$ and beyond are permitted to vary, since they land in the codeword rather than in the received word. Injectivity: $c_M = c_{M'}$ as words forces $\Lambda_M = \Lambda_{M'}$ on all of D, and both are polynomials of degree $cm < cN = n$, hence equal as polynomials, with equal root multisets, hence $M = M'$ through the fibers. The family counts are the number of choices of antipodal classes. □

Corollary 11.7 (explicit witnesses at the deployed rows)**.** *At the KoalaBear row of Corollary 5.4 ($n = 2^{21}$, $k = 2^{20}$, $\mathbb{B} = \mathbb{F}_p$, $q = p^6$) and at the Chebyshev line-round row of Corollary 7.13*(a) *($n = 2^{21}$, $k = 2^{20}$, $q = p'^4$, $p' = 2^{31}-1$, target 2^{-100}):*

(i) *(MCA, gap 2^{-8}) With $c = 2^{12}$, $m = k/c + 2 = 258$, $K = k+1$: the explicit word $u = X^{k+2^{13}}|_D$ (resp. $u = T_c^{258}|_D$) has list at agreement $k + 2^{13}$ at least $\binom{256}{129} \geq 2^{251}$, exceeding $q/k + 1 < 2^{166}$ (resp. $< 2^{104}$); by Theorem 2.6 with $\eta = \frac{1}{2}$, exactly as in Theorem 4.1, $\varepsilon_{\mathrm{ca}}(C,\delta) > \frac{1}{2k}(1 - \frac{n}{q})$ on the integer-admissible part of $[\frac{1}{2} - 2^{-8}, \frac{1}{2})$ and $\varepsilon_{\mathrm{mca}}$ obeys the same bound throughout it.*

(ii) *(lists, gap 2^{-7}) With $c = 2^{14}$, $m = k/c + 1 = 65$, $K = k$: for any $t \in Q$ the explicit word $u = \left(X^{k+2^{14}} - t\,X^k\right)|_D$ (resp. its T_c analogue) has list at agreement $k + 2^{14}$ at least $\binom{63}{32} = 916312070471295267 > 2^{59}$, exceeding $2^{-128}q < 2^{58}$ (resp. $2^{-100}q < 2^{24}$): the first staircase step of the grand list challenge, at its original printed edge $\frac{1}{2} - 2^{-7}$, is certified by a fully explicit word.*

All four inequalities are verified exactly in integer arithmetic.

Proof. The hypotheses of Theorem 11.6 hold with $N = 512$ resp. 128: parity as displayed, $cm \leq K - 1 + 2c$ (with equality in (i)), and the Chebyshev scale conditions $2c \mid M = 2^{21}$. The threshold comparisons are the displayed integer inequalities; the MCA conversion and monotone assembly are as quoted. □

The words above certify lists explicitly, and Theorem 2.6 converts large lists into CA failure — but by contradiction, without exhibiting the failing pair. The explicit-witness question Question 22.1, however, asks for the *pair*. The proof of Theorem 2.6 is more generous than its statement: it converts through the completely explicit simple-pole lines $\big(U/(x-\alpha),\, -1/(x-\alpha)\big)$ — precisely the residue-line normal form with extension denominator that Corollary 6.2 predicts — and its collision count can be read as a lower bound holding for *most* poles.

Theorem 11.8 (explicit certifying pairs, up to the choice of a pole)**.** *At the KoalaBear row of Corollary 5.4, let $u := X^{k+2^{13}}|_D$ be the pure power word of Corollary 11.7*(i)*, let $L_0 :=$*

$\lceil (q-n)/k \rceil < 2^{166}$, *and fix any explicit subfamily of* L_0 *of its listed codewords* $P_M \in \mathbb{F}[X]_{<k+1}$ (*e.g. the lexicographically least* L_0 *antipodal-class sets;* $L_0 \le \binom{256}{129}$). *For* $\alpha \in \Omega := \mathbb{F} \setminus D$ *define the explicit pair*

$$f_\alpha(x) := \frac{u(x)}{x-\alpha}, \qquad g_\alpha(x) := -\frac{1}{x-\alpha} \qquad (x \in D).$$

Then for at least half of all $\alpha \in \Omega$, *the pair* (f_α, g_α) *has at least*

$$\left\lceil \frac{q-n}{3k} \right\rceil$$

distinct CA-bad slopes at radius $\delta = \frac{1}{2} - 2^{-8}$, *namely among the explicit values* $\{P_M(\alpha)\}$; *its CA-bad-slope density exceeds* $2^{-21.6} > 2^{-22}$ *there and at every larger integer-admissible radius, and its MCA-bad-slope density exceeds the same bound throughout* $[\frac{1}{2} - 2^{-8}, \frac{1}{2})$. *All but at most* p *of these poles have* $\alpha \notin \mathbb{B}$, *and for those the pair is genuinely* $\mathbb{F}$*-valued, in accordance with Corollary 6.2. The same construction on the circle line-round row gives density* $> 2^{-21.6} > 2^{-23}$ *for at least half of all poles.*

Proof. Each P_M agrees with u on at least $a = k + 2^{13}$ points, so, as in the proof of Theorem 2.6, for every $\alpha \in \Omega$ and every M the slope $P_M(\alpha)$ carries the point $f_\alpha + P_M(\alpha) g_\alpha$ to within radius $1 - a/n = \frac{1}{2} - 2^{-8}$ of C, while the pair is farther than every $\delta < 1 - k/n = \frac{1}{2}$ from $C^{\equiv 2}$ by the $(X-\alpha)G + 1$ argument there; so every distinct value $P_M(\alpha)$ is a CA-bad slope at every integer-admissible $\delta \in [\frac{1}{2} - 2^{-8}, \frac{1}{2})$, hence (as failing CA is failing MCA, Fact 2.4) an MCA-bad slope throughout that interval. It remains to count distinct values for many α at once. Distinct sets M give distinct polynomials P_M (Theorem 11.6), and $P_M - P_{M'}$ is nonzero of degree at most k, so the total number of colliding pairs over all $\alpha \in \Omega$ is at most $k\binom{L_0}{2}$. Let $x^* := \lceil 2k\binom{L_0}{2}/(q-n) \rceil$; by Markov, fewer than $(q-n)/2$ poles have collision count at least x^*, so at least half have collision count at most $x^* - 1$. For such α, Cauchy–Schwarz as in Theorem 2.6 gives

$$M(\alpha) \;\ge\; \frac{L_0^2}{L_0 + 2(x^*-1)} \;\ge\; \frac{L_0}{1 + 2k(L_0-1)/(q-n)} \;>\; \left\lceil \frac{q-n}{3k} \right\rceil - 1,$$

so, since $M(\alpha)$ is an integer, $M(\alpha) \ge \lceil (q-n)/(3k) \rceil$. The last comparison and $\lceil (q-n)/3k \rceil / q > 2^{-21.6}$ are verified exactly in integer arithmetic at both rows. The subfield count is $|\Omega \cap \mathbb{B}| \le |\mathbb{B}| = p$. The obstruction is projective, not merely affine. If λg_α were $\mathbb{B}$-valued for some $\lambda \in \mathbb{F}^\times$, then for two distinct $x, y \in D$ the quotient of its values would give

$$\frac{y-\alpha}{x-\alpha} \in \mathbb{B}.$$

Writing this quotient as $r \in \mathbb{B}$, the case $r = 1$ would force $x = y$, while $r \ne 1$ gives $\alpha = (rx - y)/(r-1) \in \mathbb{B}$, a contradiction. Hence no common scalar makes the pair $\mathbb{B}$-valued, so it lies outside the inert class of Lemma 6.1, as Corollary 6.2 requires of any pair with this density. □

Remark 11.9 (what remains of the explicit-witness question Question 22.1). Theorem 11.8 answers the explicit-witness question Question 22.1 at $\delta = \frac{1}{2} - 2^{-8}$, on the upper half of the stated band, in the predicted normal form, and for both densities the problem asks about; the non-constructiveness has collapsed from a received word — an object of $n \log_2 q$ bits with no known certificate — to the choice of one field element in an explicit majority subset of $\mathbb{F} \setminus D$. Two residues remain. First, a certified *individual* pair: derandomizing the pole. Second, the left half

of the band: at gap exactly 2^{-7} the explicit family at scale 2^{13} has only $\binom{128}{65} < 2^{125}$ members, short of the $\approx 2^{166}$ the counting needs, so explicitness there — and throughout the widened band of Corollary 10.9, where the floors are pigeonhole-selected deep prefixes — would require explicit families with prescribed deep prefixes, i.e. explicit subsets of μ_N with many prescribed elementary symmetric functions. We record this as the surviving form of the explicit-witness question Question 22.1.

11.4 The failure profile: the deep-point ceiling

Theorem 2.6 was used throughout with $\eta = \frac{1}{2}$, giving error floors of $\frac{1}{2k}(1 - \frac{n}{q})$. The dichotomy improves when the list is far above threshold, and the improvement is exactly quantifiable.

Proposition 11.10 (optimized failure profile)**.** *Let $C = \mathrm{RS}[\mathbb{F}, D, k]$, let δ be integer-admissible ($\lfloor \delta n \rfloor \le n - k - 1$, $q > n$), and put $T := \frac{1}{k}\big(1 - \frac{n}{q}\big)$. If some received word has $\big|\Lambda(C^+, \delta, u)\big| \ge 1 + \kappa q T$ for a real $\kappa > 0$, then*

$$\varepsilon_{\mathrm{ca}}(C, \delta) \;\ge\; \frac{\kappa}{\kappa + 1}\, T.$$

Proof. Write $x := \varepsilon_{\mathrm{ca}}(C, \delta)$; if $x \ge T$ there is nothing to prove, so assume $x < T$ and set $\eta := x/T \in [0, 1)$. Then $x = \eta T$, so Theorem 2.6 applies and gives $L \le \lceil qx/(1-\eta) \rceil \le qx/(1-\eta) + 1$, i.e. $\kappa q T \le L - 1 \le qx/(1 - x/T)$. Rearranging, $\kappa(T - x) \le x$, i.e. $x \ge \kappa T/(\kappa + 1)$. □

Corollary 11.11 (deployed profile at the first staircase step)**.** *At the KoalaBear row, the fiber floor of Lemma 7.2 at scale 2^{13} gives $L \ge \lceil \binom{256}{130}/p \rceil > 2^{220}$ at agreement $k + 2^{14}$, hence $\kappa \ge 2^{54}$ and*

$$\varepsilon_{\mathrm{ca}}\big(C, \delta\big) \;\ge\; \big(1 - 2^{-54}\big)\, \frac{1}{k}\Big(1 - \frac{n}{q}\Big)$$

for every integer-admissible $\delta \in [\frac{1}{2} - 2^{-7}, \frac{1}{2} - 2^{-21}]$, with the same lower bound for $\varepsilon_{\mathrm{mca}}(C, \delta)$ throughout $[\frac{1}{2} - 2^{-7}, \frac{1}{2})$: twice the printed constant of Corollary 5.4, saturating the dichotomy's ceiling to relative error 2^{-54}. On the circle line-round row the same floor gives $\kappa \ge 2^{116}$ and the factor $1 - 2^{-116}$.

Proof. $\kappa = \lfloor (L-1)k/(q-n) \rfloor \ge 2^{54}$ (resp. 2^{116}) is verified exactly; $\kappa/(\kappa+1) \ge 1 - 2^{-54}$; the list persists at every larger radius, and Fact 2.4 with Lemma 2.5 extend as usual. □

Remark 11.12 (the ceiling, and what keeps the error-one question Question 22.2 open). Two ceilings frame the error-one question Question 22.2. First, the route ceiling: as $\kappa \to \infty$ the profile of Proposition 11.10 tends to $T = \frac{1}{k}(1 - \frac{n}{q})$ and no list, however large, certifies more through Theorem 2.6 — error beyond $\approx \frac{1}{k}$ requires strengthening the dichotomy itself, which is exactly the unnamed problem following the error-one question Question 22.2. Second, at the specific radius $1 - \rho - \frac{1}{64}$ of the error-one question Question 22.2 over *prime* fields $2^{150} \le p \le 2^{256}$, the present floors certify nothing at all: the only subfield is $\mathbb{B} = \mathbb{F}_p$, so every prefix frontier obeys $g^* \le 1/\beta = 1/\log_2 p \le \frac{1}{150} < \frac{1}{64}$, and the radius lies strictly outside the entropy–subfield envelope. (There is also no confinement barrier there: Lemma 6.1 caps $\mathbb{B}$-valued pairs at $|\mathbb{B}|/q$, which is vacuous when $\mathbb{B} = \mathbb{F}$.) So the error-one question Question 22.2 is genuinely beyond both mechanisms of this paper: it needs either non-prefix floors at macroscopic gaps over prime fields, or conversions past $\frac{1}{k}$. The certified state after this section: error at least $(1 - 2^{-54})\frac{1}{k}(1 - \frac{n}{q})$ at gap 2^{-7}, error at least $\frac{1}{2k}(1 - \frac{n}{q})$ at the widened gaps, and error one still open everywhere.

11.5 Certificate closure of the two-sided sandwich

Proposition 9.23 supplied the initial certificate format and Table 3 the formal core. The pincer and the refinements above add new verdict types. This subsection closes the certificate grammar over the statements used by the two-sided sandwich, prints the deployed certificates, and isolates the single imported proximity-gap input.

Definition 11.13 (certificate grammar, second edition)**.** Extend Definition 9.21 by the following verdict types, each a tuple of integers together with a theorem tag:

(V5) **Prefix-floor (PF) / rational-floor (RF).** A scale datum (polynomial or Chebyshev c, or rational type (a, e)), integers (m, w, K), and the single inequality $\binom{N}{m} > |\mathbb{B}|^w \cdot \Theta$, $\Theta \in \{q/k+1,\ \varepsilon^* q\}$: mark the band $[1 - am/n, 1 - \rho)$ MCA-unsafe (via Propositions 11.1 and 10.8, Theorem 2.6) resp. $1 - am/n$ list-unsafe.

(V6) **Explicit-head (XH).** Integers (c, m, K) with $cm \le K - 1 + 2c$ and the inequality $\binom{N/2-\epsilon}{\lfloor m/2 \rfloor} > \Theta$, $\epsilon \in \{0, 1\}$ by parity: as (V5), with the received word printed in closed form (Theorem 11.6).

(V7) **Mutual-from-correlated (MC).** An integer r with $2r \le w_{\min} - 1$ and a CA-verdict handle: transfer a CA-safe verdict at radius r/n to MCA-safe with error $\max\{\varepsilon_{\mathrm{CA}}, r/q\}$ (Theorem 10.1). The CA handle is either **HJ** — integers (r, L_2) with $n-2r > 0$, $(n-2r)^2 > (k-1)n$, $(L_2+1)\big((n-2r)^2-(k-1)n\big) > n(n-2r-k+1)$, and $(1+(r+1)L_2)/q \le \varepsilon^*$, certifying the self-contained half-Johnson CA bound of Theorem 10.4 — or **CH** — the inequality $2r \le n - k$ with error $n/q \le \varepsilon^*$, tagged IMPORT: [BCIKS20, Thm. 4.1] (Corollary 10.7).

(V8) **Profile (PR).** A floor handle and an integer κ with $L \ge 1 + \kappa q T$: certify error $\ge \frac{\kappa}{\kappa+1} T$ (Proposition 11.10).

(V9) **Identity-prefix list (IP).** Integers (n, m, K, p, q, t) with $D \subseteq \mathbb{F}_p$ and

$$\binom{n}{m} 2^t > p^{m-K} q :$$

mark agreement m list-unsafe at target 2^{-t} by Lemma 14.1 and Proposition 14.2.

(V10) **Flexible-budget MCA (FB).** Integers (n, m, k, p, q, t), $w = m - k - 1$, satisfying

$$\binom{n}{m} > p^w \left\lfloor \frac{q}{2^t} \right\rfloor, \qquad \binom{\lfloor q/2^t \rfloor + 1}{2} k < q - p, \qquad 2^t \nmid q :$$

mark agreement m MCA-unsafe at target 2^{-t} by Corollary 15.1 and Proposition 15.2.

Theorem 11.14 (certificate closure of the two-sided sandwich)**.** *With the grammar of Proposition 9.23 extended by Definition 11.13:*

(i) (*Soundness*) *Every verdict of Definition 11.13 is implied by its cited theorem once its integer inequalities hold; Theorem 9.22 extends verbatim.*

(ii) (*Completeness for the printed results*) *Every clause of Theorem 10.12, of Table 5, of Corollary 11.2, of Corollary 11.5, and of Corollaries 11.7 and 11.11 is certified by a certificate of Proposition 9.23 or Definition 11.13; the deployed certificates are printed in Table 7.*

(iii) (*Complexity*) *A binomial certificate of type* (V5), (V6), (V9), *or* (V10) *is checked with* $O(N)$ *multiplications of integers of* $O(N + \log q)$ *bits (the binomial by its product formula, the power* $|\mathbb{B}|^w$ *by repeated squaring); every other certificate with* $O(1)$ *arithmetic operations on integers of* $O(\log q + \log n)$ *bits. A full row is checked in* $O(n)$ *big-integer multiplications.*

(iv) (*Import isolation*) *Exactly one imported statement occurs anywhere in the grammar: the* IMPORT *tag of the CH handle. Deleting it degrades only the safe edge, from* $(1-\rho)/2$ *to* $\max\{(1-\rho)/3,$ *half-Johnson*$\}$, *and leaves every other verdict intact. The remaining statements — Table 3 extended by the rows of Table 6 — quantify over explicitly bounded finite sets and use only the listed elementary ingredients; their dependency closure is import-free.*

Proof. (i) is quotation: each verdict's hypotheses are the hypotheses of its theorem, specialized to integers. (ii) is inspection of the quoted proofs against Table 7: each proof consists of its printed integer inequalities plus the cited assembly steps, all of which are verdicts or core statements. (iii): $\binom{N}{m} = \prod_{j=1}^{m} \frac{N-m+j}{j}$ uses $2m \le 2N$ multiplications and divisions of integers of at most $N + O(\log N)$ bits; thresholds and powers are $O(\log w)$ squarings of $O(w \log |\mathbb{B}| + \log q)$-bit integers, and $w \log |\mathbb{B}| = O(N)$ in every certificate above; the remaining verdicts compare $O(1)$ products. (iv): the only IMPORT tag is CH; Theorem 10.12 lists its conditional clause separately, and the self-contained clauses cite only Theorems 9.6, 10.1 and 10.4 on the safe side. The new core rows are in Table 6; each proof above is elementary in the stated ingredients, as the reader may verify line by line. □

Added core statement	Proof ingredients
Theorem 10.1 (mutual from correlated)	minimum weight, two-point interpolation
Theorem 10.4 (half-Johnson CA)	pair-list double counting, convexity
Proposition 10.8, Proposition 11.1	symmetric functions, graded pigeonhole
Theorem 11.6 (explicit heads)	antipodal pairing, no pigeonhole
Theorem 11.8 (explicit pairs)	Theorem 2.6's collision count, Markov
Proposition 11.10 (optimized profile)	one-line rearrangement of Theorem 2.6
Propositions 14.2 and 15.2	exact prefix pigeonhole, flexible-budget conversion
Lemma 11.4 (stereographic model)	conic parametrization, basis change
Definition 11.13, Theorem 11.14	finite case analysis over the above

Table 6: The formal finite core, continued: statements added by the pincer and the refinements, extending Table 3. Each row quantifies over explicitly bounded finite sets; the dependency order is top to bottom, below the rows of Table 3. The one import (CH) is not a core row.

Remaining residues. The refinements above reduce the remaining landscape to explicit residues. The genus-one and circle-code questions are reduced to the same normalized safe-side package after the rational-floor and stereographic transports of Corollaries 11.2 and 11.5. The explicit-witness question is resolved at the level of closed-form words and residue-line pairs in Corollary 11.7 and Theorem 11.8, with the remaining scope stated in Remark 11.9. The error-one profile is not settled here: Corollary 11.11 raises the certified profile to $(1 - 2^{-54})\frac{1}{k}(1 - \frac{n}{q})$, while Remark 11.12 explains why reaching $1 - o(1)$ requires new input. The frontier residue is the active package stated in Section 12.1: quotient/prefix flatness, the base-field-normalized split-pencil census, and primitive shift-pair control. The certificate material above is finite and

Row / verdict	Certificate tuple		Exact comparison
KB MCA, FB	(m,w,t) $(1116047,67470,128)$	=	$\binom{2^{21}}{1116047} > p^{67470}\lfloor p^6/2^{128}\rfloor$; $\binom{\lfloor p^6/2^{128}\rfloor+1}{2}2^{20} < p^6-p$; $2^{128}\nmid p^6$
KB list, IP	(m,w,t) $(1116046,67470,128)$	=	$\binom{2^{21}}{1116046}2^{128} > p^{67470}p^6$
M31 MCA, FB	(m,w,t) $(1116023,67446,100)$	=	$\binom{2^{21}}{1116023} > (p')^{67446}\lfloor (p')^4/2^{100}\rfloor$; $\binom{\lfloor (p')^4/2^{100}\rfloor+1}{2}2^{20} < (p')^4-p'$; $2^{100}\nmid (p')^4$
M31 list, IP	(m,w,t) $(1116022,67446,100)$	=	$\binom{2^{21}}{1116022}2^{100} > (p')^{67446}(p')^4$
KB MCA, XH	$(c,m) = (2^{12},258)$, u $X^{k+2^{13}}\vert_D$	=	$\binom{256}{129} > q/k+1$
KB list, XH	$(c,m) = (2^{14},65)$		$\binom{63}{32} > 2^{-128}q$
M31 MCA, XH	$(c,m) = (2^{12},258)$, u $T_c^{258}\vert_D$	=	$\binom{256}{129} > q/k+1$
M31 list, XH	$(c,m) = (2^{14},65)$		$\binom{63}{32} > 2^{-100}q$
Both MCA rows, MC	$r = 2^{19}$		$2r \le n-k = 2^{20}$
KB safe, HJ	$(r,L_2) = (307121,1001282)$		$(1+(r+1)L_2)/q < 2^{-147}$
KB safe, CH IMPORT	$r = 2^{19}$		$n/q < 2^{-164}$
M31 safe, CH IMPORT	$r = 2^{19}$		$n/q < 2^{-102}$
KB profile, PR	$\kappa = 2^{54}$		$\lceil\binom{256}{130}/p\rceil > 1+2^{54}qT$
i-free circle, RF	$(m,w) = (2^{20}+2^{15}+1, 2^{16})$		$\binom{2^{21}}{m} > (p')^{2^{16}}2^{256}$

Table 7: Selected exact certificates for the deployed rows (Definition 11.13; $p = 2^{31}-2^{24}+1$, $p' = 2^{31}-1$, and $T = \frac{1}{k}(1-\frac{n}{q})$). The first four rows are the active unsafe endpoints used in Theorem 10.12 and Table 5; the earlier PF endpoints are intermediate certificates only. The two IMPORT rows are optional half-distance clauses and are the only non-self-contained entries.

checkable; a proof-assistant transcription would be implementation work rather than a new mathematical ingredient.

12 Frontier package: statements, certificates, and remaining inputs

This part of the paper contains the finite certificate calculus and a candidate safe-side certificate interface. The status is deliberately separated into proved certificates, imported generic safe anchors, and three residual buckets whose exhaustive coverage and parameter coalescing are not proved here.

component	status	role
prefix floors, simple-pole conversion, and deployed integer certificates	proved, with exact integer arithmetic	supply the last unsafe agreements in the deployed rows
deep-regime MCA, tangent accounting, and MCA-from-CA pincer	proved, elementary	supply unconditional safe anchors and remove support-mismatch cells in the covered region
half-distance CA proximity gap	imported where explicitly marked	gives the generic half-distance safe edge used only in the broad sandwich
quotient/prefix flatness	residual input	controls worst prefix fibers and quotient recursions beyond the entropy–subfield envelope
base-field-normalized split-pencil census	residual input	controls the primitive interior aperiodic support census after slope elimination
primitive shift-pair control	residual input	controls the primitive part of the exact second-moment ledger

The active finite deployed frontier is the following adjacent-pair prediction:

row	proved unsafe a_0	candidate safe $a_0 + 1$	deficit floor (bits)
KoalaBear MCA	1116047	1116048	> 22.1
KoalaBear list	1116046	1116047	> 22.0
Mersenne-31 MCA	1116023	1116024	> 3.2
Mersenne-31 list	1116022	1116023	> 3.0

The unsafe entries are theorems, proved by the identity-prefix floor together with the flexible-budget conversion. The adjacent safe entries are conjectural. These deficits are not separate allowances for separate arguments: the complete paid and residual upper numerator must be summed before comparison with the row budget.

The candidate residual certificate uses three buckets:

structured quotient/prefix pullback control,
base-field-normalized interior split-pencil control,
local primitive shift-pair control.

The rank-one, balanced-core, tangent, and collision sections below are reductions toward these components, not additional assumptions imposed in parallel.

12.1 Active interface for the remaining inputs

Every proposed closing statement is to be read through the upper-certificate template of Section 21. The paper does not claim that the three buckets are already an exhaustive compiler. The first input prices the boundary prefix and quotient-pullback cells with the necessary mean-plus-one model. The second prices the interior balanced profiles with the base-field floor of Proposition 20.1. The third is local, rather than merely averaged: it bounds the primitive same-prefix neighbours of a fixed support, the quantity required for a max-fiber or per-line bound. The exact second-moment identity of Theorem 19.5 records the average of these local degrees but does not by itself replace the local input.

Remark 12.1 (active closure interface). Let $R(n) \to \infty$ with $R(n) = o(n)$, and let $L(n)$ dominate the multiplicative losses and the number of residual cells in the asymptotic upper ledger. The reserve form is consumed only under

$$\log L(n) = o\big(R(n) \log |\mathbb{B}|\big).$$

In particular, polynomial loss is compatible with $R \gg \log n$; an unspecified $e^{o(n)}$ loss is not. At a finite adjacent row, a complete certificate would have to satisfy the single integer inequality displayed in Remark 21.5 after proving coverage and coalescing. No unspecified factor is promoted to a finite certificate.

13 Threshold certificate compilers

13.1 Integer staircases and corridors

Definition 13.1 (integer numerator staircase)**.** Fix an agreement interval

$$I = \{a_{\min}, a_{\min} + 1, \dots, a_{\max}\}$$

and a nonincreasing function

$$N : I \longrightarrow \mathbb{Z}_{\geq 0}.$$

For MCA or line-MCA rows, $N(a)$ is the number of bad sampled parameters at agreement at least a; the closed integer Hamming radius is $r = n - a$. For list rows one may use the same agreement coordinate, so the list numerator is nonincreasing in a. Given a denominator Q and a target ε^*, put

$$B_* := \lfloor \varepsilon^* Q \rfloor .$$

An agreement a is *safe* if $N(a) \leq B_*$ and *unsafe* if $N(a) > B_*$.

Theorem 13.2 (staircase localization)**.** *Let $N : I \to \mathbb{Z}_{\geq 0}$ be nonincreasing. Exactly one of the following holds:*

(i) *every $a \in I$ is safe;*

(ii) *every $a \in I$ is unsafe;*

(iii) *there is a unique first safe agreement $a_* \in I$ such that the safe set is $\{a \in I : a \geq a_*\}$ and the unsafe set is $\{a \in I : a < a_*\}$.*

In the interior case, when $a_* > a_{\min}$, one has

$$N(a_* - 1) > B_* \geq N(a_*).$$

Proof. The predicate $N(a) \leq B_*$ is monotone nondecreasing in a: if it holds at a, then for $a' \geq a$ we have $N(a') \leq N(a) \leq B_*$. A monotone Boolean function on a finite linearly ordered set is either identically false, identically true, or has a unique first true point. The displayed inequalities are exactly the assertion that the point just before the first true point is false while the first true point is true. □

Corollary 13.3 (closed-radius endpoint convention)**.** *In the interior case of Theorem 13.2, the largest safe closed integer radius is*

$$r_{\text{safe}} = n - a_*,$$

and the first unsafe closed integer radius above it is $r_{\text{safe}} + 1$*. As a real closed-ball threshold under the convention* $a(\delta) = \lceil (1-\delta)n \rceil$ *of Section 8, the supremum of safe radii is*

$$\frac{n - a_* + 1}{n},$$

and this endpoint is not attained when $N(a_* - 1) > B_*$*.*

Proof. The integer radius corresponding to agreement a is $n - a$, so the first safe agreement a_* gives the largest safe integer radius $n - a_*$. A real radius δ is safe exactly when $\lceil (1-\delta)n \rceil \geq a_*$, which holds if and only if $(1-\delta)n > a_* - 1$, i.e. for all $\delta < (n - a_* + 1)/n$. At $\delta = (n - a_* + 1)/n$ the agreement requirement is $a_* - 1$, which is unsafe in the interior case. □

Theorem 13.4 (corridor from lower and upper certificates)**.** *Suppose that* $L(a) \leq N(a) \leq U(a)$ *are certified lower and upper bounds on the same staircase. Define*

$$A_{\text{unsafe}} := \{a \in I : L(a) > B_*\}, \qquad A_{\text{safe}} := \{a \in I : U(a) \leq B_*\}.$$

If an interior first safe agreement a_* *exists, then*

$$1 + \max A_{\text{unsafe}} \leq a_* \leq \min A_{\text{safe}},$$

with the conventions $1 + \max \varnothing = a_{\min}$ *and* $\min \varnothing = a_{\max}$*.*

Proof. If $L(a) > B_*$, then $N(a) > B_*$, so a is unsafe and the first safe point is strictly larger than a. Taking the largest such certified unsafe point gives the lower bound. If $U(a) \leq B_*$, then $N(a) \leq B_*$, so a is safe and the first safe point is at most a. Taking the smallest such certified safe point gives the upper bound. The displayed conventions are exactly the empty-set cases. □

Remark 13.5 (relation to the main interface). Proposition 8.1 is the one-step case of Theorems 13.2 and 13.4: a matched pair of certificates $L(a_0) > B_* \geq U(a_0 + 1)$ pins $a_* = a_0 + 1$ exactly. In the conditional theorems Theorems 8.2 and 8.3, the hypothesized band inputs enter precisely as the upper staircase U, and the unsafe floors of Section 10 enter as the lower staircase L; Theorem 13.4 is the exact bookkeeping of the sandwich between them at a fixed budget.

Theorem 13.6 (coarse steepness compiler)**.** *Let* $J \subseteq I$ *be an interval of consecutive agreements, let* $M(a) > 0$ *be a positive model count, and suppose a row packet proves*

$$\frac{M(a)}{E} \leq N(a) \leq EM(a) \qquad (E \geq 1)$$

for every $a \in J$*. For a single adjacent pair* $a, a+1 \in J$*, if*

$$\frac{M(a)}{M(a+1)} > E^2,$$

then the two ambiguity intervals

$$\left[\frac{M(a)}{E}, EM(a)\right] \quad \text{and} \quad \left[\frac{M(a+1)}{E}, EM(a+1)\right]$$

are disjoint, and a fixed budget B_* can lie in at most one of them.

If the same separation holds for every adjacent pair in J, then the ambiguity intervals are pairwise disjoint in their natural order. Consequently at most one agreement in J can remain undecided by the coarse envelope for a fixed budget B_*; every other agreement has a certified safe or unsafe sign.

Proof. For the adjacent pair, the displayed inequality is equivalent to

$$\frac{M(a)}{E} > EM(a+1),$$

which is precisely the assertion that the lower end of the a-interval is above the upper end of the $(a+1)$-interval. If $B_* < M(b)/E$, then $N(b) \geq M(b)/E > B_*$ and b is certified unsafe. If $B_* \geq EM(b)$, then $N(b) \leq EM(b) \leq B_*$ and b is certified safe. Thus only budgets inside the ambiguity interval for b can be undecided.

When the adjacent separation holds throughout J, induction gives

$$\frac{M(a)}{E} > EM(a+1) \geq \frac{M(a+1)}{E} > EM(a+2) \geq \cdots$$

for consecutive points, using $E \geq 1$, so all ambiguity intervals are pairwise disjoint. A single integer budget can belong to at most one member of a pairwise disjoint collection of intervals, and the sign criterion of the earlier paragraph certifies all remaining agreements. □

13.2 Paid ledger cells

Definition 13.7 (high-agreement tangent exact cell)**.** For a Reed–Solomon row of length n and dimension $k < n$, put $r(A) = n - A$ and

$$R_{\tan} := \left\lfloor \frac{n-k}{3} \right\rfloor.$$

The high-agreement tangent paid cell is

$$\mathrm{Paid}^{\mathrm{hi}}_{\tan}(n,k,A) := r(A) + 1$$

when $0 \leq r(A) \leq R_{\tan}$, and is unavailable outside this range unless a separate tangent/common-line theorem is supplied.

Proposition 13.8 (exact tangent cell)**.** *Let $C = \mathrm{RS}[F, D, k]$ with $|D| = n$ and $k < n$. If $A = n - r$ and*

$$3r \leq n - k,$$

then the maximum over received lines of the number of finite support-wise MCA-bad slopes at closed agreement threshold A is exactly $r + 1 = n - A + 1$, and the same maximum for no-loss CA-bad slopes is also exactly $r + 1$. Hence Definition 13.7 is an exact paid cell throughout the high-agreement range.

Proof. The upper bound is Theorem 9.6 applied at radius r/n: the Reed–Solomon minimum Hamming weight is $w = n - k + 1$, so the hypothesis $3r \leq w - 1$ there is exactly $3r \leq n - k$, and the theorem bounds the number of CA- or MCA-bad finite slopes of every received line by $r + 1$.

For the matching lower witness, choose a set

$$T = \{t_0, \dots, t_r\} \subseteq D$$

of size $r+1$ (possible since $r+1 \le n$) and choose distinct field elements $\gamma_0, \dots, \gamma_r \in F$ (possible since $|F| > n \ge r+1$ for $D \subseteq F^\times$, and $|F| \ge n \ge r+1$ in general). Define two received words by

$$f_2(t_i) = 1, \qquad f_1(t_i) = -\gamma_i, \qquad f_1(x) = f_2(x) = 0 \quad (x \in D \setminus T).$$

For the slope γ_i, the word $f_1 + \gamma_i f_2$ vanishes at t_i and off T, hence agrees with the zero codeword on

$$S_i := D \setminus (T \setminus \{t_i\}), \qquad |S_i| = n - r = A.$$

We first record the two degree inequalities used below. If $r = 0$, then $n - r - 1 = n - 1 \ge k$ and $n - 2r - 1 = n - 1 \ge k$. If $r \ge 1$, then $r + 1 \le 3r \le n - k$ gives $n - r - 1 \ge k$, and $2r \le 3r - r \le n - k - r \le n - k - 1$ gives $n - 2r - 1 \ge k$.

Support-wise badness. Suppose $P_1, P_2 \in F[X]_{<k}$ satisfy $P_1 = f_1$ and $P_2 = f_2$ on S_i. On $D \setminus T \subseteq S_i$ both received words vanish, and $|D \setminus T| = n - r - 1 \ge k$, so each P_j has at least k roots and is the zero polynomial. But $t_i \in S_i$ and $f_2(t_i) = 1 \ne 0$, a contradiction. Hence the pair (f_1, f_2) is not simultaneously explained on S_i, while the point $f_1 + \gamma_i f_2$ is explained there by the zero codeword: each γ_i is support-wise MCA-bad at agreement threshold A, and the $r+1$ slopes γ_i are distinct.

No-loss badness. Let $S \subseteq D$ be any support with $|S| \ge A = n - r$ and suppose $P_1, P_2 \in F[X]_{<k}$ explain (f_1, f_2) on S. Since $|T| = r+1$, inclusion–exclusion gives $|S \cap T| \ge |S| + |T| - n \ge 1$, and

$$|S \cap (D \setminus T)| \ge |S| - |T| \ge n - 2r - 1 \ge k.$$

Hence both P_j vanish at $\ge k$ points and are identically zero, contradicting $f_2(t) = 1$ at any point $t \in S \cap T$. So no support of size $\ge A$ simultaneously explains the pair, i.e. the pair is no-loss far at radius r/n; since each $f_1 + \gamma_i f_2$ is explained on S_i, the same $r+1$ slopes are no-loss CA-bad. Combining the witness with the upper bound of Theorem 9.6 proves exactness. □

Remark 13.9 (tangent-cell calibration). The exactness makes the high-agreement staircase pointwise computable: in the tangent range the exact numerator is $N(A) = n - A + 1$, which is strictly decreasing, so Theorem 13.2 localizes with a one-step transition. As a worked integer example, take the calibration row $n = 512$, $k = 256$ with sampling denominator $Q = 17^{32}$ and target 2^{-128}. Exact arithmetic gives $B_* = \lfloor 17^{32}/2^{128} \rfloor = 6$ (17^{32} has 131 bits), and the tangent range is $r \le R_{\tan} = 85$, i.e. $A \ge 427$. Inside this range, $N(A) = n - A + 1 \le 6$ first holds at $a_* = 507$: agreement 506 is unsafe with $N = 7 > 6$ and agreement 507 is safe with $N = 6 \le 6$. The largest safe closed integer radius for the tangent-range staircase is $r_{\text{safe}} = 5$, and the real closed-ball supremum of Corollary 13.3 is $6/512$, not attained. This is a calibration example for the compiler, not a deployed row.

Definition 13.10 (quotient paid status cell)**.** Let $\mathcal{C}$ be a finite set of declared quotient fiber sizes $c \mid n$ and let A be an agreement threshold. The conservative safe-sum support count is

$$U_{\text{sum}}(n, A, \mathcal{C}) := \sum_{c \in \mathcal{C}} \sum_{B=A}^{n} \binom{n/c}{\lfloor B/c \rfloor} \binom{n - c\lfloor B/c \rfloor}{B - c\lfloor B/c \rfloor},$$

the quotient-remainder count of Section 9: for each c and each size B, choose $\lfloor B/c \rfloor$ complete c-fibers among the n/c fibers and then a residual set of size $B - c\lfloor B/c \rfloor$ among the remaining $n - c\lfloor B/c \rfloor$ points. The quotient paid cell at $(n, A, \mathcal{C})$ is status-valued:

$$\mathrm{Paid}_{\mathrm{quot}} = \begin{cases} \mathrm{EXACT\_IMAGE}, & \text{an image-lcm certificate as in Theorem 3.36,} \\ \mathrm{EXACT\_SUPPORT}, & \text{an exact block-profile union as in Theorem 3.24,} \\ \mathrm{SAFE\_SUM}(U_{\mathrm{sum}}), & \text{otherwise.} \end{cases}$$

Unsafe quotient lower cells should carry zone tags, for example norm-exact, collision-not settled here, or cap-floor. A collision-open zone is interval-valued and must not be silently replaced by a point estimate.

Proposition 13.11 (safe-sum validity)**.** *Let $C = \mathrm{RS}[F, D, k]$ and $A \geq k$. If a quotient packet declares that every finite bad slope in the cell is witnessed on at least one support counted by the quotient-remainder model above, then $U_{\mathrm{sum}}(n, A, \mathcal{C})$ is a valid conservative upper numerator for the union of the declared quotient finite-slope cells at agreement at least A, for support-wise MCA-bad and for no-loss CA-bad slopes alike.*

Proof. The displayed binomial product counts, for each declared c and size B, the supports consisting of $\lfloor B/c \rfloor$ complete c-fibers plus a residual set outside those fibers; summing over $B \geq A$ and $c \in \mathcal{C}$ is a union bound over the declared support family, which may double-count supports with several quotient descriptions but cannot undercount them. By Lemma 3.16, each fixed support pays at most one finite bad slope, in either badness sense. The declared coverage hypothesis then converts the support union bound into a finite-slope numerator. This is the safe-sum instantiation of Proposition 3.19; when the exact union or the image-lcm coalescing is available, the sharper grades of Definition 13.10 replace it. □

Proposition 13.12 (extension-pole conversion cell)**.** *Let $B \subseteq F$ be the generated field inside the line field, with $q_{\mathrm{gen}} = |B|$ and $q_{\mathrm{line}} = |F|$. If $q_{\mathrm{gen}} = q_{\mathrm{line}}$, then there is no proper extension-only line-parameter cell:*

$$\mathrm{Paid}_{\mathrm{ext}}^{\mathrm{only}} = 0.$$

Assume $q_{\mathrm{gen}} < q_{\mathrm{line}}$, put $m = q_{\mathrm{line}} - q_{\mathrm{gen}}$, and let $\mathcal{P}$ be a base list of size $L \geq 1$. Suppose a pole-conversion certificate assigns to every $P \in \mathcal{P}$ and every pole $\alpha \in F \setminus B$ a finite extension parameter $\lambda_\alpha(P)$ such that:

(i) *for each fixed pole α, distinct values among $\{\lambda_\alpha(P) : P \in \mathcal{P}\}$ give distinct extension-only bad finite line parameters for the lifted row;*

(ii) *for every $P \neq P'$ in $\mathcal{P}$,*

$$\#\{\alpha \in F \setminus B : \lambda_\alpha(P) = \lambda_\alpha(P')\} \leq \kappa.$$

Then some pole produces at least

$$\mathrm{ExtPole}(q_{\mathrm{line}}, q_{\mathrm{gen}}, \kappa, L) := \left\lceil \frac{L(q_{\mathrm{line}} - q_{\mathrm{gen}})}{q_{\mathrm{line}} - q_{\mathrm{gen}} + \kappa(L-1)} \right\rceil$$

distinct extension-only bad finite line parameters. When the conversion is evaluation at the pole, one may take $\kappa = k$ for a degree-$\leq k$ auxiliary list and $\kappa = k - 1$ for an ordinary degree-$< k$ Reed–Solomon list. For safe-side extension charts over B, a closed parameter set of degree Δ and dimension e contributes at most $\Delta q_{\mathrm{gen}}^{e}$ affine B-parameters, by Theorem 3.68.

Proof. If $B = F$, then $F \setminus B = \varnothing$, so no line parameter is extension-only. Now assume $m > 0$. For a pole α, let M_α be the number of distinct values among $\{\lambda_\alpha(P) : P \in \mathcal{P}\}$, let $n_{\alpha,\lambda}$ be the multiplicity of the value λ, and let $C_\alpha := \sum_\lambda \binom{n_{\alpha,\lambda}}{2}$ be the number of colliding unordered pairs at α. Summing the pairwise certificate (ii) over unordered pairs of list elements gives

$$\sum_{\alpha \in F \setminus B} C_\alpha \le \kappa \binom{L}{2},$$

so some pole α has $C_\alpha \le \kappa\binom{L}{2}/m$. At this pole, Cauchy–Schwarz gives

$$L^2 = \Big(\sum_\lambda n_{\alpha,\lambda}\Big)^2 \le M_\alpha \sum_\lambda n_{\alpha,\lambda}^2 = M_\alpha(L + 2C_\alpha) \le M_\alpha\Big(L + \frac{\kappa L(L-1)}{m}\Big),$$

hence

$$M_\alpha \ge \frac{Lm}{m + \kappa(L-1)}.$$

Since M_α is an integer, the displayed ceiling is certified, and condition (i) converts the distinct pole values into distinct extension-only bad line parameters. For the evaluation conversion, two distinct polynomials of degree $\le k$ (resp. $< k$) agree at most k (resp. $k-1$) times on $F \setminus B$, giving the stated κ. The safe-side clause is Theorem 3.68 verbatim. □

Remark 13.13 (relation to the main extension-pole floors). Proposition 2.11 and Corollary 2.12 realize the conversion certificate concretely: the conversion is evaluation of the base list at a simple pole $\alpha \in F \setminus B$, condition (i) holds with genuinely extension-valued witness lines, and condition (ii) holds with $\kappa = k$ there. The proof of Corollary 2.12 internally derives the denominator $m + k(L-1)$ before weakening it to $m + kL$ for display; Proposition 13.12 promotes the sharper denominator to the official compiler value. The coarser denominator $q_{\text{line}} - q_{\text{gen}} + \kappa L$ remains a valid conservative fallback and may be printed when only the weaker collision certificate is available.

13.3 Budget windows and bounded quotient census

Theorem 13.14 (budget windows and dodge selection)**.** *Let N be the actual integer numerator for a fixed object, and suppose a certificate gives $L \le N \le K$. Put $B(q) = \lfloor q/2^{128} \rfloor$. Then*

$$B(q) < L \Rightarrow \textit{certified unsafe}, \qquad B(q) \ge K \Rightarrow \textit{certified safe for this object},$$

and the only unresolved budgets are

$$L \le B(q) < K.$$

Equivalently, the unresolved field-size interval is

$$L2^{128} \le q \le K2^{128} - 1.$$

If the lower proof is $L = K - g$, then the largest tolerable missing mass while still proving unsafety is

$$g_{\max} = K - B(q) - 1.$$

Proof. The safety condition for the fixed object is $N \le B(q)$. If $B(q) < L \le N$, then $N > B(q)$ and the object certifies unsafety. If $N \le K \le B(q)$, then the object is safe under this certificate. The complement is precisely $L \le B(q) < K$. The floor definition turns this double inequality into the displayed field-size interval. Finally, if $L = K - g$, unsafety requires $K - g > B(q)$, equivalently $g < K - B(q)$, and the largest such integer is $K - B(q) - 1$. □

For a 2-power quotient order N', let $n_1 = N'/2$. The characteristic-zero antipodal quotient count used by the quotient census of [Cho26b] is

$$A_2(N', \ell') := \sum_{\substack{u \ge 0,\ t = \ell' - 2u \ge 0 \\ u \le n_1 - t}} \binom{n_1}{t} 2^t.$$

Lemma 13.15 (rate-1/2 census identity)**.** *For every even $N' \ge 2$, with $n_1 = N'/2$,*

$$A_2(N', n_1 + 1) = \frac{3^{n_1} - 1}{2}.$$

Proof. At $\ell' = n_1 + 1$ the summation constraints reduce to a single parity class. Indeed $t = n_1 + 1 - 2u$, so $t \equiv n_1 + 1 \pmod 2$; the constraint $u \le n_1 - t$ becomes $u \le n_1 - (n_1 + 1 - 2u) = 2u - 1$, i.e. $u \ge 1$, which is equivalent to $t \le n_1 - 1$; and $t \ge 0$ closes the range. Hence

$$A_2(N', n_1 + 1) = \sum_{\substack{0 \le t \le n_1 - 1 \\ t \equiv n_1 + 1 \pmod 2}} \binom{n_1}{t} 2^t = \sum_{\substack{0 \le t \le n_1 \\ t \not\equiv n_1 \pmod 2}} \binom{n_1}{t} 2^t,$$

where the last step notes that $t = n_1$ has the excluded parity, and $t = n_1 + 1$ is out of range. From $\sum_t \binom{n_1}{t} 2^t = 3^{n_1}$ and $\sum_t \binom{n_1}{t} (-2)^t = (-1)^{n_1}$, the odd-$t$ class sums to $(3^{n_1} - (-1)^{n_1})/2$ and the even-t class to $(3^{n_1} + (-1)^{n_1})/2$. If n_1 is even the class $t \not\equiv n_1$ is the odd class and $(-1)^{n_1} = 1$; if n_1 is odd it is the even class and $(-1)^{n_1} = -1$. In both cases the sum is $(3^{n_1} - 1)/2$. □

Theorem 13.16 (bounded quotient census, exact arithmetic certificate)**.** *Let $B_{\max} = 2^{128} - 1$. In the relaxed even-scale model, where N' ranges over the even integers with $\rho N' \in \mathbb{Z}$, the first admissible scales at which $A_2(N', \rho N' + 1) > B_{\max}$ are*

ρ	1/2	1/4	1/8	1/16
N'_{relaxed}	164	176	248	384.

For dyadic quotient orders $N' = 2^a$, the corresponding first scales are

ρ	1/2	1/4	1/8	1/16
N'_{dyadic}	256	256	256	512.

Thus the exact integer quotient endgame at the 2^{-128} budget sees only bounded quotient scales, independent of the deployed blocklength, once the row has been reduced to this canonical quotient-count model.

Proof. The count $A_2(N', \ell')$ is an explicit finite integer sum, so the statement is an exact arithmetic certificate. For each fixed ρ, the certificate evaluates $A_2(N', \rho N' + 1)$ by exact integer arithmetic at *every* admissible scale up to and including the printed first crossing, confirms $A_2 \le B_{\max}$ at every admissible scale strictly below the crossing, and confirms $A_2 > B_{\max}$

at the crossing; no floating-point estimate and no unproved monotonicity is used. Writing $\mathrm{bits}(X) = \lfloor \log_2 X \rfloor + 1$, so that $\mathrm{bits}(X) \le 128 \iff X \le B_{\max}$, the immediate predecessor and crossing values are

ρ	N'_{prev}	bits A_2	N'_{first}	bits A_2
1/2	162	128	164	129
1/4	172	127	176	130
1/8	240	127	248	131
1/16	368	124	384	129

ρ	N'_{prev}	bits A_2	N'_{first}	bits A_2
1/2	128	101	256	202
1/4	128	95	256	190
1/8	128	68	256	135
1/16	256	87	512	172

for the relaxed even model (left) and the dyadic model (right). At rate 1/2 the scanned values agree with the closed form of Lemma 13.15, which is strictly increasing in N', so the rate-1/2 rows also follow from the identity alone. □

Definition 13.17 (dyadic planted quotient profile)**.** For $n = 2^m$, $k = \rho n$, and integer slack $\sigma = \lceil \eta n \rceil$, define

$$K_Q(n,k,\sigma) := \max_{\substack{M=2^e \mid \gcd(n,k) \\ \sigma < M \\ k/M \le n/M - 1}} \binom{n/M - 1}{k/M},$$

with value 0 if the active set is empty.

Proposition 13.18 (bounded active quotient order)**.** *If $\sigma \ge n/256$, then every active scale in K_Q has quotient order $N = n/M < 256$; because N is dyadic in Definition 13.17, this means $N \le 128$. If $\sigma \ge n/128$, then every active scale has $N < 128$, hence dyadically $N \le 64$.*

Proof. Activity requires $\sigma < M = n/N$, hence $N < n/\sigma$. The two displayed strict bounds follow from the two lower bounds on σ, and the dyadic constraint improves the strict inequalities to $N \le 128$ and $N \le 64$ respectively. □

Proposition 13.19 (polynomial residual absorption)**.** *Fix a candidate scale and suppose a future theorem bounds an unpaid residual contribution by*

$$n^C \Phi(n,q)$$

for some explicit proxy Φ. If the verified entropy calculation gives

$$-\log_2\bigl(n^C \Phi(n,q)\bigr) > 0,$$

then the residual integer contribution is zero.

Proof. The residual contribution is a nonnegative integer bounded above by a real number strictly smaller than 1. The only such integer is 0. □

14 Identity-scale frontier and quotient base cases

The quotient floors of the main text work at every divisor scale. At the identity scale $c = 1$ no quotient remains, and the pigeonhole is simply over the leading coefficients of ordinary split locators. This terminal scale gives the base unsafe certificates; the MCA rows are then strengthened further by the flexible-budget conversion in Proposition 15.2.

14.1 The identity prefix floor

Lemma 14.1 (identity-scale prefix floor)**.** *Let $\mathbb{B} \subseteq \mathbb{F}$, let $D \subseteq \mathbb{B}$ have $|D| = n$, and let $1 \le K \le n$. For every integer m with $K \le m \le n$, put $w = m - K$. Then there is a $\mathbb{B}$-valued received word $U : D \to \mathbb{B}$ such that*

$$\left|\Lambda(\mathrm{RS}[\mathbb{F}, D, K], 1 - m/n, U)\right| \ge \left\lceil \frac{\binom{n}{m}}{|\mathbb{B}|^{\,w}} \right\rceil .$$

Consequently the graded-prefix route of Proposition 10.8 extends to the identity scale $c = 1$.

Proof. For each m-subset $M \subseteq D$, write

$$\Lambda_M(X) = \prod_{b \in M} (X - b) = X^m + \sum_{j=1}^{m} (-1)^j e_j(M) X^{m-j}.$$

For $z = (z_1, \dots, z_w) \in \mathbb{B}^w$, define

$$U_z := \Big(X^m + \sum_{j=1}^{w} z_j X^{m-j}\Big)\Big|_D .$$

The map $M \mapsto ((-1)^j e_j(M))_{1 \le j \le w}$ has codomain $\mathbb{B}^w$, so one fiber contains at least $\binom{n}{m}/|\mathbb{B}|^w$ subsets. Fix such a prefix value z^*, and for each M in the fiber put

$$c_M := U_{z^*} - \Lambda_M|_D = -\sum_{j=w+1}^{m} (-1)^j e_j(M) X^{m-j}\Big|_D .$$

Since $m - w - 1 = K - 1$, the word c_M is a codeword of $\mathrm{RS}[\mathbb{F}, D, K]$. On the m points of M, the locator vanishes, so c_M agrees with U_{z^*} on at least m positions.

It remains only to check injectivity. If $c_M = c_{M'}$, then the difference of the two displayed polynomials has degree at most $K - 1 < n$ and vanishes on all points of D, hence is the zero polynomial. Therefore the coefficients $e_j(M)$ and $e_j(M')$ agree for all $j > w$; the prefix-fiber condition gives agreement for $j \le w$. The two monic locator polynomials are equal, and hence $M = M'$. □

14.2 Identity-prefix certificates

Throughout this subsection $n = 2^{21}$, $k = 2^{20}$. For the KoalaBear row, $p = 2^{31} - 2^{24} + 1$ and $q = p^6$; for the Mersenne-31 line-round circle row, $p' = 2^{31} - 1$ and $q = (p')^4$. The MCA rows use the floor at $K = k + 1$ and then the deep-point conversion to the dimension-k code, exactly as in Corollary 10.9 and Remark 10.10; the list rows use $K = k$.

Proposition 14.2 (identity-scale list frontier, exact)**.** *The exact identity-scale list certificates are:*

row / challenge	c	m	w	*certified unsafe range*	*approx. pass / next deficit*	$\log_2$
KoalaBear, list	1	1116046	67470	$\delta \in [\frac{490553}{1048576}, 1-\rho) = [0.4678278\ldots, 1-\rho)$	+9.164 / −22.011	
Mersenne-31 line-round, list	1	1116022	67446	$\delta \in [\frac{490565}{1048576}, 1-\rho) = [0.4678392\ldots, 1-\rho)$	+28.113 / −3.073	

Each printed list value of m passes by exact integer comparison, and $m+1$ fails by exact integer comparison. The MCA rows use the same identity-prefix floor, but their final finite frontier is the stronger flexible-budget certificate of Proposition 15.2.

Proof. The construction is Lemma 14.1 with $K = k$. The displayed list floors exceed the corresponding list budgets by the exact integer inequalities

$$\binom{n}{1116046} 2^{128} > p^{67470}p^6, \qquad \binom{n}{1116022} 2^{100} > (p')^{67446}(p')^4.$$

In both rows the inequality with m replaced by $m+1$ and w by $w+1$ is false. The edge fractions are $1 - m/n$. □

Corollary 14.3 (deployed sandwich edges)**.** *At the active deployed targets,*

$$\begin{aligned} \delta^*_{C,\mathrm{KB}}(2^{-128}) &\le 981105/2^{21}, & \delta^*_{C,\mathrm{M31}}(2^{-100}) &\le 981129/2^{21}, \\ \delta^*_{\mathrm{list},\mathrm{KB}}(2^{-128}) &\le 490553/2^{20}, & \delta^*_{\mathrm{list},\mathrm{M31}}(2^{-100}) &\le 490565/2^{20}. \end{aligned}$$

The first line is supplied by the flexible-budget MCA certificates and the second by the identity-prefix list certificates. These are the upper endpoints used in Theorem 10.12 and Table 5.

Proof. The two list inequalities are Proposition 14.2; the two MCA inequalities are Proposition 15.2. Monotonicity converts each exact unsafe agreement into the displayed threshold upper bound. □

14.3 The frontier conjecture

Definition 14.4 (frontier function, unchanged)**.** For $\rho \in (0,1)$ and $\beta > 0$, put

$$g^*(\rho,\beta) = \sup\{g \in (0, 1-\rho] : \mathrm{H}_2(\rho+g) \ge \beta g\}.$$

For KoalaBear, $\beta = \log_2(2^{31} - 2^{24} + 1) = 30.988685\ldots$ and

$$1 - \rho - g^*(\rho,\beta) = 0.4678266,\ 0.7225056,\ 0.8557947,\ 0.9251114$$

for $\rho = 1/2, 1/4, 1/8, 1/16$. For Mersenne-31, the corresponding values are

$$0.4678383,\ 0.7225161,\ 0.8558023,\ 0.9251165.$$

At rate $1/2$, the identity-scale deployed gaps are within $2.7 \cdot 10^{-6}$ of the KoalaBear MCA ceiling, $1.2 \cdot 10^{-6}$ of the KoalaBear list ceiling, $1.5 \cdot 10^{-6}$ of the Mersenne-31 MCA ceiling, and $1.0 \cdot 10^{-6}$ of the Mersenne-31 list ceiling. Larger-block extrapolations from floating-point scans should be recorded as numerical evidence unless separately certified by exact integer comparisons.

Lemma 14.5 (transversality of the entropy–subfield crossing)**.** *For fixed $\rho \in (0,1)$ and $\beta > 0$, put*

$$F_{\rho,\beta}(g) := \mathrm{H}_2(\rho+g) - \beta g.$$

Then $\{g \in [0, 1-\rho] : F_{\rho,\beta}(g) \ge 0\} = [0, g^*]$, the endpoint $g^* = g^*(\rho,\beta)$ lies in $(0, 1-\rho)$, and

$$c_{\rho,\beta} := -F'_{\rho,\beta}(g^*) = \beta - \log_2 \frac{1-\rho-g^*}{\rho+g^*} > 0.$$

Consequently, for every $R = o(n)$,

$$nF_{\rho,\beta}\left(g^* + \frac{R}{n}\right) = -c_{\rho,\beta}R + O\left(\frac{R^2}{n}\right).$$

For the deployed rates $\rho \in \{1/2, 1/4, 1/8, 1/16\}$ and every $\beta \geq 30$, $c_{\rho,\beta} \geq \beta - \log_2 15 > 0.86\beta$. Thus a multiplicative loss $L(n)$ with $\log L(n) = o(R(n)\beta)$ is negligible relative to the reserve separation.

Proof. The function $F_{\rho,\beta}$ is strictly concave, with $F_{\rho,\beta}(0) = \mathrm{H}_2(\rho) > 0$ and $F_{\rho,\beta}(1-\rho) = -\beta(1-\rho) < 0$. Its superlevel set is therefore the interval $[0, g^*]$ with a unique interior boundary point. Concavity gives $F'_{\rho,\beta}(g^*) \leq (F_{\rho,\beta}(g^*) - F_{\rho,\beta}(0))/g^* < 0$, proving $c_{\rho,\beta} > 0$. Taylor's theorem gives the displayed expansion. Finally, $\mathrm{H}'_2(x) = \log_2((1-x)/x)$ is decreasing and $\rho + g^* \geq \rho$, so $\mathrm{H}'_2(\rho + g^*) \leq \mathrm{H}'_2(\rho) \leq \log_2 15$ on the four deployed rates. □

Conjecture 14.6 (frontier conjecture, asymptotic and finite forms)**.** Prove or refute the asymptotic safe-side match

$$\delta_C^*(\varepsilon^*) = 1 - \rho - g^*(\rho, \log_2 |\mathbb{B}|) + o(1)$$

for smooth multiplicative and circle line-round rows at the deployed challenge targets. The finite deployed form predicts that the first safe agreement is one more than the best exact unsafe agreement supplied by Proposition 15.2 and the list entries of Proposition 14.2:

row	proved unsafe agreement a_0	first safe $a_0 + 1$	deficit floor (bits)
KoalaBear MCA	1116047	1116048	> 22.1
KoalaBear list	1116046	1116047	> 22.0
Mersenne-31 MCA	1116023	1116024	> 3.2
Mersenne-31 list	1116022	1116023	> 3.0

with the real closed-ball endpoint convention of Corollary 13.3.

Remark 14.7 (closure mechanisms after $c = 1$). The two "free" mechanisms in the terminal-scale scan should be kept, but with the endpoint changed. Subgroup-invariant subset selection at a nontrivial scale is absorbed into the composite scale exactly as before; iterating this can only move toward the identity prefix floor, where no nontrivial quotient remains. The $c = 2$, $\sigma = 1$ planted hybrid is also dominated by the identity row at the deployed parameters: its best net agreement is at most the $c = 2$ agreement minus one, whereas Proposition 14.2 improves the $c = 2$ agreement by $5, 2, 3, 2$ steps in the four rows. Thus planting remains useful as a structural lower construction, but it is not the deployed frontier certificate once identity prefixes are allowed.

14.4 Proved base cases of the identity-scale collision problem

At the identity scale the census behind Lemma 14.1 runs over the deployed domain itself, and for the KoalaBear rows that domain is a coset of the order-2^{21} 2-adic subgroup of $\mathbb{F}_p^\times$ with $2^{21} > \lceil\sqrt{p}\rceil = 46160$. The prefix-fiber counts are therefore in Weil range, and their base cases are provable unconditionally. Throughout, $e_p(x) := \exp(2\pi i x/p)$ and $B := \lceil\sqrt{p}\rceil$.

Theorem 14.8 (bounded-prefix equidistribution over large cosets, with planted cores)**.** *Let p be prime, $H \le \mathbb{F}_p^\times$ of order N, $d = (p-1)/N$, and $S = \alpha H$ a coset. Let $P \subseteq S$ be a fixed set of size $\ell \ge 0$, let $1 \le w < N - \ell$ satisfy $wd < p$, and let $\ell \le m \le N$, $m' := m - \ell$. For $\vec{z} \in \mathbb{F}_p^w$ put*

$$\mathcal{N}_w^P(m, \vec{z}) \;:=\; \#\big\{A :\; P \subseteq A \subseteq S,\; |A| = m,\; \big(e_1(A), \dots, e_w(A)\big) = \vec{z}\,\big\}.$$

Then for every $\vec{z}$,

$$\Big|\,\mathcal{N}_w^P(m, \vec{z}) - \frac{1}{p^w}\binom{N-\ell}{m'}\Big| \;\le\; p^{w/2}\binom{wB + \ell + m' - 1}{m'},$$

and for $w = 1$ the factor $p^{1/2}$ may be dropped.

Proof. Step 1 (removing the core). Write $A = P \cup A'$ with $A' \subseteq S' := S \setminus P$, $|A'| = m'$. From $\Lambda_A = \Lambda_P \Lambda_{A'}$,

$$e_j(A) = e_j(A') + \sum_{i=1}^{\min(j,\ell)} e_i(P)\, e_{j-i}(A') \qquad (1 \le j \le w),$$

so the map $(e_1, \dots, e_w)(A') \mapsto (e_1, \dots, e_w)(A)$ is affine and unipotent triangular over $\mathbb{F}_p^w$, hence a bijection. Each fiber of the planted census is therefore a fiber of the plain census on S' at a transformed value, and it suffices to prove the displayed bound for $\mathcal{N}_w(m', \cdot)$ on S'.

Step 2 (Fourier reduction). By inversion on $\mathbb{F}_p^w$, $\mathcal{N}_w(m', \vec{z}) = p^{-w} \sum_{\vec{t}} e_p(-\vec{t} \cdot \vec{z}) E(\vec{t})$ with $E(\vec{t}) = \sum_{|A'|=m'} e_p(\sum_j t_j e_j(A'))$; the term $\vec{t} = 0$ gives the main term, so it suffices to bound $|E(\vec{t})|$ for $\vec{t} \ne 0$.

Step 3 (Newton conversion). Fix $\vec{t} \ne 0$ and $j^* = \max\{j : t_j \ne 0\}$. Since $j \le w < p$, the Newton identities express each e_j as a polynomial in the power sums with $e_j = (-1)^{j-1} p_j / j + R_j(p_1, \dots, p_{j-1})$, so $\sum_j t_j e_j(A') = Q(P(A'))$ for $P(A') = (p_1, \dots, p_{j^*})(A')$ and a polynomial Q that is linear in its last variable with constant coefficient $(-1)^{j^*-1} t_{j^*}/j^* \ne 0$. Fourier inversion on $\mathbb{F}_p^{j^*}$ gives

$$E(\vec{t}) = p^{-j^*} \sum_{\vec{s}} \widehat{Q}(\vec{s})\, G(\vec{s}), \qquad \widehat{Q}(\vec{s}) = \sum_{\vec{v}} e_p(Q(\vec{v}) - \vec{s} \cdot \vec{v}), \quad G(\vec{s}) = \sum_{|A'|=m'} e_p(\vec{s} \cdot P(A')),$$

and $\widehat{Q}(0) = 0$ because the inner sum over the last coordinate of $\vec{v}$ is a complete nontrivial character sum.

Step 4 (power sums on the punctured coset). For $\vec{s} \ne 0$, $G(\vec{s}) = e_{m'}\big((e_p(g_s(a)))_{a \in S'}\big)$ with $g_s(x) = \sum_{i \le j^*} s_i x^i$ of degree $i^* \le w$; its power sums are $\pi_j = \sum_{a \in S'} e_p(j g_s(a))$, $1 \le j \le m'$. Since $\deg g_s \le w < N - \ell = |S'|$ and g_s is nonconstant, and $j \not\equiv 0 \pmod p$, the coset sum obeys Weil's bound through the power parametrization $S = \{\alpha y^d : y \in \mathbb{F}_p^\times\}$:

$$\Big|\sum_{a \in S} e_p(j g_s(a))\Big| = \frac{1}{d}\Big|\sum_{y \in \mathbb{F}_p} e_p\big(j g_s(\alpha y^d)\big) - e_p(j g_s(0^*))\Big| \le \frac{1}{d}\big[(i^* d - 1)\sqrt{p} + 1\big] \le w\sqrt{p} \le wB,$$

where 0^* abbreviates the $y = 0$ evaluation point and $\deg_y = i^* d \le wd < p$ ensures applicability. Removing the ℓ points of P costs at most ℓ, so $|\pi_j| \le wB + \ell$ for all j.

Step 5 (Newton domination). The recursion $m' e_{m'} = \sum_{i=1}^{m'} (-1)^{i-1} e_{m'-i} \pi_i$ is majorized coefficientwise by the recursion of $[u^r](1-u)^{-(wB+\ell)}$, whence $|G(\vec{s})| \le \binom{wB+\ell+m'-1}{m'}$ for every $\vec{s} \ne 0$.

Step 6 (Parseval). $\sum_{\vec{s}} |\widehat{Q}(\vec{s})| \le p^{j^*/2}\big(\sum_{\vec{s}} |\widehat{Q}(\vec{s})|^2\big)^{1/2} = p^{3j^*/2}$, so $|E(\vec{t})| \le p^{j^*/2}\binom{wB+\ell+m'-1}{m'}$, and summing the $p^w - 1$ nonzero characters against p^{-w} gives the theorem. For $w = 1$ no linearization is needed: $E(t) = e_{m'}\big((e_p(ta))_{a\in S'}\big)$ directly, with $|\pi_j| \le \sqrt{p} + \ell$. □

The standard entropy sandwich gives a quick coarse orientation: for $0 < m < M$, taking $\lambda = m/M$ in $1 = (\lambda + (1-\lambda))^M \ge \binom{M}{m}\lambda^m(1-\lambda)^{M-m}$ gives $\binom{M}{m} \le 2^{M\mathrm{H}_2(m/M)}$, while the largest of the $M+1$ binomial probabilities gives $\binom{M}{m} \ge 2^{M\mathrm{H}_2(m/M)}/(M+1)$. The sharper displayed integer bounds are obtained from the exact binomial ratios in the theorem, with outward-rounded interval evaluation of their logarithms; equivalently, each may be certified by an exact integer power-of-two comparison. The accompanying verifier is a cross-check, not the proof.

Corollary 14.9 (the head of the identity staircase is exact)**.** *For the KoalaBear rows of Proposition 14.2 ($N = 2^{21}$, $d = 1016$, $B = 46160$, $\ell = 0$):*

(a) *at the deployed subset sizes $m = 1116047$ (MCA) and $m = 1116046$ (list), Theorem 14.8 is nonvacuous for exactly $w \le 21$, with margins*

$$> 1810769, \quad > 1620294, \quad > 1462136, \quad \ldots, \quad > 63883, \quad > 11951$$

bits at $w = 1, 2, 3, \ldots, 20, 21$ for the MCA row (the list row agrees within 4 bits)*, and failure at $w = 22$ by at least* 38283 *bits;*

(b) *at the head rungs $m = K + w$ it is nonvacuous for exactly $w \le 22$* (margin at least +32393 bits at $w = 22$, failure by at least 14136 bits at $w = 23$)*;*

(c) *consequently every prefix fiber in these ranges equals $\binom{N}{m}p^{-w}(1+\varepsilon)$ with $|\varepsilon| \le 2^{-\text{margin}}$: the pigeonhole step of Lemma 14.1 is* exact *at head depths, two-sidedly and with no polynomial factor, giving the first exactly priced strata of the identity staircase;*

(d) *the planted census is priced identically with wB replaced by $wB + \ell$, which costs about $\log_2(1 + m/B) \approx 4.65$ bits of margin per planted point: cores of size $\ell \le 2^{11}$ preserve every rung $w \le 21$, and the $w = 1$ rung tolerates ℓ up to $\approx 2^{18}$, so the planted lower constructions of Remark 14.7 are two-sided at head depths.*

Not claimed: any depth beyond the printed w_0 at these subset sizes; and the U-twisted per-word census at general received words (the fibers priced here are the ones entering the floor and the witness words it constructs; Proposition 14.13 makes that pricing exact)*. The circle rows are treated in Theorem 14.15 below.*

Proposition 14.10 (composite directions collapse to coarser scales)**.** *Let $e \mid N$, $e > 1$, and let $\vec{s} \ne 0$ be supported on indices divisible by e, so that $g_s(x) = h(x^e)$ with $\deg h \le w/e$. Then, with $S_e := \{a^e : a \in S\}$ a coset of the order-N/e subgroup,*

$$\sum_{\substack{A\subseteq S\\ |A|=m}} e_p\Big(\sum_{a\in A} g_s(a)\Big) \;=\; [x^m]\prod_{b\in S_e}\big(1 + x\,e_p(h(b))\big)^e.$$

Proof. The generating function $\sum_m x^m \sum_{|A|=m} e_p(\sum_{a\in A} g_s(a))$ equals $\prod_{a\in S}(1 + xe_p(g_s(a)))$. The value $g_s(a) = h(a^e)$ is constant on each fiber of $a \mapsto a^e$, and each $b \in S_e$ has exactly e preimages in S, so the product collapses as displayed. □

Two consequences. First, Theorem 14.8 never assumed primitivity: for composite $\vec{s}$ the reduced-scale power sums still obey $e \cdot |\pi_j^{\mathrm{red}}| \le e \cdot (w/e)\sqrt{p} = w\sqrt{p}$ by the same Weil argument at the coarser coset, so the head-depth pricing above covers all directions. Second, beyond the head, composite directions genuinely carry the character mass of their coarser scale (quantified in Remark 14.12), so any strengthening uniform in the direction must be graded along the divisor lattice — the same absorption structure as Remark 14.7.

Remark 14.11 (the corrected collision target). The proof of Theorem 14.8 pays one Weil cost $\sqrt{p}$ per power sum, and this is exactly what stops it at $w_0 = 21$–22 while the frontier rows of Proposition 14.2 sit at depth $w \approx 67466$. Computation rules out the tempting strengthening: a uniform character bound $|E(\vec{t})| \le \binom{N}{m} p^{-\kappa w}$ over all $\vec{t} \neq 0$ is false, in two distinct ways. Composite directions carry coarse-scale mass by Proposition 14.10, so no single-scale exponent can hold uniformly. And even over primitive directions, maxima are pinned near the Parseval level: $\sum_{\vec{t}} |E(\vec{t})|^2$ equals p^w times the collision count of the prefix map *exactly*, so square-root cancellation per slot is the typical size, while the passage from character maxima to fiber bounds sums p^w signed terms and would need $\kappa \ge 1$. Character maxima therefore face a square-root barrier at depth, and the correct remaining target is the max-fiber inequality itself, at every primitive scale in the dense regime $\binom{N}{m} p^{-w} \ge 1$:

$$\max_{\vec{z}} \mathcal{N}_w(m, \vec{z}) \;\le\; \mathrm{poly}(n) \cdot \binom{N}{m} p^{-w},$$

with composite scales priced through their reductions. Its consequences must be read against Remark 14.18: a polynomial-loss max-fiber input prices only the boundary prefix bucket. An asymptotic closure additionally requires the complete normalized residual certificate and the reserve conditions of Remark 21.5; none of this decides the printed adjacent pairs of Conjecture 14.6, whose next-row deficits exceed 22.1, 22.0, 3.2, and 3.0 bits, respectively. Deciding a printed pair requires the constant-factor form — max within $(1 + o(1))$ of the mean, with all constants displayed — which is precisely the regime the exact censuses of Remark 14.12 probe, and which Corollary 14.9 proves outright at head depths.

Remark 14.12 (computational calibration; experimental). The measurements behind Remark 14.11 use two exact instruments. (i) Because the Newton map $(e_1, \dots, e_w) \leftrightarrow (p_1, \dots, p_w)$ is a bijection on value vectors for $w < p$, the power-sum prefix map has identical fibers and elementwise-additive phases, so ratios $E(\vec{s})/\binom{N}{m}$ are computable at every depth by the normalized recursion $q_j^{(i)} = (1 - \frac{j}{i}) q_j^{(i-1)} + \frac{j}{i} \varepsilon_i q_{j-1}^{(i-1)}$: a convex combination of unimodularly weighted values, hence $|q| \le 1$ throughout and the update is non-expansive, so floating error cannot amplify; spot values were reproduced to all printed digits by exact arithmetic in $\mathbf{Z}[x]/(x^p - 1)$. (ii) Group-algebra dynamic programming gives exact joint fiber censuses at small scales. Findings, on subgroups with $N/\sqrt{p} \approx 3$–4 (the deployed identity scale has $N/\sqrt{p} \approx 45$, the favourable direction): primitive character maxima decay toward the Parseval level (measured exponents per slot $2.26, 1.13, 0.80, 0.67, 0.53, 0.46, 0.43, 0.39, 0.31$ at depths 2–10 for $p = 257$, $N = 64$, $m = 34$), while directions forced onto index multiples of $e = 2, 4, 8$ jump to 4.2×10^{-6}, 1.7×10^{-4}, 1.6×10^{-2} at depth 8 — the divisor hierarchy of Proposition 14.10 in plain sight. The exact censuses show the max-fiber target holding with room where the ledger operates: max-to-mean ratios $1.0012,\ 1.21,\ 2.67$ at depths 1–3 for $(p, N, m) = (17, 16, 8)$; $1.0022,\ 1.21,\ 4.10$ for $(41, 20, 10)$; $1 + 3 \times 10^{-12}$, $1 + 9 \times 10^{-9}$ at depths 1–2 for $(257, 64, 34)$, minima equally pinned. Ratios grow only at the Poisson boundary, where the mean fiber falls to $O(1)$ and the maximum of p^w near-independent cells is logarithmic (maximum 11 against mean

2.7 over 41^3 cells, matching the Poisson tail); the deployed crossings sit in the ultra-dense bulk at mean $\approx q/k$. These tables are calibration evidence only and enter no proof; the falsifiable prediction is that any violation of the max-fiber target must place a super-polynomial fiber inside the dense bulk of a primitive scale.

14.5 Circle rows, exact witness lists, and quotient scales

Three additions complete the base-case picture: the identity witness lists are *exactly* the prefix fibers, so the census statements above are list-size statements about the code; the circle rows acquire the same proved base cases through a Weil bound on the conic, with doubled constants; and the quotient census cells inherit a per-scale table.

Proposition 14.13 (witness lists are exactly the prefix fibers)**.** *In the setting of Lemma 14.1, for* every $z \in \mathbb{B}^w$*,*

$$\Lambda\big(\mathrm{RS}[\mathbb{F}, D, K],\, 1 - m/n,\, U_z\big) \;=\; \{\, c_M :\; M \text{ in the prefix fiber of } z \,\},$$

so the list size equals the fiber size exactly; moreover no codeword agrees with U_z on more than m points, so U_z is list-empty at every strictly smaller radius.

Proof. The inclusion $\supseteq$ and the distinctness of the c_M are Lemma 14.1. Conversely, let $c \in \mathrm{RS}[\mathbb{F}, D, K]$ agree with U_z on at least m points, and let $\hat{c} \in \mathbb{F}[X]$, $\deg \hat{c} < K$, be its polynomial. Then

$$P(X) := X^m + \sum_{j=1}^{w} z_j X^{m-j} - \hat{c}(X)$$

is monic of degree exactly m (as $\deg \hat{c} \le K - 1 = m - w - 1 < m - w$) and vanishes at every agreement point, so it has at least m roots in D; having degree m, it equals $\Lambda_{M'}$ for the set $M' \subseteq D$ of its m distinct roots. Comparing coefficients in degrees $m-1, \dots, m-w$ places M' in the prefix fiber of z, and comparing the remaining coefficients gives $\hat{c} = -\sum_{j>w} (-1)^j e_j(M') X^{m-j}$, i.e. $c = c_{M'}$. Since $P \neq 0$ has at most m roots, agreement above m is impossible. □

Combined with Theorem 14.8 and Corollary 14.9 (the sign twist $z_j \mapsto (-1)^j z_j$ is a coordinatewise bijection of the value space, so fiber counts transfer verbatim), every identity witness word at KoalaBear head depths $w \le w_0$ has list size $\binom{n}{m} p^{-w}(1+\varepsilon)$ — all p^w of them, not only the pigeonhole maximizer.

Lemma 14.14 (power-sum bound on twin-coset x-domains)**.** *Let $p \equiv 3 \pmod 4$, let $\mathbb{U}$ be the circle group of order $p+1$, let $\mathcal{D} = gH \cup g^{-1}H$ be a twin coset ($H \le \mathbb{U}$ of order M, $g^2 \notin H$), and let $D = \chi(\mathcal{D}) \subseteq \mathbb{F}_p$, $|D| = M$, $d' = (p+1)/M$. Then for every nonconstant $f \in \mathbb{F}_p[X]$ with $\deg f \cdot d' < p$,*

$$\Big|\sum_{x \in D} e_p(f(x))\Big| \;\le\; 2 \deg f \cdot \sqrt{p}.$$

Proof. Since $\mathcal{D}$ contains no self-inverse element, χ is exactly two-to-one on it. Therefore

$$\sum_{x \in D} e_p(f(x)) = \frac{1}{2} \sum_{u \in \mathcal{D}} e_p(f(\chi(u))).$$

It suffices to bound each coset sum by $2\deg f\sqrt{p}$. Fix the coset gH. As $\mathbb{U}$ is cyclic and $d' \mid p+1$, the map $v \mapsto gv^{d'}$ carries $\mathbb{U}$ onto gH with every fiber of size d', so

$$\sum_{u\in gH} e_p(f(\chi(u))) = \frac{1}{d'}\sum_{v\in\mathbb{U}} e_p\big(R(v)\big), \qquad R(v) := f\big(\chi(gv^{d'})\big).$$

Writing $x_g := \chi(g)$, $y_g := (g-g^{-1})/(2i) \in \mathbb{F}_p$ and using $u^a \pm u^{-a}$ Chebyshev identities, the angle-addition formula gives

$$\chi(gv^{d'}) = x_g\, T_{d'}(x_v) - y_g\, y_v\, W_{d'}(x_v),$$

where $T_{d'}, W_{d'} \in \mathbb{F}_p[X]$ have degrees d' and $d'-1$ (with $v^a - v^{-a} = (v-v^{-1})W_a(x_v)$). Hence R is an $\mathbb{F}_p$-polynomial function on the affine conic $x^2+y^2=1$, whose $\mathbb{F}_p$-points are exactly the $p+1$ elements of $\mathbb{U}$: the conic has no rational points at infinity, since $x^2+y^2=0$ has no nonzero solution over $\mathbb{F}_p$. Over $\overline{\mathbb{F}_p}$ the conic is a projective line with coordinate u, and near its two (conjugate, non-rational) points at infinity $u \in \{0,\infty\}$ the expansion $\chi(gv^{d'}) = \frac{g}{2}v^{d'}(1+O(v^{-2d'}))$ shows R has exactly two poles, each of order $\deg f \cdot d'$ with nonvanishing leading coefficient. R is nonconstant. The d'-power map on the cyclic group $\mathbb{U}$ has image H, not all of $\mathbb{U}$. Hence $v \mapsto gv^{d'}$ covers gH. The map χ is injective on gH: if $\chi(gh_1) = \chi(gh_2)$, then either $h_1 = h_2$ or $g^2h_1h_2 = 1$, and the second alternative would give $g^2 \in H$, contrary to the twin-coset hypothesis. Thus $\chi(gH)$ has $M = (p+1)/d'$ elements. Since $\deg f\, d' < p$, one has $\deg f < M$, so constancy of R would make the nonconstant polynomial $f-c$ vanish at more than $\deg f$ points, a contradiction. Its pole orders lie in $[1,p)$, so R is not of the form h^p-h+c, whose pole orders are multiples of p. The standard Weil–Bombieri additive-character bound on a complete nonsingular curve [Sch76, Ch. 6] now gives

$$\Big|\sum_{v\in\mathbb{U}} e_p(R(v))\Big| \le \big(2\cdot 0 - 2 + 2 + 2\deg f\cdot d'\big)\sqrt{p} = 2\deg f\cdot d'\sqrt{p},$$

and dividing by d' bounds the coset sum. The coset $g^{-1}H$ is identical. □

Theorem 14.15 (circle version, with planted cores)**.** *Let $D = \chi(\mathcal{D})$ be a twin-coset x-domain as in Lemma 14.14, $P \subseteq D$ with $|P| = \ell \ge 0$, $1 \le w < M-\ell$ with $wd' < p$, $\ell \le m \le M$, $m' = m-\ell < p$, and $B = \lceil\sqrt{p}\rceil$. Then for every $\vec{z} \in \mathbb{F}_p^w$,*

$$\Big|\,\mathcal{N}_w^P(m,\vec{z}) - \frac{1}{p^w}\binom{M-\ell}{m'}\Big| \;\le\; p^{w/2}\binom{2wB+\ell+m'-1}{m'},$$

and for $w=1$ the factor $p^{1/2}$ may be dropped.

Proof. Steps 1–3, 5, and 6 of the proof of Theorem 14.8 are domain-agnostic and apply verbatim. In Step 4, the power sums are $\pi_j = \sum_{x\in D\setminus P} e_p(jg_s(x))$ with jg_s nonconstant of degree $i^* \le w$ and $j \not\equiv 0 \pmod p$; Lemma 14.14 bounds the sum over D by $2i^*\sqrt{p} \le 2wB$, and removing the ℓ points of P costs at most ℓ. The domination step then runs with $2wB+\ell$ in place of $wB+\ell$. □

Corollary 14.16 (circle head rungs; all four rows priced)**.** *For the Mersenne-31 line-round rows of Proposition 14.2 ($M = 2^{21}$, $d' = 2^{10}$, $B = 46341$, $\ell = 0$):*

(a) *at the deployed sizes $m = 1116023$ (MCA) and $m = 1116022$ (list), Theorem 14.15 is nonvacuous for exactly $w \le 10$, with margin lower bounds > 1619021, > 1322221, …, > 169804, > 60240 bits at $w = 1,\dots,10$ and failure at $w = 11$ by more than 41998 bits;*

(b) *at the head rungs $m = K + w$ it is nonvacuous for exactly $w \le 11$* (margin at least +28846 at $w = 11$)*;*

(c) *by Proposition 14.13, on all four deployed rows the identity witness lists at head depths are exactly $\binom{M}{m} p^{-w}(1+\varepsilon)$, $|\varepsilon| \le 2^{-\mathrm{margin}}$; the doubled Weil constant halves the proved depth relative to the KoalaBear rows but leaves the $w = 1$ margin above 1.6×10^6 bits;*

(d) *planted cores cost about $\log_2(1 + m/(2B)) \approx 3.7$ bits of margin per point on these rows.*

All margins use the same exact-ratio and outward-rounded interval certification as above.

Corollary 14.17 (per-scale table for the bounded quotient census)**.** *At the KoalaBear quotient scales, with $N_c = 2^{21}/c$ and the per-scale certificate sizes m_c, Theorem 14.8 prices the scale-c prefix census with proved depths*

c	2	4	8	16	32
m_c	558019	279007	139501	69748	34871
$w_0(c)$	10	5	2	1	—
$w = 1$ *margin (bits)*	> 810165	> 331057	> 111115	> 18230	< -14581

so the census cells of the budget windows at scales $c \le 16$ carry exactly priced heads, while the theorem is vacuous from $c = 32$ onward ($N_c/\sqrt{p} \approx 1.42$: inside Weil range but outside the entropy margin*). For the circle rows, dyadic quotients of twin cosets are again twin cosets by Lemma 7.7, so Theorem 14.15 applies at every quotient scale with (M, d') replaced by $(M/c,\, 2^{10}c)$ and correspondingly halved reach.*

14.6 Finite certificates versus the historical two-input shorthand

Remark 14.18 (finite certificate versus the historical two-input shorthand)*.* Let a_0 be one of the four active unsafe agreements in Conjecture 14.6. An exact finite one-step theorem follows only from a complete upper certificate

$$U(a_0 + 1) \le B_* < L(a_0),$$

where L is the proved lower staircase, U is the full upper numerator for the same row, and B_* is the challenge budget. Under that assumption Proposition 8.1 and Theorems 13.2 and 13.4 gives $a_* = a_0 + 1$.

The older shorthand "(Q)+(A)" is retained only as a proof-route grouping. A polynomial-loss prefix max-fiber estimate supplies the boundary (Q) bucket, but an asymptotic closure additionally needs a complete normalized residual certificate: the corrected interior base-field census, local primitive shift-pair control, exhaustive bucket coverage, and support-to-parameter coalescing. If those data are supplied and their combined loss $L(n)$ satisfies $R(n) \to \infty$, $R(n) = o(n)$, and

$$\log L(n) = o\big(R(n) \log |\mathbb{B}|\big),$$

then the reserve absorbs the loss. None of these asymptotic statements decides a printed adjacent pair without explicit constants and the single summed integer comparison of Remark 21.5.

15 Flexible-budget conversion and the MCA frontier

The deep-point conversion has so far been used only at its threshold operating point: Theorem 2.6 converts a list of size exceeding $q/k+1$ into $\varepsilon_{\mathrm{ca}} > \frac{1}{2k}(1-n/q)$, and every deployed MCA certificate (Corollaries 5.4 and 10.9 and Proposition 14.2) paid the full threshold $q/k+1 \approx 2^{166}$ even though the challenge budgets are only $\lfloor 2^{-128}q \rfloor \approx 2^{58}$ (KoalaBear) and $\lfloor 2^{-100}q \rfloor \approx 2^{24}$ (Mersenne-31). The counting inside the proof of Theorem 2.6 is in fact operating-point-free; extracting it recovers the difference.

Corollary 15.1 (flexible-budget deep-point conversion)**.** *Let $C = \mathrm{RS}[\mathbb{F}, D, k]$, $C^+ = \mathrm{RS}[\mathbb{F}, D, k+1]$, $q = |\mathbb{F}| > n$, and let m be an integer agreement with $k+1 \le m \le n$. Suppose some received word $U : D \to \mathbb{F}$ has*

$$\big|\Lambda\big(C^+,\, 1-m/n,\, U\big)\big| \;\ge\; L_0 \qquad \textit{with} \qquad \binom{L_0}{2} k \;<\; q-n.$$

Then some simple-pole line (f_α, g_α), $f_\alpha = U/(x-\alpha)$, $g_\alpha = -1/(x-\alpha)$, $\alpha \in \mathbb{F} \setminus D$, has at least L_0 distinct CA-bad — hence, by Fact 2.4, support-wise MCA-bad — finite slopes at agreement threshold m. Consequently

$$\varepsilon_{\mathrm{ca}}\big(C, 1-m/n\big) \;\ge\; \frac{L_0}{q}, \qquad \varepsilon_{\mathrm{mca}}\big(C,\delta\big) \;\ge\; \frac{L_0}{q} \quad \textit{for every } \delta \in \big[1-\tfrac{m}{n},\, 1\big)$$

by Lemma 2.5. If moreover α is restricted to $\mathbb{F} \setminus (D \cup \mathbb{B})$, the same conclusion holds with $q-n$ replaced by $q-|\mathbb{B}|$ in the hypothesis, and the certifying line is not $\mathbb{B}$-rational.

Proof. This is the count inside the proof of Theorem 2.6, stopped before the threshold specialization. Let $P_1, \dots, P_{L_0} \in \mathbb{F}[X]_{<k+1}$ be distinct listed polynomials. For $i \ne j$ the difference $P_i - P_j$ is a nonzero polynomial of degree at most k, so $\sum_{\alpha\in\Omega} \#\{\,\{i,j\} : P_i(\alpha) = P_j(\alpha)\,\} \le \binom{L_0}{2} k < |\Omega|$ for $\Omega = \mathbb{F} \setminus D$; hence some $\alpha \in \Omega$ has *no* collisions, and the L_0 values $P_i(\alpha)$ are pairwise distinct. As in the proof of Theorem 2.6, each value $z_i = P_i(\alpha)$ is a slope at which $f_\alpha + z_i g_\alpha = (P_i(x) - P_i(\alpha))/(x-\alpha)$ is explained by C on the agreement set of P_i with U (size $\ge m \ge k+1$), while no polynomial of degree $< k$ agrees with g_α on more than k points; so all L_0 slopes are CA-bad at radius $1-m/n$ and support-wise MCA-bad with the same witnesses. The probability statements follow from Definition 2.2 with γ uniform in $\mathbb{F}$. For the last sentence run the average over $\Omega' = \mathbb{F} \setminus (D \cup \mathbb{B})$. The line is not projectively $\mathbb{B}$-rational: if some $\lambda \in \mathbb{F}^\times$ made λg_α $\mathbb{B}$-valued on D, then for distinct $x, y \in D$ the ratio $g_\alpha(x)/g_\alpha(y) = (y-\alpha)/(x-\alpha)$ would lie in $\mathbb{B}$; the two-point calculation of Corollary 2.12 would force $\alpha \in \mathbb{B}$, a contradiction. □

Combining Corollary 15.1 with the identity-prefix floor Lemma 14.1 at $K = k+1$, the MCA-row unsafe certificate becomes the pure fiber-versus-budget comparison

$$\binom{n}{m} \;>\; p^{\,m-k-1} \left\lfloor \frac{q}{2^t} \right\rfloor, \qquad t = \text{row target exponent},$$

of exactly the same shape as the list-row certificates — the threshold $q/k+1$ has disappeared. Optimizing m under exact integer verification moves the MCA frontier.

Proposition 15.2 (moved MCA frontier, exact)**.** *With $n = 2^{21}$, $k = 2^{20}$, and the deployed rows of Proposition 14.2:*

row / challenge	m_1	$w = m_1 - k - 1$	*certified unsafe range*	*approx. pass / next deficit* $\log_2$
KoalaBear, MCA, $\varepsilon^* = 2^{-128}$	1116047	67470	$\delta \in [\frac{981105}{2097152}, \frac{1}{2}) = [0.4678273\ldots, \frac{1}{2})$	+8.978 / −22.197
Mersenne-31 line round, MCA, $\varepsilon^* = 2^{-100}$	1116023	67446	$\delta \in [\frac{981129}{2097152}, \frac{1}{2}) = [0.4678387\ldots, \frac{1}{2})$	+27.927 / −3.259

Specifically, the exact integer comparisons

$$\binom{n}{1116047} > p^{67470}\left\lfloor \frac{p^6}{2^{128}} \right\rfloor, \qquad \binom{n}{1116023} > (p')^{67446}\left\lfloor \frac{(p')^4}{2^{100}} \right\rfloor$$

hold, and both fail with m_1 *replaced by* $m_1 + 1$ *and* w *by* $w + 1$*; the auxiliary admissibility* $\binom{\lfloor \varepsilon^* q \rfloor + 1}{2} k < q - n$ *holds exactly on both rows (indeed with* $q - n$ *replaced by* $q - |\mathbb{B}|$*), and* $\varepsilon^* q \notin \mathbb{Z}$*, so unsafety is strict. Hence* $\varepsilon_{\mathrm{mca}} > \varepsilon^*$ *on the printed ranges: the MCA rows are certified at these stronger agreements, and in Corollary 14.3 the imported sandwich upper edges improve to*

$$\delta^*_C(2^{-128}) \le \frac{981105}{2097152} \quad \textit{(KoalaBear)}, \qquad \delta^*_C(2^{-100}) \le \frac{981129}{2097152} \quad \textit{(Mersenne-31)}.$$

The two list rows are unchanged. Both new edges remain strictly below the entropy–subfield envelope of Definition 14.4: the deployed MCA gaps now sit within 7.2×10^{-7} *(KoalaBear) and* 4.9×10^{-7} *(Mersenne-31) of the respective ceilings* $1 - \rho - g^*$*.*

Proof. Apply Corollary 15.1 with $L_0 = \lfloor \varepsilon^* q \rfloor + 1$ and $U = U_{z^*}$ the identity witness word of Lemma 14.1 at $K = k+1$ and $m = m_1$; by Proposition 14.13 its C^+-list at radius $1 - m_1/n$ has size $|\mathrm{Fib}| \ge \left\lceil \binom{n}{m_1} p^{-w} \right\rceil \ge L_0$, which is the first displayed comparison ($\lceil x \rceil \ge B+1 \iff x > B$). The admissibility hypothesis is the second displayed comparison. Then $\varepsilon_{\mathrm{mca}}(C, \delta) \ge L_0/q > \varepsilon^*$ on the printed range. For the circle row the floor and the witness classification run verbatim on the twin-coset x-domain ($D \subseteq \mathbb{B} = \mathbb{F}_{p'}$ is all that is used), and the conversion is field-agnostic. Failure at $m_1 + 1$ is the printed failing comparison. All six integer comparisons (two passes, two failures, two admissibilities) were verified by exact arithmetic; the bit margins are orientation values. The envelope positions are $g = 67471/2097152 = 0.03217268 < g^*(\frac{1}{2}, \log_2 p) = 0.03217339$ and $g = 67447/2097152 = 0.03216124 < g^*(\frac{1}{2}, \log_2 p') = 0.03216172$. □

Corollary 15.3 (updated adjacent pairs)**.** *In Conjecture 14.6 and Remark F.1, the predicted adjacent pairs become*

row	*proved unsafe* a_0	*conjectured first safe* $a_0 + 1$	*deficit floor (bits)*
KoalaBear MCA	1116047	1116048	> 22.1
KoalaBear list	1116046	1116047	> 22.0
Mersenne-31 MCA	1116023	1116024	> 3.2
Mersenne-31 list	1116022	1116023	> 3.0

with the endpoint convention of Corollary 13.3. Two structural remarks. First, at equal m *the MCA comparison is exactly one factor of* p *easier than the list comparison (depth* $m - k - 1$ *against* $m - k$*), while one agreement step costs* $\log_2 p - \log_2 \frac{n-m}{m+1} > \log_2 p$ *bits for* $m > n/2$*; hence* $a_0^{\mathrm{list}} \le a_0^{\mathrm{MCA}} \le a_0^{\mathrm{list}} + 1$ *always in this regime, with the right equality realized on both deployed*

rows: the unsafe edges interleave as list $<$ MCA *one step apart. Second, the two Mersenne-31 rows are now the tight rows (deficits only just above* 3.0 *and* 3.2 *bits); any finite safe-side input must be exact-constant there.*

Proof. Immediate from Proposition 15.2 and the unchanged list rows of Proposition 14.2; the one-factor-of-p relation is the identity $w_{\mathrm{MCA}}(m) = w_{\mathrm{list}}(m) - 1$. □

16 Dual-weighted syndrome calculus

The exact-agreement machinery of Section 3.4 manipulates window sums $\sum_x y(x)L(x)x^m$. Over a multiplicative coset the $m = 0$ row of the unweighted sum retains a constant-term defect, which the main text absorbs by the shifted-window convention. We adopt instead the dual-weighted normalization, which is exact in the unshifted window over *every* domain D and makes the dictionary below uniform. The two conventions differ by the fixed diagonal change $y \mapsto v \cdot y$ (a code equivalence onto the dual-twisted generalized Reed–Solomon partner) together with an invertible triangular change of window rows; all rank, kernel and cofactor statements used from the main text are invariant under such changes, and Theorem B.19 is stated for arbitrary affine pencils, so it applies to the realization below verbatim.

Lemma 16.1 (residue pairing)**.** *Let $D \subseteq \mathbb{F}$ with $|D| = n \geq 1$ and set $v_x := 1/P'_D(x)$ for $x \in D$ (well defined: $P'_D(x) = \prod_{y \in D\setminus\{x\}}(x - y) \neq 0$). Then for every $h \in \mathbb{F}[X]$ with $\deg h \leq n - 1$,*

$$\sum_{x \in D} v_x \, h(x) \;=\; [X^{\,n-1}]\, h.$$

In particular the sum vanishes whenever $\deg h \leq n - 2$.

Proof. The Lagrange interpolant $\sum_{x \in D} h(x)\, \ell_{D\setminus\{x\}}(X)/P'_D(x)$ has degree $\leq n - 1$ and agrees with h on the n points of D; since $\deg h \leq n - 1$, the difference has n roots and degree $\leq n - 1$, hence is zero. Each $\ell_{D\setminus\{x\}}$ is monic of degree $n - 1$, so comparing coefficients of X^{n-1} gives the identity. □

Definition 16.2 (weighted syndromes, annihilator spaces)**.** For a word $y\colon D \to \mathbb{F}$ put $s_m(y) := \sum_{x \in D} v_x \, y(x)\, x^m$ $(m \geq 0)$. For $L = \sum_i \ell_i X^i \in \mathbb{F}[X]$ and a window length $\tau \geq 1$ we say L *annihilates y in window τ*, and write $L \in \mathcal{A}_\tau(y)$, if

$$\sum_{x \in D} v_x \, y(x)\, L(x)\, x^m \;=\; \sum_i \ell_i \, s_{m+i}(y) \;=\; 0 \qquad (0 \leq m \leq \tau - 1).$$

$\mathcal{A}_\tau(y)$ is an $\mathbb{F}$-subspace of $\mathbb{F}[X]$. Restricted to $\deg L \leq j$ the condition reads $H_{\tau,j+1}\big(s(y)\big)\,\tilde{\ell} = 0$ with the Hankel matrix built from $(s_0(y), \ldots, s_{\tau+j-1}(y))$; with $u := \big(s_m(f)\big)_m$ and $v := \big(s_m(g)\big)_m$ this realizes the exact-agreement pencils of Section 3.4 for the line (f, g).

Lemma 16.3 (annihilator–tangent dictionary)**.** *Let L be split over D with root set T, $\gamma := \deg L = |T|$, and let $\tau \geq 1$ satisfy $\tau + \gamma \leq n$. Then for every word y,*

$$L \in \mathcal{A}_\tau(y) \quad\Longleftrightarrow\quad y \in \mathrm{RS}[\mathbb{F}, D, n - \tau - \gamma] \;\oplus\; \mathbb{F}^T,$$

where $\mathbb{F}^T$ denotes the words supported on T and $\mathrm{RS}[\mathbb{F}, D, \kappa] := \{p|_D : \deg p < \kappa\}$ (the zero space if $\kappa \leq 0$). The sum on the right is direct.

Proof. ($\supseteq$) If y is supported on T then $y(x)L(x) = 0$ for all $x \in D$. If $y = p|_D$ with $\deg p \le n-\tau-\gamma-1$, then for $0 \le m \le \tau-1$ the polynomial $h := p\,L\,X^m$ has degree $\le (n-\tau-\gamma-1)+\gamma+(\tau-1) = n-2$, so $\sum_x v_x y(x)L(x)x^m = [X^{n-1}]h = 0$ by Lemma 16.1.

(Rank.) The linear map $y \mapsto \big(\sum_x v_x y(x)L(x)x^m\big)_{m<\tau}$ has, in the standard basis of $\mathbb{F}^D$, the column at x equal to $v_x L(x)\cdot(x^m)_{m<\tau}$. The columns over T vanish; over $E := D\setminus T$ the scalars $v_x L(x)$ are nonzero, so the matrix is a column-scaled $\tau\times(n-\gamma)$ Vandermonde on distinct nodes and has rank $\min(\tau, n-\gamma) = \tau$.

(Dimensions.) The kernel has dimension $n-\tau$. The right side has dimension $(n-\tau-\gamma)+\gamma = n-\tau$ once the sum is direct: a word in the intersection is $p|_D$ vanishing on E, i.e. a polynomial of degree $\le n-\tau-\gamma-1 < n-\gamma = |E|$ with $|E|$ roots, hence zero. The containment of a subspace of equal dimension forces equality. □

At the critical window $\tau = n-k-\gamma$ the dictionary reads: *a split annihilator of degree γ certifies agreement with a codeword of C off its root set.* Shorter windows certify agreement with higher-degree polynomials; this graded form is used in Proposition 17.8.

Lemma 16.4 (window cost of multiplication)**.** *If $L \in \mathcal{A}_\tau(y)$ and $A \in \mathbb{F}[X]$ with $d := \deg A \le \tau-1$, then $AL \in \mathcal{A}_{\tau-d}(y)$.*

Proof. $\sum_x v_x y(AL)(x)x^m = \sum_c A_c \sum_x v_x y L(x)x^{m+c} = 0$ whenever $m+d \le \tau-1$. □

Lemma 16.5 (gcd window)**.** *Let $L_1, L_2 \in \mathcal{A}_\tau(y)$ be nonzero with $\deg L_i \le j$, let $G := \gcd(L_1, L_2)$ and $\gamma := \deg G$, and suppose $\tau \ge j-\gamma$. Then $G \in \mathcal{A}_{\tau-(j-\gamma-1)}(y)$ (with the convention that for $L_2 \mid L_1$ this is the trivial statement $G \in \mathcal{A}_\tau(y)$).*

Proof. Write $L_i = GC_i$ with C_1, C_2 coprime. If C_2 is constant then G is a scalar multiple of L_2 and the claim is trivial. Otherwise Bézout gives A, B with $AC_1+BC_2 = 1$, $\deg A < \deg C_2 \le j-\gamma$ and $\deg B < \deg C_1 \le j-\gamma$; hence $G = AL_1+BL_2$ with $\deg A, \deg B \le j-\gamma-1$, and Lemma 16.4 plus linearity give $G \in \mathcal{A}_{\tau-(j-\gamma-1)}(y)$. □

Lemma 16.6 (small annihilator)**.** *If $\dim\big(\mathcal{A}_\tau(y) \cap \mathbb{F}[X]_{\le j}\big) \ge d+1$ for some $0 \le d \le j$, then $\mathcal{A}_\tau(y)$ contains a nonzero element of degree $\le j-d$.*

Proof. $\mathbb{F}[X]_{\le j-d}$ has codimension d in $\mathbb{F}[X]_{\le j}$, so a subspace of dimension $\ge d+1$ meets it nontrivially. □

Remark 16.7 (support-side companions). The functionals $S_r(h,T) = \sum_{x\in T} h(x)x^r/\Lambda_T'(x)$ of Lemmas E.9 and F.11 are the *support-side* instance of the same residue identity as Lemma 16.1: there the domain is T itself, the weights are $1/\Lambda_T'$, and window annihilation tests interpolability on T. The dictionary Lemma 16.3 is the domain-side dual: the domain is all of D, the weights are $1/P_D'$, and window annihilation against a split locator L tests membership in $\mathrm{RS}[\mathbb{F}, D, n-\tau-\gamma] \oplus \mathbb{F}^{Z(L)}$. The two calculi meet in the rank-one census: the twisted vectors $(u(x)/\Lambda_T'(x))\mathbf{1}_T$ of Remark F.13 are supported window data in the support-side calculus, while the pencil realization used in Section 17 is domain-side. Statements proved in either calculus transfer along the diagonal code equivalence recorded above.

17 Split rigidity for kernel sections

The residual branches of the exact-agreement atlas are populated by kernel sections — one-parameter families $W(Z, X)$ annihilating the line's syndromes identically in the slope variable — whose normalized locators pass the split test identically. We show these branches collapse: the locator family is constant, and the line is a tangent pair. The engine is an elementary rigidity statement for divisors with constant coefficients.

Lemma 17.1 (relative algebraic closure)**.** *$\mathbb{F}$ is relatively algebraically closed in $\mathbb{F}(Z)$: if $c \in \mathbb{F}(Z)$ satisfies $m(c) = 0$ for some nonzero $m \in \mathbb{F}[Y]$, then $c \in \mathbb{F}$.*

Proof. Write $c = a/b$ with $a, b \in \mathbb{F}[Z]$ coprime, b monic, and $m = \sum_{i \le d} m_i Y^i$ with $m_d \neq 0$, $d \geq 1$ (if $d = 0$ then m is a nonzero constant and has no root, so $d \geq 1$). Multiplying $m(c) = 0$ by b^d gives $m_d a^d = -b \sum_{i<d} m_i a^i b^{d-1-i}$, so $b \mid m_d a^d$; by coprimality $b \mid m_d$, hence $b = 1$ and $c = a \in \mathbb{F}[Z]$. If $\deg_Z a \geq 1$ then, letting i_0 be the largest index with $m_{i_0} \neq 0$ and $i_0 \geq 1$, the polynomial $m(a(Z)) = \sum_i m_i a(Z)^i$ has a nonzero leading term of degree $i_0 \deg_Z a$ (the degrees $i \deg_Z a$ are distinct), so $m(a) \not\equiv 0$, a contradiction; if $i_0 = 0$ then $m(a) = m_0 \neq 0$, likewise a contradiction. Hence $\deg_Z a = 0$ and $c \in \mathbb{F}$. □

Lemma 17.2 (constant-divisor rigidity)**.** *Let $P \in \mathbb{F}[X]$ be nonzero and let $L \in \mathbb{F}(Z)[X]$ be monic with $L \mid P$ in $\mathbb{F}(Z)[X]$. Then $L \in \mathbb{F}[X]$ and $L \mid P$ in $\mathbb{F}[X]$.*

Proof. Over $\overline{\mathbb{F}}(Z)$ the polynomial P factors as $\mathrm{lc}(P) \prod_r (X - \xi_r)$ with all $\xi_r \in \overline{\mathbb{F}}$ (they are roots of P, hence algebraic over $\mathbb{F}$, hence constant). Since $\overline{\mathbb{F}}(Z)[X]$ is a unique factorization domain and L is a monic divisor of P there, $L = \prod_{r \in I} (X - \xi_r)$ for some index set I; its coefficients lie in $\overline{\mathbb{F}}$. They also lie in $\mathbb{F}(Z)$, hence in $\mathbb{F}$ by Lemma 17.1. Finally, division with remainder of P by L in $\mathbb{F}[X]$ must reproduce the (unique) division in $\mathbb{F}(Z)[X]$, whose remainder is zero; so $L \mid P$ in $\mathbb{F}[X]$. □

Theorem 17.3 (split kernel sections are tangent-borne)**.** *Fix an exact agreement a for $C = \mathrm{RS}[\mathbb{F}, D, k]$, and set $j = n - a$, $t = a - k \geq 1$. Let (f, g) be a received line and let $W(Z, X) = \sum_{i=0}^{j} W_i(Z) X^i \in \mathbb{F}[Z][X]$ be nonzero with*

$$\sum_{x \in D} v_x \, (f + Zg)(x) \, W(Z, x) \, x^m \;\equiv\; 0 \quad \text{in } \mathbb{F}[Z] \qquad (0 \leq m \leq t - 1), \tag{4}$$

equivalently $\big(H_t(u) + Z H_t(v)\big) \big(W_0(Z), \ldots, W_j(Z)\big)^{\top} \equiv 0$ with u, v the weighted syndrome sequences of f, g. Let $\gamma := \deg_X W$ over $\mathbb{F}(Z)$ (so $W_\gamma \not\equiv 0$) and let $\widehat{L} := W / W_\gamma \in \mathbb{F}(Z)[X]$ be the monic normalization. Assume

$$\widehat{L} \text{ divides } P_D \text{ in } \mathbb{F}(Z)[X].$$

Then:

- *(i) $\widehat{L} = L^* \in \mathbb{F}[X]$ is a (constant) monic divisor of P_D, hence split over D; write $T^* := Z(L^*)$, $|T^*| = \gamma$.*
- *(ii) $f, g \in \mathrm{RS}[\mathbb{F}, D, n - t - \gamma] \oplus \mathbb{F}^{T^*}$; write $f = c_f + e_f$, $g = c_g + e_g$ with $\deg c_f, \deg c_g < n - t - \gamma$ and e_f, e_g supported on T^*.*

(iii) At every agreement threshold a' with $n - a' \le t$, the line (f, g) has at most $\max(\gamma, 1)$ MCA-bad finite slopes. If $c_f, c_g \in C$, every such slope is a tangent ratio $-e_f(x)/e_g(x)$ with $x \in T^*$, $e_g(x) \neq 0$, and there are at most γ of them; otherwise there is at most one such slope, namely the unique z (if any) with $c_f + zc_g \in C$.

(iv) If $c_f, c_g \in C$, the pair (f, g) is column-close to C at radius γ/n, hence has no CA-bad slopes at any radius pair $(\delta, \delta_{\text{int}})$ with $\delta_{\text{int}} \ge \gamma/n$.

Proof. (i) is Lemma 17.2 applied to $P := P_D$ (squarefree with root set D, so any monic divisor is split over D).

(ii) Dividing (4) by the unit $W_\gamma \in \mathbb{F}(Z)^\times$ and using $\widehat{L} = L^*$ constant,

$$0 = \sum_x v_x (f + Zg)(x) L^*(x) x^m = \Big(\sum_x v_x f(x) L^*(x) x^m \Big) + Z \Big(\sum_x v_x g(x) L^*(x) x^m \Big) \quad \text{in } \mathbb{F}(Z)$$

for $0 \le m \le t - 1$; both coefficients of this degree-≤ 1 polynomial in Z vanish, i.e. $L^* \in \mathcal{A}_t(f) \cap \mathcal{A}_t(g)$. Since $t + \gamma \le t + j = n - k \le n$, Lemma 16.3 with $\tau = t$ gives (ii); note $\kappa := n - t - \gamma \ge n - t - j = k \ge 1$.

(iii) Fix a' with $r' := n - a' \le t$ and let z be an MCA-bad finite slope at threshold a', with witness support S, $|S| \ge a'$, and explaining codeword $c \in C$ for the point $f + zg$ on S. On $S \setminus T^*$, which has size $\ge (n - r') - \gamma \ge n - t - \gamma = \kappa$, both c and $c_f + zc_g$ are restrictions of polynomials of degree $< \kappa$ (using $C \subseteq \mathrm{RS}[\mathbb{F}, D, \kappa]$) agreeing at $\ge \kappa$ points, hence equal as polynomials:

$$c \;=\; c_f + zc_g. \tag{5}$$

Suppose first $c_f, c_g \in C$. Explanation of the *pair* on S by codewords $p_1, p_2 \in C$ forces, by the same κ-point argument on $S \setminus T^*$, $p_1 = c_f$ and $p_2 = c_g$; hence the pair is explained on S if and only if e_f, e_g both vanish on $S \cap T^*$. Since z is MCA-bad there is a witness S on which the pair is *not* explained, so some $x \in S \cap T^*$ has $(e_f(x), e_g(x)) \neq (0, 0)$; and (5) restricted to S gives $e_f(x) + ze_g(x) = 0$ at that x. If $e_g(x) = 0$ this is impossible; if $e_g(x) \neq 0$ then $z = -e_f(x)/e_g(x)$. Thus every bad slope is one of at most $|T^*| = \gamma$ tangent ratios. (If $\gamma = 0$ then $e_f = e_g = 0$, $f, g \in C$, and no slope is ever MCA-bad.)

Suppose now that $c_f \notin C$ or $c_g \notin C$. By (5), badness of z forces $c_f + zc_g \in C$. If $c_g \in C$ this forces $c_f \in C$, a contradiction, so there is no bad slope. If $c_g \notin C$, then two distinct bad slopes $z \neq z'$ would give $(z - z')c_g \in C$, again a contradiction; so there is at most one bad slope. This proves (iii).

(iv) If $c_f, c_g \in C$ then (f, g) differs from the codeword pair (c_f, c_g) only on T^*, so its column distance to $C \times C$ is at most γ/n; a pair that is δ_{int}-close for $\delta_{\text{int}} \ge \gamma/n$ is not δ_{int}-far, hence supports no CA-bad slope at $(\delta, \delta_{\text{int}})$ by definition. □

Corollary 17.4 (collapse of the identically split top chart at deficiency one)**.** *In the setting of Theorem B.19 at deficiency one ($t = j$, realized by the weighted pencil $M(Z) = H_t(u) + ZH_t(v)$ of Definition 16.2), assume the top chart is nonempty ($c_j \not\equiv 0$) and that case (b) holds, i.e. the split test passes identically. Then there is a fixed j-set $T^* \subseteq D$ such that f and g each agree with a codeword of C off T^*; the line has at most j MCA-bad finite slopes at every agreement threshold $a' \ge a$, each a tangent ratio; and it has no CA-bad slopes at any radius pair $(\delta, \delta_{\text{int}})$ with $\delta_{\text{int}} \ge j/n$. In particular the residual named in case (b) of Theorem B.19 is closed.*

Proof. For the $t \times (j+1)$ pencil $M(Z)$ let $c(Z) = \big(c_0(Z), \dots, c_j(Z)\big)$ be the cofactor vector, $c_i = (-1)^i \det M^{(i)}(Z)$ with $M^{(i)}$ the maximal minor omitting column i. Then $M(Z)c(Z) \equiv 0$ identically: appending any row of $M(Z)$ on top of $M(Z)$ produces a $(j+1) \times (j+1)$ matrix with a repeated row, whose Laplace expansion along the top row is the corresponding entry of $M(Z)c(Z)$. Set $W(Z,X) := \sum_i c_i(Z)X^i$; on the top chart $\deg_X W = j = \gamma$ and $\widehat{L} = W/c_j$. Case (b) states that the pseudo-division remainders of $c_j^{\Delta} P_D$ by W vanish identically, i.e. $c_j^{\Delta} P_D = QW$ in $\mathbb{F}[Z][X]$; since c_j is a unit in $\mathbb{F}(Z)$ this is exactly $\widehat{L} \mid P_D$ in $\mathbb{F}(Z)[X]$. Theorem 17.3 applies with $\gamma = j$; here $n - t - \gamma = n - 2j = k$, so $c_f, c_g \in C$ automatically and parts (iii)–(iv) give the stated conclusions for all a' with $n - a' \le t = j$, i.e. all $a' \ge a$. □

Corollary 17.5 (identically singular chains with split sections)**.** *In the setting of Lemma 9.8 (any $t \ge j$), suppose some polynomial kernel section $W(Z,X)$ of the pencil, of X-degree $\gamma \le j$ over $\mathbb{F}(Z)$, has identically split monic normalization (equivalently, $\widehat{L} \mid P_D$ in $\mathbb{F}(Z)[X]$). Then the conclusions of Theorem 17.3 hold; in particular the line has at most $\max(\gamma, 1)$ MCA-bad finite slopes at every threshold a' with $n - a' \le t$.*

Proof. A kernel section satisfies (4) by definition of the pencil realization; apply Theorem 17.3. □

Remark 17.6 (quantitative identically-split criterion). Let $W(Z,X)$ be as in Theorem 17.3 with $e_Z := \deg_Z W$ and top coefficient W_γ. Pseudo-dividing P_D by W in X takes $n - \gamma + 1$ steps, each multiplying the running remainder by W_γ and subtracting a coefficient multiple of W, so every remainder coefficient lies in $\mathbb{F}[Z]$ with degree at most $(n - \gamma + 1)e_Z$. Consequently, if the specialized monic normalization $W(z,\cdot)/W_\gamma(z)$ is split over D for more than $(n - \gamma + 2)\, e_Z$ values of z (the extra e_Z absorbing the roots of W_γ), then it is split identically and Theorem 17.3 applies. In words: a kernel section cannot pass the split test at superlinearly many slopes without collapsing to a tangent pair.

We now reduce the remaining deficiency-one residual, class (0) (identically rank-deficient pencils), using the calculus of Section 16.

Lemma 17.7 (small kernel sections in class (0))**.** *Let $t = j$ and suppose the pencil $M(Z)$ has rank $\le j - 1$ over $\mathbb{F}(Z)$. Then there is a nonzero $W(Z,X) \in \mathbb{F}[Z][X]$ with $\deg_X W \le j - 1$, $\deg_Z W \le j - 1$, satisfying (4).*

Proof. The $\mathbb{F}(Z)$-kernel of $M(Z)$ has dimension ≥ 2 inside $\mathbb{F}(Z)^{j+1}$, hence meets the hyperplane $\{w_j = 0\}$ nontrivially; equivalently, the $t \times j$ subpencil $M'(Z)$ of the first j columns has rank $\rho \le j - 1$ over $\mathbb{F}(Z)$ and a nontrivial kernel. Choose a $\rho \times \rho$ minor of M' that is not identically zero and solve the pivot coordinates by Cramer's rule with one free coordinate set to that minor; the resulting kernel vector has polynomial entries, each a signed $\rho \times \rho$ or bordered minor of an affine pencil, hence of Z-degree $\le \rho \le j - 1$. □

Proposition 17.8 (class-(0) dichotomy at deficiency one)**.** *Let $t = j$, let the pencil of the line (f, g) be identically rank-deficient (class (0)), and let W be as in Lemma 17.7, with $\gamma_W := \deg_X W \le j - 1$ and $e_Z := \deg_Z W \le j - 1$. Consider the MCA-bad finite slopes of (f, g) at exact agreement a, i.e. those with a witness support of size exactly a, and stratify each bad slope z by $\gamma_z := \deg \gcd\big(\ell_T, W(z,\cdot)\big)$, where T is the co-support of a witness (so $\ell_T \in \mathcal{A}_t(f + zg)$ by Lemma 16.3) and z avoids the $\le e_Z$ roots of the top coefficient and the $\le e_Z$ zeros of $W(z,\cdot)$. Then:*

(a) *(*coprime stratum, $\gamma_z = 0$*) if* $\big(s_0(f), s_0(g)\big) \neq (0,0)$*, at most one bad slope lies in this stratum: the unique* z *with* $s_0(f) + z\, s_0(g) = 0$*.*

(b) *(*section-split stratum, $W(z,\cdot) \mid \ell_T$*) these slopes number at most* $(n - \gamma_W + 2)\, e_Z \leq (n+1)(j-1)$*, unless the monic normalization of* W *is identically split, in which case Corollary 17.5 bounds the line's bad slopes by* $\max(\gamma_W, 1) \leq j - 1$ *outright.*

(c) *(*divisor-chain residual, $1 \leq \gamma_z < \gamma_W$*) each such slope carries a split* $G_z := \gcd(\ell_T, W(z,\cdot))$ *of degree* γ_z *with* $G_z \in \mathcal{A}_{\gamma_z+1}(f + zg)$ *and* $G_z \mid W(z,\cdot)$*; by the dictionary,* $f + zg$ *agrees with a polynomial of degree* $< n - 2\gamma_z - 1$ *off* $Z(G_z)$*. We make no polynomial claim for this stratum; it is the named residual.*

Lines with $s_0(f) = s_0(g) = 0$ *(equivalently, both interpolants of degree* $\leq n-2$*) form the second named residual.*

Proof. For a bad slope z at exact agreement a with witness support S and $T = D \setminus S$, $|T| = j$: the dictionary (Lemma 16.3, $\tau = t$, $\gamma = j$, $n - t - j = k$) converts the explanation into $\ell_T \in \mathcal{A}_t(f + zg)$. Specializing (4) at z gives $W(z,\cdot) \in \mathcal{A}_t(f + zg)$, nonzero for the admitted z, of degree $\leq j - 1$. Lemma 16.5 with $\tau = t = j$ and degree cap j places G_z in $\mathcal{A}_{\gamma_z+1}(f + zg)$. If $\gamma_z = 0$ then $1 \in \mathcal{A}_1(f + zg)$, i.e. $s_0(f + zg) = s_0(f) + z s_0(g) = 0$, giving (a). If $\gamma_z = \deg W(z,\cdot)$, i.e. $W(z,\cdot) \mid \ell_T$, then the monic normalization of $W(z,\cdot)$ is split over D; by Remark 17.6 either this happens for at most $(n - \gamma_W + 2)e_Z$ slopes or the section is identically split and Corollary 17.5 applies — giving (b). In the remaining stratum, G_z divides the split ℓ_T, hence is split, and the window $\gamma_z + 1$ instance of the dictionary ($\kappa = n - (\gamma_z + 1) - \gamma_z$) gives the stated agreement statement — (c). □

Remark 17.9 (closure of the identically-split residual of slope elimination). Theorem F.12(d) bounds the slope count of any exact-agreement syndrome pencil $M(Z) = M_0 + ZM_1$ by a rank-one census over split candidates Λ_A, *except* on the residual branch $M_0\Lambda_A = M_1\Lambda_A = 0$ "as in Theorem B.19". For a fixed split Λ_A this branch is a common support in the sense of Proposition E.10 and is already harmless by Theorem F.12(b). The genuinely open case was the *family* version — kernel sections $W(Z, X)$ that are identically split without any single Z-independent split annihilator being given — and that case is closed by Theorem 17.3: any rational identically split kernel section is constant, so the branch collapses to the fixed-candidate case, at every deficiency admitting a rational kernel section, with the tangent count $\leq \max(\gamma, 1)$. In particular the "known unconditional strata" list of Remark F.13 may now cite deficiency one *without* the residual caveat: by Corollary 17.4, deficiency-one charts are eliminant-or-tangent outright.

18 Identity-scale aperiodic slope floors and the refuted band conjecture

Lemma 14.1 produces list-size floors at the identity periodicity scale, and Proposition 14.13 shows the witness lists are exactly the prefix fibers. We now transport these statements to *per-line, per-slope* form: for the pole lines of Theorem 2.6, the threshold-m witness supports of every MCA-bad slope are exactly fiber members, and for $\gcd(m, n) = 1$ every witness is aperiodic. The consequence is that Remark 9.11, which asks for a q-independent polynomial bound on aperiodically witnessed slopes throughout the band $k + 2n/N < a < n - \lfloor (n-k)/3 \rfloor$, *fails* below

the entropy–subfield envelope. This upgrades two prior gestures to theorems: the informal band-edge move recorded in Remark 10.11, and the genericity heuristic in the preamble of Section E.1 ("generic supports are aperiodic except for a negligible quotient-periodic subfamily") — at odd m the correct statement is that *all* witnesses are aperiodic, with no exceptional subfamily, so no periodic bucket can absorb any part of the floor.

Definition 18.1 (prefixes, fibers, prefix words, pole lines)**.** Let $B \subseteq \mathbb{F}$ be a subfield with $D \subseteq B$, and fix $k+1 \le m \le n$; put $K := k+1$ and $w := m - K \ge 0$. For an m-subset $M \subseteq D$ let $\Lambda_M := \ell_M$ and define the depth-w prefix $\Phi_w(M) := \big([X^{m-h}]\Lambda_M\big)_{h=1,\dots,w} \in B^w$. For $z \in B^w$ set $\mathrm{Fib}_w(z) := \Phi_w^{-1}(z)$ and

$$U_z(X) := X^m + \textstyle\sum_{h=1}^{w} z_h X^{m-h}.$$

For $\alpha \in \mathbb{F} \setminus D$ the *pole line* at (z, α) is the line spanned by the words

$$f_\alpha(x) := \frac{U_z(x)}{x-\alpha}, \qquad g_\alpha(x) := \frac{-1}{x-\alpha} \qquad (x \in D),$$

i.e. the pole lines of Theorem 2.6 instantiated at the prefix word U_z.

Lemma 18.2 (pole-line witness classification)**.** *Fix $z \in B^w$, $\alpha \in \mathbb{F} \setminus D$, and write $(f, g) := (f_\alpha, g_\alpha)$.*

(i) For every $\zeta \in \mathbb{F}$ and every $S \subseteq D$ with $|S| \ge m$: the word $f + \zeta g$ is explained by C on S if and only if

$$|S| = m, \qquad S \in \mathrm{Fib}_w(z), \qquad \Lambda_S(\alpha) = U_z(\alpha) - \zeta.$$

(ii) No S with $|S| \ge k+1$ admits a simultaneous explanation of the pair (f, g) by C; in particular the pair is column-far at radius $1 - m/n$.

Consequently the MCA-bad finite slopes of (f, g) at agreement threshold m are exactly

$$\big\{\, U_z(\alpha) - \Lambda_S(\alpha) \;:\; S \in \mathrm{Fib}_w(z) \,\big\},$$

each of them CA-bad at radii $(1-m/n,\, 1-m/n)$ as well; and every *threshold-m witness support of every such slope is a member of* $\mathrm{Fib}_w(z)$.

Proof. (i, $\Rightarrow$) Let $c = P_c|_D$ with $\deg P_c \le k-1$ explain $f + \zeta g$ on S; on S this reads $U_z(x) - \zeta = (x-\alpha)P_c(x)$. Set $\Psi := U_z - \zeta - (X-\alpha)P_c$. Since $\deg\big((X-\alpha)P_c\big) \le k \le m-1$, Ψ is monic of degree m, and it vanishes on S; hence $|S| \le m$, so $|S| = m$ and $\Psi = \Lambda_S$. The coefficients of Ψ in degrees $m-1, \dots, m-w = k+1$ are unaffected by the degree-$\le k$ correction, so they equal those of U_z, i.e. $\Phi_w(S) = z$. Evaluating at α kills the correction term and gives $\Lambda_S(\alpha) = U_z(\alpha) - \zeta$.

(i, $\Leftarrow$) Let $S \in \mathrm{Fib}_w(z)$ and $\zeta := U_z(\alpha) - \Lambda_S(\alpha)$. Then $\Theta := U_z - \Lambda_S$ has degree $\le m - w - 1 = k$ (the monic terms and the depth-w prefix cancel), and $\zeta = \Theta(\alpha)$. The polynomial $P_c := \big(\Theta - \Theta(\alpha)\big)/(X-\alpha)$ has degree $\le k-1$, and for $x \in S$ (where $\Lambda_S(x) = 0$): $U_z(x) - \zeta = \Theta(x) - \Theta(\alpha) = (x-\alpha)P_c(x)$, i.e. $(f + \zeta g)(x) = P_c(x)$; so $P_c|_D$ explains $f + \zeta g$ on S, $|S| = m$.

(ii) If $G \in \mathbb{F}[X]$, $\deg G \le k-1$, has $G|_D = g$ on some S with $|S| \ge k+1$, then $(X-\alpha)G(X)+1$ has degree $\le k$, vanishes on S ($\ge k+1$ points), hence is identically zero; but its value at α is 1 — a contradiction. Since a simultaneous explanation of the pair requires an explanation of g, none exists on any S with $|S| \ge k+1$; in particular no pair of codewords agrees with (f, g) on $\ge m > k$ columns, so the pair is column-far at radius $1 - m/n$.

(Consequences.) A slope ζ is MCA-bad at threshold m iff some S, $|S| \geq m$, carries an explanation of the point but not of the pair; by (ii) the pair is never explained at these sizes, so badness is equivalent to the existence of an explanation of the point, which by (i) is equivalent to $\zeta \in U_z(\alpha) - \Lambda_\bullet(\alpha)\big(\mathrm{Fib}_w(z)\big)$, with witnesses exactly the fiber members mapping to ζ. CA-badness at $(1 - m/n, 1 - m/n)$ follows since the point is $(1 - m/n)$-close by (i) and the pair is $(1 - m/n)$-far by (ii) (cf. Fact 2.4). □

Remark 18.3 (relation to the exact witness lists). Lemma 18.2 is the pole-twisted form of Proposition 14.13: the map $c \mapsto (X - \alpha)c + \zeta$ is a bijection from degree-$< k$ explanations of $f_\alpha + \zeta g_\alpha$ on S to $\mathrm{RS}[\mathbb{F}, D, k+1]$-explanations c' of U_z on S with $c'(\alpha) = \zeta$, and under this bijection the classification above is the classification of Proposition 14.13 refined by the evaluation-at-α fibration $c_M \mapsto c_M(\alpha) = U_z(\alpha) - \Lambda_M(\alpha)$. Part (ii) (the pair is never explained) is what the fibration adds: it is exactly the far condition of Definition 2.2, and it is what makes the transported floor land in $\varepsilon_{\mathrm{mca}}$ rather than in list size.

Theorem 18.4 (identity-scale aperiodic slope floor)**.** *Let $D \subseteq \mathbb{B}$ be either a smooth multiplicative coset $D = \theta H$, $H = \mu_n(\mathbb{B})$, or a twin-coset Chebyshev x-domain of size n as in Theorem 9.2, and let $k + 1 \leq m \leq n$ satisfy $\gcd(m, n) = 1$. Let $z^* \in B^w$ maximize $|\mathrm{Fib}_w(\cdot)|$, so that*

$$N_{\mathrm{fib}} := |\mathrm{Fib}_w(z^*)| \;\geq\; \left\lceil \binom{n}{m} |B|^{-w} \right\rceil,$$

and set $N_0 := \min\big(N_{\mathrm{fib}}, \lceil (q-n)/k \rceil\big)$, where $q = |\mathbb{F}| > n$. Then there exists $\alpha \in \mathbb{F} \setminus D$ such that the pole line (f_α, g_α) at (z^*, α) has at least

$$\tfrac{1}{2} N_0 \Big(1 - \tfrac{k}{q-n}\Big)$$

distinct MCA-bad finite slopes at agreement threshold m, each also CA-bad at radii $(1-m/n, 1-m/n)$, and such that every *threshold-m witness support of every one of these slopes is aperiodic ($c(S) = 1$ in the sense of Definition 9.1). In particular all these slopes lie in the aperiodic bucket of the stabilizer partition (Theorem 9.2).*

Proof. By Lemma 18.2 every threshold-m witness of every bad slope of such a line is a member of $\mathrm{Fib}_w(z^*)$. For any m-subset $S \subseteq D$, Theorem 9.2(i) gives $c(S) \mid n$ and exhibits S as a union of complete scale-$c(S)$ fibers; every scale-c fiber of D has exactly c points (H_c-orbits in the multiplicative case, uniform T_c-fibers on the twin-coset x-domain by Lemma 7.8), so $c(S)$ divides $|S| = m$. Hence $c(S) \mid \gcd(m, n) = 1$: every witness is aperiodic.

For the count, fix any subfamily $\mathcal{G} \subseteq \mathrm{Fib}_w(z^*)$ with $|\mathcal{G}| = N_0$. For distinct $S, S' \in \mathcal{G}$ the polynomial $\Lambda_S - \Lambda_{S'}$ is nonzero of degree $\leq m - w - 1 = k$ (monic terms and prefixes cancel), so $\Lambda_S(\alpha) = \Lambda_{S'}(\alpha)$ for at most k values $\alpha \in \mathbb{F}$. Summing over $\alpha \in \Omega := \mathbb{F} \setminus D$,

$$\sum_{\alpha \in \Omega} \#\{\text{colliding pairs at } \alpha\} \;\leq\; \binom{N_0}{2} k,$$

so some $\alpha \in \Omega$ has at most $\binom{N_0}{2} k/(q-n)$ colliding pairs, whence the map $S \mapsto U_{z^*}(\alpha) - \Lambda_S(\alpha)$ takes at least

$$N_0 - \frac{N_0^2 k}{2(q-n)} \;\geq\; N_0\Big(1 - \tfrac{1}{2}\big(1 + \tfrac{k}{q-n}\big)\Big) \;=\; \tfrac{1}{2} N_0 \Big(1 - \tfrac{k}{q-n}\Big)$$

distinct values on $\mathcal{G}$, using $N_0 \le (q-n)/k+1$. By Lemma 18.2 each distinct value is an MCA-bad finite slope at threshold m (and CA-bad at the stated radii), with all witnesses in the fiber, hence aperiodic; membership in the aperiodic bucket follows from the bucket convention of Theorem 9.2. □

Corollary 18.5 (deployed MCA frontier floors)**.** *Let the rows be as in Proposition 15.2. At $m = 1116047$ for KoalaBear and $m = 1116023$ for Mersenne-31, the fiber pigeonhole gives*

$$N_1 := \left\lceil \binom{n}{m} p^{-w} \right\rceil$$

with $N_1 \ge 2^{66}$ and $N_1 \ge 2^{51}$ respectively. The exact comparisons $N_1(N_1-1)k < 2(q-n)$ hold, so the zero-collision average in the proof of Corollary 15.1 yields a pole α at which all N_1 fiber slopes are distinct. Thus, at the two MCA frontier agreements, a single received line carries at least N_1 aperiodically witnessed MCA-bad finite slopes. These agreements lie strictly inside the unnormalized band of Remark 9.11; this is the deployed obstruction that forces the normalized band problem Remark 9.12.

Proof. The two comparisons $\binom{n}{m} > p^w \lfloor \varepsilon^* q \rfloor$ of Proposition 15.2 give the stated N_1 values. The two zero-collision inequalities were included in the exact certificate of Proposition 15.2. Hence some pole has no collisions among the fiber evaluations, and Lemma 18.2 makes each value an MCA-bad slope with witnesses inside the fiber. Since both frontier agreements are odd while n is a power of two, the witnesses are aperiodic. □

Corollary 18.6 (no uniform polynomial bound below the envelope)**.** *For $\nu \to \infty$ set $n = 2^\nu$, $k = n/2$, let p_ν be the least prime $\equiv 1 \pmod n$ (so $\beta_\nu := \log_2 p_\nu = O(\log n)$ by Linnik's theorem [IK04]), $B_\nu := \mathbb{F}_{p_\nu}$, $D_\nu := \mu_n \subseteq B_\nu^\times$; let m_ν be an odd integer with $m_\nu - k \in \left[n/(8\beta_\nu),\, n/(4\beta_\nu)\right]$, and let $\mathbb{F}_\nu := \mathbb{F}_{p_\nu^{d_\nu}}$ with d_ν minimal such that $q_\nu \ge k \cdot 2^{n/8}$. Then for all large ν:*

(a) $N_{\mathrm{fib}} \ge 2^{n/2} \ge \lceil (q_\nu - n)/k \rceil$, so $N_0 = \lceil (q_\nu - n)/k \rceil$;

(b) some received line of the rate-1/2 multiplicative row has at least $2^{n/8-2}$ MCA-bad finite slopes at agreement threshold m_ν, all of whose witness supports are aperiodic;

(c) $a = m_\nu$ lies in the band of Remark 9.11 for every smoothness convention with $2n/N \le n/(16\beta_\nu)$, in particular for the finest scale $N = n$.

Consequently no polynomial $n \mapsto P(n)$, uniform over this family (and a fortiori none independent of q), bounds the number of aperiodically witnessed MCA-bad finite slopes of a received line on the band of Remark 9.11: the conjecture is false as stated below the entropy–subfield envelope.

Proof. Write $g := (m_\nu - k)/n \le 1/(4\beta_\nu) \le 1/4$. The elementary bound $H_2(\frac{1}{2}+x) \ge 1 - 4x^2$ on $[0, \frac{1}{4}]$ (the difference vanishes at $x = 0$ and has positive derivative $8x - \log_2 \frac{1+2x}{1-2x}$ there) gives

$$\begin{aligned}
\log_2 N_{\mathrm{fib}} &\ge n H_2(m_\nu/n) - \log_2(n+1) - w\beta_\nu \\
&\ge n(1-4g^2) - ng\beta_\nu - \log_2(n+1) \\
&\ge n\left(1 - \frac{1}{4\beta_\nu^2} - \frac{1}{4}\right) - \log_2(n+1) \\
&\ge \frac{n}{2}.
\end{aligned}$$

while $\lceil (q_\nu - n)/k \rceil \le q_\nu/k \le p_\nu 2^{n/8} = 2^{n/8+O(\log n)} < 2^{n/2}$ for large ν — giving (a). Then Theorem 18.4 (with $\gcd(m_\nu, 2^\nu) = 1$ since m_ν is odd) yields at least $\frac{1}{2}\lceil (q_\nu - n)/k \rceil (1 - k/(q_\nu - n)) \ge (q_\nu - n - k)/(2k) \ge 2^{n/8-2}$ slopes for large ν — (b). For (c), $m_\nu - k \ge n/(8\beta_\nu) > 2n/N$ under the stated convention, and $m_\nu < n - \lfloor (n-k)/3 \rfloor$ since $m_\nu - k \le n/(4\beta_\nu) = o(n)$ while $n - \lfloor (n-k)/3 \rfloor - k = n/2 - \lfloor n/6 \rfloor = \Theta(n)$. The count in (b) is $2^{\Omega(n)}$, which exceeds every polynomial in n. □

Remark 18.7 (the band conjecture after the floors). Corollary 18.6 refutes Remark 9.11 as stated: no polynomial P, uniform over the row family and independent of q, can bound the aperiodically witnessed slope counts on the sub-envelope part of the band. At the deployed size the refutation is quantitative rather than formal — any P certifying the band at $n = 2^{21}$ must satisfy $P(2^{21}) > 2^{164}$. For the coefficient-normalized monomial model $P(n) = n^C$ this would require $C \ge 8$; no degree lower bound follows for an arbitrary polynomial without controlling its coefficients. Either way, a value above 2^{164} already destroys the intended consumption: Theorems 8.2 and 8.3 would inherit an aperiodic numerator of at least 2^{164} at every band agreement, far above both challenge budgets, so the conditional theorems would certify nothing. Accordingly:

(i) the operative band input is Remark 9.12, now by necessity rather than by convention; its left-edge reserve $R(n)$ absorbs exactly the identity-prefix floors proved here, and the sentence of Section E.1 assigning the sub-band $A \lesssim k + g^* n$ to the unsafe side is a theorem (Corollaries 18.5 and 18.6);

(ii) version 13.2 has replaced every operative use of the old proposal in Theorems 8.2 and 8.3 and the closing statements by Remark 9.12; the label Remark 9.11 is retained only for the historical statement refuted here, and the finite deployed target is the new adjacent value of Corollary 15.3;

(iii) the flatness form of the band input is Remark E.2, whose base-field normalization $|\mathbb{B}|^{-s}$ is confirmed — rather than refuted — by the floors: the first term of its model crosses one exactly at the envelope, which is where the floors stop.

Remark 18.8 (even m: the two list rows). At the deployed list rows m is even with $\gcd(m, n) = 2$, so Theorem 18.4 does not apply verbatim. The witness classification (Lemma 18.2) still confines all threshold-m witnesses to a single fiber, and a 2-periodic member S (stabilizer $\{\pm 1\}$, $S = -S$) has $\Lambda_S(X) = G(X^2)$ up to the coset twist, so its prefix vanishes at every odd index $h \le w$. Hence either the pigeonholed z^* has some nonzero odd-index coordinate — and then the fiber contains *no* periodic support and the floor is purely aperiodic exactly as before — or z^* is even-supported, in which case the periodic members of the fiber are precisely the scale-2 pullbacks of an image fiber at half parameters and are charged to the periodic image ledger (Theorem 9.4 and Corollary 9.3); the aperiodic residue of the fiber remains subject to the same collision count. We do not claim a floor for the list rows here; only that periodic contamination of identity-scale fibers is confined to even-supported prefixes and is quotient-borne.

Remark 18.9 (relation to the conditional theorems). Corollaries 18.5 and 18.6 do not touch the unconditional results (Theorems 2.6, 9.6 and 10.1); they show why the hypothesis consumed by Theorems 8.2 and 8.3 is the normalized Remark 9.12; the original Remark 9.11 is retained only as a refuted historical proposal (see Remark 18.7). The deployed conditional rows already sit at the envelope frontier with the recorded margins, so the substitution is a restatement of what those rows in fact require, not a weakening; see the final package in Section 21.

19 Prefix-collision rigidity and exact second moments

The floors of Section 18 and the closing accounting of Remarks 14.18 and F.31 both turn on the sizes of the prefix fibers $\mathrm{Fib}_w(z)$. Remark E.8 names, as a low-risk Q1 deliverable, the sharpening of the kernel side-distance $\delta_w = \lceil (w+1)/2 \rceil$ of Lemma E.6 and Proposition F.6 to $e \geq w+1$ by locator subtraction, together with an exact collision ledger built on it. This section pays that deliverable: the side-distance $w+1$ is proved characteristic-free, the packing cap of Proposition E.7 and Corollary F.7 is replaced with the packing exponent doubled from $\delta_w - 1 \approx w/2$ to w, and the second moment is stratified exactly with strata starting at $e = w+1$ instead of δ_w and with the top stratum identified.

Lemma 19.1 (prefix-collision rigidity)**.** *Let $D \subseteq \mathbb{F}$, $|D| = n$, and $1 \leq w < m \leq n$. If $M \neq M'$ are m-subsets of D with $\Phi_w(M) = \Phi_w(M')$, then $|M \setminus M'| = |M' \setminus M| \geq w+1$. Moreover if $|M \setminus M'| = w+1$ then $\ell_{M \setminus M'} - \ell_{M' \setminus M}$ is a nonzero constant (a* constant-shift split pair*).*

Proof. Let $R := M \cap M'$, $A := \ell_{M \setminus M'}$, $B := \ell_{M' \setminus M}$, $e := |M \setminus M'| = |M' \setminus M|$, and $G := \ell_R$ of degree $m - e$. Then $\Lambda_M - \Lambda_{M'} = G\,(A - B)$. Prefix equality (together with monicity) means $\deg(\Lambda_M - \Lambda_{M'}) \leq m - w - 1$, hence $\deg(A - B) \leq e - w - 1$. If $e \leq w$ this forces $A = B$, hence $M = M'$ (disjoint nonempty root sets cannot coincide unless $e = 0$), a contradiction; so $e \geq w+1$. If $e = w+1$ then $\deg(A - B) \leq 0$ with $A \neq B$ (their root sets are disjoint and nonempty), so $A - B$ is a nonzero constant. □

Remark 19.2 (relation to the moment-kernel bounds). Lemma 19.1 strengthens the fiber statements of Lemma E.6 and Proposition F.6(b) in two directions at once. First, the per-side distance doubles: the moment-kernel argument bounds only the *support* of $\mathbf{1}_M - \mathbf{1}_{M'} \in K_w(D)$, giving $|M \triangle M'| \geq w+1$, i.e. $e \geq \delta_w$ per side, whereas locator subtraction uses the multiplicative structure of the balanced word and gives $e \geq w+1$ per side, i.e. $|M \triangle M'| \geq 2(w+1)$. Equivalently: the *balanced-ternary* minimum weight of the moment kernel $K_w(D)$ is at least $2(w+1)$, twice its BCH-style minimum weight $w+1$ — a statement about $K_w(D)$ itself, relevant to Proposition F.6(d). Second, the hypothesis weakens: the locator argument never invokes Newton identities, so it is valid in every characteristic and at every depth w, whereas the power-sum transfer requires $w < \mathrm{char}\,\mathbb{F}$ (harmless at the deployed rows, $w \approx 6.7 \times 10^4 < 2^{31}$, but not free in general).

Corollary 19.3 (unconditional anticode fiber cap; replaces the packing bounds)**.** *For every $z \in B^w$,*

$$|\mathrm{Fib}_w(z)| \;\leq\; \binom{n}{m-w} \Big/ \binom{m}{w}.$$

This replaces Proposition E.7 and Corollary F.7, whose packing parameter is $\delta_w - 1$: at the deployed frontier rows the saving over the trivial bound $\binom{n}{m}$, in bits, improves as follows (orientation values; first line the four rows of Corollary F.7, second line the moved MCA rows of Proposition 15.2):

row	(m, w)	*saved, Corollary F.7*	*saved, this cap*	*gap to conjectured average*
KoalaBear list	$(1116046, 67470)$	213510	361177	1729629
Mersenne-31 list	$(1116022, 67446)$	213453	361085	1729741
KoalaBear MCA	$(1116047, 67470)$	213510	361177	1729629
Mersenne-31 MCA	$(1116023, 67446)$	213453	361085	1729741

The cap thus removes about 17% *of the trivial-to-average gap, roughly* 1.48×10^5 *bits more per row than the* δ_w*-packing; the remaining* $\approx 1.73 \times 10^6$ *bits per row are the flatness targets of Remarks F.2 and E.25.*

Proof. By Lemma 19.1 the fiber is a family of m-subsets with pairwise $|M \setminus M'| > w$; apply Proposition B.18 with $s = w$. The comparison numbers are $\log_2 \binom{n}{m} - \log_2\big[\binom{n}{m-s}/\binom{m}{s}\big]$ at $s = w$ versus $s = \delta_w - 1$, evaluated by the entropy sandwich as in Corollary F.7. □

Remark 19.4 (calibration at the deployed rows). Average fibers (bits): 191.61, 67.10, 114.30, 52.11 at the four rows of Proposition 14.2, and 66.91, 51.93 at the moved MCA edges; at the new safe-side rungs $a_0 + 1$ the averages are 35.7 (KoalaBear MCA) and 20.7 (Mersenne-31 MCA). Neither the δ_w-packing nor this cap decides any adjacent pair — the caps sit $\approx 1.73 \times 10^6$ bits above average — but the ledger they anchor is exact: Theorem 19.5 expresses the entire L^2 mass in closed form, so the flatness inputs of Remarks F.3 and E.25 can now be posed about a single explicitly named quantity (the shift-pair counts below) rather than about an unstructured second moment.

Theorem 19.5 (exact second-moment stratification)**.** *Let* $N_w(z) := |\mathrm{Fib}_w(z)|$. *Then*

$$\sum_{z \in B^w} N_w(z)^2 \;=\; \binom{n}{m} \;+\; \sum_{e=w+1}^{\min(m,\,n-m)} P_e, \qquad P_e \;=\; \sum_{\substack{R \subseteq D \\ |R| = m-e}} \mathrm{sp}_w\big(e;\, D \setminus R\big),$$

where for $D' \subseteq D$

$$\mathrm{sp}_w(e; D') := \#\left\{(A,B):\ \begin{array}{l} A, B \text{ are monic of degree } e \text{ and split over } D', \\ \text{their root sets are disjoint, and } \deg(A-B) \le e-w-1 \end{array}\right\}$$

counts ordered depth-w shift pairs. *In particular* $P_e = 0$ *for* $1 \le e \le w$, *and the top stratum is*

$$\mathrm{sp}_w(w+1; D') = \#\left\{(A,c):\ \begin{array}{l} A \text{ is monic of degree } w+1 \text{ and split over } D', \\ c \in \mathbb{F}^\times, \quad A - c \text{ splits over } D' \end{array}\right\}.$$

Proof. The left side counts ordered pairs (M, M') with equal prefixes. The diagonal contributes $\binom{n}{m}$. For $M \neq M'$ the map $(M, M') \mapsto (R, A, B) := (M \cap M',\, \ell_{M \setminus M'},\, \ell_{M' \setminus M})$ is a bijection onto triples with R an $(m-e)$-subset and (A,B) monic split of degree e over $D \setminus R$ with disjoint root sets, for the unique $e = |M \setminus M'|$; and as in the proof of Lemma 19.1, prefix equality is equivalent to $\deg(A-B) \le e-w-1$. Rigidity kills $1 \le e \le w$; e cannot exceed m nor $n-m$. At $e = w+1$ the condition reads $A - B = c \in \mathbb{F}^\times$ (nonzero since disjoint nonempty root sets force $A \neq B$); conversely for any split A and $c \neq 0$ with $A - c$ split, the root sets are automatically disjoint (a common root would force $c = 0$). □

Remark 19.6 (relation to the balanced-ternary collision ledger, and the exact Γ_2). Theorem 19.5 refines Proposition E.4 and Proposition F.6(c) in three ways. (i) The collision strata are proved *empty* for $\delta_w \le e \le w$: in the notation there, $T_w(e) = 0$ for all $e \le w$ (for $w < \mathrm{char}\,\mathbb{F}$ via Newton; the elementary-prefix statement is characteristic-free), so the ledger starts at $e = w+1$. (ii) The top stratum $e = w+1$ is identified: its pairs are exactly the constant-shift split pairs $(A,\ A-c)$, $c \in \mathbb{B}^\times$, both split over D — an explicitly algebraic family — so the leading correction to fiber

flatness is a single named census. (iii) In the notation of Remark F.3, the exact collision-gap value is

$$\Gamma_2(m,w) \;=\; \frac{\binom{n}{m}+\sum_{e\ge w+1}P_e}{\binom{n}{m}^2\,|\mathbb{B}|^{-w}} \;=\; \frac{1}{\bar F}+\frac{\sum_{e\ge w+1}P_e}{\binom{n}{m}\,\bar F}, \qquad \bar F := \binom{n}{m}|\mathbb{B}|^{-w},$$

so in the dense bulk ($\bar F \gg 1$) the collision gap *is* the shift-pair ledger normalized by $\binom{n}{m}\bar F$: the exact $r=2$ input to Remark F.3 and extends the anchor Proposition F.5 beyond the head depths as an identity (though not yet as a bound). A proof of Remark 19.8 below therefore bounds Γ_2 exactly, and, by Proposition F.26, fixed-r analogues of the same stratification are the asymptotic form of (Q).

Remark 19.7 (quotient prototype and the primitive part). Let $D' \supseteq \mu_n$ contain the full group and suppose $(w+1) \mid n$. Then for every $\beta \in \mu_{n/(w+1)}$ the binomial $X^{w+1}-\beta$ is split over μ_n (its root set is a full fiber of the $(w+1)$-power map), and any two distinct $\beta,\beta' \in \mu_{n/(w+1)}$ give a shift pair with $c=\beta-\beta'$; hence $\mathrm{sp}_w(w+1;\mu_n) \ge \frac{n}{w+1}\big(\frac{n}{w+1}-1\big)$. These pairs are exactly scale-$(w+1)$ pullbacks: both root sets are unions of complete fibers, i.e. periodic supports in the sense of Definition 9.1, so this mass belongs to the quotient ledger (Theorem 9.4). When $(w+1)\nmid n$ the binomial prototype is unavailable and no pullback pair of scale $w+1$ exists; shift pairs may still exist at composite scales dividing gcd-data of (e,n), and are pullbacks whenever both root sets are periodic at a common scale $c\ge 2$. We call a shift pair *primitive* if it is not a pullback at any scale $c\ge 2$.

Remark 19.8 (Primitive shift-pair ledger). For multiplicative cosets D and the ranges of Proposition 14.2, bound the primitive part of $\mathrm{sp}_w(e;D')$ for $D'=D\setminus R$: show it is dominated, up to a factor $\mathrm{poly}(n)$, by the pullback part plus the diagonal-scale term implicit in $\binom{n}{m}$. Via Theorem 19.5 and Cauchy–Schwarz this is the exact combinatorial content of second-moment flatness; it refines Lemma B.8 and Remark A.21 to the support side, and is the support-level counterpart of the alignment moment calculus (Theorems B.1 and B.4 and Corollary B.5; scope as in Remark B.7).

Remark 19.9 (flatness wiring after capg). The first-pass capg flatness problem is withdrawn as a separate item: its content is covered by Remark E.2 (asymptotic, base-field-normalized form) and Remark E.25 (finite extremality form), and this part's contribution to them is structural:

(i) the collision side of Remark E.25 now has the exact ledger of Theorem 19.5, whose sole unbounded ingredient is Remark 19.8; a proof of the latter with constant κ' bounds every fiber's excess over average via Γ_2 and Cauchy–Schwarz, though, per Remark C.7 and Proposition F.26, second-moment control alone cannot decide the printed pairs;

(ii) the finite targets sit at the *new* rungs $a_0+1 = 1116048, 1116024$ with next-row deficits exceeding 22.1 and 3.2 bits (Corollary 15.3), and the rung audit Remark E.26 is a prerequisite there exactly as before;

(iii) nothing in this part bounds $\mathrm{sp}_w(e)$ for the primitive part; that residual is named in Remark 19.8 and is a strict sub-problem of the mode-at-null input.

20 Subfield floors for split-pencil censuses

The aperiodic side of the safe programme has been reduced, in Section E.4 and the final-form subsections of Section F, to census problems whose conjectured models are normalized at the

challenge-field scale: Remarks E.11 and F.13 with $\max\bigl(1, \binom{n}{m} q^{-(w'-1)}\bigr)$, Remarks F.17 and E.18 with $\max\bigl(1, \binom{n}{m} q^{-w'}\bigr)$, and Remarks E.22 and F.22 with $\max\bigl(1, \binom{n}{\omega} q^{1-w'}\bigr)$, where in this section K is the code dimension of the census ($K = k$ for the deployed rows), $w' := m - K$, and $\omega = n - m$. These normalizations are calibrated to challenge-field-generic words and lines. The identity-prefix mechanism, however, produces *base-field* instances at every balanced profile, and their censuses sit at the $|\mathbb{B}|$-scale, which is larger than the q-scale by a factor exponential in n whenever $[\mathbb{F} : \mathbb{B}] \geq 2$. As stated, all three problems are therefore refuted; the corrections are recorded at the end of the section, and they align the census problems with the base-field normalization that Remarks E.2 and F.2 already carry.

Proposition 20.1 (subfield census floors at every balanced profile)**.** *Let $C = \mathrm{RS}[\mathbb{F}, D, K]$, $D \subseteq \mathbb{B}$, $|D| = n$, $q = |\mathbb{F}|$, $p = |\mathbb{B}|$, and fix an agreement m with $K \leq m$ and $2(m-K) \leq n-K$; write $w' = m - K$. For a profile parameter d_1 with $w' + 1 \leq d_1 \leq \lfloor (n - K + 1)/2 \rfloor$ put $m' := K - 1 + d_1$, assume $m' + d_1 \leq n$, and define*

$$A_{\mathbb{B}}(d_1) := \binom{n}{m'} p^{-(d_1-1)},$$
$$M_{\mathbb{B}}^{\mathrm{disc}}(d_1) := \binom{m'}{m} \lceil A_{\mathbb{B}}(d_1) \rceil,$$
$$M_{\mathbb{B}}^{\mathrm{soft}}(d_1) := \binom{m'}{m} \bigl(1 + A_{\mathbb{B}}(d_1)\bigr).$$

Then $M_{\mathbb{B}}^{\mathrm{disc}}(d_1) \leq M_{\mathbb{B}}^{\mathrm{soft}}(d_1)$, and:

(a) (Boundary profile, $d_1 = w' + 1$.) *For every $z \in \mathbb{B}^{w'}$ the witness word U_z of Lemma 14.1 has minimal shifted lattice degree $d_1(U_z) = w' + 1$, and $\mathrm{cen}(U_z; m) = |\Phi_{m,w'}^{-1}(z)|$; at the heaviest prefix,*

$$\mathrm{cen}(U_{z^*}; m) \;\geq\; \left\lceil \binom{n}{m} p^{-w'} \right\rceil.$$

(b) (Interior profiles, $d_1 \geq w' + 2$.) *For every $z' \in \mathbb{B}^{d_1 - 1}$ the level-m' witness word $U_{z'}^{(m')}$ has $d_1\bigl(U_{z'}^{(m')}\bigr) = d_1$ and is a balanced determinantal representation datum in the sense of Propositions E.21 and F.21. If z'^* indexes a heaviest level-m' prefix fiber, then its size-m census satisfies*

$$\mathrm{cen}\bigl(U_{z'^*}^{(m')}; m\bigr) \;\geq\; \binom{m'}{m} \left\lceil \binom{n}{m'} p^{-(d_1-1)} \right\rceil \;=\; M_{\mathbb{B}}^{\mathrm{disc}}(d_1).$$

At $d_1 = w' + 1$ the same formula reproduces (a).

(c) (Per-line R1 floors.) *Let $w' \geq 1$, $m \geq K + 1$, and let (f_α, g_α) be the pole line of U_{z^*} at any $\alpha \in \mathbb{F} \setminus (D \cup \mathbb{B})$. Then every member T of the heaviest fiber is a non-common rank-one support of (f_α, g_α) in the sense of Proposition E.10, so*

$$\#\mathrm{R1}(f_\alpha, g_\alpha; m) \;\geq\; \left\lceil \binom{n}{m} p^{-w'} \right\rceil,$$

and the line is not $\mathbb{B}$-rational, hence outside the subfield-confined stratum of Lemma 6.1 cited in Remark F.13.

Consequently, along the asymptotic family of Corollary 18.6 (where $\binom{n}{m}p^{-w'} \geq 2^{n/2-1}$ *while* $\binom{n}{m}q^{-(w'-1)} \leq 1$*), any final census model omitting the base-field floor is too small. At the deployed list rows the missing floor contributes, in bits,* 67.1 *(KoalaBear) and* 52.1 *(Mersenne-31) at the boundary profile, and* 56.0, 43.9, 57.7 *(KoalaBear,* $d_1 = w'+2, w'+3, w'+4$*), resp.* 41.0, 39.2, 57.7 *(Mersenne-31), at the first interior profiles, against challenge-field models below one. The rise at the last displayed profile is the discrete effect of one guaranteed nonempty level-*m' *prefix contributing all* $\binom{m'}{m}$ *of its size-*m *supports.*

Proof. First, $M_{\mathbb{B}}^{\mathrm{disc}} \leq M_{\mathbb{B}}^{\mathrm{soft}}$ follows from $\lceil x \rceil \leq x+1$ for $x \geq 0$.

(a) At the boundary $m' = m$ and the standing hypothesis $m' + d_1 \leq n$ gives $m + w' + 1 \leq n$, hence $m + w' < n$. Suppose a nonzero $(W, N) \in M_{U_z}$ had shifted degree at most $w' = d_1 - 1$. Then $\deg W \leq w'$, $\deg N \leq K + w' - 1 = m - 1$, and $\deg(WP_z - N) \leq m + w' < n$. This polynomial vanishes on D, so it is zero. If $W = 0$ then $N = 0$; otherwise $N = WP_z$ has degree at least m, a contradiction. Thus $d_1(U_z) = w' + 1$, and Theorem F.23 gives the census identity. The fiber lower bound is the pigeonhole of Lemma 14.1.

(b) $U_{z'}^{(m')}$ is defined by the same recipe at level m': its interpolant is the monic polynomial $P_{z'}^{(m')} = X^{m'} + \sum_{j \leq d_1 - 1} z'_j X^{m'-j}$ of degree $m' < n$. The vector $(1, P_{z'}^{(m')})$ has shifted degree $m' - K + 1 = d_1$. No nonzero lattice vector (W, N) can have shifted degree at most $d_1 - 1$: then $\deg W \leq d_1 - 1$, $\deg N \leq K + d_1 - 2 = m' - 1$, and $\deg(WP_{z'}^{(m')} - N) < n$ by $m' + d_1 \leq n$. This polynomial vanishes on D, hence is zero. If $W = 0$ then $N = 0$; otherwise $N = WP_{z'}^{(m')}$ has degree at least m', contradicting $\deg N \leq m' - 1$. Thus $d_1(U_{z'}^{(m')}) = d_1$. The reduced-basis computation of Theorem F.23 then gives $\big((1, P_{z'}^{(m')}), g_2\big)$, which is a coprime unimodular quadruple by Lemma F.20; this is the claimed determinantal datum. For the census: by Proposition 14.13 at level m', each member M' of the level-m' fiber of z' yields a codeword $c_{M'}$ of $\mathrm{RS}[\mathbb{F}, D, K]$ agreeing with $U_{z'}^{(m')}$ on exactly the m' points of M', and distinct M' give distinct codewords. Each such codeword contributes $\binom{m'}{m}$ valid size-m supports to the binomial-moment formula of Proposition E.15(c), all distinct across distinct M' since the codeword attached to a valid support is unique ($m \geq K$). The heaviest level-m' fiber has size at least $\lceil \binom{n}{m'} p^{-(d_1-1)} \rceil$ by pigeonhole.

(c) By Lemma 18.2(i) each fiber member T carries the single close slope $\zeta_T = U_{z^*}(\alpha) - \Lambda_T(\alpha)$, so by Proposition E.10 (or directly Lemma E.9) $\alpha(T) + \zeta_T \beta(T) = 0$. The support is non-common: $\beta(T) = 0$ would mean g_α extends on T to a polynomial G of degree $< K$, whence $(X - \alpha)G + 1$, of degree at most $K \leq m - 1 < |T|$, vanishes on T, forcing the zero polynomial, which contradicts its value 1 at α. Thus $\alpha(T), \beta(T)$ are proportional with $\beta(T) \neq 0$: a non-common rank-one support, and distinct T are distinct supports. Non-$\mathbb{B}$-rationality is projective, not merely affine: if λg_α were $\mathbb{B}$-valued for some $\lambda \in \mathbb{F}^\times$, the ratio at two distinct points of D would force $\alpha \in \mathbb{B}$ by the calculation in Corollary 2.12.

The deployed violation values are orientation computations of $\log_2 M_{\mathbb{B}}^{\mathrm{disc}}(d_1)$ at $(K, m) = (k, 1116046)$ resp. $(k, 1116022)$; the asymptotic claim substitutes the family parameters of Corollary 18.6, for which $q \geq k \cdot 2^{n/8}$ makes every challenge-field model term at most one while the base-field floor is at least $2^{n/2-1}$ at the boundary and at least $2^{n/4}$ at the bounded-offset interior profiles. □

Remark 20.2 (boundary profile and active split-pencil range). Theorem F.23 states that (Q) is the boundary profile $d_1 = w' + 1$ of the split pencil and the aperiodic balanced core consists of *interior* profiles; but Remarks F.22 and E.22 quantify over "balanced profile $d_1 \geq w' + 1$",

which includes the boundary and hence includes (Q) itself with the wrong (q-scale) model. Part (a) above identifies the boundary contribution. The intended reading — boundary delegated to (Q), split-pencil census over $d_1 \geq w' + 2$ — is adopted in the restatement below; part (b) shows that the interior model must also be base-field-normalized.

Remark 20.3 (Base-field-normalized split-pencil model). Let (W_1, N_1, W_2, N_2) be a determinantal representation of Λ_D as in Proposition E.21 with interior balanced profile

$$w' + 2 \leq d_1 \leq \lfloor (n - K + 1)/2 \rfloor, \qquad m' = K - 1 + d_1.$$

Write

$$A_{\mathbb{B}}(d_1) = \binom{n}{m'} |\mathbb{B}|^{-(d_1-1)}, \quad M_{\mathbb{B}}^{\mathrm{disc}}(d_1) = \binom{m'}{m} \lceil A_{\mathbb{B}}(d_1) \rceil, \quad M_{\mathbb{B}}^{\mathrm{soft}}(d_1) = \binom{m'}{m} \big(1 + A_{\mathbb{B}}(d_1)\big).$$

The first quantity after $A_{\mathbb{B}}$ is the unavoidable discrete lower floor of Proposition 20.1; the second is a smooth upper baseline satisfying $M_{\mathbb{B}}^{\mathrm{disc}} \leq M_{\mathbb{B}}^{\mathrm{soft}}$. The corrected target census is

$$\begin{aligned} &\#\big\{(A, B) : \ \deg A \leq \omega - d_1, \ \deg B \leq \omega - d_2, \\ &\qquad\qquad AW_1 + BW_2 \mid \Lambda_D, \ \deg(AW_1 + BW_2) = \omega\big\} \\ &\leq L_{\mathrm{BC}}(n, d_1) \max\left\{ M_{\mathbb{B}}^{\mathrm{soft}}(d_1), \binom{n}{\omega} q^{1-w'} \right\} + E_{\mathrm{BC}}(n, d_1), \end{aligned}$$

counting monic normalizations. Applying a ceiling only after multiplying the average prefix-fiber size is invalid: once $A_{\mathbb{B}}(d_1) < 1$, one nonempty level-m' prefix still contributes all $\binom{m'}{m}$ supports. In the reserve regime one requires $R \to \infty$, $R = o(n)$, and combined logarithmic loss $o(R \log |\mathbb{B}|)$. In a finite deployed certificate, L_{BC} and E_{BC} are explicit and are summed with every other bucket before comparison with the budget.

The same correction applies *mutatis mutandis* to Remarks F.17 and E.18: the interior subfield model is $M_{\mathbb{B}}^{\mathrm{soft}}(d_1)$ together with the challenge-field generic term, while $M_{\mathbb{B}}^{\mathrm{disc}}(d_1)$ is the unavoidable lower floor. For Remarks E.11 and F.13, the per-line model remains

$$\max\left\{ 1, \binom{n}{m} |\mathbb{B}|^{-w'}, \binom{n}{m} q^{-(w'-1)} \right\},$$

with the middle term achieved by the non-$\mathbb{B}$-rational pole lines of Proposition 20.1(c). These models are consistent with the master mean-plus-one target Remark E.2 and with the normalized band Remark 9.12.

Remark 20.4 (what the subfield floors do not say). The floors identify the correct model; they do not weaken the programme. In particular: (i) the reductions Propositions E.10, E.21 and F.21, Theorems E.16, F.12 and F.15, and Lemmas E.19 and F.19 are unconditional and untouched — only the conjectured *sizes* change; (ii) the per-line slope counts remain bounded by the per-line census, and on the certifying pole lines the slope count and the support count agree to within the collision correction of Theorem 18.4, so the corrected R1 model is tight at the floors; (iii) for genuinely challenge-field-primitive representations — no member of the pencil proportional to a $\mathbb{B}$-rational datum in the sense of Lemma 6.1 — the original q-scale models remain the plausible targets, and separating that stratum cleanly is part of Remark 20.3.

21 Frontier upper ledger and remaining inputs

The unsafe side is theorem-level: the identity-prefix floor, the list certificates of Proposition 14.2, and the flexible-budget MCA certificates of Proposition 15.2 supply the lower staircase. The adjacent safe rows remain conditional. This section states exactly what their upper certificate must contain.

Definition 21.1 (frontier budget and summed upper ledger)**.** Fix a row, a sampled-parameter denominator Q, target ε^*, agreement a, and budget

$$B_* := \lfloor \varepsilon^* Q \rfloor .$$

For a received line ℓ, let $N(\ell, a)$ be its number of bad sampled parameters at agreement at least a, and put $N(a) := \max_\ell N(\ell, a)$. A complete certificate must print an integer $U_{\mathrm{paid}}(a)$ together with theorem handles covering the tangent, already-coalesced quotient, extension-lift, common-factor, fixed-dimensional, and other discharged cells; this manuscript does not aggregate those handles into one finished numerator.

A proposed certificate must supply an explicit priority map assigning every unresolved cell to three disjoint classes: boundary prefix/quotient-pullback cells, interior balanced split-pencil cells, and primitive same-prefix cells not assigned to a pullback scale. Let their line-wise counts be $N_Q(\ell, a)$, $N_{\mathrm{BC}}(\ell, a)$, and $N_{\mathrm{sp}}(\ell, a)$. A valid frontier certificate must prove both the coverage statement

$$N(\ell, a) \le U_{\mathrm{paid}}(a) + N_Q(\ell, a) + N_{\mathrm{BC}}(\ell, a) + N_{\mathrm{sp}}(\ell, a)$$

for every line and explicit bounds $N_Q \le U_Q$, $N_{\mathrm{BC}} \le U_{\mathrm{BC}}$, $N_{\mathrm{sp}} \le U_{\mathrm{sp}}$. Thus

$$N(a) \le U_{\mathrm{total}}(a) := U_{\mathrm{paid}}(a) + U_Q(a) + U_{\mathrm{BC}}(a) + U_{\mathrm{sp}}(a).$$

Coverage is part of the inputs below; a size estimate without the corresponding bucket assignment is not a closure certificate.

Residual Input 21.2 (structured quotient and prefix pullbacks)**.** For every boundary prefix or quotient cell ν appearing in the compiler, with local parameters $(n_\nu, m_\nu, w_\nu, \mathbb{B}_\nu)$, choose a canonical support in each prefix fiber and let $P_\nu(M)$ be the number of other members related to it by a declared nontrivial pullback scale. Prove the cell-wise bound

$$1 + P_\nu(M) \le \kappa_{Q,\nu} \left(1 + \binom{n_\nu}{m_\nu} |\mathbb{B}_\nu|^{-w_\nu} \right) + E_{Q,\nu}$$

uniformly in M, together with the parameter-image coalescing that makes the sum of these cell bounds a valid $U_Q(a)$. The additive 1 is mandatory: above the crossing, a nonempty fiber can have average below one. In the finite form every $\kappa_{Q,\nu}$ and $E_{Q,\nu}$ is an explicit integer or rational interval contribution.

Residual Input 21.3 (base-field-normalized interior split pencils)**.** For every interior balanced profile $d_1 \ge w' + 2$, prove the coverage reduction to the determinantal census of Remark 20.3 and the explicit bound

$$C_{\mathrm{BC}}(d_1) \le \kappa_{\mathrm{BC},d_1} \max \left\{ M_{\mathbb{B}}^{\mathrm{soft}}(d_1), \binom{n}{\omega} q^{1-w'} \right\} + E_{\mathrm{BC},d_1}.$$

The upper numerator $U_{\mathrm{BC}}(a)$ is the sum over the compiler's interior profiles after all overlaps already paid by the priority rule are removed. The boundary profile $d_1 = w' + 1$ belongs to Residual Input 21.2, not to this input.

Residual Input 21.4 (local primitive shift-pair control)**.** For an m-subset $M \subseteq D$, let

$$\Delta^{\mathrm{prim}}_{w,e}(M) := \#\left\{ M' : \begin{array}{l} |M \setminus M'| = e,\ \Phi_w(M') = \Phi_w(M), \\ (M, M') \text{ is not a pullback at any common scale } c \geq 2 \end{array} \right\}.$$

Prove a uniform local bound

$$\sum_{e=w+1}^{\min(m,n-m)} \Delta^{\mathrm{prim}}_{w,e}(M) \leq \kappa_{\mathrm{sp}} \left(1 + \binom{n}{m} |\mathbb{B}|^{-w}\right) + E_{\mathrm{sp}}$$

for every relevant M, and prove that the remaining primitive prefix cells of the compiler inject into these local neighbours. This local statement is stronger than a bound for the global second moment: Theorem 19.5 gives the sum of the local degrees over M, whereas a max-fiber and per-line certificate requires their maximum. The finite form must give explicit constants and an exact $U_{\mathrm{sp}}(a)$.

Remark 21.5 (conditional closure template with one budget). Assume that a future certificate supplies: an explicit cell set and priority map; a proved coverage inequality of Definition 21.1; support-to-parameter coalescing for every bucket; exact integers $U_{\mathrm{paid}}, U_Q, U_{\mathrm{BC}}, U_{\mathrm{sp}}$; and the bounds requested in Residual Inputs 21.2 to 21.4. Under those assumptions the following bookkeeping implications hold.

(i) Let $R(n) \to \infty$ with $R(n) = o(n)$, and let $L(n)$ dominate every multiplicative loss, normalized additive error, and residual-cell count. If the normalized band begins at $k + g^* n + R(n)$ and

$$\log L(n) = o\big(R(n) \log |\mathbb{B}|\big),$$

then the supplied upper certificate matches the entropy–subfield frontier with reserve $R(n)$. In particular, fixed polynomial losses are absorbed by $R \gg \log n$; the phrase "$e^{o(n)}$ loss" alone is insufficient.

(ii) For a finite adjacent row with proved unsafe agreement a_0, lower certificate $L(a_0) > B_*$, and a fully instantiated upper certificate, the single comparison

$$\boxed{U_{\mathrm{paid}}(a_0+1) + U_Q(a_0+1) + U_{\mathrm{BC}}(a_0+1) + U_{\mathrm{sp}}(a_0+1) \leq B_* < L(a_0)}$$

forces the first safe agreement to be $a_0 + 1$ by Theorem 13.2 and Proposition 8.1. The row margins of Corollary 15.3 are room for this sum, not separate allowances.

This remark is a closure template, not a theorem that the required priority map, coverage, coalescing, or upper numerators exist.

The missing certificate data are therefore testable. One must define and prove exhaustive bucket coverage and support-to-parameter coalescing, then prove structured pullback coalescing with the mean-plus-one floor, prove the base-field-normalized interior split-pencil census, prove a *local* primitive shift-pair degree bound rather than only a global moment estimate, and print the resulting summed numerator at each candidate adjacent row.

22 Discussion

Values versus fibers. Above the per-rate value-set reach of [Cho26b] (roughly 2^{150}–2^{162}), the value route — showing that the slope multiset $e_1(\ell^{\wedge}Q)$ covers a 2^{-128} fraction of $\mathbb{F}$ at a fixed large prime — runs into the per-fiber collision problem ([Cho26b, Rem. "The remaining band"]): characteristic-zero locator values can collide arbitrarily badly under reduction at a worst-case prime. [Cho26b] therefore closed the band at the δ^* grade only conditionally and with weaker constants — error 2^{-42} at gap 2^{-11}, assuming $|\mathbb{F}| \geq 2n$, via quotient-core lists and the same conversion Theorem 2.6 — leaving the *error-one* grade outside the present method. The present route also never counts values, and improves each parameter of that cap: pigeonhole guarantees one heavy fiber regardless of how the values $e_1(A)$ distribute, Theorem 2.6 converts heavy fibers into correlated-agreement error directly, and the subfield pigeonhole carries the construction to extension fields with no condition relating n and $|\mathbb{F}|$. The cost relative to the error-one results of [Cho26a, Cho26b] below 2^{150} is quantitative: error $\approx \frac{1}{2k}$ instead of $1 - o(1)$, and gap 2^{-9} (uniformly 2^{-10} across all four rates) instead of $\frac{1}{64}$. For the prize question — is $\varepsilon_{\mathrm{mca}} \leq 2^{-128}$ achievable above the stated gaps? — this cost is immaterial.

Limits of the method. Hypothesis (1) needs $N\,\mathrm{H}_2(\rho) \gtrsim 2\log_2 q - \log_2 k$, so the smallest certifiable gap is

$$\frac{2}{N} \;\approx\; \frac{2\,\mathrm{H}_2(\rho)}{2\log_2 q - \log_2 k} \;\approx\; \frac{\mathrm{H}_2(\rho)}{\log_2 q},$$

and the certified error is capped at $\approx \frac{1}{k}$ by the shape of Theorem 2.6 (any $\eta < 1$). Since q enters only as the size of the field from which the line challenge γ is drawn, this is also a protocol-facing floor: enlarging the challenge field cannot push a deep-point-conversion failure below gap $\approx \mathrm{H}_2(\rho)/\log_2 q_{\mathrm{line}}$, which quantitatively bounds the "larger challenge field" fallback in certificate frameworks. The construction also needs $N \leq n$, i.e. $\log_2 q \lesssim \mathrm{H}_2(\rho)\,n/2$; this covers the entire cryptographic regime $\mathrm{poly}(n) \leq q \leq 2^{o(n)}$, but for $q \geq 2^{\mathrm{H}_2(\rho)n/2}$ the method is vacuous and the value-route caps of [Cho26a, Cho26b] are the only ones known. Within the band, the *error-one* question at gap $\frac{1}{64}$ (the error-one question Question 22.2) remains exactly as open as before: the present result bounds δ^*, not the failure profile near capacity.

Comparison with the cyclotomic sieve. The sieve of [Cho26a] produces, for infinitely many primes, failures of error $(\log p)^{-6}$ at gap $\Theta(1/\log p)$, but its prime set is given by Siegel–Walfisz and is ineffective. Corollary 5.1 trades error strength for universality: every field below 2^{256}, fully effective, and elementary given Theorem 2.6. The two results triangulate the failure landscape from different directions.

Isogeny-smooth reach and the lacunarity dichotomy. The counting side of the method extends beyond multiplicative smoothness. The fiber lemma needs only a polynomial folding map with complete uniform fibers inside the domain (Definition 7.1 and Lemma 7.2), and the divisibility hypothesis $a \mid k$ is removed by rounding, so the universal cap now covers mixed-radix multiplicative rows without integral ρN (Corollary 7.6), the x-coordinate line rounds of circle FRI over $p \equiv 3 \pmod 4$, and — through the torus uniformization, over any challenge field containing i — the bivariate circle codes themselves, at the deployed Mersenne-31/QM31 parameters with error $> 2^{-23}$ at gap 2^{-7} (Corollaries 7.12 and 7.13). The boundary of the method is now a clean dichotomy in the degree type (a, e) of the folding map: polynomial

towers ($e = 0$: powers, Chebyshev) give gap $\approx 2/N$, while elliptic isogeny towers ($e = a - 1$: ECFFT) collapse the slack-two window to gap $2/n$ (Remark 7.16). Subfield confinement and the extension-pole mechanism also carry over to circle rows (Corollary 7.14), so the deployed circle failures are again fiber-borne, genuinely extension-valued, and constructively reachable from any supplied list via extension poles.

Pressure on the packing conjecture. By the exact normal form of [Cho26b], $\varepsilon_{\mathrm{mca}}(C, \delta) = \max_t \Lambda_t^{\mathrm{NC}}(\delta)/q$, where Λ^{NC} is a noncontained residue-line packing number. Corollary 5.4 therefore forces $\max_t \Lambda^{\mathrm{NC}} \geq (q-n)/2k \approx 2^{164}$ at gap 2^{-7} on the deployed sextic — superpolynomially many noncontained residue lines at $n = 2^{21}$. This puts no direct pressure on the final MCA conjecture of [Cho26b, Conj. "Final floor- and quotient-corrected MCA asymptotic"]: at gap 2^{-7} both of its corrected-reserve hypotheses fail ($\sigma = 2^{14}$ lies below $Cn/\log_2 n \approx 2^{16.6}$, and $\sigma \log_2 q_{\mathrm{gen}} = 2^{14} \cdot 31 \approx 2^{19}$ is far below $\log_2 \binom{n}{k+\sigma} \approx 2^{21}$), so the conjectured ceiling of the form $\left(n^{1+o(1)} + 2^{(\beta/\mathrm{H}_2)\, Q_{\mathrm{H}_2}(\eta)}\right)/q$ is simply not asserted there — the packing mass extracted through Theorem 2.6 lives in the region the framework already marks as below-floor. The genuine structural question this paper raises for the framework of [Cho26b] is whether conversion-extracted fiber mass persists *above* the corrected reserve: whether, at reserves where the conjecture does apply, the fiber route still produces packings beyond the $n^{1+o(1)}$ aperiodic term plus the quotient-periodic term, which would force an additional fiber-borne term into the conjectured ceiling. This is not settled here.

Verification caveat. The main cap no longer depends on importing the Crites–Stewart theorem: Theorem 2.6 proves the needed deep-point list-to-CA conversion directly, including the integer-radius condition $f \leq n-k-1$. Theorem 2.16 remains a separate slacked fallback imported through the BCHKS/ABF route, and readers using that fallback should consult the underlying source directly. The main composition is deliberately shallow: it uses the self-contained conversion, the locator-fiber pigeonhole, and support-wise MCA monotonicity.

What remains after the frontier package. The structural ledgers, extension-line accounting, scanner, and certificate grammar are closed at the level needed for the printed unsafe certificates and for conditional safe-side compilation. After the entropy–subfield normalization, the remaining safe-side package is more precise than the band statement by itself. The active research programme is the candidate three-bucket package

quotient/prefix max-fiber flatness,
base-field-normalized interior split-pencil census,
primitive shift-pair control in the exact second-moment ledger.

They are formulated precisely in Section 21. In the asymptotic regime one also requires $R(n) \to \infty$ and $R(n) = o(n)$; a loss $L(n)$ is admissible only when $\log L(n) = o(R(n) \log |\mathbb{B}|)$. Polynomial losses therefore suffice with $R \gg \log n$, whereas a generic $e^{o(n)}$ loss need not. The finite adjacent rows require explicit constants below the printed margins of Corollary 15.3. Separate auxiliary questions remain: the error-one profile Question 22.2, the explicit deployed-witness strengthening Question 22.1, and the genus-one/circle transport problem Question 22.3.

Questions resolved by the present methods. The genus-one macroscopic-cap and circle field-of-definition questions are resolved by Corollaries 11.2 and 11.5. The explicit-witness question is resolved at the first staircase step and for simple-pole certifying pairs at gap 2^{-8} (Corollary 11.7 and Theorem 11.8); its full-gap strengthening is restated below.

Further frontier questions.

Question 22.1 (explicit deployed witnesses at the full deployed gap)**.** For the KoalaBear-sextic deployed row, exhibit explicit pairs (f, g) with MCA-bad-slope density $> 2^{-22}$ at the gap 2^{-7} certified by Corollary 5.4; on the integer-admissible subinterval, exhibit the corresponding CA-bad-slope density as well. By Corollary 6.2 such pairs exist and are not $\mathbb{B}$-rational, and the natural candidates are residue-line normal forms with denominator $E \in \mathbb{F}[X]$ not defined over $\mathbb{B}$. The present paper gives explicit received words and simple-pole certifying pairs of exactly this normal form, $E = X - \alpha$, at the first staircase step and at gap 2^{-8} (Corollary 11.7 and Theorem 11.8); the remaining task is to push the explicit pair construction to the full deployed gap, derandomize the choice of the pole α, and print the deployed-size bucket tables.

Question 22.2. Error one in the band: does $\varepsilon_{\mathrm{mca}}(C, 1 - \rho - \frac{1}{64}) = 1 - o(1)$ hold for all primes $2^{150} \le p \le 2^{256}$? This is the per-fiber collision problem proper; Corollary 5.1 caps δ^* there but says nothing above error $\approx \frac{1}{2k}$. Corollary 11.11 raises the certified profile to $(1 - 2^{-54})\frac{1}{k}$ at gap 2^{-7}, and Remark 11.12 shows the question lies strictly beyond both mechanisms of this paper; it remains outside the present methods.

Further direction: improving the conversion threshold. Improve the threshold of Theorem 2.6: any strengthening of the hypothesis range $\varepsilon_{\mathrm{ca}} \le \eta/k$ toward $\varepsilon_{\mathrm{ca}} \le \eta \cdot \mathrm{polylog}(q)/k$, or of the list conclusion, immediately tightens the certified error of Theorem 4.1 toward $1/\mathrm{poly}\log$; the gap, by contrast, is governed by the binomial in (1) and would need a denser fiber construction to shrink. Cf. Proposition 11.10: optimizing the dichotomy already attains $(1 - o(1))/k$, so passing $1/k$ is exactly what a polylog strengthening would buy.

Question 22.3 (aperiodic band on genus-one and circle tangent domains)**.** Extend the active safe-side package—quotient/prefix flatness, base-field-normalized split-pencil control, and primitive shift-pair control—to the tangent-domain Reed–Solomon models arising from genus-one isogeny towers and to the stereographic circle models over challenge fields not containing i. The macroscopic unsafe floors on these rows are supplied by Corollaries 11.2 and 11.5; the remaining task is to transport the normalized safe-side input through the rational scale map.

Acknowledgements. This paper builds on the Crites–Stewart conversion perspective [CS25], the BCHKS theorem [BCHKS25], and the locator framework of [Cho26a, Cho26b], following the synthesis of Arnon–Boneh–Fenzi [ABF26]; it proves the main conversion needed here directly, sharpens the conditional universal-cap theorem, and refines the subfield-confinement theorem of [Cho26b].

Guide to the auxiliary appendices

The main body states the final logical interface and all theorem-level additions used by version 13.2. The appendices preserve the longer L1/M1 compilers, alternative proof routes, rank-one and split-pencil reductions, and computational calibration that motivate the three candidate residual buckets. Statements marked experimental or conjectural are not consumed by any unconditional theorem.

A L1 list-side compilers and sunflower residuals

Let $D \subset F$ be an evaluation domain and let $C = \mathrm{RS}[F, D, k]$ be a Reed–Solomon code of length $n = |D|$. For a received word $U : D \to F$ and an agreement threshold s, write

$$\mathrm{ImgFib}_U^{<k}(s) := \{P \in F[X]_{<k} : \#\{x \in D : U(x) = P(x)\} \geq s\},$$

and, for an auxiliary degree bound d and a subdomain $T \subseteq D$,

$$\mathrm{ImgFib}_W^{\leq d}(s; T) := \{G \in F[X]_{\leq d} : \#\{x \in T : W(x) = G(x)\} \geq s\}.$$

These are the actual list objects used below; they count codewords rather than raw supports.

A.1 Planted quotient-core lower counts

Theorem A.1 (planted quotient-core lower count)**.** *Let $D = \alpha H \subset F^\times$ be a multiplicative coset of a cyclic subgroup H of order n, let $C = \mathrm{RS}[F, D, k]$, and suppose $M \mid \gcd(n, k)$ satisfies*

$$1 \leq \sigma < M, \qquad k/M \leq n/M - 1.$$

Then there is a received word $U : D \to F$ with at least

$$P_M(n, k) := \binom{n/M - 1}{k/M}$$

distinct codewords agreeing with U on at least $k + \sigma$ positions. Consequently

$$\max_U \left|\mathrm{ImgFib}_U^{<k}(k + \sigma)\right| \geq \max_{\substack{M \mid \gcd(n,k) \\ \sigma < M \\ k/M \leq n/M - 1}} \binom{n/M - 1}{k/M}.$$

Proof. Let $K \leq H$ be the unique subgroup of order M and let $Q := \{x^M : x \in D\}$ be the quotient image, so that the map $\pi : D \to Q$, $\pi(x) = x^M$, has fibers equal to the K-cosets inside D, each of size exactly M, and $|Q| = N := n/M$. Choose one quotient value $b_0 \in Q$ and a set $T \subseteq \pi^{-1}(b_0)$ of size σ; this is possible since $\sigma < M$. Put

$$E_T(X) := \prod_{t \in T} (X - t).$$

For every subset $A \subseteq Q \setminus \{b_0\}$ of size k/M, define

$$L_A(X) := E_T(X) \prod_{b \in A} (X^M - b).$$

Every point of the fiber $\pi^{-1}(b)$ with $b \in A$ is a root of $X^M - b$, and every point of T is a root of E_T; these point sets are disjoint since $b_0 \notin A$. Hence L_A vanishes at at least $M \cdot (k/M) + \sigma = k + \sigma$ points of D.

Since

$$\prod_{b \in A} (X^M - b) = X^k + \text{terms of degree at most } k - M,$$

and $\deg E_T = \sigma < M$, the part of L_A of degree at least k is independent of A, namely

$$U_0(X) := X^k E_T(X).$$

Thus $L_A = U_0 + R_A$ with $\deg R_A < k$. Let U be the received word $x \mapsto U_0(x)$ on D and let $P_A := -R_A \in F[X]_{<k}$. For every root $x \in D$ of L_A,

$$0 = L_A(x) = U(x) - P_A(x),$$

so P_A agrees with U on at least $k + \sigma$ positions.

It remains to check distinctness. If P_A and $P_{A'}$ give the same codeword on D, then $R_A - R_{A'}$ has degree $< k < n$ and vanishes on all n points of D, so $R_A = R_{A'}$ as polynomials, hence $L_A = L_{A'}$, and dividing by the fixed nonzero polynomial E_T gives

$$\prod_{b \in A} (X^M - b) = \prod_{b \in A'} (X^M - b).$$

The substitution homomorphism $F[Y] \to F[X]$, $G(Y) \mapsto G(X^M)$, is injective (the monomials X^{Mi} are linearly independent), so $\prod_{b \in A}(Y - b) = \prod_{b \in A'}(Y - b)$ in $F[Y]$, and unique factorization into the distinct monic linear factors gives $A = A'$. (This argument needs no separability of $X^M - b$ and is valid in every characteristic.) There are $\binom{N-1}{k/M}$ choices of A, which proves the claim. □

Corollary A.2 (complete polynomial-folding planted core)**.** *Let $H \subset F$ be finite, let $\phi \in F[X]$ have degree $M > 0$, and assume the map $\phi : H \to Q$ is onto with every fiber $\phi^{-1}(b) \cap H$ of size exactly M. Suppose*

$$M \mid k, \qquad 1 \le \sigma < M, \qquad k/M \le |Q| - 1.$$

Then there is a received word $U : H \to F$ with at least

$$\binom{|Q| - 1}{k/M}$$

distinct degree-$< k$ Reed–Solomon codewords agreeing with U on at least $k + \sigma$ positions.

Proof. Choose $b_0 \in Q$ and a subset $T \subseteq \phi^{-1}(b_0) \cap H$ of size σ. Put

$$E_T(X) := \prod_{t \in T} (X - t), \qquad \ell := k/M.$$

For every $A \subseteq Q \setminus \{b_0\}$ of size ℓ, define

$$L_A(X) := E_T(X) \prod_{b \in A} (\phi(X) - b).$$

Its roots in H include the σ planted points of T and the $M\ell = k$ points of the complete fibers over A, disjointly, hence at least $k + \sigma$ points. Since

$$\prod_{b\in A}(Y - b) = Y^{\ell} + \text{terms of degree at most } \ell - 1,$$

we have

$$L_A(X) = E_T(X)\phi(X)^{\ell} + R_A(X), \qquad \deg R_A \le \sigma + M(\ell - 1) < k.$$

Let U be the word $x \mapsto E_T(x)\phi(x)^{\ell}$ on H and put $P_A := -R_A \in F[X]_{<k}$; then $U = P_A$ at every root of L_A in H.

For distinctness, note $|H| = M|Q|$ by the complete-fiber hypothesis, and $k = M\ell \le M(|Q| - 1) < |H|$. If P_A and $P_{A'}$ define the same codeword on H, then $R_A - R_{A'}$ has degree $< k < |H|$ and vanishes on H, hence $R_A = R_{A'}$ and

$$\prod_{b\in A}(\phi(X) - b) = \prod_{b\in A'}(\phi(X) - b).$$

The substitution map $F[Y] \to F[X]$, $G(Y) \mapsto G(\phi(X))$, is injective because $\deg\phi > 0$, so $\prod_{b\in A}(Y - b) = \prod_{b\in A'}(Y - b)$ and $A = A'$. The number of choices is $\binom{|Q|-1}{\ell}$. □

Remark A.3 (relation to the main fiber lemmas and extension persistence). Theorem A.1 is the exact-count planted form of the quotient-core construction of [Cho26b]: a single received word carries the whole binomial sublist, with no pigeonhole loss and no subfield denominator. By contrast, Lemma 3.1 pigeonholes locator subsets over their slope values to extract one heavy fiber at slack two; the planted form is the right tool on the list side, where no slope average is available. Corollary A.2 is the map-smooth analogue: its complete-uniform-fiber hypothesis is exactly the condition of Definition 7.1 used by Lemma 7.2, so the planted core applies verbatim to Chebyshev and circle-derived foldings. Finally, all these list floors persist over every field extension: $\mathrm{RS}[B, D, k] \subseteq \mathrm{RS}[F, D, k]$ for $D \subseteq B \subseteq F$, so the planted sublist survives unchanged when the row is lifted, exactly as in the list-side persistence discussion of Section 6.

Corollary A.4 (list unsafe test). *For a list denominator q_{list} and target 2^{-128}, the planted quotient-core packet proves a list row unsafe whenever*

$$\max_{M}\binom{n/M - 1}{k/M} > \left\lfloor \frac{q_{\mathrm{list}}}{2^{128}} \right\rfloor,$$

where the maximum ranges over the active divisors of Theorem A.1.

Proof. By Theorem A.1, the maximum list size at agreement $k + \sigma$ is at least the displayed planted count. If that count exceeds the allowed numerator $\lfloor q_{\mathrm{list}}/2^{128} \rfloor$, the normalized list quantity exceeds 2^{-128}. □

Proposition A.5 (dyadic planted crossings). *For $n = 2^m$, $k = \rho n$, and*

$$\rho \in \{1/2,\ 1/4,\ 1/8,\ 1/16\},$$

the first dyadic quotient orders $N = n/M$ at which the planted count exceeds $2^{128} - 1$ are

ρ	1/2	1/4	1/8	1/16
N	256	256	256	512,

and the exact bit lengths of $\binom{N-1}{\rho N}$ *at the predecessor and first crossing are*

ρ	N_{prev}	bits	N_{first}	bits
1/2	128	124	256	251
1/4	128	100	256	204
1/8	128	67	256	136
1/16	256	83	512	169.

Proof. The planted count at quotient order N is $\binom{N-1}{\rho N}$, and $\text{bits}(X) \le 128 \iff X \le 2^{128}-1$; the displayed bit lengths were checked by exact integer arithmetic on binomial coefficients, not floating-point logarithms. It remains to justify that the predecessor check identifies the first dyadic crossing. Along dyadic orders the counts are monotone: from a subset $S \subseteq \{1,\dots,N-1\}$ of size ρN, form a subset of $\{1,\dots,2N-1\}$ of size $2\rho N$ by adjoining a fixed set of ρN elements of $\{N,\dots,2N-1\}$. This map is injective, so

$$\binom{N-1}{\rho N} \le \binom{2N-1}{2\rho N}.$$

Hence once the predecessor is below the budget, every smaller dyadic order is below it as well. □

Corollary A.6 (planted list budget windows)**.** *If a planted exact count is P and the certified lower count is $L \le P$, then the only unresolved list budgets are*

$$L \le \lfloor q_{\text{list}}/2^{128} \rfloor < P,$$

equivalently $L2^{128} \le q_{\text{list}} \le P2^{128}-1$.

Proof. This is Theorem 13.14 with $q = q_{\text{list}}$, $K = P$. □

A.2 Agreement-threshold Johnson bounds

Theorem A.7 (κ-intersecting agreement bound)**.** *Let Ω be a finite set of size n_0, let $W : \Omega \to F$ be a received word, and let $\mathcal{F}$ be a family of words $\Omega \to F$ such that any two distinct members of $\mathcal{F}$ agree on at most κ points of Ω. For an agreement threshold s with $0 < s \le n_0$, let L be the number of members of $\mathcal{F}$ agreeing with W on at least s points.*

(i) *If $2s > n_0 + \kappa$, then $L \le 1$.*

(ii) *If $s^2 > \kappa n_0$, then*

$$L \le \frac{n_0(s-\kappa)}{s^2 - \kappa n_0}.$$

In particular, for $C = \text{RS}[F, \Omega, k]$ and $\kappa = k-1$,

$$\left|\text{ImgFib}_W^{<k}(s)\right| \le \frac{n_0(s-k+1)}{s^2-(k-1)n_0} \qquad \text{when } s^2 > (k-1)n_0,$$

and for the auxiliary degree-$\le d$ code on a petal domain, $\kappa = d$.

Proof. Let $P_1, \dots, P_L$ be the listed members and $A_i := \{x \in \Omega : P_i(x) = W(x)\}$, so $|A_i| \geq s$ and, for $i \neq j$, $|A_i \cap A_j| \leq \kappa$ by the intersecting hypothesis (two words agreeing with W at a point agree with each other there).

(i) For $L \geq 2$, inclusion–exclusion gives $|A_1 \cap A_2| \geq 2s - n_0 > \kappa$, a contradiction.

(ii) Assume $L \geq 1$ and put $m_x := \#\{i : x \in A_i\}$ and $I := \sum_{x \in \Omega} m_x \geq Ls$. Counting incident pairs,

$$\sum_{x \in \Omega} \binom{m_x}{2} = \sum_{i<j} |A_i \cap A_j| \leq \binom{L}{2} \kappa.$$

First suppose $Ls \leq n_0$. Then $L \leq n_0/s$, and since $s \leq n_0$,

$$\frac{n_0}{s} \leq \frac{n_0(s-\kappa)}{s^2 - \kappa n_0} \iff s^2 - \kappa n_0 \leq s^2 - \kappa s \iff \kappa s \leq \kappa n_0,$$

which holds; note $s - \kappa > 0$ because $s^2 > \kappa n_0 \geq \kappa s$. Now suppose $Ls > n_0$. The function $\varphi(y) = y(y-1)$ is convex and nondecreasing on $[1, \infty)$, so Jensen and $I/n_0 \geq Ls/n_0 > 1$ give

$$n_0 \cdot \frac{Ls}{n_0}\Big(\frac{Ls}{n_0} - 1\Big) \leq n_0\, \varphi\Big(\frac{I}{n_0}\Big) \leq \sum_x m_x(m_x - 1) \leq L(L-1)\kappa.$$

Multiplying out, $Ls(Ls - n_0) \leq n_0 L(L-1)\kappa$, hence $s(Ls - n_0) \leq n_0(L-1)\kappa$, i.e.

$$L(s^2 - \kappa n_0) \leq n_0(s - \kappa),$$

which is the claim. □

Remark A.8 (relation to the main Johnson theorem). For $\kappa = k-1$ this is exactly the $m = 1$ case of Theorem 9.15, restated in agreement-threshold list coordinates; the numerator $n_0(s-k+1)$ is the sharp one and strictly improves the numerator $n_0(n_0 - k + 1)$ that a plain Cauchy–Schwarz argument yields. The abstract κ-intersecting form is recorded because the sunflower reduction below applies it with $\kappa = d$ to an auxiliary code that is not a row of the challenge.

A.3 Sunflower and petal interfaces

The statements in this subsection are compiler reductions: they become theorems about an actual sunflower decomposition only when the decomposition supplies the stated layer maps. This is the safe capf form: the paid pieces are proved, while the mixed-petal and growing-defect classification remains a named residual branch. The first proposition records one concrete situation in which the required layer injection is automatic.

Definition A.9 (concrete sunflower core layer)**.** Let $\Omega \subset F$ be a finite evaluation set for $\mathrm{RS}[F, \Omega, k]$. Fix pairwise disjoint subsets

$$Y, \quad R_0, \quad T_1, \dots, T_M \subseteq \Omega, \qquad |Y| = k - 1, \qquad |T_i| = \ell = \sigma + 1,$$

a reference polynomial $P_\star \in F[X]_{<k}$, and labels $c_1, \dots, c_M \in F$. Put

$$L_Y(X) = \prod_{y \in Y} (X - y),$$

and assume the received word $U : \Omega \to F$ satisfies

$$U = P_\star \quad \text{on } Y \cup R_0, \qquad U = P_\star + c_i L_Y \quad \text{on } T_i.$$

For $D_0 \subseteq Y$ write $d = |D_0|$, $L_{D_0}(X) = \prod_{y \in D_0}(X - y)$, and $T := \bigcup_i T_i$. The concrete (D_0, R_0) layer consists of those $P \in F[X]_{<k}$ that agree with U on at least $k + \sigma$ points of Ω, agree with U on all of $(Y \setminus D_0) \cup R_0$, and have no agreements with U outside $(Y \setminus D_0) \cup R_0 \cup T$.

Proposition A.10 (concrete sunflower layer injection)**.** *In the setting of Definition A.9, the map*

$$P \longmapsto G_P := \frac{P - P_\star}{L_{Y \setminus D_0}}$$

is well-defined and injects the concrete (D_0, R_0) layer into $F[X]_{\le d}$. Moreover G_P agrees on at least

$$a = \sigma + d + 1 - |R_0|$$

points of the petal domain T with the auxiliary word

$$U_{D_0}(x) := c_i L_{D_0}(x), \qquad x \in T_i,$$

and if the layer retains the residual set R_0, then in addition $G_P(x) = 0$ for every $x \in R_0$.

Proof. Since P and $P_\star$ agree on $Y \setminus D_0$ (both equal $U = P_\star$ there), the difference $P - P_\star$ is divisible by $L_{Y \setminus D_0}$, which has degree $k - 1 - d$; the quotient G_P therefore has degree at most $(k-1) - (k-1-d) = d$. The quotient determines $P = P_\star + G_P L_{Y \setminus D_0}$, so the map is injective.

The layer has at least $k + \sigma$ agreements in total, of which at most $|Y \setminus D_0| + |R_0| = (k - 1 - d) + |R_0|$ lie off the petal domain, by the containment condition of the definition. Hence at least

$$(k + \sigma) - (k - 1 - d) - |R_0| = \sigma + d + 1 - |R_0| = a$$

agreement points lie in T. If $x \in T_i$ is such a point, then $x \notin Y$, so $L_{Y \setminus D_0}(x) \neq 0$, and

$$G_P(x) L_{Y \setminus D_0}(x) = P(x) - P_\star(x) = U(x) - P_\star(x) = c_i L_Y(x) = c_i L_{D_0}(x) L_{Y \setminus D_0}(x);$$

cancelling the nonzero factor gives $G_P(x) = c_i L_{D_0}(x) = U_{D_0}(x)$. Finally, if $x \in R_0$, then $P(x) = U(x) = P_\star(x)$ and $x \notin Y$, so $G_P(x) = 0$. □

Definition A.11 (fixed sunflower auxiliary layer)**.** Fix a missed-core set D_0 of size d with monic locator $L_{D_0}(X) = \prod_{y \in D_0}(X - y)$, a retained residual subset R_0 of size r, pairwise disjoint petals $T_1, \dots, T_M$ of common size $\sigma + 1$ disjoint from $D_0 \cup R_0$, and labels $c_i \in F$. Let $T := \bigcup_i T_i$ and define

$$U_{D_0}(x) := c_i L_{D_0}(x), \qquad x \in T_i.$$

A fixed (D_0, R_0) sunflower layer is a finite set $\mathcal{L}(D_0, R_0)$ together with an injective map

$$\Psi : \mathcal{L}(D_0, R_0) \hookrightarrow F[X]_{\le d}$$

such that every $G = \Psi(P)$ agrees with U_{D_0} on at least

$$a := \sigma + d + 1 - r$$

points of T. Extra conditions forcing zeros on R_0 are called retained zero constraints.

Proposition A.12 (sunflower core-defect reduction)**.** *For every fixed sunflower auxiliary layer in the sense of Definition A.11,*

$$|\mathcal{L}(D_0, R_0)| \le \left|\mathrm{ImgFib}^{\le d}_{U_{D_0}}(a; T)\right|,$$

the auxiliary degree-$\le d$ list on the petal domain T at agreement $a = \sigma + d + 1 - r$. Dropping the retained zero constraints on R_0 only enlarges the right-hand side, so it is safe for upper bounds.

Proof. By definition, Ψ injects the fixed layer into the set of degree-$\le d$ polynomials agreeing with U_{D_0} on at least a points of T, which is the displayed auxiliary list. Adding zero constraints cuts out a subset of it; removing them can only increase the counted set. □

Corollary A.13 (few-petal Johnson boundary)**.** *In the notation of Proposition A.12, the auxiliary petal domain has size $|T| = M(\sigma + 1)$, and two distinct degree-$\le d$ polynomials agree on at most d points, so Theorem A.7 applies with $\kappa = d$. If $a^2 > d|T|$, the fixed (D_0, R_0) auxiliary list has size at most*

$$\frac{|T|\,(a - d)}{a^2 - d\,|T|}.$$

Equivalently, for $d > 0$, the few-petal Johnson-covered range is exactly

$$M \le \left\lfloor \frac{a^2 - 1}{d(\sigma + 1)} \right\rfloor.$$

Proof. The bound is Theorem A.7(ii) with $n_0 = |T|$, $s = a$, $\kappa = d$ (if $a > |T|$ the list is empty and the bound is vacuous). Since all quantities are integers, the strict inequality $a^2 > M(\sigma+1)d$ is equivalent to $M(\sigma + 1)d \le a^2 - 1$, which is the stated floor. □

Definition A.14 (full-petal fixed-excess counting chart)**.** A full-petal counting chart with cofactor excess $e = d - \ell$ is a parameterization of a full-petal layer with the following certificate; the certificate is for the whole layer being charged, and if several independent charts are used the bounds below must be summed over the chart index.

(i) the zero-excess layer $e = 0$ is determined by an unordered pair of petals and one scalar in F;

(ii) for every $e \ge 1$, the layer is determined by a subset of the M petals and at most $e + 1$ scalars in F.

Proposition A.15 (distinct-label zero-excess full-petal chart)**.** *Work in a concrete sunflower layer* (Definition A.9) *with $R_0 = \varnothing$, petal size $\ell = \sigma + 1$, and $d = \ell$, and assume the labels $c_1, \dots, c_M$ are pairwise distinct. Consider a residual-free full-petal sublayer: every listed codeword P agrees with U on the full union of a set $I(P)$ of complete petals, and the listedness inequality forces $|I(P)| \ge 2$. Then the map*

$$P \longmapsto (\{i, j\}, \alpha_i),$$

where $i < j$ are the two smallest indices in $I(P)$ and α_i is the scalar constructed in the proof, is injective, so the sublayer has size at most

$$\binom{M}{2}\,|F|.$$

Moreover the injection recovers the missed-core locator L_{D_0} itself, hence D_0; consequently the same bound $\binom{M}{2}|F|$ holds for the union *of these sublayers over all missed-core sets $D_0 \subseteq Y$ of size ℓ. In particular the chart supplies the zero-excess certificate of Definition A.14.*

Proof. Write $L := L_{D_0}$ and let L_i be the monic locator of the petal T_i; both are monic of degree ℓ. For a codeword P in the sublayer write, as in Proposition A.10,

$$P - P_\star = G\, L_{Y \setminus D_0}, \qquad \deg G \le d = \ell.$$

If $i \in I(P)$, then P agrees with $U = P_\star + c_i L_Y$ at every point of T_i, so by Proposition A.10 the polynomial $G - c_i L$, of degree at most ℓ, vanishes at all ℓ points of T_i; hence

$$G - c_i L = \alpha_i L_i \qquad \text{for some } \alpha_i \in F.$$

Subtracting the identities for the two smallest touched petals $i < j$,

$$(c_j - c_i)L = \alpha_i L_i - \alpha_j L_j.$$

Comparing leading coefficients of the three monic degree-ℓ locators gives

$$c_j - c_i = \alpha_i - \alpha_j, \qquad \text{so} \qquad \alpha_j = \alpha_i - (c_j - c_i).$$

Given the pair $\{i, j\}$ and the scalar α_i, the label difference $c_j - c_i \neq 0$ determines α_j, then

$$L = \frac{\alpha_i L_i - \alpha_j L_j}{c_j - c_i}$$

determines L, then $G = c_i L + \alpha_i L_i$, and finally $P = P_\star + G\, L_{Y \setminus D_0}$, where D_0 is recovered as the root set of the monic squarefree polynomial L inside Y and $L_{Y \setminus D_0} = L_Y / L$. Thus the map is injective, even jointly over the choice of D_0, and there are at most $\binom{M}{2}|F|$ values. □

Proposition A.16 (distinct-label fixed-excess full-petal chart)**.** *Work in a concrete sunflower layer* (Definition A.9) *with $R_0 = \varnothing$, petal size $\ell = \sigma + 1$, and pairwise distinct labels $c_1, \dots, c_M$. Fix $e \ge 1$ and put $d = \ell + e$. Consider the charged pairs (P, D_0), where $D_0 \subseteq Y$ has size d and P lies in a residual-free all-full sublayer with missed core D_0: there is a touched set $I(P, D_0) \subseteq [M]$ such that P agrees with U on every point of T_i for $i \in I(P, D_0)$, has no petal agreements for $i \notin I(P, D_0)$, and has no agreements outside $(Y \setminus D_0) \cup \bigcup_{i \in I(P, D_0)} T_i$. Then, for this fixed e, the number of charged pairs is at most*

$$2^M |F|^{e+1}.$$

Consequently the union of these residual-free all-full distinct-label sublayers over all missed cores D_0 also has size at most $2^M |F|^{e+1}$. Equivalently, in this distinct-label residual-free all-full regime the higher-excess certificate of Definition A.14 is automatic, not an extra hypothesis.

Proof. Let (P, D_0) be such a charged pair, put $L := L_{D_0}$, and let L_i be the monic locator of the petal T_i. As in Proposition A.10, write

$$P - P_\star = G\, L_{Y \setminus D_0}, \qquad \deg G \le d = \ell + e.$$

If $i \in I := I(P, D_0)$, full-petal agreement on T_i gives

$$G - c_i L = L_i B_i$$

for a uniquely determined polynomial B_i of degree at most e, because the left side has degree at most $\ell + e$ and vanishes on the ℓ points of T_i.

The listedness threshold is $k + \sigma = k + \ell - 1$. Since the sublayer is residual-free and all-full, all agreements are contained in the $(k-1-d)$ retained core points and the $|I|\ell$ full petal points. Hence

$$(k-1-d) + |I|\ell \geq k + \ell - 1, \qquad \text{so} \qquad (|I|-1)\ell \geq d.$$

In particular I has at least two elements. Let i_0 be the least element of I. For every $j \in I \setminus \{i_0\}$, subtracting the full-petal identities gives

$$(c_{i_0} - c_j)L + L_{i_0}B_{i_0} = L_j B_j,$$

so, since $c_{i_0} \neq c_j$,

$$L \equiv -(c_{i_0} - c_j)^{-1} L_{i_0} B_{i_0} \pmod{L_j}.$$

The petal locators are pairwise coprime, so the touched set I and the single polynomial B_{i_0} determine, by the Chinese remainder theorem, the residue class of L modulo

$$Q_I := \prod_{j \in I \setminus \{i_0\}} L_j, \qquad \deg Q_I = (|I|-1)\ell \geq d.$$

There is at most one monic polynomial of degree d in this residue class: if $\deg Q_I > d$ this is immediate because the residue representative has degree $< \deg Q_I$, and if $\deg Q_I = d$ the difference of two monic degree-d representatives has degree $< d$ while being divisible by Q_I. Thus $L = L_{D_0}$ is recovered from (I, B_{i_0}).

Once L is known, the actual charged pair recovers D_0 as the root set of the monic squarefree locator L inside Y, then recovers $L_{Y \setminus D_0} = L_Y / L$, then recovers

$$G = c_{i_0} L + L_{i_0} B_{i_0}, \qquad P = P_\star + G L_{Y \setminus D_0}.$$

Hence the map $(P, D_0) \mapsto (I, B_{i_0})$ is injective. There are at most 2^M possible touched sets and at most $|F|^{e+1}$ possible polynomials B_{i_0} of degree at most e, proving the bound. □

Remark A.17 (higher-excess full petals after the distinct-label chart). Proposition A.16 pays residual-free all-full higher-excess petals when the petal labels are pairwise distinct. For repeated labels, retained residual sets, partial or mixed petals, or charts that do not satisfy the residual-free all-full containment hypothesis, Chinese-remainder reconstructions may still certify the $2^M q^{e+1}$ chart of Definition A.14, but the reconstruction data must be printed as a certificate. The part does not infer those remaining higher-excess charts from listedness alone.

Proposition A.18 (fixed-excess full-petal compiler)**.** *For every fixed $E \geq 0$, any single certified full-petal atlas satisfying Definition A.14 has total size at most*

$$\binom{M}{2} q + 2^M \sum_{e=1}^{E} q^{e+1}, \qquad q = |F|,$$

the certified layer bounds being $\binom{M}{2} q$ at $e = 0$ and $2^M q^{e+1}$ at each $e \geq 1$.

Proof. For $e = 0$ the certificate chooses an unordered pair of petals and one scalar: at most $\binom{M}{2} q$ possibilities. For fixed $e \geq 1$ it chooses a subset of the M petals in at most 2^M ways and $e+1$ scalars in at most q^{e+1} ways. In the residual-free pairwise-distinct-label all-full case, Proposition A.16 proves exactly this certificate; in other cases it remains a printed atlas certificate. Summing over $1 \leq e \leq E$ proves the claim. □

Corollary A.19 (fixed-excess polynomiality)**.** *At the intended L1 lower cutoff, suppose $M = O(\log n)$, $q \le n^Q$ for a fixed Q, and E is fixed. Then the fixed-excess full-petal contribution certified by Proposition A.18 is polynomial in n.*

Proof. The term $\binom{M}{2}q$ is polynomial because $M = O(\log n)$ and $q \le n^Q$. For fixed $e \le E$, the term $2^M q^{e+1}$ is $n^{O(1)} \cdot n^{Q(e+1)}$, hence polynomial, and a fixed finite sum of polynomially bounded terms is polynomial. □

Remark A.20 (L1 evidence is not a theorem)*.* The small-field auxiliary scans and worst-word sweeps that motivated the sunflower split, together with the explicit growing-defect full-petal examples, should be retained only as evidence and route cuts. They support the residual decomposition but do not prove polynomial growth for the mixed-petal or growing-defect branches.

Remark A.21 (L1 primitive image-fiber reduction)*.* Prove a polynomial bound for the stabilizer-primitive part of $\mathrm{ImgFib}_U^{<k}(k+\sigma)$ once quotient-periodic codewords are charged separately and the generated-field entropy reserve and lower cutoff clear.

Remark A.22 (mixed-petal and growing-defect residuals)*.* Classify or bound mixed-petal sunflower extras in the sub-Johnson region after the few-petal Johnson layer (Corollary A.13), the distinct-label all-full fixed-excess layer (Proposition A.16), and any certified fixed-excess full-petal layers (Proposition A.18) are paid. Determine whether the full-petal branch with $d - \ell \to \infty$ remains polynomial or produces a new obstruction ledger.

B M1 residue-line tools and Conjecture-F reductions

The safe-side program asks for upper bounds on bad line parameters after tangent, quotient, and extension cells have been charged. In exact-agreement notation,

$$j = n - A, \qquad t = A - k, \qquad j + t = n - k,$$

and the M1 problem is to control aperiodic residue-line packings in this (j, t) window — the residual kernel of the broad band formulation Remark 9.12. The results below are reusable local algebra and combinatorics; they do not prove the full worst-case M1 local limit.

B.1 Split-locator moment calculus

Let $F = \mathbb{F}_q$, let $D \subset F^\times$ have size n, and let $k < A \le n$. Put $t = A - k \ge 1$ and $j = n - A$. For $R \subseteq D$ with $|R| = j$, define

$$\ell_R(X) = \prod_{r \in R} (X - r)$$

and the locator-syndrome vector of a word $w : D \to F$,

$$S_R(w) = \Big(\sum_{x \in D} w(x)\, \ell_R(x)\, x^m \Big)_{m=0}^{t-1} \in F^t.$$

The shifted convention $m = 1, \dots, t$ is equivalent on $D \subset F^\times$: it is the composite of S_R with the linear automorphism $w(x) \mapsto xw(x)$ of F^D, which preserves the uniform measure. For independent uniform words $u, v : D \to F$, call R *aligned* if $S_R(v) \ne 0$ and $S_R(u) \in F\, S_R(v)$, and let $N_A(u, v)$ be the number of aligned locators.

Theorem B.1 (exact first moment)**.**

$$\mathbb{E}N_A = \binom{n}{j}\,(1-q^{-t})\,q^{1-t}.$$

Proof. Fix R and put $E := D \setminus R$, $|E| = n - j = A$. The linear map $w \mapsto S_R(w)$ has matrix entries $\ell_R(x)x^m$; the columns indexed by $x \in R$ vanish, and on E we have $\ell_R(x) \neq 0$, so scaling the column of x by $\ell_R(x)^{-1}$ leaves the $t \times A$ Vandermonde block $(x^m)_{0\le m<t,\ x\in E}$ on distinct points. Since $A = k + t \geq t$, any t of its columns form an invertible Vandermonde matrix, so the map is surjective onto F^t. A surjective linear image of a uniform vector is uniform, and u, v are independent, so $(S_R(u), S_R(v))$ is uniform on $F^t \times F^t$. The number of pairs (a, b) with $b \neq 0$ and $a \in Fb$ is $(q^t - 1)q$, giving alignment probability $(q^t - 1)q/q^{2t} = (1 - q^{-t})q^{1-t}$; summing over the $\binom{n}{j}$ locators proves the formula. □

Theorem B.2 (two-locator joint rank)**.** *For $R, T \subseteq D$ of size j, put $c = |R \cap T|$. The linear map*

$$w \longmapsto (S_R(w), S_T(w)) \in F^t \times F^t$$

has rank

$$2t - \max(0,\, t - j + c).$$

Proof. A linear relation among the $2t$ coordinate functionals is a pair of polynomials $A_0, B_0 \in F[X]_{<t}$ such that $\sum_x w(x)\big(\ell_R(x)A_0(x) + \ell_T(x)B_0(x)\big) = 0$ for all w, i.e. such that

$$\ell_R(X)A_0(X) + \ell_T(X)B_0(X)$$

vanishes on all of D. Its degree is at most $j + t - 1 \leq n - 1 < n$, so it is the zero polynomial. Let $G = \gcd(\ell_R, \ell_T)$, the locator of $R \cap T$, and write $\ell_R = Gr$, $\ell_T = Gs$ with $\deg G = c$, $\deg r = \deg s = j - c$, and $\gcd(r, s) = 1$. The identity becomes $rA_0 + sB_0 = 0$, so $s \mid A_0$ and $r \mid B_0$: $A_0 = sH$, $B_0 = -rH$ for a polynomial H with

$$\deg H \leq t - 1 - (j - c).$$

Hence the relation space has dimension $\max(0,\, t-j+c)$, and the rank is $2t$ minus this nullity. □

Corollary B.3 (covariance support)**.** *The alignment indicators for R and T are independent whenever*

$$d(R, T) := j - |R \cap T| \geq t.$$

All covariance is supported in the radius-$(t-1)$ Johnson neighborhood.

Proof. The condition $d(R, T) \geq t$ is $t - j + |R \cap T| \leq 0$, so by Theorem B.2 the map $w \mapsto (S_R(w), S_T(w))$ is surjective onto F^{2t}; hence $S_R(w)$ and $S_T(w)$ are independent uniform vectors for each of the independent words u, v, and the four syndrome vectors are mutually independent. The alignment event for R depends only on $(S_R(u), S_R(v))$ and that for T only on $(S_T(u), S_T(v))$, so the two events are independent. □

Theorem B.4 (exact second moment)**.** *Put $h(c) := \max(0,\, t - j + c)$ and, for $0 \leq h \leq t$,*

$$P_h := \frac{q\,(q^h - 1)\,q^{2(t-h)} + q^2\,(q^{t-h} - 1)^2}{q^{4t-2h}}.$$

Then

$$\mathbb{E}N_A^2 = \sum_{c=\max(0,2j-n)}^{j} \binom{n}{j}\binom{j}{c}\binom{n-j}{j-c} P_{h(c)}.$$

Proof. Fix an ordered pair of locators (R,T) with $|R\cap T| = c$ and put $h = h(c)$; the number of such ordered pairs is $\binom{n}{j}\binom{j}{c}\binom{n-j}{j-c}$, and c ranges over $\max(0,2j-n) \le c \le j$. It suffices to show that the probability that both R and T are aligned equals P_h.

Normal form of the joint image. Let $V \subseteq F^t \oplus F^t$ be the image of $w \mapsto (S_R(w), S_T(w))$; by Theorem B.2, $\dim V = 2t-h$, and both coordinate projections $\pi_1, \pi_2 : V \to F^t$ are surjective (each single-locator map is surjective, as in the first-moment proof). Put

$$W_1 := \{a \in F^t : (a,0) \in V\}, \qquad W_2 := \{b \in F^t : (0,b) \in V\}.$$

Since π_1 is surjective with kernel $\{0\}\oplus W_2$, we get $\dim W_2 = (2t-h)-t = t-h$, and symmetrically $\dim W_1 = t-h$. The correspondence $a+W_1 \mapsto b+W_2$ for $(a,b) \in V$ is a well-defined linear isomorphism $F^t/W_1 \to F^t/W_2$ between h-dimensional spaces: well-defined because $(a,b),(a,b') \in V$ force $b-b' \in W_2$, injective and surjective by the symmetric statement and the dimension count. Choose splittings $F^t = Z_i \oplus W_i$ with $\dim Z_i = h$, identify $Z_1 \cong Z_2 \cong F^h$ through the isomorphism, and correct the W-components by a linear section of V over Z_1. In the resulting coordinates,

$$V = U_h := \{(z,p;\, z,q) : z \in F^h,\ p,q \in F^{t-h}\},$$

with the first block representing S_R and the second S_T; alignment is invariant under these blockwise invertible changes of coordinates because each acts by an invertible linear map on the corresponding syndrome vector, preserving the relation $S(u) \in F\, S(v)$ and the condition $S(v) \neq 0$.

Counting. The pair of words (u,v) maps to a uniform element of $U_h \times U_h$, of size $q^{2(2t-h)} = q^{4t-2h}$. Write the syndrome of v as $(z,p;z,p')$. If $z \neq 0$ (q^h-1 choices, with p,p' free in $q^{2(t-h)}$ ways), then both $S_R(v)$ and $S_T(v)$ are nonzero, and simultaneous alignment of u forces a common scalar: $S_R(u) = \lambda_R S_R(v)$ and $S_T(u) = \lambda_T S_T(v)$ share the first block $z_u = \lambda_R z = \lambda_T z$, so $\lambda_R = \lambda_T =: \lambda$, and each $\lambda \in F$ gives exactly one admissible syndrome of u, which lies in U_h; this contributes $q(q^h-1)q^{2(t-h)}$ aligned pairs. If $z = 0$, alignment requires $p \neq 0$ and $p' \neq 0$ ($(q^{t-h}-1)^2$ choices), the two scalars are unconstrained and each choice $(\lambda_R,\lambda_T) \in F^2$ gives a distinct admissible syndrome of u inside U_h; this contributes $q^2(q^{t-h}-1)^2$. Dividing the total by q^{4t-2h} gives P_h. □

Corollary B.5 (exact variance ledger)**.** *Let $p := (1-q^{-t})q^{1-t}$. Then*

$$\mathrm{Var}(N_A) = \sum_{c=\max(0,2j-n)}^{j} \binom{n}{j}\binom{j}{c}\binom{n-j}{j-c}\big(P_{h(c)} - p^2\big),$$

and the terms with $j-c \ge t$ vanish.

Proof. By Theorem B.1 each locator is aligned with probability p, and the Vandermonde identity $\sum_c \binom{j}{c}\binom{n-j}{j-c} = \binom{n}{j}$ decomposes $(\mathbb{E}N_A)^2 = \binom{n}{j}^2 p^2$ over the same ordered overlap classes as Theorem B.4; subtracting termwise gives the display. If $j-c \ge t$, then $h(c) = 0$ and $P_0 = q^2(q^t-1)^2/q^{4t} = p^2$, so those covariance terms are zero, in accordance with Corollary B.3. □

Corollary B.6 (dependency-degree concentration)**.** *Let*

$$D_t(n,j) := \sum_{d=0}^{\min(t-1,\,j,\,n-j)} \binom{j}{d}\binom{n-j}{d}$$

be the size of the closed radius-$(t-1)$ *Johnson ball. Then*

$$\mathrm{Var}(N_A) \le \mathbb{E}N_A \cdot D_t(n,j),$$

so $\mathbb{E}N_A \ge D_t(n,j)\, n^s$ *implies* $\mathrm{Var}(N_A)/(\mathbb{E}N_A)^2 \le n^{-s}$*. The coarser sufficient condition* $\mathbb{E}N_A \ge t\, n^{2(t-1)+s}$ *implies the same conclusion.*

Proof. Write $N_A = \sum_R I_R$ with alignment indicators I_R. By Corollary B.3, $\mathrm{Cov}(I_R, I_T) = 0$ unless $d(R,T) \le t-1$; the number of such T for fixed R is $\sum_{d\le t-1}\binom{j}{d}\binom{n-j}{d} = D_t(n,j)$ (choose the d points of R to remove and the d points outside to add). For every pair, $\mathrm{Cov}(I_R, I_T) \le \mathbb{E}(I_R I_T) \le \mathbb{E}I_R$. Hence

$$\mathrm{Var}(N_A) = \sum_{R,T}\mathrm{Cov}(I_R, I_T) \le \sum_R D_t(n,j)\,\mathbb{E}I_R = \mathbb{E}N_A \cdot D_t(n,j),$$

and dividing by $(\mathbb{E}N_A)^2$ gives the exact condition. The coarse condition follows from $D_t(n,j) \le \sum_{d=0}^{t-1} n^{2d} \le t\, n^{2(t-1)}$. □

Remark B.7 (random-word scope)*.* The moment calculus is an average over random word pairs. It does not bound worst-case pairs; those require the aperiodic M1 local limit of the broad band formulation Remark 9.12, a Conjecture-F incidence input, or exchange-rigidity input.

B.2 Conjecture-F reductions

Let K be a field, let $H \subset K$ be finite, and put $P_H(X) = \prod_{h\in H}(X-h)$. Let $\mathcal{D}_j(H)$ be the set of monic squarefree degree-j divisors of P_H, i.e. of locators $L_S(X) = \prod_{h\in S}(X-h)$ with $S \subseteq H$, $|S| = j$. Conjecture-F asks for upper bounds on $|\mathbb{P}(W) \cap \mathcal{D}_j(H)|$ for linear subspaces $W \le K[X]_{\le j}$; the lemmas below pay its reducible strata.

Lemma B.8 (common-GCD reduction)**.** *Let* $W \le K[X]_{\le j}$ *be a nonzero vector space, let* $P = \mathbb{P}(W)$*, and let* $E = P \cap \mathcal{D}_j(H)$*. If every locator in* E *is divisible by a fixed monic divisor* G *of* P_H *with* $\deg G = w$*, then division by* G *injects* E *into*

$$\mathcal{D}_{j-w}\big(H \setminus Z(G)\big)$$

inside a projective linear space of dimension at most $\dim P$*.*

Proof. Division by the fixed polynomial G is a linear injection from the subspace $W \cap GK[X]$ into $K[X]_{\le j-w}$; write W_G for its image, so $\dim \mathbb{P}(W_G) \le \dim P$. If $[L] \in E$ with L its monic representative, then L/G is monic of degree $j-w$ and, since L is squarefree with roots in H and contains all roots of G, the quotient L/G is squarefree with roots exactly $Z(L)\setminus Z(G) \subseteq H\setminus Z(G)$. Thus $[L] \mapsto [L/G]$ is the claimed injection into $\mathbb{P}(W_G) \cap \mathcal{D}_{j-w}(H \setminus Z(G))$. □

Lemma B.9 (quotient-pullback recursion)**.** *Assume $H = \mu_n$ and $\operatorname{char} K \nmid n$. If $M \mid \gcd(n, j)$, then*

$$g(Y) \longmapsto g(X^M)$$

maps $\mathcal{D}_{j/M}(\mu_{n/M})$ bijectively onto the M-periodic stratum of $\mathcal{D}_j(\mu_n)$, i.e. onto the locators whose root sets are unions of complete μ_M-cosets. Intersecting with a projective linear space descends to an intersection with a projective linear space of no larger dimension in the quotient locator space.

Proof. Because $\operatorname{char} K \nmid n$, the polynomial $X^n - 1$ is squarefree, and the M-power map $\mu_n \to \mu_{n/M}$ is surjective with all fibers of size M. A monic squarefree divisor g of $Y^{n/M} - 1$ of degree j/M pulls back to $g(X^M)$, a monic divisor of $X^n - 1$ of degree j whose root set is the union of the fibers over $Z(g)$, hence M-periodic and squarefree. Conversely, an M-periodic squarefree divisor L of $X^n - 1$ has root set a union of fibers; its image $S \subseteq \mu_{n/M}$ has size j/M, and $g := \prod_{b \in S}(Y - b)$ satisfies $g(X^M) = L$ (both are monic with the same roots inside the squarefree $X^n - 1$). This proves the bijection. For the descent, the space $\{g : g(X^M) \in W\}$ is linearly isomorphic to $W \cap K[X^M]$ by injectivity of the substitution, hence has dimension at most $\dim W$. □

Lemma B.10 (dimension-one voting)**.** *Assume $j \geq 1$. Let $P = \mathbb{P}(\operatorname{span}(L_0, L_1))$ be a projective line in $K[X]_{\leq j}$, and assume no $h \in H$ is a common zero of L_0 and L_1. Then*

$$|P \cap \mathcal{D}_j(H)| \leq \left\lfloor \frac{|H|}{j} \right\rfloor.$$

Proof. For each $h \in H$, the functional $[a : b] \mapsto aL_0(h) + bL_1(h)$ is nonzero by the common-zero hypothesis, so it has a unique kernel point on the projective line; say that h votes for this point. A locator point $[aL_0 + bL_1] \in P \cap \mathcal{D}_j(H)$ has exactly j roots in H and receives a vote precisely from each of them. There are $|H|$ votes in total, so the number of locator points is at most $|H|/j$, hence at most the floor. □

Lemma B.11 (hyperplane concurrency)**.** *Let $W \leq K[X]_{\leq j}$ be nonzero and assume no $h \in H$ is a common root of all elements of W. For each h, let*

$$E_h := \{[L] \in \mathbb{P}(W) : L(h) = 0\},$$

a projective hyperplane of $\mathbb{P}(W)$. Then $\mathbb{P}(W) \cap \mathcal{D}_j(H)$ is exactly the set of projective points lying on at least j of the hyperplanes E_h.

Proof. If $[L] \in \mathbb{P}(W) \cap \mathcal{D}_j(H)$, then L has exactly j distinct roots in H, so $[L]$ lies on the corresponding j hyperplanes. Conversely, if a nonzero $L \in W$ vanishes at j or more distinct points of H, then $\deg L \leq j$ forces L to have exactly j roots, all simple and all among those points; after monic normalization, $L \in \mathcal{D}_j(H)$. □

Theorem B.12 (projective-plane pair bound)**.** *Under the hypotheses of Lemma B.11, assume $j \geq 2$ and $\dim \mathbb{P}(W) = 2$. Then*

$$|\mathbb{P}(W) \cap \mathcal{D}_j(H)| \leq \frac{\binom{|H|}{2}}{j - 1},$$

and if the evaluation lines E_h are pairwise distinct, the sharper bound with denominator $\binom{j}{2}$ holds.

Proof. Here the hyperplanes E_h are lines in the projective plane $\mathbb{P}(W)$; distinct lines meet in exactly one point. Group the E_h into coincidence classes. No class has multiplicity $\geq j$: if $E_{h_1} = \cdots = E_{h_j}$ for distinct $h_1, \ldots, h_j$, this common line is a two-dimensional projective space of polynomials of degree $\leq j$ vanishing at the j fixed points $h_1, \ldots, h_j$; but that vanishing space is spanned by the single locator $\prod_i (X - h_i)$, hence at most zero-dimensional projectively — a contradiction.

Now let $[L]$ be a locator point and let $m_1, \ldots, m_s$ be the multiplicities of the line classes through $[L]$ among the j lines E_h, $h \in Z(L)$, so $\sum_i m_i = j$ and each $m_i \leq j - 1$; in particular $s \geq 2$. Call a pair $\{h, h'\} \subseteq Z(L)$ a *cross pair* if $E_h \neq E_{h'}$. Fixing any one class, of multiplicity m, the number of cross pairs at $[L]$ is at least $m(j - m)$, which for integers $1 \leq m \leq j - 1$ is minimized at the endpoints with value $j - 1$. Each cross pair $\{h, h'\}$ determines the unique intersection point $E_h \cap E_{h'}$, so it can be charged to at most one locator point. Since there are at most $\binom{|H|}{2}$ pairs of points of H in total, the number of locator points is at most $\binom{|H|}{2}/(j-1)$. When all E_h are pairwise distinct, every locator point is incident to at least j distinct lines, hence carries at least $\binom{j}{2}$ cross pairs, giving the sharper denominator. □

Theorem B.13 (fixed-dimensional Conjecture-F bound)**.** *Let $W \leq K[X]_{\leq j}$ have vector-space dimension $d + 1$ with $0 \leq d \leq j$, and assume W is gcd-trivial on H, i.e. no point of H is a common root of all elements of W. Then*

$$|\mathbb{P}(W) \cap \mathcal{D}_j(H)| \leq \binom{|H|}{d}.$$

More generally, if $C_0 \subseteq H$ is the common root set of W on H and $d \leq j - |C_0|$, then

$$|\mathbb{P}(W) \cap \mathcal{D}_j(H)| \leq \binom{|H| - |C_0|}{d}.$$

Proof. First assume gcd-triviality and fix a total order on H. For a locator point $[L]$ with root set $R \subseteq H$, $|R| = j$, the vanishing subspace

$$W(-R) := \{G \in W : G|_R = 0\}$$

equals $K \cdot L$: it contains L, and any $G \in W \subseteq K[X]_{\leq j}$ vanishing at the j points of R is a scalar multiple of the degree-j locator of R.

Process the roots of R in the fixed order, imposing one vanishing condition at a time, and record the roots at which the dimension strictly drops; call this set $S_0 \subseteq R$. At every stage the current subspace equals $W(-R')$ for the set R' of processed roots, and a skipped root leaves the subspace unchanged, so by induction the current subspace equals $W(-S')$ for the set S' of *kept* roots so far; in particular $W(-S_0) = W(-R) = K \cdot L$ after all of R is processed. Each vanishing condition drops the dimension by at most one, and the total drop is $(d+1) - 1 = d$, so $|S_0| = d$ exactly. The assignment $[L] \mapsto S_0$ is therefore a map into the d-subsets of H, and it is injective because S_0 determines $W(-S_0) = K \cdot L$, i.e. the point $[L]$. This gives the bound $\binom{|H|}{d}$.

For the general version, every element of W vanishes on C_0 and has degree $\leq j$, so $W \subseteq G_{C_0} K[X]$ for $G_{C_0} := \prod_{h \in C_0} (X - h)$; moreover every locator in the intersection vanishes on C_0, hence is divisible by G_{C_0}. By Lemma B.8, division by G_{C_0} injects the intersection into a gcd-trivial instance on $H \setminus C_0$ with locator degree $j - |C_0|$ and the same vector-space dimension $d+1$ (gcd-trivial because a new common root in $H \setminus C_0$ would have been a common root of W). The first part applies whenever $d \leq j - |C_0|$ and gives $\binom{|H| - |C_0|}{d}$. □

Remark B.14 (what remains to be settled). Theorem B.13 shows that the primitive Conjecture-F core is a dimension-growing problem: common-root, quotient-periodic, dimension-one, and projective-plane strata are paid by Lemmas B.8 to B.10 and Theorem B.12, and every fixed-dimensional stratum is paid by Theorem B.13, but the growing-dimensional incidence theorem needed for the aperiodic band is not settled here and is part of the broad band formulation Remark 9.12.

B.3 Quotient seam arithmetic and Hankel tools

Lemma B.15 (GAP-2 quotient seam). *Let n, k, A be fixed and put $j = n - A$, $t = A - k$, and $r = n - k$, so $j + t = r$. If $M \mid \gcd(n, k)$, then $M \mid r$, and hence*

$$M \mid j \quad \Longleftrightarrow \quad M \mid t.$$

On such buckets all quotient parameters $n/M, k/M, A/M, j/M, t/M$ are integral. If $M \mid n$ and $M \mid j$ but $M \nmid k$, then the support descends but the dimension and window do not: a non-rate-preserving seam.

Proof. Since $M \mid n$ and $M \mid k$, we have $M \mid n - k = r = j + t$, so $j \equiv -t \pmod{M}$, proving the equivalence, and then $A = n - j$ and t descend whenever j does. If $M \mid n$ and $M \mid j$ but $M \nmid k$, then $A = n - j$ is divisible by M while k/M and $t/M = (A - k)/M$ are not integral. □

For a finite ordered set X of nonzero field elements and $m \geq 1$, write $V_m(X)$ for the matrix with rows indexed by $x \in X$, columns $0 \leq r < m$, and entries x^r.

Proposition B.16 (Hankel factorization). *Let $H \subseteq F^\times$ be finite. For a word $w : H \to F$, define the rectangular syndrome Hankel block*

$$H_{t,s}(w)_{a,b} := \sum_{x \in H} w(x) x^{a+b} \qquad (0 \leq a < t,\ 0 \leq b < s).$$

Then

$$H_{t,s}(w) = V_t(H)^{\mathsf{T}} \operatorname{diag}\big(w(x)\big) V_s(H).$$

If $X = \{x_0, \ldots, x_{m-1}\}$ is a set of distinct nonzero points and

$$A_h(X) := V_m(X)^{\mathsf{T}} \operatorname{diag}(x^h : x \in X)\, V_m(X),$$

then

$$\det A_h(X) = \det\big(V_m(X)\big)^2 \prod_{x \in X} x^h \ \neq 0.$$

Proof. The (a, b) entry of the triple product is $\sum_x x^a w(x) x^b$, the defining Hankel entry. The determinant identity is multiplicativity of the determinant together with $\det(V_m(X)^{\mathsf{T}}) = \det V_m(X)$, and the Vandermonde determinant on distinct points is nonzero. □

Proposition B.17 (Johnson exchange mixing). *Let $1 \leq j \leq n - 1$, and let $J(n, j)$ be the Johnson graph on j-subsets of an n-set, with degree $d = j(n - j)$. If $\mathcal{A}$ is a family of j-subsets of density δ and T_1 is obtained from a uniform T_0 by one uniform exchange step, then*

$$\Pr[T_0 \in \mathcal{A},\ T_1 \in \mathcal{A}] \leq \delta^2 + \Big(1 - \frac{n}{d}\Big)\delta(1 - \delta),$$

and point dictators $\mathcal{A} = \{T : x_0 \in T\}$ attain equality.

Proof. Let P be the normalized adjacency operator of $J(n,j)$ and write $1_{\mathcal{A}} = \delta\mathbf{1} + f$ with $f \perp \mathbf{1}$ and $\|f\|_2^2 = \delta(1-\delta)$ under the uniform measure. The eigenvalues of the Johnson graph are classically

$$\theta_i = (j-i)(n-j-i) - i, \qquad 0 \le i \le \min(j, n-j),$$

the association-scheme spectrum of $J(n,j)$. For $i < \min(j, n-j)$,

$$\theta_i - \theta_{i+1} = (j-i)(n-j-i) - (j-i-1)(n-j-i-1) + 1 = n - 2i > 0,$$

so the sequence is strictly decreasing and the largest nontrivial normalized eigenvalue is $\lambda_1 = \theta_1/d = 1 - n/d$. Hence

$$\Pr[T_0, T_1 \in \mathcal{A}] = \langle 1_{\mathcal{A}}, P 1_{\mathcal{A}} \rangle = \delta^2 + \langle f, Pf \rangle \le \delta^2 + \lambda_1 \|f\|_2^2,$$

which is the display; only the largest nontrivial eigenvalue is needed for this upper bound. The centered indicator of a point dictator lies in the $i = 1$ eigenspace, so equality is attained there. □

Proposition B.18 (anticode packing)**.** *Let $0 \le s \le j \le n$. If a family $\mathcal{F}$ of j-subsets of an n-set satisfies $|A \setminus B| > s$ for all distinct $A, B \in \mathcal{F}$, then*

$$|\mathcal{F}| \binom{j}{s} \le \binom{n}{j-s}.$$

Proof. For each $A \in \mathcal{F}$, consider its $(j-s)$-subsets, $\binom{j}{j-s} = \binom{j}{s}$ of them. These collections are disjoint across $\mathcal{F}$: a common $(j-s)$-subset of $A \ne B$ would give $|A \cap B| \ge j - s$, hence $|A \setminus B| \le s$, contrary to hypothesis. There are at most $\binom{n}{j-s}$ subsets of size $j-s$ in total. □

B.4 SPI deficiency-one chart

Let

$$M(Z) = M_0 + Z M_1$$

be a $t \times (j+1)$ matrix pencil with affine-linear entries over F, and assume $t = j$. For $0 \le i \le j$, let $M_i(Z)$ be the maximal minor obtained by deleting column i, put

$$c_i(Z) = (-1)^i M_i(Z), \qquad L_Z(X) = \sum_{i=0}^{j} c_i(Z) X^i,$$

and recall that whenever $\operatorname{rank} M(z) = j$, the cofactor vector $(c_0(z), \dots, c_j(z))$ is nonzero and spans $\ker M(z)$. In the syndrome application, $M(Z)$ is the exact-agreement syndrome pencil of Section 3.4 in the deficiency-one shape $t = j$, the kernel vector is the coefficient vector of a candidate locator of degree $\le j$, and a slope z *passes the split test* if the normalized kernel locator divides $X^n - 1$; as in Lemma 9.8, passing the split test is necessary for badness at the exact agreement but not sufficient.

Theorem B.19 (deficiency-one eliminant or residual)**.** *Assume H is the full root set of $X^n - 1$ and* $\operatorname{char} F \nmid n$.

(0) *If every maximal minor of $M(Z)$ vanishes identically, the pencil is identically rank-deficient; this chart is a named residual branch, to be classified by the polynomial-kernel chain of Lemma 9.8, and the statements below do not apply to it.*

Otherwise fix a maximal minor that is not identically zero. Every finite slope z then lies in exactly one of the classes:

(i) *rank drop: all maximal minors vanish at z; this set is contained in the root set of the chosen minor, a nonzero polynomial of degree at most t;*

(ii) *low-degree full-rank chart:* $\operatorname{rank} M(z) = j$ *and* $c_j(z) = 0$*; the Cramer locator has degree $< j$ and is charged to a contained higher-agreement branch;*

(iii) *top chart: $c_j(z) \neq 0$; the Cramer locator is proportional to a monic degree-j locator, and z passes the split test if and only if $L_z(X) \mid X^n - 1$.*

If $c_j \equiv 0$, the top chart is empty and every full-rank slope lies in class (ii)*. Assume $c_j \not\equiv 0$ and perform pseudo-division over $F[Z]$:*

$$c_j(Z)^{\Delta}\,(X^n - 1) = Q(Z, X)\,L_Z(X) + R(Z, X), \qquad \Delta = n - j + 1,$$

with $\deg_X R < j$ *and* $\deg_Z R_m \leq (n - j + 1)t$ *for every X-coefficient $R_m(Z)$ of R. Then exactly one of the following holds:*

(a) *some R_m is not identically zero; then the polynomial $c_j(Z)R_m(Z)$, of degree at most*

$$t + (n - j + 1)t,$$

is a finite eliminant whose root set contains every top-chart slope passing the split test together with every slope of classes (i) *and* (ii)*;*

(b) *all R_m vanish identically; then every top-chart slope passes the split test, and the chart must be retained as a named residual branch rather than discarded.*

Proof. Class (0) and the trichotomy are immediate from the definitions: each maximal minor is the determinant of a $j \times j$ matrix with affine-linear entries in Z, hence a polynomial of degree at most $j = t$; if all vanish at z then $\operatorname{rank} M(z) < j$, and otherwise the kernel is one-dimensional and spanned by the nonzero cofactor vector, whose top coordinate $c_j(z)$ decides between (ii) and (iii). In class (iii), L_z has degree exactly j; since $X^n - 1$ is squarefree with root set H, its monic normalization is H-split if and only if it divides $X^n - 1$, which is the split test.

For the pseudo-division, dividing the degree-n polynomial $X^n - 1$ by the degree-j polynomial L_Z with leading coefficient $c_j(Z)$ performs $\Delta = n - j + 1$ elimination steps, each multiplying the running remainder by $c_j(Z)$ once and subtracting a $c_i(Z)$-multiple of a shift of L_Z; since every c_i has Z-degree at most t, each step raises the Z-degree of the remainder coefficients by at most t, giving $\deg_Z R_m \leq \Delta t = (n - j + 1)t$ from the degree-0 start.

In case (a), let z be a top-chart slope passing the split test. Substituting $Z = z$, the left-hand side $c_j(z)^{\Delta}(X^n - 1)$ is divisible by L_z, and so is $Q(z, X)L_z(X)$; hence $L_z(X) \mid R(z, X)$, and $\deg_X R(z, X) < j = \deg L_z$ forces $R(z, X) = 0$, so $R_m(z) = 0$ for every m, in particular for the chosen nonzero one. Slopes of classes (i) and (ii) satisfy $c_j(z) = 0$ (at rank drop, all minors vanish, in particular M_j). Hence all these slopes are roots of $c_j R_m$, whose degree is at most $t + (n - j + 1)t$. In case (b), substituting any top-chart slope gives $c_j(z)^{\Delta}(X^n - 1) = Q(z, X)L_z(X)$ with $c_j(z) \neq 0$, so L_z divides $X^n - 1$: the split test is identically passed on the top chart, which is then a genuine residual family. □

Remark B.20 (relation to the main eliminant machinery, and calibration). Theorem B.19 complements the certificate framework of Section 3.4: Theorem 3.59 and Theorem 3.73 bound chart contributions *given* a nonzero eliminant, while Theorem B.19 *produces* one for every deficiency-one chart with explicit degree, or else names the residual branch (identically rank-deficient pencils, handled by Lemma 9.8, and identically split top charts). As calibration: for $n = 512$, $k = 256$, $A = 384$, one has $t = j = 128$, and the global one-parameter cap of case (a) is

$$t + (n - j + 1)t = 128 + 385 \cdot 128 = 49408,$$

a bounded eliminant problem. The theorem does not classify higher-deficiency SPI charts; those remain inside the broad band formulation Remark 9.12.

C Frontier prize and quotient flatness

The identity-scale floor of Lemma 14.1 and Proposition 14.2 changes the shape of the remaining problem. The unsafe side is no longer the main source of leverage: the best printed lower certificate is already the entropy prefix floor, and the gap between the deployed rows and the asymptotic ceiling of Definition 14.4 is at the scale of a few parts in 10^{-6}. The prize-winning statement is therefore a safe-side theorem. It must prove that no upper mechanism beats the identity-prefix average except through structures already paid by the capf ledger.

Remark C.1 (First-form frontier formulation). Prove the asymptotic identity-prefix frontier

$$\delta_C^*(\varepsilon^*) = 1 - \rho - g^*(\rho, \log_2 |\mathbb{B}|) + o(1)$$

for the smooth multiplicative and circle line-round rows at the deployed challenge budgets. In finite deployed form, prove the one-step safe side at the four adjacent candidates of Conjecture 14.6: namely, prove safety at agreements

$$1116048, \quad 1116047, \quad 1116024, \quad 1116023$$

for the KoalaBear MCA, KoalaBear list, Mersenne-31 MCA, and Mersenne-31 list rows respectively, while the preceding agreements remain unsafe by Proposition 15.2 for MCA and Proposition 14.2 for lists.

Remark C.2 (separation between asymptotic and finite prizes). The asymptotic prize and the finite adjacent prize should be kept separate. A factor n^C in any residual component costs $C \log_2 n$ bits; at $n = 2^{21}$ this is $21C$ bits, larger than several adjacent fail margins in Proposition 14.2 and Remark 14.18. Polynomial losses are compatible with an asymptotic closure only after the complete active certificate, coverage, coalescing, and the reserve conditions of Remark 21.5 have been supplied. A printed adjacent pair needs exact constants, integer rounding, and the single summed budget.

C.1 First split of the safe side

The first useful decomposition separates prefix/quotient fibers from the aperiodic residue. This split is falsifiable and is used by the subsequent reductions, but it is not the final active interface; after slope elimination and the split-pencil normalization it is read through Section 12.1.

Remark C.3 (First-form quotient/prefix formulation). Let $D \subseteq \mathbb{B} \subseteq \mathbb{F}$, $|D| = n$, $K = \rho n$, $m = K + w$, and define the identity prefix map

$$\Phi_{m,w} : \binom{D}{m} \longrightarrow \mathbb{B}^w, \qquad M \longmapsto \big((-1)^h e_h(M)\big)_{1 \le h \le w}.$$

After quotient-periodic components and declared paid scales have been separated, prove the primitive discrete-flatness estimate

$$\max_{z \in \mathbb{B}^w} |\Phi_{m,w}^{-1}(z)| \le L_Q(n) \left(1 + \frac{\binom{n}{m}}{|\mathbb{B}|^w}\right).$$

The additive 1 is indispensable once the average fiber is below one. In an asymptotic closure L_Q is consumed only with a reserve satisfying $R \to \infty$, $R = o(n)$, and $\log L_Q = o(R \log |\mathbb{B}|)$. The finite deployed version must replace L_Q by an explicit multiplier and additive error whose exact bit cost fits into the summed one-step upper ledger at $a_0 + 1$.

Remark C.4 (historical first-form aperiodic umbrella). An earlier organization placed every unpaid primitive aperiodic M1/L1 configuration into one umbrella term and sought

$$N_{\mathrm{ap}}(A) \le L_A(n) N_{\mathrm{model}}(A).$$

This remains useful as proof-route shorthand only. To instantiate it above the entropy–subfield envelope, the model must include the corrected discrete/soft base-field census, the proof must control the local primitive same-prefix contribution, and the compiler must prove exhaustive coverage and support-to-parameter coalescing. An exact deployed version is therefore not a single unspecified integer U_A: it is the explicit sum $U_{\mathrm{BC}} + U_{\mathrm{sp}}$ together with any coverage correction, as required by Definition 21.1. The active status statement is Remark 21.5.

Remark C.5 (historical (Q)+(A) implication). If the historical umbrella (A) were proved in a form that already included the corrected base-field floors, local primitive control, exhaustive coverage, and parameter coalescing, then (Q)+(A) would merely repackage the active certificate. In that case the asymptotic loss is absorbable only for $R \to \infty$, $R = o(n)$, and $\log L = o(R \log |\mathbb{B}|)$, while a finite adjacent theorem still requires

$$U_{\mathrm{paid}}(a_0 + 1) + U_Q(a_0 + 1) + U_{\mathrm{BC}}(a_0 + 1) + U_{\mathrm{sp}}(a_0 + 1) \le B_* < L(a_0).$$

This is not an additional closure theorem and supplies none of the missing certificate data.

C.2 Collision-gap formulation of (Q)

Let $N_m := \binom{n}{m}$, $Q_w := |\mathbb{B}|^w$, and

$$S_{m,w} := \min\{N_m, Q_w\}.$$

For the probability distribution $\mu_{m,w}$ of $\Phi_{m,w}(M)$ on a uniformly random m-subset, define the discrete-normalized collision gaps

$$\Gamma_r(m,w) := S_{m,w}^{r-1} \sum_{z \in \mathbb{B}^w} \mu_{m,w}(z)^r, \qquad r \ge 2.$$

A distribution uniform on $S_{m,w}$ attainable cells has $\Gamma_r = 1$, and $\Gamma_r \ge 1$ for every prefix distribution because its support has at most $S_{m,w}$ points. This normalization agrees with the usual codomain-uniform one while $N_m \ge Q_w$, and switches to the injective baseline once $N_m < Q_w$.

Remark C.6 (Fixed-moment collision hierarchy). In the primitive frontier window, prove a quantitative hierarchy strong enough to imply the first-form quotient input Remark C.3. A useful target is

$$\Gamma_r(m,w) \le n^{Cr} \quad \text{for all } 2 \le r \le c \log S_{m,w},$$

with constants uniform in the row, or an equivalent tail estimate implying

$$\max_z \mu_{m,w}(z) \le n^C S_{m,w}^{-1} \le n^C \left(N_m^{-1} + Q_w^{-1}\right).$$

For exact deployment, print the finite moment or direct-tail constants and show that the induced discrete max-fiber bound fits inside the summed adjacent ledger.

Remark C.7 (why fixed low moments are not enough). A bound on Γ_r for fixed r proves useful bulk equidistribution but does not by itself imply the max-fiber estimate. Indeed

$$\max_z \mu_{m,w}(z) \le \Gamma_r(m,w)^{1/r} S_{m,w}^{-(1-1/r)}.$$

Thus polynomial control at the effective discrete baseline requires moment order comparable to $\log S_{m,w}/\log n$, or a separate tail or inverse theorem. This is why the collision-gap programme must measure the whole joint tail, not only the second moment.

C.3 Stable Fourier measurement protocol

The normalized elementary-symmetric recursion gives a numerically stable way to measure Fourier coefficients of the prefix distribution without ever forming the enormous coefficients e_m.

Proposition C.8 (non-expansive normalized recursion). *Let $\varepsilon_1, \dots, \varepsilon_N \in \mathbb{C}$ satisfy $|\varepsilon_i| \le 1$. Define*

$$q_j^{(i)} := \frac{1}{\binom{i}{j}} \sum_{1 \le r_1 < \cdots < r_j \le i} \varepsilon_{r_1} \cdots \varepsilon_{r_j}, \qquad 0 \le j \le i,$$

with $q_0^{(i)} = 1$. Then, for $1 \le j \le i$,

$$q_j^{(i)} = \left(1 - \frac{j}{i}\right) q_j^{(i-1)} + \frac{j}{i}\, \varepsilon_i q_{j-1}^{(i-1)}.$$

In particular $|q_j^{(i)}| \le 1$ for all i, j.

Proof. The elementary symmetric recursion is

$$e_j^{(i)} = e_j^{(i-1)} + \varepsilon_i e_{j-1}^{(i-1)}.$$

Dividing by $\binom{i}{j}$ and using

$$\frac{\binom{i-1}{j}}{\binom{i}{j}} = 1 - \frac{j}{i}, \qquad \frac{\binom{i-1}{j-1}}{\binom{i}{j}} = \frac{j}{i}$$

gives the formula. If $|q_j^{(i-1)}|, |q_{j-1}^{(i-1)}| \le 1$, the displayed formula is a convex combination of two complex numbers of modulus at most one. The claim follows by induction. □

For a frequency vector $\tau = (\tau_1, \dots, \tau_w)$ over a prime field and an additive character ψ, put

$$\varepsilon_x = \psi\left(\sum_{r=1}^{w} \tau_r x^r\right), \qquad x \in D.$$

Then

$$q_m^{(n)}(\varepsilon_x : x \in D) = \frac{1}{\binom{n}{m}} \sum_{|M|=m} \psi\left(\sum_{r=1}^{w} \tau_r \sum_{x \in M} x^r\right),$$

which is the normalized Fourier coefficient of the power-sum prefix distribution. When $w < \operatorname{char} \mathbb{F}$, Newton identities give a triangular bijection between the first w power sums and the first w elementary symmetric prefixes, so the same measurement probes the first-form quotient input Remark C.3.

Remark C.9 (Weil barrier and the reason to measure joint cancellation). The per-power-sum Weil bound is the wrong endpoint for (Q). It pays a visible $\sqrt{p}$ cost per active power-sum direction, and the bounded-prefix base-case theorem Theorem 14.8 shows exactly how far that method reaches at head depths. The computations described in Remark 14.12 indicate that the second-moment refinement over a full period returns the same square-root scale: individual character maxima are controlled by the wrong statistic. The normalized recursion should therefore be used to measure joint tails and collision gaps, not to claim a uniform character bound strong enough to imply (Q) by itself.

C.4 Heuristic picture and falsification tests

The current heuristic is as follows. In the dense primitive bulk, the prefix map behaves like an almost-random constant-weight syndrome map. Quotient-periodic directions account for the visible heavy Fourier modes, and those are precisely the directions absorbed into coarser scales by Proposition 14.10. After those scales are removed, the remaining primitive fibers should have only polynomial max-to-mean inflation. On the M1 side, the split-locator ledger of Theorems B.1 and B.4 and Corollary B.5 says that random residual correlations are supported only in a bounded Johnson neighborhood; the missing safe-side theorem is an inverse statement that every worst-case excess correlation has a paid structural cause.

A practical falsification programme should therefore print, for each tested scale:

(i) max-to-mean prefix fiber ratios in the dense bulk, separated by primitive and composite frequency support;

(ii) Γ_r values for increasing r, with enough growth in r to test the max-fiber prediction rather than only the second moment;

(iii) normalized-recursion Fourier tail plots or tables, with divisor-lattice buckets for composite directions;

(iv) the exact paid scale to which every visible structured spike is absorbed.

A counterexample to the active certificate programme should then be classifiable: a super-polynomial primitive prefix fiber; an oversized corrected interior split-pencil census; an oversized local primitive shift-pair degree; or a defect in the proposed bucket coverage or support-to-parameter coalescing.

D Safe-side proof routes

The frontier package leaves explicit upper inputs, formalized in Sections 12 and 21. This section records proof routes and intermediate reductions for those inputs. The symbols (A) and (B/Q) are used locally for the first aperiodic and quotient split; the authoritative status statement is the conditional three-bucket certificate template of Section 12.1.

D.1 First-form contracts

First-form aperiodic contract. Fix a multiplicative or circle-line row $C = \mathrm{RS}[\mathbb{F}, D, k]$, an agreement value A, and write

$$j = n - A, \qquad t = A - k, \qquad j + t = n - k.$$

For a received affine line $f + zg$, let

$$\mathrm{Bad}_{\mathrm{ap}}^{\mathrm{M1}}(f, g; A)$$

be the set of finite slopes z admitting a support-wise noncontained explanation at agreement at least A after the tangent cell, declared quotient cells, extension-pole cells, common-GCD strata, quotient-pullback strata, and printed fixed-dimensional Conjecture-F strata have been paid. Every remaining witness support is required to be aperiodic in the sense of Definition 9.1.

As a historical umbrella target, an asymptotic (A) theorem would prove a row-uniform polynomial bound

$$\sup_{f,g} \left|\mathrm{Bad}_{\mathrm{ap}}^{\mathrm{M1}}(f, g; A)\right| \le n^{C_A}$$

throughout the frontier band of the broad band formulation Remark 9.12, with C_A independent of q. The deployed finite theorem must be sharper: for each printed adjacent candidate $a_0 + 1$, it must supply an explicit integer

$$U_A(a_0 + 1) \ge \sup_{f,g} \left|\mathrm{Bad}_{\mathrm{ap}}^{\mathrm{M1}}(f, g; a_0 + 1)\right|$$

whose proof is then decomposed into the active exact numerators and included in

$$U_{\mathrm{paid}}(a_0 + 1) + U_Q(a_0 + 1) + U_{\mathrm{BC}}(a_0 + 1) + U_{\mathrm{sp}}(a_0 + 1) \le B_*.$$

Thus the symbol U_A is only historical shorthand for proved residual components; it is not an additional active bucket.

First-form quotient contract. For the quotient-remainder prefix map of Definition 3.6, write

$$\Phi_{c,m,s}^{K} : \mathcal{X}_{c,m,s} \longrightarrow \mathbb{B}^{w_c(s, A_0 - K)}, \qquad \mu(\xi) = \left|{\Phi_{c,m,s}^{K}}^{-1}(\xi)\right|,$$

where $\mathcal{X}_{c,m,s}$ is the finite set of quotient-remainder supports counted by $M_{c,m,s}$. In the identity-scale case of Lemma 14.1, this is simply

$$\Phi_{1,m,0}^{K}(M) = \big((-1)^h e_h(M)\big)_{1 \le h \le m-K}, \qquad |M| = m.$$

The quotient-fiber upper theorem should bound the heaviest prefix fiber

$$H_{c,m,s}^{K} := \max_{\xi} \mu(\xi)$$

from above, not only from below by pigeonhole. The asymptotic grade may allow

$$H^K_{c,m,s} \le n^{C_Q}\left(1+\frac{M_{c,m,s}}{|\operatorname{im}\Phi^K_{c,m,s}|}\right)$$

or the corresponding mean-plus-one ambient-box model when an image-size theorem is also supplied. This loses only $C_Q \log_2 n$ bits and is therefore usable away from an $O(\log n)$ reserve.

The deployed finite grade must instead print constants. For each adjacent candidate it should give a certified inequality of the form

$$H^K_{c,m,s} \le \Lambda_Q(c,m,s,K)\left(1+\frac{M_{c,m,s}}{|\operatorname{im}\Phi^K_{c,m,s}|}\right)+E_Q(c,m,s,K)$$

with exact integer rounding and with

$$\log_2\Lambda_Q+\log_2\left(1+\frac{E_Q}{1+M_{c,m,s}/|\operatorname{im}\Phi^K_{c,m,s}|}\right)$$

small enough to fit inside the printed adjacent fail margin after the paid, interior-census, and local shift-pair terms have been added. Since the smallest deployed adjacent margin in Proposition 14.2 and Remark 14.18 is only a few bits, an unspecified poly(n) multiplier is not a finite certificate.

D.2 The normalized aperiodic input

A1. Deterministic local-limit upgrade from the split-locator moment ledger. The moment ledger of Theorems B.1 and B.4 and Corollaries B.5 and B.6 gives the correct random-word model but not a worst-case theorem. The proposed upgrade is an inverse theorem.

A1.1 Convert every bad slope to at least one co-support locator L_T with $|T| = j$ satisfying the fixed-locator Hankel equation

$$H_{t,j}(\mathrm{Syn}(f+zg))\boldsymbol{\ell}_T = 0.$$

Use a deterministic tie-breaking rule so that each slope pays one canonical primitive locator after all paid branches are removed.

A1.2 Prove a density dichotomy for the chosen locator set inside the Johnson graph on j-subsets: either it has random-scale local intersections, in which case the dependency-degree calculation forces a polynomial number of distinct slopes, or it has too many short Johnson exchanges.

A1.3 In the second case, apply the exchange identities of Proposition B.17 and the anticode obstruction of Proposition B.18. Long exchange chains should force one of the already paid structures: common GCD, quotient-pullback periodicity, tangent containment, extension-pole containment, or an identically deficient SPI chart.

A1.4 The deliverable is a deterministic replacement for the random-pair variance sentence: after paid structures are removed, every worst-case pair behaves locally like the random split-locator model up to $n^{O(1)}$ exceptional choices.

The advantage of this route is that it reuses the strongest capf calculations already printed. The risk is the inverse step: one must prove that excess local correlation has a structural cause, not merely that it is rare on average.

A2. Growing-dimensional primitive Conjecture-F incidence. For a fixed slope z, the locator equation cuts a projective linear space

$$\mathbb{P}(W_z) \subseteq \mathbb{P}(K[X]_{\le j})$$

and the explanations at that slope are precisely the squarefree locator points in

$$\mathbb{P}(W_z) \cap \mathcal{D}_j(D)$$

after the usual locator normalization. The fixed-dimensional and projective-plane cases are already paid in Theorems B.12 and B.13 and Lemma B.10; the missing part is the primitive growing-dimensional case.

A proof can be organized as follows.

A2.1 Use Lemmas B.8, B.9 and B.15 to remove common-root and quotient-periodic components before the incidence theorem is invoked. This ensures that the incidence problem is genuinely primitive.

A2.2 Prove a Hilbert-function or polynomial-method bound for

$$|\mathbb{P}(W) \cap \mathcal{D}_j(D)|$$

when W has growing dimension but contains no paid subspace forced by quotient pull-back, large common divisor, or low-dimensional projection.

A2.3 Apply the bound uniformly to the pencil W_z as z varies. If too many z give nonempty primitive intersections, eliminate z from the minors defining W_z; the resulting identically singular or identically split charts are precisely the named SPI residuals of Theorem B.19, and the non-identical charts contribute by their printed degree caps.

A2.4 The target output is either

$$\#\{z : \mathbb{P}(W_z) \cap \mathcal{D}_j(D)_{\mathrm{prim}} \neq \varnothing\} \le n^{C_A}$$

for asymptotic closure, or an exact degree-sum table for the deployed adjacent rows.

This route is the cleanest conceptual proof of (A) because it attacks exactly the frontier sentence in the proof guide: the growing-dimensional primitive incidence problem inside the broad band formulation Remark 9.12.

A3. Exchange-rigidity and Johnson-spectral route. Instead of bounding each projective incidence independently, one can exploit the fact that many bad slopes produce many locators with constrained exchanges. The proof target is a rigidity theorem of the following form: a large primitive aperiodic locator family in the M1 band either expands in the Johnson graph like a random family, or its adjacent exchanges preserve enough locator-syndrome data to force quotient periodicity.

The concrete steps are:

A3.1 Define the exchange graph on canonical bad co-supports, connecting T and T' when $|T \triangle T'| = 2$.

A3.2 Use Propositions B.16 and B.17 to express equality or near-equality of contributed slopes across an edge as an explicit low-rank condition on the two exchanged points.

A3.3 Use a Johnson-graph spectral or container estimate to show that a super-polynomial primitive family contains many exchange paths unless it is contained in a small anticode.

A3.4 Use Proposition B.18 to pay the anticode case; use the exchange identities along paths to force a periodic stabilizer, a common divisor, or an SPI residual in the remaining case.

This route is attractive for finite certificates because the spectral/container constants can be stated as explicit integer inequalities on the deployed Johnson graph.

A4. Finite SPI atlas route. For the four adjacent deployed rows, a full asymptotic primitive incidence theorem may be more than is needed. A finite route is to promote Theorem B.19 into a chart atlas with a degree ledger:

$$\mathrm{Bad}_{\mathrm{ap}}^{\mathrm{M1}}(f,g;a_0+1) \subseteq \bigcup_{\mathfrak{c}\in\mathfrak{C}_{\mathrm{ap}}} Z(E_{\mathfrak{c}}) \cup \mathfrak{R}_{\mathrm{named}}.$$

Here $E_{\mathfrak{c}}$ is an explicit eliminant in the slope parameter and every named residual branch is either paid by a prior theorem or assigned, by a proved priority map, to one of the active residual ledgers. The atlas contribution is therefore an exact integer

$$U_{\mathrm{atlas}}(a_0+1) = \sum_{\mathfrak{c}} \deg E_{\mathfrak{c}} + |\mathfrak{R}_{\mathrm{named}}|,$$

with no asymptotic notation. Before it can enter a one-step certificate, every term of U_{atlas} must be charged to U_Q, U_{BC}, U_{sp}, or an explicit coverage correction; it is not a fourth unnamed bucket.

D.3 Quotient-fiber equidistribution

Q0. Collision-gap measurement before proof. Before committing to any proof architecture, measure the primitive prefix distribution using the non-expansive recursion of Proposition C.8. The required outputs are not individual character maxima alone, but the joint quantities of the collision-gap programme Remark C.6: high moments Γ_r, max-to-mean ratios, and divisor-lattice decompositions of the Fourier tail. The experiment is useful only if every table records whether the frequency is primitive or collapses to a coarser quotient scale in the sense of Proposition 14.10. A table of single-character Weil ratios should be treated as calibration, not as evidence for a finite adjacent certificate.

The theoretical target suggested by the measurements is an inverse theorem: if the primitive dense-bulk max fiber is larger than $n^C\binom{n}{m}|\mathbb{B}|^{-w}$, then the responsible support family has a quotient stabilizer, a common block structure, or a low-dimensional exceptional incidence already paid by the capf ledger. Without such an inverse theorem, the normalized recursion remains an audit instrument rather than a proof of (Q).

Q1. Constant-weight syndrome local central limit theorem. The identity prefix map is a constant-weight syndrome map: an m-subset of D is sent to the first $w = m-K$ elementary-symmetric coefficients of its locator. More generally, Definition 3.6 gives the quotient-remainder version. The desired theorem is a local central limit theorem for these fibers.

The proof route is:

Q1.1 Use additive-character inversion over $\mathbb{B}^w$:

$$\mu(\xi) = |\mathbb{B}|^{-w} \sum_{\alpha \in \mathbb{B}^w} \psi(-\alpha \cdot \xi) \sum_{X \in \mathcal{X}_{c,m,s}} \psi\big(\alpha \cdot \Phi^K_{c,m,s}(X)\big).$$

The trivial character gives the average fiber size.

Q1.2 Bound every nontrivial symmetric exponential sum in the last display. The needed estimate is not merely cancellation on average over ξ; it must be uniform in the character α and in the deployed window of m, w.

Q1.3 Convert the character-sum bound into an explicit multiplier Λ_Q and additive error E_Q. For asymptotic closure one may allow $\Lambda_Q = n^{O(1)}$; for deployed adjacent closure the multiplier and additive term must be evaluated in bits and compared against the printed fail margins.

This route has the clearest finite audit trail: every loss is visible as a character-sum constant.

Q2. Algebraic-monodromy route for elementary-symmetric fibers. The map from squarefree locators to their high-degree prefix coefficients is a finite morphism after the degree and support model are fixed. One can try to prove that, after quotient decompositions are removed, this morphism has full primitive monodromy and hence nearly uniform fibers.

The proof package would contain:

Q2.1 a geometric irreducibility theorem for the primitive fiber product of two prefix fibers;

Q2.2 a branch-locus bound for exceptional prefix values;

Q2.3 a finite-field point-count theorem with constants, specialized to the multiplicative and circle-line domains;

Q2.4 an explicit treatment of quotient-periodic components, so that the theorem does not re-count cells already handled by Theorem 3.36.

This is conceptually strong, but the finite deployed use requires effective constants. A qualitative Lang–Weil or monodromy theorem would only give an asymptotic (B/Q) input unless all constants are printed.

Q3. Switch-chain stability and inverse equidistribution. There is also a purely combinatorial route. Run the usual switch chain on quotient-remainder supports: replace one selected point or fiber by one unselected point or fiber while keeping the agreement size fixed. A very heavy prefix fiber is stable under unexpectedly many switches. The inverse statement should say that such stability forces one of the paid quotient structures.

Q3.1 Prove an edge-isoperimetric inequality for the support model $\mathcal{X}_{c,m,s}$.

Q3.2 Show that a prefix fiber whose size exceeds the average by more than the permitted multiplier must have many internal switches.

Q3.3 Translate internal switches into polynomial identities among elementary symmetric coefficients. Long chains of such switches force either a quotient stabilizer, a common block structure, or a low-dimensional exceptional family.

Q3.4 Pay those forced structures using the quotient image ledger, leaving only a uniformly distributed primitive remainder.

This route aligns well with the exchange-rigidity strategy for (A) and may permit one combined inverse theorem: excess mass in either the M1 local-limit problem or the quotient prefix fibers must be quotient-periodic.

Q4. Exact finite quotient-ledger certificate. For the finite adjacent theorem, the useful output is a certificate rather than a broad asymptotic theorem. For each row and each candidate agreement $a_0 + 1$, print a quotient-fiber ledger containing:

Q4.1 the exact map $\Phi^K_{c,m,s}$ or its identity-scale specialization;

Q4.2 the denominator used, either $|\operatorname{im}\Phi|$ or the ambient box $|\mathbb{B}|^w$;

Q4.3 a proved upper bound for $H^K_{c,m,s}$, with all constants and rounding;

Q4.4 the resulting quotient numerator $U_Q(a_0 + 1)$ after parameter-image coalescing;

Q4.5 a bit-margin line showing

$$U_{\mathrm{paid}}(a_0 + 1) + U_Q(a_0 + 1) + U_{\mathrm{BC}}(a_0 + 1) + U_{\mathrm{sp}}(a_0 + 1) \le B_*.$$

If the line contains a hidden n^C or an unspecified absolute constant, it should be downgraded to an asymptotic certificate and not used for the adjacent deployed pair.

D.4 Proof order

The lowest-risk development order is:

(1) Prove the boundary quotient/prefix theorem first, even with a printed polynomial multiplier, to test the identity-scale formalism above the entropy–subfield crossing.

(2) Prove the corrected interior split-pencil census with both $M^{\mathrm{disc}}_{\mathbb{B}}$ and $M^{\mathrm{soft}}_{\mathbb{B}}$ visible, and separately prove the local primitive shift-pair bound.

(3) Define the priority map, prove exhaustive bucket coverage, and prove support-to-parameter coalescing. Only then combine the asymptotic losses under $R \to \infty$, $R = o(n)$, and $\log L = o(R \log |\mathbb{B}|)$.

(4) Return to finite deployment with exact $U_Q, U_{\mathrm{BC}}, U_{\mathrm{sp}}, U_{\mathrm{paid}}$ and run the one-step compiler of Theorems 13.2 and 13.4. Finite closure is claimed only after their exact sum is below B_* at $a_0 + 1$.

Remark D.1 (what not to promote). The following statements should remain visibly weaker than the finite adjacent theorem: random-word M1 concentration without a worst-case inverse theorem; support-family quotient counts without parameter-image coalescing; quotient-fiber equidistribution with an unspecified polynomial multiplier; and SPI charts whose identically deficient branches are not paid or named. Each of these is useful evidence or an asymptotic component, but none is by itself an exact deployed upper certificate.

E Rank-one and split-pencil reduction of the safe side

This section refines the safe-side package. The first reduction has the form

$$\boxed{(\mathrm{Q}) + \mathrm{R1}}$$

where (Q) is the prefix-fiber census and R1 is a rank-one support census obtained by eliminating the slope parameter. Thus R1 is not a proof of (A) by itself; rather, after the paid tangent, quotient, extension-pole, common-GCD, quotient-pullback, fixed-dimensional, sunflower, and SPI cells are removed, R1 is the concrete support-counting statement that implies the remaining aperiodic-band input. The additional lattice reduction below sharpens this one step further: R1 is reduced to a balanced two-generator lattice census, and then to a split-pencil problem for determinantal representations of the domain polynomial. Thus the intermediate aperiodic target becomes a split-pencil census; the final closing statement later adds the base-field normalization and the primitive shift-pair ledger.

E.1 Band normalization and the master flatness target

The identity-prefix floor changes how the broad band formulation Remark 9.12 should be read. Below the entropy–subfield envelope, the identity-scale construction produces exponentially many generic support witnesses, and generic supports are aperiodic in the sense of Definition 9.1 except for a negligible quotient-periodic subfamily. Therefore the aperiodic band input should be normalized with its left edge at the entropy frontier plus a reserve.

Remark E.1 (developmental restatement of the normalized band). The operative definition is Remark 9.12. The identity-prefix floor explains its left edge: below $k + g^*(\rho, \beta)n$ the first term in the split-locator model is exponentially large, while a reserve $R \to \infty$ with $R = o(n)$ pays a multiplicative loss only when $\log L = o(R \log |\mathbb{B}|)$. The finite target remains the exact adjacent value $a_0 + 1$ with every term included in the integer ledger.

The common geometric object behind (Q) and R1 is split-locator flatness. Let $\mathcal{D}_j(D)$ denote the monic squarefree degree-j locators whose roots lie in D.

Remark E.2 (master split-locator flatness). Find an explicit polynomial P such that, for every affine subspace $\mathcal{A}$ of the monic degree-j polynomial space over $\mathbb{B}$ of codimension s, after quotient and common-divisor components have been removed,

$$|\mathcal{A} \cap \mathcal{D}_j(D)| \le P(n) \left(\binom{n}{j} |\mathbb{B}|^{-s} + 1 \right).$$

The prefix-fiber input (Q) is the case where $\mathcal{A}$ fixes the top prefix of a locator. The aperiodic/rank-one input is the case where $\mathcal{A}$ is cut out by the support interpolation equations below.

The first term in Remark E.2 crosses one at exactly the entropy–subfield envelope. A polynomial-loss version would price the relevant support intersections once the reserve dominates the polynomial cost, but it would not alone prove the frontier: local primitive control, exhaustive coverage, and support-to-parameter coalescing are still required. Even a complete asymptotic certificate cannot certify a one-step deployed pair without exact constants.

E.2 The Q-side census: moment kernels and packing

Assume in this subsection that $w < \operatorname{char} \mathbb{F}$, so Newton identities identify the first w elementary-symmetric prefixes with the first w power sums. For m-subsets define

$$\Psi_{m,w}(M) = \left(\sum_{x \in M} x^h \right)_{1 \le h \le w}.$$

Define the moment kernel

$$K_w(D) = \left\{ f : D \to \mathbb{B} : \sum_{x \in D} f(x) x^h = 0 \text{ for } 1 \le h \le w \right\}.$$

Proposition E.3 (moment-kernel dictionary)**.** *Let $M, M' \subseteq D$ be m-subsets. If $\Psi_{m,w}(M) = \Psi_{m,w}(M')$, then*

$$f = 1_M - 1_{M'}$$

belongs to $K_w(D)$, takes values in $\{-1, 0, 1\}$, and has equally many $+1$ and -1 coordinates. Conversely, if $f \in K_w(D)$ is $\{-1, 0, 1\}$-valued with $|f^{-1}(1)| = |f^{-1}(-1)| = e$, then for every common core $C \subseteq D \setminus \operatorname{supp}(f)$ of size $m - e$, the two sets

$$M = C \cup f^{-1}(1), \qquad M' = C \cup f^{-1}(-1)$$

have the same first w power sums.

Proof. The equality $\Psi_{m,w}(M) = \Psi_{m,w}(M')$ says precisely that

$$\sum_{x \in D} (1_M(x) - 1_{M'}(x)) x^h = 0 \qquad (1 \le h \le w),$$

which is the definition of membership in $K_w(D)$. The values and balance are immediate from the fact that M and M' have the same size. The converse is the same identity read backwards after adding an arbitrary common core on which the two indicators agree.

□

Proposition E.4 (balanced-ternary collision ledger)**.** *Assume that the moment map defining $K_w(D)$ has rank w over the base field. Then $K_w(D)$ has dimension $n - w$, and the second moment of the power-sum prefix fibers is the weighted balanced-ternary enumerator*

$$\sum_z |\Psi_{m,w}^{-1}(z)|^2 = \binom{n}{m} + \sum_{e \ge \delta_w} \binom{n-2e}{m-e} T_w(e),$$

where $T_w(e)$ is the number of ordered pairs (A, A') of disjoint e-subsets of D satisfying

$$\sum_{x \in A} x^h = \sum_{x \in A'} x^h \qquad (1 \le h \le w),$$

equivalently the number of $\{0, \pm 1\}$-valued words in $K_w(D)$ with exactly e positive and e negative coordinates. Consequently Γ_2 is exactly a balanced-ternary weight enumerator of the moment kernel.

Proof. The rank claim follows because the Vandermonde minor on any w distinct nonzero points of D is nonsingular. For the moment identity, count ordered pairs (M, M') in one prefix fiber. If $M = M'$, there are $\binom{n}{m}$ pairs. Otherwise write

$$A = M \setminus M', \qquad A' = M' \setminus M,$$

so $|A| = |A'| = e$, the two sets are disjoint, and the common core $C = M \cap M'$ has size $m - e$. Equality of power-sum prefixes is exactly the displayed condition on (A, A'), because the common core cancels. Conversely, any such (A, A') and any common core $C \subseteq D \setminus (A \cup A')$ of size $m - e$ give a colliding ordered pair $(C \cup A, C \cup A')$. By Lemma E.6, nontrivial collisions have $e \geq \delta_w$. The number of possible cores is $\binom{n-2e}{m-e}$, giving the formula. □

Remark E.5 (certified size of the packing improvement). At the four deployed frontier rows, the all-depth packing bound of Proposition E.7 improves the trivial bound $\binom{n}{m}$ by about $2.13 \cdot 10^5$ bits, while still remaining about $1.88 \cdot 10^6$ bits above the conjectured average fiber. Thus it is a genuine unconditional (Q) theorem, but the finite adjacent problem still requires the high-moment or extremality input of Remark E.25.

Lemma E.6 (BCH-style distance of the moment kernel)**.** *If $0 \neq f \in K_w(D)$ is supported on fewer than $w+1$ points, then this is impossible. Consequently two distinct sets in one $\Psi_{m,w}$-fiber differ on at least*

$$\delta_w := \left\lceil \frac{w+1}{2} \right\rceil$$

points on each side.

Proof. Let $S = \operatorname{supp}(f)$, $|S| = s \leq w$. The equations $\sum_{x \in S} f(x) x^h = 0$ for $1 \leq h \leq s$ form a Vandermonde system after multiplying the rows by the nonzero diagonal $x \mapsto x$ and using that the elements of D are distinct and nonzero in the smooth rows. The only solution is $f|_S = 0$, contradiction. If two fiber sets differ by e points on each side, the difference vector has support $2e$. The first assertion gives $2e \geq w + 1$, hence $e \geq \delta_w$. □

Proposition E.7 (all-depth packing bound for prefix fibers)**.** *For every value z of the prefix map and every m, w,*

$$|\Psi_{m,w}^{-1}(z)| \leq \frac{\binom{n}{m-\delta_w+1}}{\binom{m}{\delta_w-1}}, \qquad \delta_w = \left\lceil \frac{w+1}{2} \right\rceil.$$

The same bound holds for the elementary-prefix map $\Phi_{m,w}$ when $w < \operatorname{char} \mathbb{F}$.

Proof. Let $\mathcal{F}$ be one fiber. By Lemma E.6, distinct $M, M' \in \mathcal{F}$ satisfy $|M \cap M'| \leq m - \delta_w$. Therefore no $(m-\delta_w+1)$-subset of D can be contained in two different members of $\mathcal{F}$. Counting pairs (R, M) with $M \in \mathcal{F}$, $R \subseteq M$, and $|R| = m - \delta_w + 1$, gives

$$|\mathcal{F}| \binom{m}{\delta_w - 1} \leq \binom{n}{m - \delta_w + 1}.$$

This is the displayed inequality. The elementary-prefix statement follows from the triangular Newton equivalence for $w < \operatorname{char} \mathbb{F}$. □

Remark E.8 (what the packing bound does and does not do). The packing bound is a genuine all-depth (Q) theorem and should be kept. It is still far above the conjectured average fiber in the deployed frontier window, so it is not a finite adjacent certificate. The stronger locator-subtraction rigidity for equal elementary prefixes often gives the sharper side-distance $e \geq w+1$; turning that sharpening into an exact second-moment or high-moment ledger is one of the low-risk Q1 deliverables.

E.3 Slope elimination and the rank-one support census

Let $T \subseteq D$ have size m, and set $w = m - K$. Write

$$\Lambda_T(X) = \prod_{x \in T} (X - x).$$

For a word $h : T \to \mathbb{F}$, define the interpolation functionals

$$S_r(h, T) = \sum_{x \in T} \frac{h(x)x^r}{\Lambda_T'(x)}, \qquad 0 \leq r \leq w - 1.$$

Lemma E.9 (support interpolation test)**.** *A word $h : T \to \mathbb{F}$ is the restriction to T of a polynomial of degree $< K$ if and only if*

$$S_r(h, T) = 0 \qquad (0 \leq r \leq w - 1).$$

Proof. Let P_T be the unique interpolating polynomial of degree $< m$ with $P_T(x) = h(x)$ on T. For every integer d, Lagrange interpolation applied to X^d shows

$$\sum_{x \in T} \frac{x^d}{\Lambda_T'(x)} = 0 \quad (0 \leq d \leq m - 2), \qquad \sum_{x \in T} \frac{x^{m-1}}{\Lambda_T'(x)} = 1.$$

If $\deg P_T < K$, then for $0 \leq r \leq w - 1 = m - K - 1$, every monomial of $X^r P_T(X)$ has degree at most $m - 2$, so $S_r(h, T) = 0$.

Conversely, write $P_T = \sum_{i=0}^{m-1} c_i X^i$. The equation $S_0 = 0$ has leading term c_{m-1} and lower-degree terms vanish by the displayed identity, so $c_{m-1} = 0$. Having shown $c_{m-1} = \cdots = c_{m-r} = 0$, the equation $S_r = 0$ has leading remaining term c_{m-1-r}, and all lower terms vanish. Induction for $r = 0, \ldots, w - 1$ gives $c_i = 0$ for every $i \geq K$. Thus $\deg P_T < K$. □

For a received affine line (u, v), put

$$\alpha(T) = \big(S_r(u, T)\big)_{0 \leq r < w}, \qquad \beta(T) = \big(S_r(v, T)\big)_{0 \leq r < w}.$$

Call T a *rank-one support* for (u, v) if the two vectors $\alpha(T), \beta(T) \in \mathbb{F}^w$ span a subspace of dimension at most one. It is a *common support* if $\alpha(T) = \beta(T) = 0$.

Proposition E.10 (slope elimination)**.** *Fix (u, v) and a support T of size m. A finite slope z has $(u + zv)|_T$ extending to a degree-$< K$ polynomial if and only if*

$$\alpha(T) + z\beta(T) = 0.$$

If T is not common, then it determines at most one such slope. Consequently, after common-support/tangent cells are paid, the number of bad finite slopes is bounded by the number of non-common rank-one supports.

Proof. By linearity of the functionals and Lemma E.9, the restriction $(u + zv)|_T$ extends to degree $< K$ exactly when

$$S_r(u, T) + zS_r(v, T) = 0 \qquad (0 \le r < w),$$

which is the vector equation. If $\beta(T) = 0$, then either $\alpha(T) = 0$, the common case, or no slope works. If $\beta(T) \neq 0$, then a solution, if it exists, is forced by any nonzero coordinate of $\beta(T)$, and the vector equation is possible exactly when $\alpha(T)$ and $\beta(T)$ are proportional. Thus a non-common support determines at most one slope. Choosing one canonical non-common support for each unpaid bad slope gives an injection from slopes to non-common rank-one supports. □

Remark E.11 (Rank-one support census after slope elimination). After all paid cells and common supports have been removed, an intermediate primitive target is

$$\#\{T \subseteq D : \ |T| = m, \ T \text{ non-common rank-one for } (u, v)\} \le L_{\mathrm{R1}}(n) \max\left(1, \binom{n}{m} q^{-(w-1)}\right),$$

uniformly over primitive affine lines (u, v). This challenge-field support census is not the active all-strata model: a usable asymptotic theorem must also incorporate the corrected base-field floors and the local primitive shift-pair ledger, with $R = o(n)$ and $\log L_{\mathrm{R1}} = o(R \log |\mathbb{B}|)$. The finite deployed form must replace L_{R1} by exact constants and an integer numerator assigned through the proved bucket compiler.

Remark E.12 (scope of the rank-one reduction). The quotient/prefix input (Q) and the rank-one support census R1 form an intermediate support-counting decomposition. By Proposition E.10, a proved R1 bound injects unpaid primitive slopes into non-common rank-one supports. It does *not* by itself prove the active frontier: the support model must still be corrected by the base-field-normalized interior floor, the local primitive shift-pair term, exhaustive bucket coverage, and support-to-parameter coalescing. A valid finite certificate has the authoritative form

$$U_{\mathrm{paid}}(a_0 + 1) + U_Q(a_0 + 1) + U_{\mathrm{BC}}(a_0 + 1) + U_{\mathrm{sp}}(a_0 + 1) \le B_* < L(a_0),$$

as in Remark 21.5.

Lemma E.13 (pair descent for close slopes)**.** *Fix an affine line (u, v) with $v \notin C$, and fix an agreement $A > n/2$. Call a finite slope z close if $u + zv$ agrees with some codeword c_z on at least A points. If $z_1 \neq z_2$ are close, then*

$$v_{12} := \frac{c_{z_1} - c_{z_2}}{z_1 - z_2} \in C$$

agrees with v on at least $2A - n$ points. Coloring each close pair $\{z_i, z_j\}$ by v_{ij}, every color class is a clique. Inside a color class with color v', the close slopes are exactly close slopes of the derived line

$$\big(u - c_{z_0} + z_0 v', \ v - v'\big)$$

at the same agreement, for any fixed member z_0 of the clique.

Proof. Let T_i be an agreement set of size at least A for $u + z_i v$ and c_{z_i}. Since $A > n/2$, the intersection $T_1 \cap T_2$ has size at least $2A - n$. On this intersection,

$$(z_1 - z_2)v = c_{z_1} - c_{z_2},$$

which proves the first assertion. If $v_{12} = v_{23} = v'$, then

$$c_{z_1} - c_{z_3} = (z_1 - z_2)v' + (z_2 - z_3)v' = (z_1 - z_3)v',$$

so $v_{1,3} = v'$, and color classes are cliques. For a fixed clique member z_0, the relation $c_{z_i} = c_{z_0} + (z_i - z_0)v'$ gives

$$u + z_i v - c_{z_i} = (u - c_{z_0} + z_0 v') + z_i(v - v'),$$

so the agreement sets are unchanged after passing to the derived line. □

Remark E.14 (why pair descent alone does not close the band). If two distinct slopes z_1, z_2 are close on agreement supports of size at least A, then subtracting the two explanations shows that the direction word v agrees with a codeword on at least $2A - n$ positions. This is useful only when that descended agreement is itself in a controlled list-decoding range. At the KoalaBear adjacent MCA value $A = 1116048$, one has

$$2A - n = 134944, \qquad 3A - 2n = -846160,$$

so the first-order descent is far too weak to eliminate the underdetermined band. This is precisely why the rank-one support census is the right replacement for a naive descent argument.

E.4 Lattice resolution of R1: near-rational slopes, balanced core, and split pencils

The rank-one support census admits a second reduction. For one received word $U : D \to \mathbb{F}$, support counting can be expressed as the split locus of a rank-two $\mathbb{F}[X]$-lattice. Its near-rational stratum resolves exactly; outside that stratum every surviving support belongs to an explicit balanced two-generator family. Finally, the divisibility condition in that family is automatic, so the terminal aperiodic input is a split-pencil problem for determinantal representations of Λ_D. Throughout this subsection K denotes the code dimension, m is the agreement support size, $w = m - K$, and $\omega = n - m$. Let $\widehat{U}$ be the degree-$\leq n-1$ interpolant of U on D, and put $\Lambda_D(X) = \prod_{x \in D}(X - x)$.

Proposition E.15 (the support census as a lattice locus)**.** *Define*

$$M_U := \{(W, N) \in \mathbb{F}[X]^2 : \ W(x)U(x) = N(x) \text{ for all } x \in D\},$$

and give $\mathbb{F}[X]^2$ the shifted degree

$$\operatorname{wdeg}(W, N) := \max\big(\deg W, \deg N - (K - 1)\big).$$

Then:

(a) *M_U is a free rank-two $\mathbb{F}[X]$-module with basis $(1, \widehat{U}), (0, \Lambda_D)$. A shifted weak-Popov basis g_1, g_2 has predictable degrees, and if its row degrees are $d_1 \leq d_2$, then $d_1 + d_2 = n - K + 1$.*

(b) *The number $\operatorname{cen}(U; m)$ of valid agreement supports of size m is exactly*

$$\#\{(W, N) \in M_U : \ W \text{ monic}, \ \deg W = \omega, \ \operatorname{wdeg}(W, N) \leq \omega, \ W \mid \Lambda_D, \ W \mid N\}.$$

Each element gives the codeword $c = N/W$ and the support $T = D \setminus Z(W)$.

(c) *Equivalently,*

$$\operatorname{cen}(U;m)=\sum_{c\in C}\binom{\operatorname{agr}(U,c)}{m},$$

the m-th binomial moment of the agreement profile of U.

Proof. For $(W,N)\in M_U$, the difference $N-W\widehat{U}$ vanishes on every point of D, hence is a multiple of Λ_D. Thus $(W,N)=W(1,\widehat{U})-C(0,\Lambda_D)$ for some $C\in\mathbb{F}[X]$, and the two displayed generators form a basis. Since $\mathbb{F}[X]$ is Euclidean, shifted row reduction gives a weak-Popov basis; the predictable-degree property is the standard leading-position argument for weak-Popov rows, and the sum of the shifted row degrees equals the shifted degree of the determinant Λ_D, namely $n-K+1$.

If T is a valid support with codeword c, take $W=\Lambda_{D\setminus T}$ and $N=Wc$. Then $(W,N)\in M_U$, W is monic of degree ω, satisfies $\operatorname{wdeg}(W,N)\le\omega$, and $W\mid\Lambda_D,N$. Conversely, such an element gives a polynomial $c=N/W$ of degree $<K$, because $\operatorname{wdeg}(W,N)\le\omega$ in the census range, and the identity $W(U-c)=0$ on D shows that $U=c$ on $D\setminus Z(W)$, a support of size m. Since $m\ge K$, the codeword attached to a valid support is unique. Counting by supports or by codewords gives the binomial-moment formula. □

Theorem E.16 (near-rational dichotomy)**.** *Assume $2w\le n-K$, and let $d_1=d_1(U)$ be the minimal shifted degree of a nonzero element of M_U.*

(i) *If $d_1\le w$, then either U is within Hamming distance exactly d_1 of a unique codeword c, the minimal vector is the error-locator pair $(\Lambda_E,\Lambda_E c)$ with $E=\operatorname{Disagr}(U,c)$, and*

$$\operatorname{cen}(U;m)=\binom{n-d_1}{m},$$

or $\operatorname{cen}(U;m)=0$.

(ii) *If $d_1\ge w+1$, every census element lies in the two-generator family*

$$Ag_1+Bg_2,\qquad \deg A\le\omega-d_1,\quad \deg B\le\omega-d_2,$$

whose parameter space has $\mathbb{F}$-dimension $\omega-w+1$.

Proof. Let $g_1=(W_1,N_1)$, $g_2=(W_2,N_2)$ be a shifted weak-Popov basis with row degrees $d_1\le d_2$. If $d_1\le w$, then $d_2=n-K+1-d_1\ge n-K+1-w=\omega+1$. Predictable degrees imply that any census element of shifted degree ω must be a multiple Ag_1. Hence the census is empty unless W_1 is, up to scalar, a monic squarefree divisor of Λ_D, and unless N_1/W_1 is a polynomial c of degree $<K$. In the nonempty case c agrees with U off the roots of W_1, and minimality forces the disagreement set to be exactly those roots; otherwise the true error-locator pair would have smaller shifted degree. The codeword is unique because two codewords both within distance w of U would differ in at most $2w\le n-K$ positions, below the Reed–Solomon distance $n-K+1$. The valid supports are exactly the m-subsets of the $n-d_1$ agreement points, yielding the displayed binomial. If any of the required split/divisibility/codeword conditions fails, no census element exists.

If $d_1\ge w+1$, both rows can appear. Predictable degrees bound the multipliers by $\deg A\le\omega-d_1$ and $\deg B\le\omega-d_2$. The number of coefficients in A,B is $(\omega-d_1+1)+(\omega-d_2+1)=2\omega+2-(n-K+1)=\omega-w+1$. □

Corollary E.17 (per-line reduction to the balanced core)**.** *Let (u, v) be an affine line at agreement m, write $w = m - K$, and assume*

$$n - 3w \geq m \qquad \textit{equivalently} \qquad m \leq \frac{n + 3K}{4}.$$

If two finite slopes z_1, z_2 have $d_1(u + z_i v) \leq w$ and nonzero census, then (u, v) has a common codeword-pair explanation on at least $n-3w \geq m$ points and hence is not MCA-bad at agreement m. Consequently, outside at most one near-rational slope,

$$N_{\text{MCA-bad}}(u, v; m) \leq 1 + \#\{z : \ d_1(u + zv) \geq w + 1, \ \operatorname{cen}(u + zv; m) > 0\}.$$

The condition holds at all four deployed adjacent rows. Without it, the near-rational calculation alone does not imply an agreement-m explanation.

Proof. By Theorem E.16, a slope with $d_1 \leq w$ and nonzero census has $u + z_i v = c_i + \eta_i$ with $\operatorname{wt}(\eta_i) \leq w$. For two such slopes,

$$v = \frac{c_1 - c_2}{z_1 - z_2} + \eta_v, \qquad \operatorname{wt}(\eta_v) \leq 2w,$$

and then u is within distance $3w$ of a codeword. Hence the pair has a common codeword-pair explanation on at least $n - 3w$ positions, which is at least m by hypothesis. Therefore no slope is MCA-bad; otherwise at most one near-rational slope exists, and every remaining close slope belongs to the balanced stratum. □

Remark E.18 (Balanced-core reduction). For every word U with $d_1(U) \geq w + 1$, prove

$$\#\left\{(A, B) : \begin{array}{l} \deg A \leq \omega - d_1, \quad \deg B \leq \omega - d_2, \\ W_{A,B} \mid \Lambda_D, \quad W_{A,B} \mid N_{A,B} \end{array}\right\} \leq e^{o(n)} \max\left\{1, \binom{n}{m} q^{-w}\right\},$$

where $(W_{A,B}, N_{A,B}) = Ag_1 + Bg_2$. This statement implies the twisted support census for every word, hence Remark E.11. The family has dimension $\omega - w + 1$, and its split condition is the same divisor variety that appears in Theorem B.19, now with a global two-generator origin.

Lemma E.19 (automatic divisibility on split supports)**.** *If $(W, N) \in M_U$ and W is squarefree with all roots in D, then $W \mid N$.*

Proof. For every root x of W, the defining relation gives $N(x) = W(x)U(x) = 0$. Since W is squarefree, vanishing at all roots of W is equivalent to divisibility by W. □

Lemma E.20 (unimodularity of lattice bases)**.** *For any $\mathbb{F}[X]$-basis $(W_1, N_1), (W_2, N_2)$ of M_U, one has* $\gcd(W_1, W_2) = 1$ *and*

$$W_1 N_2 - W_2 N_1 = \gamma \Lambda_D$$

for some $\gamma \in \mathbb{F}^\times$.

Proof. Because $(1, \widehat{U}), (0, \Lambda_D)$ is also a basis, the change-of-basis matrix has determinant in $\mathbb{F}^\times$. Hence the determinant of $(W_1, N_1), (W_2, N_2)$ is a nonzero scalar multiple of Λ_D. Also $(1, \widehat{U}) = A_0(W_1, N_1) + B_0(W_2, N_2)$ for some $A_0, B_0 \in \mathbb{F}[X]$; taking first coordinates gives $1 = A_0 W_1 + B_0 W_2$, so W_1, W_2 are coprime. □

Proposition E.21 (determinantal representation of the adversary)**.** *Every* $\mathbb{F}[X]$*-basis* $(W_1, N_1), (W_2, N_2)$ *of an interpolation lattice* M_U *gives a determinantal representation*

$$W_1 N_2 - W_2 N_1 = \gamma \Lambda_D, \qquad \gamma \in \mathbb{F}^\times, \qquad \gcd(W_1, W_2) = 1.$$

Conversely, any quadruple satisfying the displayed identity and coprimality defines a word U *on* D *by*

$$U(x) = N_i(x)/W_i(x)$$

using any index i *for which* $W_i(x) \neq 0$*, and the two pairs form a basis of the resulting lattice* M_U*. Therefore the balanced aperiodic adversary may be studied as a degree-capped family of determinantal representations of* Λ_D*; the shifted weak-Popov profile* (d_1, d_2)*,* $d_1 + d_2 = n - K + 1$*, is a normalization imposed after this representation is chosen.*

Proof. The forward direction is Lemma E.20. For the converse, coprimality ensures that at least one of $W_1(x), W_2(x)$ is nonzero at each $x \in D$. If both are nonzero, the determinant identity, evaluated at x, gives $W_1(x)N_2(x) = W_2(x)N_1(x)$, because $\Lambda_D(x) = 0$; hence the two ratios agree and U is well defined. If $W_1(x) = 0$, then $W_2(x) \neq 0$, and the same identity gives $N_1(x) = 0$, so $(W_1, N_1) \in M_U$; the argument for the second row is identical. The two rows generate a full-rank submodule of M_U with determinant $\gamma\Lambda_D$, the same determinant as M_U up to a unit, hence the generated submodule equals M_U. Shifted row reduction can then be applied to choose the weak-Popov basis and its degree profile without changing the underlying determinantal parent. □

Remark E.22 (Split-pencil reduction)*.* Let (W_1, N_1, W_2, N_2) be a determinantal representation of Λ_D as in Proposition E.21, with balanced profile $d_1 \geq w + 1$. Prove

$$\#\left\{(A,B): \begin{array}{ll} \deg A \leq \omega - d_1, & \deg B \leq \omega - d_2, \\ AW_1 + BW_2 \mid \Lambda_D, & \deg(AW_1 + BW_2) = \omega \end{array}\right\} \leq e^{o(n)} \max\left(1, \binom{n}{\omega} q^{1-w}\right),$$

counting monic normalizations. By Lemma E.19, this is the full balanced-core census: the condition $W \mid N$ has been eliminated, and the remaining split condition is the single congruence $\Lambda_D \equiv 0 \pmod{AW_1 + BW_2}$. For a smooth multiplicative coset this is $X^n - \beta \equiv 0 \pmod{AW_1 + BW_2}$; the circle rows replace $X^n - \beta$ by the corresponding twin-coset domain polynomial. Thus the terminal aperiodic input is an explicit split-pencil problem.

Remark E.23 (active reading of the split-pencil reduction)*.* The split-pencil reduction above is a proof reduction, not the final census scale. The active split-pencil input is the base-field-normalized interior census of Residual Input 21.3; the boundary profile is delegated to quotient/prefix flatness, and the primitive shift-pair ledger is kept as a separate input.

Remark E.24 (scope of the intermediate split-pencil package)*.* The algebraic reductions Lemma E.19, Propositions E.10 and E.21, and Corollary E.17 show that the balanced primitive support problem factors through a two-generator split-pencil census. They do not establish the final scale, exhaustive coverage, or parameter coalescing. The active all-strata certificate must instead use the corrected base-field-normalized numerator of Residual Input 21.3, keep the primitive shift-pair numerator of Residual Input 21.4 separate, and fit all terms into the single budget of Remark 21.5. In particular, an unspecified polynomial or $e^{o(n)}$ split-pencil loss does not decide any printed adjacent pair.

E.5 Finite extremality and audits

The finite adjacent pair is not a soft local-limit problem. The printed fail margins are constant-sized; an unspecified factor n^C costs $21C$ bits at $n = 2^{21}$ and therefore exceeds three of the four deployed adjacent margins. The finite version should be reduced to two exact extremality statements.

Remark E.25 (mode-at-null / exchange-compression for (Q)). Let $N_w(z) = |\Phi_{m,w}^{-1}(z)|$ at the identity scale, with quotient rungs separated. Prove either

$$N_w(z) \le N_w(0) \qquad \text{for all } z,$$

or an exchange-compression theorem implying the same deployed max-fiber bound after quotient-pulled-back fibers are charged to their rungs. Then prove the exact finite mean-plus-one inequality

$$N_w(0) \le \kappa \left(1 + \binom{n}{m} |\mathbb{B}|^{-w}\right)$$

with κ and the additive unit included in the printed row budget.

Remark E.26 (Finite-rung audit). For each of the four deployed adjacent values $a_0 + 1$, run the exact integer scanner over every divisor/rung and every slack profile used by the quotient ledger. A periodic support count is far above budget at support level, so the quotient bucket must be paid at image level by descent and lcm coalescing. The audit must print per-rung bit losses; if any rung is tight or inverted, the one-step finite conjecture is threatened from the periodic side before (Q) or R1 is even invoked.

E.6 Updated proof ladder and risk register

The work queue is now ordered by leverage.

(1) **Immediate.** Use the normalized left edge; run the rung-margin audit; prove and use the slope-elimination reduction; keep the proved anchors Corollaries 14.9, 14.16 and 14.17 and Propositions 14.10 and 14.13 as proved base cases for (Q).

(2) **Low-to-medium risk.** Build the exact prefix-collision ledger using Propositions E.3 and E.7; stratify collisions by common core, difference size, and quotient scale. This gives a typical-fiber theorem and calibrates the high-moment hierarchy, but not max-fiber control by itself.

(3) **Main Q effort.** Prove heavy-fiber symmetry descent: few heavy fibers plus the twist action $M \mapsto \zeta M$ force quotient-pulled-back prefix values, reducing worst-case flatness to a null-fiber or quotient-rung census.

(4) **Main aperiodic effort.** Prove the split-pencil census of the split-pencil reduction Remark E.22. This would imply the balanced-core census of the balanced-core reduction Remark E.18, then R1, then the normalized aperiodic input. Johnson-exchange rigidity, bounded-deficiency SPI induction, and determinantal monodromy now all target this single two-generator pencil.

(5) **Finite prize.** Replace all polynomial losses by exact constants. The Mersenne-31 list row, with only a few bits of room, should be treated as requiring exact extremality rather than a soft analytic estimate.

item	risk	failure meaning
band normalization, slope elimination, rung audit	low	editorial or arithmetic correction
prefix collision ledger and packing	low–medium	technical refinement, not frontier failure
heavy-fiber symmetry descent for (Q)	medium	defect recursion may lose constants; fatal only for finite one-step
split-pencil / balanced-core census	high	a true primitive aperiodic split-locator mechanism may exist
mode-at-null / exchange-compression	open	finite adjacent theorem remains conditional

The strategic conclusion is that the Proximity Prize has a clean form. Asymptotically, prove split-locator flatness at the entropy–subfield boundary. Finitely, prove exact prefix-fiber extremality, base-field-normalized split-pencil bounds, and primitive shift-pair control at the four adjacent rows. A counterexample cannot fail silently: it must either produce a super-polynomial primitive prefix fiber, refuting (Q), or a super-polynomial primitive split-pencil family, refuting Remark E.22 and hence R1.

F Unified safe-side frontier package

After Proposition 14.2, the unsafe side is essentially finished: the identity-prefix floor reaches the entropy transition $\mathrm{H}_2(\rho+g)=\beta g$ directly, quotient scales and planting are subordinate mechanisms, and the deployed rows sit within 3×10^{-6} of the ceiling $g^*(\rho,\beta)$. The prize is now a safe-side theorem: prove that nothing beats the identity floor. This section states that goal as a package of named conjectures, registers which of its base cases are already theorems, proves two new unconditional statements — a rigidity/dictionary result for the prefix map, with the first nontrivial all-depth fiber bound, and the pair-descent interface for the aperiodic band — and records why the band resists every first-order mechanism.

F.1 The conjecture package

Remark F.1 (Frontier prize formulation). With $\beta=\log_2|\mathbb{B}|$, prove

$$\delta_C^*(\varepsilon^*) \;=\; 1-\rho-g^*(\rho,\beta)+o(1), \qquad g^*(\rho,\beta)=\sup\{g:\mathrm{H}_2(\rho+g)\ge\beta g\},$$

for the deployed smooth multiplicative and circle line-round rows. The finite deployed form is the prediction of the four adjacent pairs of Conjecture 14.6: the first safe agreement is a_0+1 for the unsafe agreements $a_0=1116047$, 1116046, 1116023, 1116022.

Remark F.2 (Boundary quotient/prefix formulation). For $m=(\rho+g)n$, $K=\rho n$, $w=m-K$, let

$$\Phi_{m,w}:\binom{D}{m}\to\mathbb{B}^w, \qquad M\mapsto\big(e_1(M),\dots,e_w(M)\big)$$

be the prefix map (equivalently the power-sum prefix, by Newton). The identity floor is $\max_z|\Phi_{m,w}^{-1}(z)|\ge\lceil\binom{n}{m}|\mathbb{B}|^{-w}\rceil$. Prove the matching discrete upper bound

$$\max_z\big|\Phi_{m,w}^{-1}(z)\big|\le L_Q(n)\left(1+\binom{n}{m}|\mathbb{B}|^{-w}\right)$$

throughout the normalized band, with quotient-periodic structure absorbed at its own scale as in Proposition 14.10. The asymptotic form is consumed only when $R \to \infty$, $R = o(n)$, and $\log L_Q = o(R \log |\mathbb{B}|)$; the finite form uses explicit constants in the summed deployed ledger at $a_0 + 1$.

Remark F.3 (Fixed-moment collision hierarchy). Let $N_m = \binom{n}{m}$, $Q_w = |\mathbb{B}|^w$, $S_{m,w} = \min\{N_m, Q_w\}$, and let $\mu_{m,w}$ be the distribution of $\Phi_{m,w}$ on a uniform m-subset. For $r \geq 2$ define

$$\Gamma_r(m, w) := S_{m,w}^{r-1} \sum_z \mu_{m,w}(z)^r.$$

Then $\Gamma_r \geq 1$, equality holds for the effective uniform baseline, and $\Gamma_r^{1/r}$ increases to $R_{m,w} := S_{m,w} \max_z \mu_{m,w}(z)$. While $N_m \geq Q_w$, Parseval recovers the usual codomain-normalized formula; once $N_m < Q_w$, the normalization switches to the injective baseline. Prove $\Gamma_r(m, w) \leq n^{C_r}$ for each fixed r in the band and quantify the growth of C_r. The exact conversion is

$$R_{m,w} \leq \inf_{r \geq 2} \Gamma_r(m, w)^{1/r} S_{m,w}^{1/r}.$$

Thus polynomial discrete max-fiber control requires moment order about $\log_2 S_{m,w} / \log_2 n$ at a finite row; second-moment control alone cannot close the adjacent problem.

Remark F.4 (historical first-form aperiodic input). An earlier organization grouped every unpaid aperiodic configuration into one term and sought an inequality

$$N_{\text{aper}}(a) \leq L(n) N_{\text{model}}(a).$$

This is retained only to explain the subsequent slope-elimination reductions. It is not an active closure statement: the model must include the base-field floors of Proposition 20.1, the primitive same-prefix term must be controlled locally, and an asymptotic loss is usable only when $R = o(n)$ and $\log L = o(R \log |\mathbb{B}|)$. The authoritative status statement is the conditional certificate template of Section 21.

F.2 Proved anchors of the package

Proposition F.5 (base cases are theorems)**.** *At the deployed rows, for every prefix depth $w \leq w_0$ of Corollaries 14.9 and 14.16 ($w_0 = 21$–22 KoalaBear, 10–11 Mersenne-31) and every finite integer $r \geq 2$,*

$$1 \leq \Gamma_r(m, w) \leq \left(1 + 2^{-\text{margin}}\right)^{r-1},$$

with the printed margins. At the max-fiber endpoint,

$$1 \leq R_{m,w} := S_{m,w} \max_z \mu(z) \leq 1 + 2^{-\text{margin}}.$$

At these head depths $N_m \geq Q_w$, so $S_{m,w} = Q_w$ and the effective normalization agrees with the original one. The exact witness-list, composite-scale, and quotient-cell claims are those of Propositions 14.10 and 14.13 and Corollary 14.17.

Proof. The lower bound is the power-mean inequality on a support of size at most $S_{m,w}$. For the upper bound put $R_{m,w} := S_{m,w} \max_z \mu(z) \leq 1 + 2^{-\text{margin}}$ by the two-sided pricing. Then

$$\Gamma_r = S_{m,w}^{r-1} \sum_z \mu(z)^r \leq R_{m,w}^{r-1} \sum_z \mu(z) = R_{m,w}^{r-1}.$$

The separate max-fiber statement is the displayed bound for $R_{m,w}$; the remaining claims are the quoted results. □

F.3 The moment kernel: rigidity, dictionary, and an all-depth bound

The following unconditional statements convert the L^2 layer of Remark F.2 into a coding-theoretic object and make the inverse strategy precise. Assume $0 \notin D$ (true for the multiplicative rows; for circle x-domains, $0 \in D$ only if $\pm i$ lies in the twin coset, in which case the statements hold verbatim on $D \setminus \{0\}$).

Proposition F.6 (moment-kernel dictionary)**.** *Let* $K_w := \{f : D \to \mathbb{B} : \sum_{x \in D} f(x)x^j = 0,\ 1 \le j \le w\}$. *Then:*

(a) *K_w is a linear code of dimension $n - w$ and minimum Hamming weight at least $w + 1$;*

(b) *(fiber rigidity) distinct members $M \ne M'$ of one $\Phi_{m,w}$-fiber satisfy $|M \triangle M'| \ge w + 1$; equivalently $|M \cap M'| \le m - \delta$ with $\delta := \lceil (w + 1)/2 \rceil$;*

(c) *(collision stratification) with $T_w(\ell)$ the number of ordered pairs of disjoint ℓ-subsets A, A' with $p_j(A) = p_j(A')$ for $j \le w$ — equivalently the number of $\{0, \pm 1\}$-valued words of K_w with ℓ entries of each sign —*

$$\sum_z \left|\Phi^{-1}_{m,w}(z)\right|^2 \;=\; \binom{n}{m} + \sum_{\ell \ge \delta} \binom{n - 2\ell}{m - \ell} T_w(\ell);$$

(d) *consequently Γ_2 is a weighted balanced-ternary weight enumerator of K_w, and an exponentially heavy fiber forces exponentially many balanced ternary codewords of K_w concentrated on the fiber's members.*

Proof. (a) The defining map $V : f \mapsto (\sum f(x)x^j)_{j \le w}$ has, on any $s \le w$ coordinates $x_1, \dots, x_s \in D$, the top minor $\det(x_i^j)_{1 \le i,j \le s} = (\prod_i x_i) \cdot \prod_{i<i'}(x_{i'} - x_i) \ne 0$, since the x_i are distinct and nonzero; so V is injective on words of support size $\le w$ and surjective, giving both claims. (b) $f = \mathbf{1}_M - \mathbf{1}_{M'} \in K_w$ is nonzero of weight $|M \triangle M'|$; apply (a); the weight is even and $\ge w + 1$, hence $\ge 2\delta$. (c) Stratify pairs by $A = M \setminus M'$, $A' = M' \setminus M$ (sizes equal, say ℓ, with $\ell \ge \delta$ or $\ell = 0$ by (b)); prefix matching is $p_j(A) = p_j(A')$ for $j \le w$ since the common part cancels, i.e. $\mathbf{1}_A - \mathbf{1}_{A'} \in K_w$; the common part is an arbitrary $(m - \ell)$-subset of the remaining $n - 2\ell$ points. (d) Restate (c). For the last claim, fix one member M_0 of a fiber of size S. The map $M \mapsto \mathbf{1}_M - \mathbf{1}_{M_0}$ is injective on the other $S - 1$ members and produces distinct nonzero balanced ternary words of K_w. □

Corollary F.7 (unconditional all-depth fiber bound)**.** *For every m, w and every z,*

$$\left|\Phi^{-1}_{m,w}(z)\right| \;\le\; \binom{n}{m - \delta + 1} \Big/ \binom{m}{\delta - 1}, \qquad \delta = \lceil (w + 1)/2 \rceil .$$

At the four deployed frontier rows ($w = 67466, 67470, 67444, 67446$) this improves the trivial bound $\binom{n}{m}$ by 213501, 213510, 213448, 213453 bits respectively — the first nontrivial fiber bound at full frontier depth — while remaining about 1.877×10^6 bits above the conjectured average: an exact measure of the distance from rigidity alone to Remark F.2.

Proof. By rigidity, two distinct fiber members share at most $m - \delta$ points, so no $(m - \delta + 1)$-subset lies in two members; each member contains $\binom{m}{m-\delta+1} = \binom{m}{\delta-1}$ of the $\binom{n}{m-\delta+1}$ available. The bit counts are certified by the entropy sandwich as before. □

F.4 The (A) interface: pair descent and band obstructions

Lemma F.8 (pair descent)**.** *Fix a line (u, v) with $v \notin C$, an agreement $a > n/2$, and call a finite slope z* close *if $u + zv$ agrees with some codeword c_z on at least a points. If $z_1 \neq z_2$ are close, then $v_{12} := (c_{z_1} - c_{z_2})/(z_1 - z_2)$ is a codeword (possibly zero) agreeing with v on at least $2a - n$ points. Color each pair of close slopes by its codeword v_{ij}; then each color class is a clique, and the close slopes of a clique in color v' are, slope for slope, close slopes of the derived line $(u - c_{z_0} + z_0 v', v - v')$ at the same agreement, for any fixed clique member z_0. Consequently, with $\lambda := \#\{v' \in C : \ v' \text{ agrees with } v \text{ on } \geq 2a - n \text{ points}\}$,*

$$L(u, v; a) \ \leq \ 1 + \lambda \cdot \max_{v'} L\big(u - c_{z_0} + z_0 v', \, v - v'; a\big).$$

Proof. Subtracting the two agreement relations on $T_1 \cap T_2$ ($|T_1 \cap T_2| \geq 2a - n$) gives $(z_1 - z_2)v = c_{z_1} - c_{z_2}$ there, whence the first claim. For transitivity, if $v_{12} = v_{23} = v'$ then $c_{z_1} - c_{z_3} = (z_1 - z_2)v' + (z_2 - z_3)v' = (z_1 - z_3)v'$, so $v_{1,3} = v'$: color classes are cliques. Within the clique of v', $c_{z_i} = c_{z_0} + (z_i - z_0)v'$, so $u + z_i v - c_{z_i} = (u - c_{z_0} + z_0 v') + z_i(v - v')$, and the agreement set is unchanged. Every close slope other than a fixed one shares an edge, hence a color, with it; counting cliques by color gives the displayed inequality. □

Remark F.9 (why the band is the band)*.* The inventory of first-order mechanisms against Remark F.4, at rate $\rho = \frac{1}{2}$ and the deployed sizes: (i) r-wise differences that eliminate both u and v need $r \geq 3$ and bite only when $3(n - a) < n - k$, i.e. $a > \frac{5}{6}n$ — the deep regime already covered. (ii) Lemma F.8 is unconditional, but its hypothesis-free use needs the direction list to be empty or small at agreement $2a - n$, and $2a - n$ exceeds the direction word's own entropy frontier only when $a \geq n(1 + \rho + g^*)/2 \approx 0.76609\, n$; below that, generic directions carry astronomically many codewords at agreement $2a - n$, and adversarial directions sit near the code outright. (iii) Subtracting the near codeword descends to a direction of weight at most $2(n - a)$, after which every agreement account passes through $|T_i| - 2(n - a) \geq 3a - 2n$, which is *negative* throughout the band — at the KoalaBear adjacent pair, $3a - 2n = -846160$ while $2a - n = 134944$. Thus every pencil-type first-order mechanism degenerates identically at the quantity $3a - 2n$, and the band $a/n \in (\rho + g^*, \frac{5}{6})$ admits no unconditional exclusion by descent alone. This is the precise sense in which (A) is a named input rather than a missing lemma: the viable routes are stability classification of near-extremal locator families, growing-dimensional incidence bounds for the residual Conjecture-F strata, and bounded-deficiency SPI induction, each of which must produce structure that the paid compilers then price.

Remark F.10 (experimental hierarchy values are not certificate inputs)*.* The Fourier coefficients $\widehat{\mu}_{m,w}(\tau)$ are computable on small instances by the non-expansive recursion of Remark 14.12, and such calculations were useful for choosing conjectures. No experimental census value is used in a theorem, in the active unsafe certificates, or in the conditional closure template. The falsifiable asymptotic question remains whether $\log_2 \Gamma_r(m, w) = O_r(\log n)$ for fixed r across the normalized band, while the finite adjacent rows require either direct max-fiber control or the much deeper row-dependent hierarchy of Remark F.27.

F.5 Slope elimination: reducing the aperiodic side to a support census

The following results turn the first-form aperiodic input into an explicit support census. They eliminate the slope quantifier unconditionally, remove the tangent stratum from the inequality

outright, and subsume the counting content of every bounded- or growing-deficiency SPI chart. The census is then refined further by the split-pencil and base-field-normalized formulations.

Lemma F.11 (interpolation functionals)**.** *Let $T \subseteq D$, $|T| = m$, $w = m - K$, and for a word $f : D \to \mathbb{F}$ put*

$$S_r(f,T) \;:=\; \sum_{x\in T} \frac{f(x)\,x^r}{\Lambda_T'(x)}, \qquad 0 \le r \le w-1.$$

Then $f|_T$ extends to a polynomial of degree $< K$ if and only if $S_r(f,T) = 0$ for all $r < w$. The functionals are $\mathbb{F}$-linear in f, and $\mathbb{B}$-valued whenever f is $\mathbb{B}$-valued (since $T \subseteq D \subseteq \mathbb{B}$).

Proof. The interpolant is $I_f(X) = \sum_{x\in T} f(x)\Lambda_T(X)/((X-x)\Lambda_T'(x))$, of degree $\le m-1$, and extension of degree $< K$ means its top w coefficients vanish. From $\prod_{y\in T\setminus x}(1+uy) = \big(\sum_i e_i(T)u^i\big)\big(\sum_r (-x)^r u^r\big)$ one gets $e_j(T\setminus x) = \sum_{r\le j}(-1)^r e_{j-r}(T)\,x^r$, so the coefficient of X^{m-1-j} in I_f is $(-1)^j\sum_{r\le j}(-1)^r e_{j-r}(T)\,S_r(f,T)$: a unitriangular combination of $S_0,\dots,S_j$ with T-symmetric coefficients. Vanishing for $j = 0,\dots,w-1$ is therefore equivalent to $S_0 = \cdots = S_{w-1} = 0$. □

Theorem F.12 (slope elimination)**.** *Fix a line (u,v) over the challenge field, an agreement $a = m > n/2$, and $w = m - K$. Call T* rank-one *if the vectors $\alpha(T) := (S_r(u,T))_{r<w}$ and $\beta(T) := (S_r(v,T))_{r<w}$ are proportional and not both zero, with $\beta(T) \neq 0$; each rank-one support determines the single slope $z(T)$ with $\alpha(T) = -z(T)\beta(T)$. Then:*

(a) *a slope z is close at support T (i.e. $u+zv$ interpolable on T) if and only if $\alpha(T) = -z\,\beta(T)$; supports with $\beta = 0 \neq \alpha$ carry no slope, and supports with $\alpha = \beta = 0$* (common *supports*) *carry every slope;*

(b) *if some m-support is common, the pair (u,v) has correlated agreement at density m/n on it, and* no *slope is MCA-bad at radius $1 - m/n$;*

(c) *in every case,*

$$N_{\text{MCA-bad}}(u,v;\,m) \;\le\; \#\mathrm{R1}(u,v;m) \;:=\; \#\{\textit{rank-one supports}\};$$

(d) (*pencil dictionary*) *for any exact-agreement syndrome pencil $M(Z) = M_0 + ZM_1$ of* any *deficiency, parametrizing split kernel candidates by their root sets A: per candidate, $M_0\Lambda_A + zM_1\Lambda_A = 0$ has at most one solution z unless $M_0\Lambda_A = M_1\Lambda_A = 0$ (an identically-split residual branch as in Theorem B.19); so the chart's slope count is bounded by the same rank-one census with columns $(M_0\Lambda_A, M_1\Lambda_A)$, at every deficiency.*

Proof. (a) By Lemma F.11 applied to $u+zv$ and $\mathbb{F}$-linearity, closeness at T is $\alpha(T)+z\beta(T) = 0$, and the trichotomy is immediate. (b) A common support gives polynomials c_u, c_v of degree $< K$ with $u|_T = c_u|_T$, $v|_T = c_v|_T$; then $(u+zv)|_T = (c_u + zc_v)|_T$ for every z: correlated agreement of the pair at density m/n, which by definition excludes MCA-badness at that radius. (c) If a common support exists, the left side is zero by (b). Otherwise every MCA-bad slope is close, every close slope has an m-support inside its agreement set, no support is common, supports with $\beta = 0$ carry no slope by (a), and distinct slopes have distinct rank-one supports since $z(T)$ is a function of T. (d) The pencil is affine-linear in Z, so per split candidate the condition is $\mu_0 + z\mu_1 = 0$ with $\mu_i = M_i\Lambda_A$: one slope, or none, or all (the residual branch). □

Remark F.13 (Rank-one support census after slope elimination). Prove, for every line at band agreements,

$$\#\mathrm{R1}(u,v;m) \;\le\; e^{o(n)}\cdot\max\Big(1,\;\binom{n}{m}\,q^{-(w-1)}\Big),$$

with the finite form fitting under the deployed budget at a_0+1. By Theorem F.12 and the paid compilers, this rank-one census implies the first-form aperiodic input. It is an intermediate algebraic reduction; the final active split-pencil input is Residual Input 21.3. The former challenge-field-only calibration is withdrawn because it omits the discrete base-field floors of Proposition 20.1. Structure available to attacks remains: the columns are $\mathbb{B}$-valued on $\mathbb{B}$-valued coordinate words (Lemma F.11), so (R1) is a base-field point count on an explicit determinantal incidence variety, and the twisted vectors $(u(x)/\Lambda_T'(x))\mathbf{1}_T$, $(v(x)/\Lambda_T'(x))\mathbf{1}_T$ live in the order-w moment kernel of Proposition F.6. Lemma F.8 bounds the clique structure of the image of $z(\cdot)$; (R1) bounds its domain. Known unconditional strata: $\mathbb{B}$-rational lines (confinement: at most $p+1$ slopes), and deficiency-one pencils, where Theorem B.19 gives a polynomial eliminant with no census needed.

F.6 The lattice resolution of the census: near-rational dichotomy and the balanced core

The rank-one census, and with it input (A), now admits a second unconditional reduction: the twisted census of *every* word is the split locus of an explicit rank-two $\mathbb{F}[X]$-lattice, its near-rational stratum resolves exactly, and per line at most one slope survives outside the balanced core. Throughout, $\omega := n-m$, $\widehat{U}$ denotes the degree-$\le n-1$ interpolant of the word U on D, and we assume the band condition $2w \le n-K$.

Proposition F.14 (the census as a shifted lattice locus)**.** *Let*

$$M_U := \{(W,N)\in\mathbb{F}[X]^2 :\; W(x)U(x)=N(x) \text{ for all } x\in D\},$$

and put

$$\mathrm{wdeg}(W,N) := \max\big(\deg W, \deg N-(K-1)\big).$$

Then:

(a) *M_U is a free rank-two $\mathbb{F}[X]$-module with basis $(1,\widehat{U})$ and $(0,\Lambda_D)$. A basis in shifted weak Popov form has predictable degrees, and its row degrees $d_1\le d_2$ satisfy $d_1+d_2=n-K+1$.*

(b) *If $\omega=n-m$, then the agreement-support census is exactly*

$$\begin{aligned}\mathrm{cen}(U;m) &:= \#\{T\subseteq D :\; |T|=m,\; U|_T\in \mathrm{RS}[\mathbb{F},T,K]\}\\ &= \#\left\{(W,N)\in M_U :\; \begin{array}{ll} & W \text{ is monic},\quad \deg W=\omega,\\ & \mathrm{wdeg}(W,N)\le\omega,\\ & W\mid\Lambda_D,\quad W\mid N\end{array}\right\}.\end{aligned}$$

(c) *Equivalently,*

$$\mathrm{cen}(U;m)=\sum_{c\in\mathrm{RS}[\mathbb{F},D,K]}\binom{\mathrm{agr}(U,c)}{m}.$$

The shifted-degree cap in (b) *is part of the census: without it one counts high-degree congruent representatives rather than degree-$< K$ codeword explanations.*

Proof. For (a), any $(W, N) \in M_U$ satisfies $N \equiv W\widehat{U} \pmod{\Lambda_D}$, hence $N = W\widehat{U} - C\Lambda_D$ for some $C \in \mathbb{F}[X]$. Thus $(W, N) = W(1, \widehat{U}) - C(0, \Lambda_D)$, and the displayed pair is a basis. Its determinant is Λ_D, of shifted degree $n - (K-1)$, so a shifted weak Popov basis has predictable degrees and row-degree sum $n - K + 1$.

If T is a valid support and $c \in \mathrm{RS}[\mathbb{F}, D, K]$ is the unique codeword agreeing with U on T (uniqueness follows from $m \geq K$), put $W = \Lambda_{D\setminus T}$ and $N = Wc$. Then $(W, N) \in M_U$, W is monic of degree ω, $W \mid \Lambda_D$, $W \mid N$, and $\mathrm{wdeg}(W, N) \leq \omega$ because $\deg c < K$.

Conversely, suppose (W, N) satisfies the conditions in (b). Since $W \mid N$, write $N = Wc$. The shifted-degree cap gives $\deg N \leq \omega + K - 1$, hence $\deg c < K$. The identity defining M_U gives $W(U - c) = 0$ on D, so $U = c$ on $T = D \setminus Z(W)$, and $|T| = m$. This is a valid agreement support. The two constructions are inverse, and the binomial-moment identity follows by grouping the valid supports by the codeword that explains them. □

Theorem F.15 (near-rational dichotomy)**.** *Let $d_1 = d_1(U)$ be the minimal shifted degree of $M_U \setminus \{0\}$.*

(i) *If $d_1 \leq w$, then either U lies within Hamming distance exactly d_1 of a* unique *codeword c, the minimal vector is the error-locator pair $(\Lambda_E, \Lambda_E \hat{c})$ with $E = \mathrm{Disagr}(U, c)$, and*

$$\mathrm{cen}(U; m) \;=\; \binom{n - d_1}{m} \qquad \textit{exactly;}$$

or $\mathrm{cen}(U; m) = 0$.

(ii) *If $d_1 \geq w+1$, every census element lies in the two-parameter family $\{Ag_1 + Bg_2 : \deg A \leq \omega - d_1,\ \deg B \leq \omega - d_2\}$, of $\mathbb{F}$-dimension exactly $\omega - w + 1$.*

Proof. If $d_1 \leq w$ then $d_2 = n - K + 1 - d_1 \geq n - K + 1 - w = \omega + 1$, so by predictable degrees every census element (shifted degree ω) is $A\, g_1$ with $g_1 = (W_1, N_1)$ of shifted degree d_1. Then $W_1 \mid W \mid \Lambda_D$, so cen $= 0$ unless W_1 is (up to scalar) monic squarefree and D-split; and $W \mid N \iff W_1 \mid N_1$ by cancellation, so cen $= 0$ unless $c := N_1/W_1$ is a polynomial. If $\deg N_1 - (K-1) > \deg W_1$, predictable degrees give $\mathrm{wdeg}(Ag_1) = \deg A + d_1$ with $\deg W = \deg A + \deg W_1 < \omega$ at the cap, so no element reaches $\deg W = \omega$: cen $= 0$. Otherwise $\deg W_1 = d_1$, $\deg c \leq K - 1$: a codeword, agreeing with U off the d_1 roots of W_1. Minimality forces $\mathrm{wt}(U - c) = d_1$ with disagreement set exactly $\mathrm{roots}(W_1)$: a smaller disagreement set would make the true error-locator pair a nonzero lattice element of smaller shifted degree. Uniqueness of c: two codewords within distance w of U differ in at most $2w \leq n - K$ positions, below the code distance $n - K + 1$. Finally, valid supports are exactly the m-subsets of the agreement set: containment follows from $W_1 \mid W$, and every m-subset $T \subseteq D \setminus E$ arises from $A = \Lambda_{(D\setminus E)\setminus T}$, which satisfies all constraints. Part (ii) is predictable degrees with both caps active, of total dimension $2\omega + 2 - (d_1 + d_2) = \omega - w + 1$. □

Corollary F.16 (per-line reduction to the balanced core; intermediate form)**.** *Let (u, v) be a line at agreement m, write $w = m - K$, and assume $n - 3w \geq m$. If two finite slopes have $d_1(u + zv) \leq w$ with nonzero census, then the pair has a common codeword-pair explanation on at least $n - 3w \geq m$ points, so no slope is MCA-bad. Hence*

$$N_{\text{MCA-bad}}(u, v; m) \leq 1 + \#\{z :\ d_1(u + zv) \geq w + 1,\ \mathrm{cen}(u + zv; m) > 0\}.$$

This is an intermediate reduction toward the active certificate template; it does not supply bucket coverage or parameter coalescing.

Proof. By Theorem F.15(i), two near-rational slopes give $u + z_i v = c_i + \eta_i$ with $\operatorname{wt} \eta_i \le w$. Subtracting gives a codeword approximation to v with error at most $2w$, and substituting back gives one for u with error at most $3w$. Their common agreement has size at least $n - 3w \ge m$. If there are not two such slopes, at most one lies outside the balanced stratum; slopes with zero census are not close at agreement m. □

Remark F.17 (balanced core, primitive challenge-field form). Prove, for every word U with $d_1(U) \ge w + 1$ at band agreements,

$$\#\left\{(A,B): \begin{array}{ll} \deg A \le \omega - d_1, & \deg B \le \omega - d_2, \\ W_{A,B} \mid \Lambda_D, & W_{A,B} \mid N_{A,B} \end{array}\right\} \le e^{o(n)} \max\left\{1, \binom{n}{m} q^{-w}\right\},$$

where $(W_{A,B}, N_{A,B}) = Ag_1 + Bg_2$. This intermediate statement would imply the twisted census bound for every word and hence the historical first-form input (A). It does not by itself close the frontier: the active template also requires the corrected base-field normalization, local primitive shift-pair control, exhaustive bucket coverage, and parameter coalescing. The family is explicit, of dimension $\omega - w + 1$ (913642 at the KoalaBear adjacent pair, 913686 at Mersenne-31); the split condition $W \mid \Lambda_D$ is the same divisor variety as in Theorem B.19, now with a canonical global origin: the basis (g_1, g_2) arises from one shifted-Popov reduction of $\big((1, \widehat{U}), (0, \Lambda_D)\big)$, and along a pencil $u + zv$ it varies algebraically in z, giving every SPI chart a common parent. The common-GCD and sunflower paid strata appear inside the coordinates as $\gcd(A, B)$ and common-factor loci. Development-time small-instance checks are not part of the proof or of the release certificate.

Remark F.18 (status of the intermediate balanced-core form). Remark F.17 is the challenge-field primitive model suggested by the shifted-lattice dichotomy. It is not the all-strata final census. The boundary profile belongs to the quotient/prefix problem, and the interior profiles must include the base-field floor terms isolated in Proposition 20.1. The final active version is therefore Residual Input 21.3, equivalently the base-field-normalized split-pencil census of Remark 20.3.

F.7 Split-pencil form of the balanced core

Three further unconditional statements put the balanced core into an explicit intermediate shape. Their proofs, not development-time small-instance checks, are the justification used here.

Lemma F.19 (divisibility is automatic). *If $(W, N) \in M_U$ and W is squarefree with all roots in D, then $W \mid N$.*

Proof. Every root x of W lies in D, where $N(x) = W(x)U(x) = 0$; as W is squarefree, $W \mid N$. □

Lemma F.20 (unimodularity). *For any basis $(W_1, N_1), (W_2, N_2)$ of M_U: $\gcd(W_1, W_2) = 1$ and $W_1 N_2 - W_2 N_1 = \gamma \Lambda_D$ with $\gamma \in \mathbb{F}^\times$.*

Proof. Expanding $(1, \widehat{U}) = A_0 g_1 + B_0 g_2$ in the first coordinate gives $1 = A_0 W_1 + B_0 W_2$. The cross determinant vanishes at every $x \in D$ (both products equal $W_1(x) W_2(x) U(x)$), so Λ_D divides it; its degree is at most n and it is the basis determinant, a unit multiple of Λ_D. □

Proposition F.21 (the adversary is a determinantal representation)**.** *The shifted-reduced bases of the lattices M_U are precisely the quadruples (W_1, N_1, W_2, N_2) satisfying*

$$W_1N_2 - W_2N_1 = \gamma\,\Lambda_D, \qquad \gamma \in \mathbb{F}^\times, \quad \gcd(W_1, W_2) = 1,$$

with the shifted-degree profile (d_1, d_2), $d_1 + d_2 = n - K + 1$: the word is recovered pointwise by $U(x) = N_i(x)/W_i(x)$ for either i with $W_i(x) \neq 0$ (some i works at every x by coprimality, and the two ratios agree by the identity). Consequently the entire aperiodic input is a statement about degree-capped determinantal representations of Λ_D: for the KoalaBear rows, of $X^n - \beta$.

Proof. The forward direction is Lemma F.20. Conversely, given such a quadruple, define U pointwise as stated; at a root x of W_1, the identity gives $W_2(x)N_1(x) = 0$ with $W_2(x) \neq 0$, so $N_1(x) = 0$ and $(W_1, N_1) \in M_U$, likewise (W_2, N_2). The submodule they generate has full rank and determinant $\gamma\Lambda_D$, equal to that of M_U, hence equals M_U, and the degree profile makes the basis reduced. □

Remark F.22 (Primitive split-pencil formulation). Let (W_1, N_1, W_2, N_2) be a determinantal representation of $\Lambda_D = X^n - \beta$ as in Proposition F.21 with balanced profile $d_1 \geq w + 1$. Prove

$$\#\big\{(A, B) :\ \deg A \leq \omega - d_1,\ \deg B \leq \omega - d_2,\ AW_1 + BW_2 \,\big|\, X^n - \beta\big\} \ \leq\ e^{o(n)} \max\Big(1,\ \binom{n}{\omega} q^{1-w}\Big),$$

counting monic normalizations of exact degree ω. By Lemma F.19 this is the full census condition — the divisibility $W \mid N$ has been eliminated — and the split condition is the single congruence $X^n \equiv \beta \pmod{AW_1 + BW_2}$. This intermediate problem implies Remark F.17, hence Remark F.13, hence the first-form aperiodic input; the active final problem is the base-field-normalized version of Residual Input 21.3. The circle rows replace $X^n - \beta$ by the twin-coset domain polynomial. Internal structure, fully proved: the det-degenerate directions $AB' - A'B = 0$ organize the solutions into *rays* $r \cdot (A_0, B_0)$, each ray corresponding to a single list codeword c with census contribution exactly $\binom{\mathrm{agr}(U,c)}{m}$, and distinct primitive members obey $\gcd(f_0, f_0') \mid A_0B_0' - A_0'B_0$, which is precisely the classical bound of $K - 1$ on pairwise codeword agreement in disguise. The split-pencil problem is therefore exactly the m-th binomial moment of the agreement profile, carried by an explicit two-generator pencil — the form in which partitioning, stability, or monodromy methods can now be aimed at a single algebraic family.

F.8 The (Q) stratum: divisor form and unification of the package

Applying the lattice machinery to the witness words themselves identifies (Q) as the boundary stratum of the split-pencil formulation: the package rests on *one* problem. Assume throughout the band profile condition $2m \leq n+K-1$, i.e. $w+1 \leq \omega$ (deployed: $2m \approx 2.23\times10^6 \leq 3.15\times10^6$).

Theorem F.23 ((Q) is the extreme profile of the split pencil)**.** *Fix $z \in \mathbb{B}^w$ and the witness word U_z, so $\widehat{U_z} = P_z := X^m + \sum_{j\leq w} z_j X^{m-j}$, and write $X^n = P_zQ_z + R_z$ with $\deg R_z < m$. Then $(1, P_z) \in M_{U_z}$ has shifted degree $w + 1$, and $g_2 := (Q_z,\ \beta - R_z) \in M_{U_z}$ has shifted degree ω. Exactly one of:*

(i) *$d_1(U_z) \leq w$, and the fiber is the exactly classified branch of Theorem F.15: $\binom{n-d_1}{m}$ for a unique near codeword, or empty;*

(ii) $d_1(U_z) = w+1$, *the pair* $\big((1,P_z),g_2\big)$ *is a reduced basis of profile* $(w+1,\omega)$, *and*

$$\big|\Phi_{m,w}^{-1}(z)\big| \;=\; \#\big\{A:\ \deg A \le \omega-w-1,\quad Q_z + A \;\big|\; X^n-\beta\big\}.$$

In particular (Q) *is exactly the prescribed-top-coefficient divisor problem for* $X^n-\beta$ *— count monic degree-*ω *divisors whose top* $w+1$ *coefficients equal those of the truncated quotient* Q_z *— and coincides with the split-pencil formulation Remark F.22 at the boundary profile* $d_1 = w+1$. *Inputs* (Q) *and* (A) *are therefore strata of one problem:* (Q) *the extreme profile, the aperiodic balanced core the interior profiles.*

Proof. $\deg P_z = m < n$ gives $\widehat{U_z} = P_z$ and $\mathrm{wdeg}(1,P_z) = m-(K-1) = w+1$. For g_2: at $x \in D$, $Q_z(x)P_z(x) = x^n - R_z(x) = \beta - R_z(x)$, so $g_2 \in M_{U_z}$; its shifted degree is $\max(\omega, \deg R_z - K+1) \le \max(\omega, w) = \omega$. The determinant of the pair is $(\beta - R_z) - Q_zP_z = \beta - X^n = -\Lambda_D$, a unit multiple, so the pair is a basis; the shifted degrees sum to $(w+1)+\omega = n-K+1$, the first row is N-led (as $w+1>0$) and the second W-led (as $\deg R_z - K + 1 \le w < \omega$), so the basis is reduced with profile $(w+1,\omega)$, provided $d_1 = w+1$; since $(1,P_z)$ shows $d_1 \le w+1$, the only alternative is $d_1 \le w$, which is branch (i) by Theorem F.15. In branch (ii), census elements have shifted degree ω and W-degree exactly ω: writing them as $A(1,P_z) + Bg_2$, predictable degrees force B scalar and nonzero, normalized to 1 by monicity, and $\deg A \le \omega - (w+1)$; so $W = Q_z + A$, whose top $w+1$ coefficients are those of Q_z, and by Lemma F.19 the split condition $W \mid X^n - \beta$ is the entire census condition. Conversely every monic degree-ω divisor with those top coefficients is $Q_z + A$. Finally, the prefix of Λ_T and the prefix of $W = \Lambda_{D\setminus T}$ determine each other unitriangularly through $\Lambda_T W = X^n - \beta$, so prescribing z is prescribing the top of W. □

Lemma F.24 (orbit constancy)**.** *For* $\zeta \in \mu_n$, *the twisted action* $z_j \mapsto \zeta^j z_j$ *preserves fiber sizes:* $T \mapsto \zeta^{-1}T$ *is a bijection of fibers, so all witness lists are constant along the* $p^w/\!\sim$ *orbits.*

Proof. $\Lambda_{\zeta^{-1}T}(X) = \zeta^{-m}\Lambda_T(\zeta X)$ has j-th prefix coordinate ζ^j times that of Λ_T, and $\zeta^{-1}D = D$. □

Remark F.25 (equivalent forms, and a hardness calibration). Three further faces of the same object, the first two proved. (1) *Additive form:* for $w < p$ the truncated logarithm is a group isomorphism $(1+u\mathbb{F}_p[u]/u^{w+1})^\times \cong (\mathbb{F}_p^w, +)$, under which the prefix of a divisor becomes the sum $\sum_{\xi\in S} v_\xi$ of the moment-curve vectors $v_\xi = (\xi, \xi^2/2, \ldots, \xi^w/w)$ over the root set: (Q) is balanced subset-sum equidistribution on a moment curve, at near-full entropy ($\binom{n}{m} \approx p^w \cdot 2^{192}$). (2) *Fourier accounting:* Parseval pins typical character sums at $\binom{N}{m}p^{-w/2}$, so fiber control requires square-root cancellation *again* in the signed sum over the p^w frequencies — two independent layers, which is what near-full-entropy equidistribution means quantitatively. (3) *Calibration:* (Q) is the function-field analogue of equidistribution of the divisors of a fixed integer in residue classes; the proved head depths $w \le w_0$ correspond to the classical regime of polylogarithmic moduli, while the band depth $w \approx 67466$ corresponds to moduli of size $(\#\text{divisors})^{1-o(1)}$, open even over $\mathbf{Z}$. Development-time small-instance checks of these identities are not used in any theorem or release certificate.

F.9 Fixed moments and the monodromy wall

Two final results delimit (Q) exactly: the asymptotic form reduces to *fixed*-moment collision bounds, and the strongest existing analytic framework is shown, by a proved expansion and exact arithmetic, to stop precisely one cohomological layer short.

Proposition F.26 (moment sandwich)**.** *For all $r \geq 2$, let*

$$R_{m,w} := S_{m,w} \max_z \mu_{m,w}(z), \qquad \theta_{m,w} := \log_2 S_{m,w}.$$

Then

$$\frac{\log_2 \Gamma_r}{r-1} \leq \log_2 R_{m,w} \leq \frac{\log_2 \Gamma_r + \theta_{m,w}}{r}.$$

Consequently, when $\theta_{m,w} = O(n)$, the asymptotic discrete-flatness bound $R_{m,w} \leq e^{o(n)}$ holds if and only if $\Gamma_r(m,w) \leq e^{o(n)}$ for every fixed r. Deciding a finite row with margin Δ bits by moments alone requires order on the scale $\theta_{m,w}/\Delta$.

Proof. The first inequality follows from $\Gamma_r \leq R_{m,w}^{r-1}$. For the second,

$$\max_z \mu(z) \leq \left(\sum_z \mu(z)^r \right)^{1/r} = \Gamma_r^{1/r} S_{m,w}^{-(r-1)/r},$$

so $R_{m,w} \leq (\Gamma_r S_{m,w})^{1/r}$. For the equivalence, fix an arbitrarily small $\varepsilon > 0$, take a sufficiently large constant r, and use $\theta_{m,w} = O(n)$ in the right inequality; the converse is the left inequality. □

Remark F.27 (finite moment orders at the deployed rows). At the active adjacent rows of Corollary 15.3, the bit scale $w\beta$ is about $2.09 \cdot 10^6$ bits. Thus the finite moment depth needed for a direct moment-to-max conversion depends on the row margin:

row	w	Δ bits	$w\beta/\Delta$
KoalaBear MCA	67470	22.197	$9.42 \cdot 10^4$
KoalaBear list	67470	22.011	$9.50 \cdot 10^4$
Mersenne-31 MCA	67446	3.259	$6.42 \cdot 10^5$
Mersenne-31 list	67446	3.073	$6.80 \cdot 10^5$.

Thus the finite adjacent rows require either direct constant-factor fiber control or a finite hierarchy reaching the displayed order; a uniform slogan such as "order 10^5" is not canonical for the Mersenne-31 rows.

Thus one possible proof route for the boundary-prefix bucket rests on *fixed-order* collision counts — for each constant r, the number of r-tuples of degree-ω divisors of $X^n - \beta$ with pairwise equal top-$(w+1)$ coefficients must not exceed $e^{o(n)} \binom{n}{m}^r S_{m,w}^{-(r-1)}$ — together with an interior split-pencil bound whose actual loss is checked against the reserve condition $\log L = o(R \log |\mathbb{B}|)$. The notation $e^{o(n)}$ alone does not supply that check. Fixed moments are exact algebraic objects (via Proposition F.6, weighted counts of $\{0, \pm 1\}$-configurations in the moment kernel), the most attackable known form of the input.

Proposition F.28 (character expansion of the split-divisor indicator)**.** *Let $D = \alpha H \subseteq \mathbb{F}_p^\times$, where H has order n, and put $r = (p-1)/n$. Let $H^\perp$ be the r multiplicative characters of $\mathbb{F}_p^\times$ that are trivial on H, and extend every one of them, including the trivial character, by $\chi(0) := 0$. Then for every $x \in \mathbb{F}_p$,*

$$\Phi_D(x) := \mathbf{1}_D(x) = \frac{1}{r} \sum_{\chi \in H^\perp} \chi(x\alpha^{-1}).$$

If D is the root set of $X^n - \beta$ and W is monic, define

$$I_D(W) := \begin{cases} \prod_{i=1}^{d} \Phi_D(\xi_i), & \text{if } W \text{ splits squarefree over } \mathbb{F}_p \text{ with roots } \xi_1, \dots, \xi_d, \\ 0, & \text{otherwise.} \end{cases}$$

Then $\mathbf{1}\{W \mid X^n - \beta\} = I_D(W)$. For a split squarefree polynomial W of degree d, choosing any ordering $\xi_1, \dots, \xi_d$ of its roots gives

$$\prod_{i=1}^{d} \Phi_D(\xi_i) = r^{-d} \sum_{(\chi_1,\dots,\chi_d) \in (H^\perp)^d} \prod_{i=1}^{d} \chi_i(\xi_i \alpha^{-1}).$$

The full sum is invariant under reordering the roots; equivalently, grouping the character tuples by permutation orbits produces well-defined symmetric factorization functions. For the deployed KoalaBear domain, $r = 1016$.

Proof. For $x \neq 0$, the first identity is character orthogonality on $\mathbb{F}_p^\times / H$; for $x = 0$, both sides are zero by the stated extension. A monic polynomial divides $X^n - \beta$ exactly when it is squarefree, splits, and every root lies in D; if it does not split squarefree, both the divisor indicator and $I_D(W)$ are zero by definition, while a zero root makes the product zero because $\beta \neq 0$. Expanding the product of the character sums gives the third identity. Permuting the roots permutes the full set of character tuples, so the total is independent of the chosen ordering. □

Remark F.29 (experimental monodromy route). By Theorem F.23, a prefix fiber can therefore be written as a short-interval sum of the symmetric twisted factorization functions above. This is a possible analytic route to prefix flatness, but the number of character assignments is enormous and termwise triangle inequalities lose far more than the deployed budget. Small-instance checks used during development do not enter any theorem or release certificate.

Remark F.30 (the exact wall, located). The expansion has $\approx 1016^\omega$ tuples, $\log_2 \approx 9{,}799{,}978$ bits, and the maximal conceivable per-tuple estimate — full square-root cancellation over the interval — saves $\frac{h+1}{2} \log_2 p \approx 14{,}156{,}266$ bits against an interval of 28.3 million bits. The triangle inequality over tuples therefore lands $\approx 23{,}956{,}053$ bits above the 192-bit target: no per-term theorem, however deep, can finish (Q). What remains is cancellation *among* the twisted sums — a statement about the cohomology of the whole character-tuple family over the prefix space, not of its fibers. The theory is otherwise tight at both ends: at the deep end $\omega \lesssim 50$ the trivial interval bound already meets the target (the regime of Theorem 9.6), and at the head $w \leq w_0$ the character route closes it (Corollary 14.9); the band is exactly the stretch where neither end reaches.

F.10 Conditional closing statement

Remark F.31 (historical Q/A grouping and the active template). The labels (Q) and (A) describe an earlier two-way grouping: boundary prefix fibers versus the remaining aperiodic census. They are useful proof-route coordinates, but they are not a complete closure theorem. The boundary bucket may be attacked through fixed-order collision bounds, and the interior aperiodic bucket through the balanced split-pencil reductions above; neither supplies the local primitive same-prefix maximum, an exhaustive priority assignment, or support-to-parameter coalescing.

Accordingly, the four adjacent-pair predictions follow only from a fully instantiated certificate of Section 21, including the corrected discrete/soft base-field models and the single summed finite budget. Failure of a candidate certificate may still locate a heavy prefix fiber, a large interior split-pencil census, a large local primitive shift-pair degree, or a coverage/coalescing defect. The row-dependent moment orders of Remark F.27 are calibration for one possible boundary proof route, not a substitute for the complete certificate.

References

[PP26] The Proximity Prize. The Proximity Prize: $1M Reed–Solomon Challenge. Preliminary release, 2026. https://proximityprize.org/.

[ABF26] G. Arnon, D. Boneh, and G. Fenzi. Open problems in list decoding and correlated agreement. IACR Cryptology ePrint Archive, Report 2026/680, 2026. https://eprint.iacr.org/2026/680.

[BCHKS25] E. Ben-Sasson, D. Carmon, U. Haböck, S. Kopparty, and S. Saraf. On proximity gaps for Reed–Solomon codes. IACR Cryptology ePrint Archive, Report 2025/2055, 2025. https://eprint.iacr.org/2025/2055.

[BCIKS20] E. Ben-Sasson, D. Carmon, Y. Ishai, S. Kopparty, and S. Saraf. Proximity gaps for Reed–Solomon codes. In *Proc. 61st IEEE Symposium on Foundations of Computer Science (FOCS)*, 2020. Extended manuscript: IACR Cryptology ePrint Archive, Report 2020/654.

[IK04] H. Iwaniec and E. Kowalski. *Analytic Number Theory.* American Mathematical Society Colloquium Publications, Vol. 53, 2004.

[Sch76] W. M. Schmidt. *Equations over Finite Fields: An Elementary Approach.* Lecture Notes in Mathematics, Vol. 536. Springer, 1976.

[BCKL21] E. Ben-Sasson, D. Carmon, S. Kopparty, and D. Levit. Elliptic curve fast Fourier transform (ECFFT) part I: Fast polynomial algorithms over all finite fields. arXiv:2107.08473, 2021.

[Cho26a] P. Chojecki. Capacity-edge obstructions to Reed–Solomon mutual correlated agreement over smooth multiplicative domains. Manuscript, 2026.

[Cho26b] P. Chojecki. Slack, quotient cores, and the entropy gap for smooth-domain Reed–Solomon codes: list fibers, cyclotomic rigidity, Galois-amplified collisions, and the corrected slack mutual-correlated-agreement theory. Manuscript, merged corrected edition, 2026.

[CS25] E. Crites and A. Stewart. On Reed–Solomon proximity gaps conjectures. IACR Cryptology ePrint Archive, Report 2025/2046, 2025. https://eprint.iacr.org/2025/2046.

[HLP24] U. Haböck, D. Levit, and S. Papini. Circle STARKs. IACR Cryptology ePrint Archive, Report 2024/278, 2024. https://eprint.iacr.org/2024/278.