

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

Przemek Chojceki
ulam.ai

June 27, 2026

Abstract

We study Reed–Solomon codes evaluated on smooth multiplicative domains: subgroups $H \leq \mathbb{F}_q^\times$ of power-of-two order n , at fixed rate $\rho = k/n$, the domains of deployed proximity-testing protocols. This merged corrected edition consolidates two complementary developments into a single document with an explicit proof-status ledger: the list/entropy route based on locator fibers, quotient cores, characteristic-zero rigidity, and Galois-amplified reduction collisions (Part I); and the slack mutual-correlated-agreement (MCA) theory based on the exact bad-slope calculus, failure ladders, density positive results, direction arithmetic, quotient floors, and residue-line normal forms (Part II).

On the list side, the negative conclusions are unconditional: the coefficient-pigeonhole mechanism gives exponential lists below the ambient entropy scale, and quotient cores force superpolynomial lists at the deployed dyadic rates — more generally along any quotient-divisible scale $M \mid k$ — whenever $\sigma = o(n/\log n)$, refuting the one-scale ambient- q entropy-gap prediction even over prime fields. Any polynomial-list theorem must therefore use a corrected reserve of the form $\tau \gtrsim \max\{\tau^*(\rho, q_D), c_{\rho, B}/\log n\}$, with subfield variants governed by the field of definition. The positive structural results are: arbitrary list size reduces to locator fibers; over characteristic zero, equal prefix fibers among subsets of 2-power roots of unity are exactly quotient-periodic; and in prime fields $p \equiv 1 \pmod{n}$ a Galois-amplified norm argument shows that for $p > \exp(Cn \log n/\sigma)$ every monomial-prefix finite-field collision is already a characteristic-zero quotient collision, so monomial-prefix list sizes are polynomial at the corrected reserve for split primes in the quasi-polynomial range $p > \exp(O((\log n)^2))$.

On the MCA side, the no-slack form of the principle is false, and the corrected slack theory at radius $1 - \rho - t/N$ separates proved theorems from conditional analytic inputs and open inverse problems. The unconditional spine is the exact slack characterization of the canonical line $x^{k+T} + zx^k$, the quotient locator and coset collapse, the cyclotomic and Fermat rigidity lemmas, the two-moment square-root construction, and the dyadic descent theorem. The free-pool results are stated as conditional consequences of a precise subgroup exponential-sum input. On the positive side the safe-slack bookkeeping uses the effective parameter $C_{\text{eff}} = \sigma \log_2 n/n$; the quotient-exact floor shows the closed-form quotient count $A(N', \ell') = 2^{\beta(\rho)N'(1-o(1))}$ survives at every prime above the quotient norm threshold $p > (2\ell')^{N'/2}$, giving a genuine MCA lower bound from the quotient profile at large primes and a concrete cap for several deployed-rate challenge parameters. The earlier all-bases, all non-dyadic-direction theorem remains removed (its rigidity step is false for arbitrary bases) and replaced by a Newton-obstruction criterion and the valid zero-base theorem.

A transfer part extends the mechanisms beyond the smooth multiplicative prime-field setting. A subfield-confinement theorem decouples the two grand challenges over deployed extension fields: B -valued lines have all MCA-bad slopes in B , deflating every prime-field witness to density

$|B|/|\mathbb{F}| < 2^{-154}$ at the deployed sextic parameters, while quotient-core lists persist verbatim. The quotient-core and slack-one mechanisms transfer to Dickson–Chebyshev quotient chains, hence to circle-group x -coordinate Reed–Solomon constituents at Mersenne-31 parameters; the $\{2, 3\}$ -smooth relation module and set decomposition supply the mixed-radix rigidity input; and a torsion-coset template theorem, conditional on an imported effective torsion-coset bound (Laurent, Aliev–Smyth), replaces the retracted arbitrary-base claim at the level of characteristic-zero template structure. Finally, conditional on an imported list-to-agreement conversion of Crites–Stewart, composing quotient-core lists with the conversion caps the achievable MCA threshold at challenge dimensions: for challenge codes with $n \leq q/2$ over any field of size at most 2^{256} , the error is at least 2^{-42} already at gap 2^{-11} . The sharpened, final form of this cap — gap 2^{-9} (2^{-10} at rate $1/16$), error 2^{-86} with no condition relating n and q , and extension fields included — is proved in the companion cap paper [2], which this edition cross-references throughout.

The document closes with an integrated frontier: the final threshold statements are isolated as finite-field local-limit conjectures for locator fibers and residue lines, calibrated on both sides by the same two-scale, floor-corrected reserve. We do not claim the full polynomial-field list theorem or the all-line MCA positive theorem as proved.

This strengthened version also promotes the stable high-agreement tangent material from the experimental notes into the main theory: the affine, no-loss CA, and projective-line numerators are exactly $r + 1$ in the range $r \leq (n - k)/3$, interleaved lists are unique for $r \leq (n - k)/2$, and degree- d finite power-curve ledgers have a theorem-backed $d(r + 1)$ upper envelope with exactness under a split moving-root hypothesis. Thus the high-agreement boundary layer is now solved, and the remaining frontier is explicitly the near-capacity quotient-profile and aperiodic local-limit problem.

Contents

1	Introduction	4
I	List entropy, quotient cores, and Route 1	7
2	List decoding as a locator-fiber problem	7
3	Two lower-bound scales	8
3.1	The entropy pigeonhole scale	8
3.2	The quotient-core scale	9
4	Characteristic-zero inverse quotient theorem	11
5	Galois-amplified finite-field collision sieve	13
6	The corrected positive conjectures	15
7	Why pairwise/BCH information is insufficient	16
II	Slack mutual correlated agreement: the corrected theory	16
8	The list-side ledger, and the MCA comparison point	17
9	Definitions, and the two delimitations	17

10 The slack calculus	18
10.1 The exact characterization	18
10.2 Quotients, and the collapse of symmetric witnesses	19
10.3 Certified slack thresholds, and curve transfer	20
11 The two-moment certificate, and fullness at the root barrier	21
12 The failure ladder, I: the free pool	22
13 The failure ladder, II: descent, and the dichotomy	24
14 The rigidity theory at every slack	25
15 Positive theorems: the density regime	26
16 The direction-arithmetic of slack	28
17 The exact regime, and the closed-form count	30
18 The boundary, and the conjecture	32
18.1 The exact normal form	32
18.2 The obstruction, the open core, and the conjecture	33
19 The quotient-exact floor, and the necessity of the quotient profile	34
19.1 The quotient-exact floor	35
19.2 Necessity of the quotient profile for MCA	36
19.3 A concrete cap for challenge-size prime fields	38
20 Solved high-agreement boundary layer	39
21 Summary, and open problems	42
 III Transfers, caps, and repairs	 42
22 A universal field-size cap via list-to-agreement conversion	43
23 Subfield confinement over extension fields	45
24 Chebyshev domains and circle codes	46
25 Torsion-coset repairs: $\{2, 3\}$-smooth rigidity and arbitrary-base templates	48
26 Additive domains, dilation invariance, and further transfers	50
 IV Integrated frontier and synthesis	 51
27 Integrated frontier and final conjectural asymptotics	51
28 Conclusion	53

1 Introduction

Proximity testing protocols [3, 4] fold pairs of oracles f, g into random combinations $u_z = f + zg$ and rest their soundness on a mutual correlated agreement principle: if u_z agrees with a Reed–Solomon codeword on a large support, then f and g should themselves agree with codewords on that support, except for a controlled set of slopes. Capacity is $1 - \rho$ at rate ρ ; the Johnson radius is $1 - \sqrt{\rho}$, below which proximity gaps and correlated agreement are known [27, 26]; the difficult regime lies between them, especially for the smooth multiplicative domains used by FFT-based protocols.

The no-slack, up-to-capacity formulation is false on smooth domains [1]. Thus the right object is MCA with an explicit slack reserve. At a quotient scale $N \mid n$ and slack t , the relevant radius is

$$1 - \rho - \frac{t}{N},$$

and the central question becomes: how many slopes are bad, and which mechanisms produce them?

This merged edition consolidates the list/entropy route and the slack-MCA theory under a single corrected ledger: it keeps the parts that are proved from the algebraic setup and downgrades the parts that require additional fixed-prime equidistribution input. Part I develops the list side (locator fibers, the two lower-bound scales, characteristic-zero rigidity, and the Galois-amplified collision sieve); Part II develops the corrected slack-MCA theory; Part III develops transfers of the failure mechanisms beyond smooth multiplicative prime-field domains (extension fields, additive and Chebyshev/circle x -coordinate domains), a conditional universal field-size cap, and the torsion-coset repair of the arbitrary-base theory; Part IV states the integrated frontier. On the MCA side, the contributions are as follows.

1. *Exact slack calculus.* The bad-slope set of the canonical slack line $x^{k+T} + zx^k$ is identified exactly as a multi-symmetric image of the evaluation set (Theorem 10.2). The quotient locator (Lemma 10.3) and the coset-collapse proposition (Proposition 10.6) explain why subgroup-symmetric slack witnesses reduce to a one-condition ladder. Cyclotomic and digit rigidity (Theorems 10.7 and 10.8) show that the remaining asymmetric witnesses are genuinely modular rather than characteristic-zero phenomena.
2. *Failure mechanisms.* The Dias da Silva–Hamidoune theorem gives certified slack-one failure at the square-root barrier (Corollary 10.11), while the two-moment construction gives full coverage of (e_1, e_2) under explicit square-root hypotheses (Theorem 11.2 and Corollary 11.3). The dyadic descent theorem is retained with all divisibility assumptions made explicit (Theorem 13.1). The free-pool ladder is stated conditionally on a precise subgroup exponential-sum estimate (Theorem 12.2); where the cited estimates supply that estimate, the stated ranges follow, but the theorem is not used as a black-box proof of stronger fixed-prime equidistribution.
3. *Characteristic-zero rigidity and direction arithmetic.* The inverse quotient theorem (Theorem 14.1) proves that vanishing at dyadic frequencies forces quotient invariance. Consequently, zero-base non-dyadic directions have no characteristic-zero nonzero slopes (Theorem 16.2), while dyadic directions reduce to lower-level direction-one problems. The earlier all-bases non-dyadic theorem is replaced by the Newton-obstruction criterion of Theorem 16.5; arbitrary bases require extra equations before the rigidity theorem can be applied.

4. *Positive density, quotient floors, and stable results.* At deepest-level slack σ , the correct normalization is

$$C_{\text{eff}} = \frac{\sigma \log_2 n}{n}.$$

With this bookkeeping, the canonical line has a no-threshold descent floor of order $\rho(1 - \rho)(\log_2 n / C_{\text{eff}})^2$ at every prime and a density ceiling for at least half the admissible primes in explicit windows ([Theorem 15.1](#)). The quotient-exact floor of [Proposition 19.1](#) upgrades the floor to $n^{\beta(\rho)/C_{\text{eff}}(1-o(1))}$ at every prime above the quotient norm threshold. Every monomial base in the canonical direction obeys the same quotient-class cap, with a sharpened collision term ([Theorem 15.2](#)). In the stable range $p > (2s)^{n/2}$, the canonical-line and zero-base direction predictions are exact at every prime — exactness is claimed only for these proved strata, not for arbitrary bases — and the closed-form canonical count remains valid ([Theorems 17.1](#) and [17.2](#)).

5. *Boundary and conjecture.* The exact normal form ([Theorem 18.3](#)) identifies the residue-line packing number that equals MCA error after maximizing over line data. The denominator-closure statement is kept as a coordinate normal form, not a quantitative bound ([Proposition 18.5](#)). The quotient-profile lower bound ([Theorem 19.4](#)) calibrates the necessary quotient reserve for polynomial MCA bounds at large primes, while the final conjecture remains a per-fiber packing conjecture with a separate conditional failure half ([Conjecture 18.8](#)).

The merged theorem package, with proof status, is as follows.

Theorem 1.1 (Integrated theorem package, with status). *For smooth multiplicative domains H of order $n = 2^m$ at fixed rate ρ :*

- (a) (Proved.) *Below the entropy scale $\tau^*(\rho, q)$, explicit monomial-prefix received words have exponentially many nearby codewords ([Theorem 3.2](#) and [Corollary 3.4](#)).*
- (b) (Proved.) *Independently of q , quotient cores produce at least $\binom{n/M-1}{k/M}$ codewords at agreement $k+\sigma$ whenever $M \mid k$, $k/M \leq n/M-1$, and $\sigma < M$; consequently, at the deployed dyadic rates (and along any scale with $M \mid k$), polynomial-in- n list decoding requires $\sigma = \Omega(n/\log n)$, and the one-scale ambient- q entropy-gap prediction is false even over prime fields ([Theorem 3.5](#) and [Corollary 3.6](#)).*
- (c) (Proved.) *Over characteristic zero, equal prefix moments among subsets of μ_n are exactly quotient-periodic; characteristic-zero prefix fibers are polynomial once $\sigma \geq Cn/\log n$ ([Theorem 4.2](#) and [Corollary 4.3](#)).*
- (d) (Proved.) *If $p \equiv 1 \pmod{n}$ and $p > \exp(C_1 n \log(2n)/\sigma)$, every finite-field monomial-prefix collision is quotient-periodic; monomial-prefix received words have polynomial lists at $\sigma \geq C_0 n / \log_2 n$ for split primes in the quasi-polynomial range $p > \exp(O((\log n)^2))$ ([Theorem 5.2](#) and [Corollary 5.3](#)).*

- (e) (Proved or conditional, as marked.) *For MCA, the canonical-line bad-slope set at slack T is exactly the multi-symmetric image $\mathcal{B}_T$ (Theorem 10.2), and subgroup-symmetric witnesses collapse onto the one-condition ladder (Proposition 10.6). Unconditional failure mechanisms are the certified slack-one rung (Corollary 10.11), two-moment fullness at the root barrier (Theorem 11.2 and Corollary 11.3), and dyadic descent (Theorem 13.1); the polynomial-range free-pool ladder is conditional on a precise subgroup exponential-sum input (Theorem 12.2). The quotient-exact floor survives at every prime above the quotient norm threshold (Proposition 19.1), the quotient profile is necessary for polynomial MCA slope bounds at large primes (Theorem 19.4), and the deployed-rate challenge cap holds below the per-rate prime-field reach (Proposition 19.7). All-line positive MCA is equivalent, up to the tangent floor, to residue-line packing (Theorem 18.3).*
- (f) (Proved, with imported black boxes as marked.) *Transfers and caps: MCA-bad slopes of B -valued lines over an extension $\mathbb{F} \supseteq B$ are confined to B , deflating the prime-field witness theory to density at most $|B|/|\mathbb{F}|$ while list obstructions persist (Theorem 23.1 and Corollaries 23.2 and 23.3); the canonical locator stratum on $\mathbb{F}_2$ -affine domains is confined to a single additive coset (Proposition 26.1); quotient cores and the slack-one stratum transfer to Dickson–Chebyshev and circle-group x -coordinate domains, with error one at Mersenne-31 parameters for those Reed–Solomon constituents (Theorem 24.2, Proposition 24.3, and Corollary 24.4); the $\{2, 3\}$ -smooth relation module gives the mixed-radix rigidity input, and arbitrary-base slabs reduce to at most $C(\sigma, s)$ torsion-coset template families under the stated effective torsion-coset import (Theorems 25.2 and 25.5 and Corollary 25.6); and, conditional on the imported list-to-agreement conversion of [25], smooth-domain challenge codes at the deployed rates with $2^{11} \mid n$, $n \leq q/2$, $k \leq 2^{40}$, and $q \leq 2^{256}$ have $\varepsilon_{\text{mca}} > 2^{-42}$ at gap 2^{-11} (Theorem 22.3 and Corollary 22.4); the sharpened form — gap 2^{-9} , error 2^{-86} , no field-size condition — is [2].*
- (g) (Conjectural.) *The full polynomial-field positive list theorem and the floor-corrected all-line MCA theorem are isolated as finite-field local-limit statements with quotient exceptions explicitly separated (Conjectures 6.1 to 27.2, 18.8 and 3.9).*

Setting and notation. p is prime (q a general prime power where stated); D or $H \leq \mathbb{F}_q^\times$ is a multiplicative subgroup of order n , *smooth* meaning n is a power of two; $C = \text{RS}[\mathbb{F}_q, D, k] = \{(P(x))_{x \in D} : \deg P < k\}$ with $k = \rho n$, $\rho \in (0, 1)$ fixed. For a finite set A in a field, $L_A(X) = \prod_{a \in A} (X - a) = X^{|A|} + \sum_{j \geq 1} (-1)^j e_j(A) X^{|A|-j}$ and $p_j(A) = \sum_{a \in A} a^j$; for $h \geq 0$, $h^\wedge A$ is the set of sums of h distinct elements of A . H_2 is binary entropy. Divisibility hypotheses ($\rho N \in \mathbb{Z}$, $t \mid N$,

$t \mid \rho N$, $\sigma \mid k$, and the like) are stated where they are used.

Part I

List entropy, quotient cores, and Route 1

2 List decoding as a locator-fiber problem

Every received word $y : H \rightarrow \mathbb{F}_q$ has a unique interpolant $U \in \mathbb{F}_q[X]$ of degree $< n$ on H . Write $\text{List}(y, \delta)$ for the list of codewords within relative radius δ of y , and $\text{List}(C, \delta) = \max_y |\text{List}(y, \delta)|$ for the worst-case list size. A codeword P agrees with y on a set $S \subseteq H$ if and only if

$$U(X) \equiv P(X) \pmod{L_S(X)}.$$

Since $\deg P < k$, this gives the following exact reduction.

Definition 2.1 (Locator fibers). For $k < s \leq n$ and $\deg U < n$, define

$$\mathcal{F}_U(s) = \{S \subseteq H : |S| = s, \deg(U \bmod L_S) < k\}.$$

For monomial-prefix data $c = (c_1, \dots, c_\sigma) \in \mathbb{F}_q^\sigma$, $s = k + \sigma$, define

$$\Phi_\sigma(S) = (e_1(S), \dots, e_\sigma(S)), \quad S \in \binom{H}{s}.$$

Proposition 2.2 (Arbitrary-word fiber upper bound). *For every received word y with interpolant U ,*

$$|\text{List}(y, 1 - s/n)| \leq |\mathcal{F}_U(s)|.$$

Consequently

$$\text{List}(C, 1 - s/n) \leq \max_{\deg U < n} |\mathcal{F}_U(s)|.$$

Proof. Each codeword P agreeing with y on at least s coordinates has an s -subset S of its agreement set. Then $U \equiv P \pmod{L_S}$, so $S \in \mathcal{F}_U(s)$. For a fixed S , the residue $U \bmod L_S$ is unique, hence the corresponding P is unique. Choosing one such S for each listed codeword injects the list into $\mathcal{F}_U(s)$. $\square$

Proposition 2.3 (Monomial-prefix words are exact prefix fibers). *Let $s = k + \sigma$ and*

$$U_c(X) = X^s + \sum_{j=1}^{\sigma} (-1)^j c_j X^{s-j}.$$

Then the map

$$S \longmapsto U_c - L_S$$

is a bijection from $\Phi_\sigma^{-1}(c)$ to the list of degree- $< k$ codewords agreeing with U_c on at least s points.

Thus

$$|\text{List}(U_c, 1 - s/n)| = |\Phi_\sigma^{-1}(c)|.$$

Proof. If $\Phi_\sigma(S) = c$, then the top σ coefficients of U_c and L_S agree, so $P_S = U_c - L_S$ has degree $< s - \sigma = k$. Also $P_S(x) = U_c(x)$ for $x \in S$. Conversely, if P has degree $< k$ and agrees with U_c on an s -set S , then $U_c - P$ is monic of degree s and vanishes exactly on S , hence $U_c - P = L_S$, forcing $\Phi_\sigma(S) = c$. Injectivity follows because L_S determines S . $\square$

Thus the positive half of list decoding can be stated without coding language:

$$\max_{\deg U < n} |\mathcal{F}_U(k + \sigma)| \quad \text{should be polynomial above the corrected reserve.}$$

3 Two lower-bound scales

The reference scale for the first mechanism is the classical entropy gap.

Definition 3.1 (Entropy gap). For $0 < \rho < 1$ and $q \geq 2$, $\tau^*(\rho, q)$ is the unique $\tau \in (0, 1 - \rho)$ with $\tau \log_2 q = H_2(\rho + \tau)$; as $q \rightarrow \infty$ with ρ fixed, $\tau^*(\rho, q) = (H_2(\rho) + o(1))/\log_2 q$.

3.1 The entropy pigeonhole scale

Theorem 3.2 (Coefficient pigeonhole lower bound). *Let $1 \leq \sigma \leq n - k$ and $s = k + \sigma$. There exists a monomial-prefix word U_c such that*

$$|\text{List}(U_c, 1 - s/n)| \geq \frac{1}{q^\sigma} \binom{n}{s}.$$

Proof. The map $\Phi_\sigma : \binom{H}{s} \rightarrow \mathbb{F}_q^\sigma$ has $\binom{n}{s}$ inputs and at most q^σ outputs. Some fiber has size at least $q^{-\sigma} \binom{n}{s}$. Apply [Proposition 2.3](#). $\square$

Corollary 3.3 (Generated-field pigeonhole). *If $H \subseteq B$ for a subfield $B \leq \mathbb{F}_q$ of size q_D , then Φ_σ has image in B^σ — its coordinates $e_j(S)$ are elementary symmetric functions of subsets of $H \subseteq B$ — so the pigeonhole improves to*

$$|\text{List}(U_c, 1 - s/n)| \geq \frac{1}{q_D^\sigma} \binom{n}{s}$$

for some monomial-prefix word U_c with coefficients in B . This is the form consumed by the necessary half of [Conjecture 3.9](#), which states the entropy floor at the generated field; in locator-fiber form, the same one-line strengthening is the subfield pigeonhole of [\[2, Lem. “locator fibers are lists”\]](#).

Proof. Identical to [Theorem 3.2](#) with $\mathbb{F}_q^\sigma$ replaced by B^σ : $e_j(S) \in B$ for every $S \subseteq H$, so Φ_σ has at most q_D^σ outputs. $\square$

Corollary 3.4 (Failure below the entropy gap). *If $q = 2^{o(n)}$, $\sigma = \tau n$, and*

$$\sigma \log_2 q \leq (1 - \varepsilon) \log_2 \binom{n}{k + \sigma},$$

then some received word has list size at least

$$2^{\varepsilon n H_2(\rho+\tau)-o(n)}.$$

In particular, this gives $2^{\Omega(n)}$ -size lists for every fixed multiplicative gap below the entropy scale, $\tau \leq (1 - \varepsilon_0)\tau^*(\rho, q)$ with $\varepsilon_0 > 0$ fixed; for $\varepsilon_0 \rightarrow 0$ the displayed pigeonhole bound still applies but may be subexponential.

Proof. By [Theorem 3.2](#), some monomial-prefix word has list size at least

$$q^{-\sigma} \binom{n}{k+\sigma} = 2^{\log_2 \binom{n}{k+\sigma} - \sigma \log_2 q} \geq 2^{\varepsilon \log_2 \binom{n}{k+\sigma}},$$

and $\log_2 \binom{n}{k+\sigma} = n H_2(\rho + \tau) - o(n)$. For the last sentence, set $h(\tau) = H_2(\rho + \tau)$ and $g(\tau) = \tau \log_2 q / h(\tau)$ on $(0, 1 - \rho)$. Concavity of h gives $h(\tau) - \tau h'(\tau) \geq h(0) = H_2(\rho) > 0$, so g is strictly increasing, and $g(\tau^*) = 1$ by the definition of τ^* . If $\tau \leq (1 - \varepsilon_0)\tau^*$ with fixed $\varepsilon_0 > 0$, then concavity again gives $h((1 - \varepsilon_0)\tau^*) \geq (1 - \varepsilon_0)h(\tau^*) + \varepsilon_0 H_2(\rho)$, whence

$$g(\tau) \leq g((1 - \varepsilon_0)\tau^*) \leq \frac{(1 - \varepsilon_0)h(\tau^*)}{(1 - \varepsilon_0)h(\tau^*) + \varepsilon_0 H_2(\rho)} \leq 1 - c, \quad c := \frac{\varepsilon_0 H_2(\rho)}{1 + \varepsilon_0 H_2(\rho)} > 0,$$

uniformly in q , using $h \leq 1$. Thus $\sigma \log_2 q \leq (1 - c)n H_2(\rho + \tau)$, and the displayed bound is at least $2^{c n H_2(\rho+\tau)-o(n)} = 2^{\Omega_{\rho, \varepsilon_0}(n)}$: on $[\rho, \rho + \tau^*]$ the entropy is at least $\min\{H_2(\rho), H_2(\rho + \tau^*)\}$, and $H_2(\rho + \tau^*) = \tau^* \log_2 q$ cannot be small for fixed ρ , since otherwise $\rho + \tau^* \rightarrow 1$ would force $\tau^* \log_2 q \geq \tau^* \rightarrow 1 - \rho > 0$, a contradiction. $\square$

3.2 The quotient-core scale

The entropy scale is not the only obstruction. Smooth subgroups have quotient structure that random domains do not have.

Theorem 3.5 (Quotient-core list obstruction). *Let $H \leq \mathbb{F}_q^\times$ have order n . Let $K \leq H$ have order M , put $N = n/M$, and assume $M \mid k$ and $k/M \leq N - 1$. If $1 \leq \sigma < M$, then there is a received word $Y : H \rightarrow \mathbb{F}_q$ with at least*

$$\binom{N-1}{k/M}$$

codewords agreeing with Y on $k + \sigma$ coordinates. Equivalently,

$$\text{List}(\text{RS}[\mathbb{F}_q, H, k], 1 - \rho - \sigma/n) \geq \binom{n/M - 1}{k/M}.$$

Proof. Choose one K -coset $C_0 \subset H$ and choose $T \subset C_0$ with $|T| = \sigma$. Let

$$L_T(X) = \prod_{t \in T} (X - t), \quad Y(X) = X^k L_T(X).$$

For every subset A of the quotient $H/K \setminus \{C_0\}$ of size k/M , let $U_A \subset H$ be the union of the K -cosets in A . Since each K -coset is a fiber of $x \mapsto x^M$, its locator has the form $X^M - \alpha$. Hence

$$L_A(X) = \prod_{x \in U_A} (X - x) = \prod_{\alpha \in A'} (X^M - \alpha) = X^k + c_1 X^{k-M} + c_2 X^{k-2M} + \dots.$$

Thus $X^k - L_A(X)$ has degree at most $k - M$. Define

$$P_A(X) = Y(X) - L_T(X)L_A(X) = L_T(X)(X^k - L_A(X)).$$

Since $\deg L_T = \sigma < M$, we have $\deg P_A < k$, so P_A is a codeword. On $T \cup U_A$, either L_T or L_A vanishes, so $P_A = Y$. The agreement set has size $\sigma + k$. Distinct A give distinct locators and distinct codewords. $\square$

Corollary 3.6 (The one-scale ambient- q conjecture is false). *Fix a deployed rate $\rho = 2^{-a}$ and $\varepsilon > 0$. There are infinitely many smooth prime-field domains $H \leq \mathbb{F}_p^\times$ with $|H| = n$ and*

$$p = 2^{\Theta(\sqrt{n})}$$

such that at radius

$$1 - \rho - (1 + \varepsilon + o(1))\tau^*(\rho, p)$$

the list size is at least

$$p^{1/(2(1+\varepsilon)) - o(1)} = 2^{\Theta(\sqrt{n})},$$

which is superpolynomial in n .

Proof. Take $n = 2^m$. By the prime number theorem in arithmetic progressions in a Siegel–Walfisz range [15], there are primes $p \equiv 1 \pmod{n}$ with $\log_2 p = \Theta(\sqrt{n})$ for all large n along a suitable sequence. Let $H \leq \mathbb{F}_p^\times$ have order n . Since p is prime, there is no proper subfield obstruction.

Let $\sigma = \lceil (1 + \varepsilon)\tau^*(\rho, p)n \rceil$. Because $\tau^*(\rho, p) = (H_2(\rho) + o(1))/\log_2 p$, we have $\sigma = \Theta(\sqrt{n})$. Let M be the least power of two larger than σ . For large n , $M \mid k = \rho n$ and $M \leq 2\sigma$. Theorem 3.5 gives

$$L \geq \binom{n/M - 1}{k/M}.$$

By Stirling,

$$\log_2 L \geq \frac{n}{M} H_2(\rho) - o(n/M) \geq \frac{H_2(\rho) + o(1)}{2(1 + \varepsilon)\tau^*(\rho, p)} = \left(\frac{1}{2(1 + \varepsilon)} - o(1) \right) \log_2 p.$$

$\square$

Remark 3.7 (Corrected reserve). For a list bound n^B , the quotient-core family at $\sigma = Cn/\log_2 n$ has size about $n^{H_2(\rho)/C}$ on the grid. Thus a theorem with list bound n^B must have $C \gtrsim H_2(\rho)/B$ at the second scale. For unspecified polynomial list size, the necessary scale is $\sigma = \Omega(n/\log n)$; for fixed exponent B , the constant matters.

Definition 3.8 (Quotient-core profile). For a smooth domain H of order n and $k = \rho n$,

$$\mathcal{Q}_H(\eta) := \max_{\substack{M \mid \gcd(n, k), \lceil \eta n \rceil < M \\ k/M \leq n/M - 1}} \log_2 \binom{n/M - 1}{k/M}$$

($-\infty$ over the empty set): the logarithmic size of the largest quotient-core list still active at gap η . Correspondingly, every quotient contribution of the form $2^{(\beta(\rho)/H_2(\rho))\mathcal{Q}_H(\eta)(1+o(1))}$ below is read as zero when the defining maximum is empty.

Conjecture 3.9 (Quotient-profile corrected list entropy gap). *Let $C_n = \text{RS}[\mathbb{F}_{q_n}, H_n, \rho n]$ with $\text{poly}(n) \leq q_n \leq 2^{o(n)}$ and $q_{D,n}$ the generated-field size, and fix a list target n^B . Necessary for $\text{List}(C_n, 1 - \rho - \eta_n) \leq n^B$ is*

$$\eta_n \geq (1 - o(1)) \tau^*(\rho, q_{D,n}), \quad \mathcal{Q}_{H_n}(\eta_n) \leq B \log_2 n + O(\log \log n)$$

(*Corollary 3.3* and *Theorem 3.5*); the conjectural positive half is that, after fixed slack in both inequalities, the conditions are sufficient. On dyadic domains at the deployed rates the scale form is $\eta_n \gtrsim \max\{\tau^*(\rho, q_{D,n}), \Theta_{\rho,B}(1/\log n)\}$.

4 Characteristic-zero inverse quotient theorem

We now prove that, over characteristic zero, the quotient-core phenomenon is the entire prefix-fiber structure for 2-power roots. This is the toric Manin–Mumford/torsion-coset part of the program (compare the torsion-coset theorem [7] and the theory of vanishing sums of roots of unity [5, 6]), but for 2-primary roots it has an elementary proof.

Let $n = 2^m$, let $\zeta = e^{2\pi i/n}$, and identify $\mu_n = \langle \zeta \rangle$ with the complex n -th roots of unity. For $h : \mu_n \rightarrow \mathbb{Z}$, define

$$\hat{h}(j) = \sum_{x \in \mu_n} h(x) x^j \in \mathbb{Z}[\zeta].$$

Lemma 4.1 (Antipodal identity). *Let $n' = 2^r$ with $r \geq 1$, and let $g : \mu_{n'} \rightarrow \mathbb{Z}$. If*

$$\sum_{y \in \mu_{n'}} g(y) y = 0,$$

then $g(y) = g(-y)$ for every $y \in \mu_{n'}$.

Proof. Write $y = \zeta_{n'}^a$. The sum becomes

$$\sum_{a=0}^{n'/2-1} (g(\zeta_{n'}^a) - g(-\zeta_{n'}^a)) \zeta_{n'}^a.$$

The powers $1, \zeta_{n'}, \dots, \zeta_{n'}^{n'/2-1}$ are linearly independent over $\mathbb{Q}$ because $\Phi_{n'}(X) = X^{n'/2} + 1$. Hence every coefficient difference is zero. $\square$

Theorem 4.2 (Inverse quotient theorem over $\mathbb{C}$). *Let $n = 2^m$, $1 \leq \sigma < n$, and let M_0 be the least power of two with $M_0 > \sigma$. Let $S, T \subseteq \mu_n$ have the same cardinality and satisfy*

$$e_j(S) = e_j(T) \quad (1 \leq j \leq \sigma)$$

over $\mathbb{C}$. Then $S \triangle T$ is a union of cosets of the subgroup $\mu_{M_0} \leq \mu_n$. Equivalently, $\mathbf{1}_S - \mathbf{1}_T$ is constant on μ_{M_0} -cosets. Conversely, every such quotient rearrangement preserves all e_j for $1 \leq j < M_0$.

Proof. By Newton identities in characteristic zero, equality of $e_1, \dots, e_\sigma$ and equality of cardinalities imply equality of the first σ power sums. Thus, for $h = \mathbf{1}_S - \mathbf{1}_T$,

$$\widehat{h}(j) = 0 \quad (1 \leq j \leq \sigma).$$

It is enough to use the dyadic frequencies. Let $M_0 = 2^a$. Since $2^{a-1} \leq \sigma$, we have

$$\widehat{h}(1) = \widehat{h}(2) = \dots = \widehat{h}(2^{a-1}) = 0$$

at the listed dyadic powers.

We prove by induction that vanishing at $1, 2, \dots, 2^i$ in this dyadic sense forces invariance under $\mu_{2^{i+1}}$. For $i = 0$, $\widehat{h}(1) = 0$ and [Lemma 4.1](#) give invariance under μ_2 . Suppose h is invariant under μ_{2^i} and $\widehat{h}(2^i) = 0$. The map $x \mapsto x^{2^i}$ sends μ_n onto $\mu_{n/2^i}$ with fibers the μ_{2^i} -cosets. Let $h^*(y)$ be the common value on the fiber above y . Then

$$0 = \widehat{h}(2^i) = 2^i \sum_{y \in \mu_{n/2^i}} h^*(y)y.$$

By [Lemma 4.1](#), $h^*(y) = h^*(-y)$. Pulling this identity back says that h is invariant under a primitive 2^{i+1} -st root of unity, hence under $\mu_{2^{i+1}}$. Iterating to $i = a - 1$ gives μ_{M_0} -invariance.

For the converse, if h is constant on μ_{M_0} -cosets, then

$$\widehat{h}(j) = \sum_{\alpha \in \mu_{M_0}} h(\alpha) \alpha^j \sum_{\kappa \in \mu_{M_0}} \kappa^j = 0$$

whenever $M_0 \nmid j$, in particular for $1 \leq j < M_0$. Newton identities convert the corresponding power-sum preservation into preservation of $e_1, \dots, e_{M_0-1}$ for equal-size subsets. $\square$

Corollary 4.3 (Polynomial size of characteristic-zero fibers). *Every characteristic-zero prefix fiber in $\binom{\mu_n}{s}$ has size at most*

$$2^{n/M_0} \leq 2^{n/\sigma}.$$

If $\sigma \geq Cn/\log_2 n$, then every such fiber has size at most $n^{1/C}$. With fixed density one may sharpen this to

$$2^{(n/M_0)H_2(\rho) + o(n/M_0)} \leq n^{H_2(\rho)/C + o(1)}.$$

Proof. Fix one set S_0 in the fiber. By [Theorem 4.2](#), every other set differs from S_0 only by toggling whole μ_{M_0} -cosets. There are n/M_0 such cosets. The entropy refinement counts fixed-weight choices among those quotient cosets. $\square$

5 Galois-amplified finite-field collision sieve

We now compare characteristic-zero prefix equality with equality after reduction modulo a split prime. Let $p \equiv 1 \pmod{n}$, choose $\omega \in \mathbb{F}_p^\times$ of order n , and reduce $\zeta \mapsto \omega$.

For $S, T \subseteq \mu_n$ with exponent sets $A, B \subseteq \mathbb{Z}/n\mathbb{Z}$, define the signed subset polynomial

$$F_{S,T}(X) = \sum_{a \in A} X^a - \sum_{b \in B} X^b \in \mathbb{Z}[X], \quad \deg F_{S,T} < n.$$

Equality of the first σ power sums in $\mathbb{F}_p$ is exactly

$$F_{S,T}(\omega^j) = 0 \quad (1 \leq j \leq \sigma).$$

If $p > \sigma$, this is equivalent to equality of $e_1, \dots, e_\sigma$ in $\mathbb{F}_p$.

Lemma 5.1 (Conjugate-prime amplification). *Let $F \in \mathbb{Z}[X]$ have degree $< n$, and assume $F(\zeta) \neq 0$.*

If

$$F(\omega^j) = 0$$

for every odd $j \leq \sigma$, then

$$p^{[\sigma/2]} \mid |\mathrm{N}_{\mathbb{Q}(\zeta)/\mathbb{Q}}(F(\zeta))|.$$

Proof. Let $\mathfrak{p}_\omega \subset \mathbb{Z}[\zeta]$ be the degree-one prime ideal corresponding to $\zeta \mapsto \omega$. For odd j , the automorphism $\tau_j : \zeta \mapsto \zeta^j$ is in $\mathrm{Gal}(\mathbb{Q}(\zeta)/\mathbb{Q})$. The congruence $F(\omega^j) = 0$ says $\tau_j(F(\zeta)) \in \mathfrak{p}_\omega$, equivalently

$$F(\zeta) \in \tau_j^{-1}\mathfrak{p}_\omega.$$

For distinct odd j modulo n , these are distinct prime ideals above p , each of norm p . There are $[\sigma/2]$ odd integers in $[1, \sigma]$. Their product divides the principal ideal $(F(\zeta))$, yielding the norm divisibility. $\square$

The quotient factors must be treated uniformly. For $M = 2^b \mid n$ define

$$Q_M(X) = 1 + X^{n/M} + X^{2n/M} + \dots + X^{(M-1)n/M} = \frac{X^n - 1}{X^{n/M} - 1}.$$

Divisibility by Q_M means that the coefficients of F are constant on residue classes modulo n/M , i.e. on cosets of the subgroup of order M .

Theorem 5.2 (Galois-amplified no-collision bound). *There is an absolute constant C_1 with the following property. Let $n = 2^m$, $1 \leq \sigma < n$, $p \equiv 1 \pmod{n}$, and $p > \sigma$. Suppose*

$$p > \exp\left(C_1 \frac{n \log(2n)}{\sigma}\right).$$

Then every equality

$$(e_1(S), \dots, e_\sigma(S)) = (e_1(T), \dots, e_\sigma(T)) \quad \text{in } \mathbb{F}_p^\sigma$$

for equal-size subsets $S, T \subseteq H = \langle \omega \rangle$ is quotient-periodic: after identifying H with μ_n , the symmetric difference $S \triangle T$ is a union of cosets of a subgroup of order $M > \sigma$.

Proof. Use power sums, which are equivalent because $p > \sigma$. Let $F = F_{S,T}$. If F is divisible by Q_M for some $M > \sigma$, then $Q_M(\zeta^j) = 0$ for $1 \leq j \leq \sigma$, and this is exactly a quotient-periodic characteristic-zero collision, handled by [Theorem 4.2](#).

Assume no such quotient divisor with $M > \sigma$ exists. Let M be the largest power of two such that $F = Q_M G$ with $G \in \mathbb{Z}[X]$; possibly $M = 1$. Then $M \leq \sigma$. Put $\xi = \zeta^M$, a primitive (n/M) -th root of unity. Maximality of M implies $G(\xi) \neq 0$; otherwise $\Phi_{n/M}$ would divide G and hence Q_{2M} would divide F .

For every $\ell \leq \sigma/M$, the finite-field equation at $j = M\ell$ gives

$$0 = F(\omega^{M\ell}) = Q_M(\omega^{M\ell})G((\omega^M)^\ell) = M G((\omega^M)^\ell),$$

and M is invertible modulo the odd prime p . Applying [Lemma 5.1](#) in the smaller field $\mathbb{Q}(\xi)$, with n replaced by n/M and σ replaced by $\lfloor \sigma/M \rfloor$, yields

$$p^{\lceil \sigma/(3M) \rceil} \leq |\mathrm{N}_{\mathbb{Q}(\xi)/\mathbb{Q}}(G(\xi))|.$$

The harmless factor 3 absorbs the floor and the restriction to odd ℓ .

Since the coefficients of F lie in $\{-1, 0, 1\}$ and $F = Q_M G$ with $\deg F < n$, the coefficients of G are also in $\{-1, 0, 1\}$. Hence every complex conjugate of $G(\xi)$ has absolute value at most n/M , and therefore

$$|\mathrm{N}_{\mathbb{Q}(\xi)/\mathbb{Q}}(G(\xi))| \leq (2n)^{\varphi(n/M)} \leq (2n)^{n/(2M)}.$$

Thus

$$p^{\sigma/(4M)} \leq (2n)^{n/(2M)}$$

for all large enough parameters, which implies

$$\log p \leq 2 \frac{n \log(2n)}{\sigma}$$

up to an absolute constant. Choosing C_1 larger gives a contradiction. Therefore the collision must have had a quotient divisor with $M > \sigma$. $\square$

Corollary 5.3 (Monomial-prefix upper bound in quasi-polynomial split prime fields). *Let $n = 2^m$, $p \equiv 1 \pmod{n}$, $H \leq \mathbb{F}_p^\times$ have order n , and $s = k + \sigma$. If*

$$p > \exp\left(C_1 \frac{n \log(2n)}{\sigma}\right),$$

then every monomial-prefix word has list size at most

$$2^{n/M_0} \leq 2^{n/\sigma},$$

where M_0 is the least power of two greater than σ . In particular, if $\sigma \geq C_0 n / \log_2 n$, then every monomial-prefix list has size at most $n^{1/C_0 + o(1)}$ provided

$$p > \exp(O_{C_0}((\log n)^2)).$$

Proof. By [Theorem 5.2](#), a finite-field prefix fiber coincides with a characteristic-zero quotient fiber. Apply [Corollary 4.3](#). At $\sigma \geq C_0 n / \log_2 n$, the threshold on p becomes $\exp(O_{C_0}((\log n)^2))$. $\square$

Remark 5.4 (What the Galois sieve does and does not prove). A one-prime norm sieve — the mechanism behind the stable range and the density windows of Part II — only gives useful information when $p \gg \binom{n}{s}^2 \text{poly}(n)$. [Theorem 5.2](#) is stronger by a factor $\asymp \sigma$ in the exponent because the equations at odd powers give distinct conjugate prime divisors. Nevertheless it still misses polynomial fields at the critical reserve $\sigma \asymp n/\log n$: then it needs $p > \exp(O((\log n)^2))$, while polynomial fields have $p = \exp(O(\log n))$. Closing this last factor of $\log n$ requires a genuinely finite-field local-limit theorem, not a norm-divisor argument.

6 The corrected positive conjectures

The lower bounds and the quasi-polynomial-field theorem point to the following two-scale statement. (For *random* evaluation sets, capacity-achieving list decodability is known [\[22, 23\]](#); the content here is the structured smooth-domain case, where [Corollary 3.6](#) shows the random prediction is false as stated.)

Conjecture 6.1 (Prefix local limit with quotient exceptions). *Fix $0 < \rho < 1$, $B > 0$, and $\varepsilon > 0$. There is $C = C(\rho, B, \varepsilon)$ such that the following holds for all large smooth n . Let $H \leq \mathbb{F}_q^\times$ have order n and generate the ambient field over its prime field, let $k = \rho n$, and let $s = k + \sigma$. If*

$$\sigma \geq C \frac{n}{\log_2 n} \quad \text{and} \quad \sigma \log_2 q \geq (1 + \varepsilon) \log_2 \binom{n}{s},$$

then for every $c \in \mathbb{F}_q^\sigma$,

$$\#\{S \in \binom{H}{s} : \Phi_\sigma(S) = c\} \leq n^B.$$

More sharply, before imposing the second inequality one expects the asymptotic

$$\#\Phi_\sigma^{-1}(c) = \frac{1}{q^\sigma} \binom{n}{s} + \text{Quot}_{\sigma,c}(H, s) + O(n^B),$$

where $\text{Quot}_{\sigma,c}$ is an explicit finite sum of quotient-core contributions from subgroups $K \leq H$ with $|K| > \sigma$.

Conjecture 6.2 (Arbitrary locator-fiber local limit). *Under the same hypotheses as [Conjecture 6.1](#), for every interpolant $U \in \mathbb{F}_q[X]$ with $\deg U < n$,*

$$|\mathcal{F}_U(k + \sigma)| \leq n^B.$$

The predicted pre-threshold asymptotic is the same random-codimension- σ law:

$$|\mathcal{F}_U(k + \sigma)| \approx \frac{1}{q^\sigma} \binom{n}{k + \sigma}$$

except for quotient-periodic and low-denominator structured families.

Theorem 6.3 (Conditional corrected list theorem). *If Conjecture 6.2 holds, then for generated-field smooth domains*

$$\text{List}(\text{RS}[\mathbb{F}_q, H, k], 1 - \rho - \sigma/n) \leq n^B$$

whenever

$$\sigma \geq C \frac{n}{\log n}, \quad \sigma \log_2 q \geq (1 + \varepsilon) \log_2 \binom{n}{k + \sigma}.$$

If only Conjecture 6.1 holds, the same conclusion holds for monomial-prefix received words.

Proof. The arbitrary-word claim is exactly Proposition 2.2. The monomial-prefix claim is Proposition 2.3. $\square$

Remark 6.4 (Subfield correction). If H is contained in a proper subfield $\mathbb{F}_{q_0} \subsetneq \mathbb{F}_q$, then the entropy scale is governed by the field of definition q_0 , not the ambient alphabet size q . Indeed, locator coefficients of subsets of H lie in $\mathbb{F}_{q_0}$, and the pigeonhole denominator in Theorem 3.2 is q_0^σ . Thus any ambient- q formulation over extension fields is false unless a generated-field hypothesis or a field-of-definition parameter is included.

7 Why pairwise/BCH information is insufficient

One tempting approach is to use the fact that equal prefix fibers have large Hamming distance as constant-weight codes. This gives a useful shadow but not a proof.

Lemma 7.1 (Pairwise intersection bound). *Let $S, T \in \binom{H}{k+\sigma}$ satisfy $\Phi_\sigma(S) = \Phi_\sigma(T)$. If $S \neq T$, then*

$$|S \cap T| \leq k - 1.$$

Equivalently, their constant-weight incidence vectors have Hamming distance at least $2(\sigma + 1)$.

Proof. The locators L_S and L_T have the same top σ coefficients, so

$$\deg(L_S - L_T) \leq (k + \sigma) - \sigma - 1 = k - 1.$$

But $L_S - L_T$ vanishes on $S \cap T$. If $|S \cap T| \geq k$, then $L_S = L_T$, hence $S = T$. $\square$

At $\sigma \asymp n/\log n$, constant-weight codes of length n , weight $\rho n + \sigma$, and distance 2σ can still be exponentially large by standard Gilbert–Varshamov type packings. Therefore any proof of Conjecture 6.1 must exploit the algebraic subgroup moment equations, not merely the pairwise distance pattern. This is the precise sense in which BCH minimum-distance information is too weak.

Part II

Slack mutual correlated agreement: the corrected theory

8 The list-side ledger, and the MCA comparison point

Part I established the corrected list-side ledger: the coefficient-pigeonhole and quotient-core lower bounds ([Theorems 3.2](#) and [3.5](#)), the prime-field refutation of the one-scale ambient- q entropy-gap prediction ([Corollary 3.6](#)), the quotient-core profile and the corrected list conjecture ([Definition 3.8](#) and [Conjecture 3.9](#)), the characteristic-zero inverse quotient theorem ([Theorem 4.2](#)), and the Galois-amplified quasi-polynomial split-prime positive theorem ([Corollary 5.3](#)). That ledger is the right comparison point for the MCA conjecture of [Section 18](#). Recall that $\text{List}(C, \delta)$ denotes the worst-case list size at relative radius δ ; in this part p is prime unless stated otherwise.

Remark 8.1 (MCA versus list decoding). The quotient-core construction refutes the one-scale list prediction directly. It does not by itself refute any MCA statement: it inflates the number of codewords explaining one received word, never the number of bad slopes of a line. Nevertheless the same quotient scales enter the conservative MCA formulation of [Conjecture 18.8](#) through the profile $\mathcal{Q}$. [Theorem 19.4](#) below shows that, above the quotient norm threshold and up to the explicit constant $H_2(\rho)/\beta(\rho)$, this quotient profile is genuinely necessary for polynomial MCA slope bounds. The remaining open direction is sufficiency, especially for arbitrary residue-line data at a fixed prime.

9 Definitions, and the two delimitations

Definition 9.1 (Support-wise agreement). For $C \subseteq \mathbb{F}^D$, $|D| = n$, $\delta \in [0, 1]$: a word $u \in \mathbb{F}^D$ is S -close to C at radius δ if $S \subseteq D$, $|S| \geq (1 - \delta)n$, and $u|_S = c|_S$ for some $c \in C$.

Definition 9.2 (Bad parameter and line-MCA error). For a line $u_z = f + zg$, $z \in \mathbb{F}$: the parameter z is *support-wise MCA-bad* at radius δ if there is $S \subseteq D$ with (i) u_z S -close to C at radius δ , and (ii) no pair $c_f, c_g \in C$ satisfying $f|_S = c_f|_S$ and $g|_S = c_g|_S$. Then

$$\varepsilon_{\text{mca}}(C, \delta) = \max_{f, g \in \mathbb{F}^D} \frac{1}{|\mathbb{F}|} \#\{z : z \text{ support-wise MCA-bad for } f + zg\},$$

and the threshold is

$$\delta_C^*(\varepsilon^*) = \sup\{\delta : \varepsilon_{\text{mca}}(C, \delta) \leq \varepsilon^*\}.$$

A witness at radius δ is one at every larger radius (neither condition mentions δ otherwise), so $\varepsilon_{\text{mca}}(C, \cdot)$ is nondecreasing and error one propagates upward; this is the monotonicity lemma of [\[1\]](#).

For a degree- c tuple $\mathbf{u} = (u_0, \dots, u_c)$ and $u(z) = \sum_j z^j u_j$, the parameter z is *curve-bad* if some S has $u(z)|_S \in C|_S$ while not every u_j has $u_j|_S \in C|_S$; $\varepsilon_{\text{mca}}^{(c)}$ is the corresponding error.

The no-slack form of the up-to-capacity principle — $\varepsilon_{\text{mca}}(C, \delta) \leq \text{negl}(q)$ for every $\delta < 1 - \rho$ — fails on smooth domains: error one holds on the interval $[1 - \rho - 1/N, 1]$ whenever some divisor $N \mid n$ has $\rho N \in \mathbb{Z}$ and $\rho(1 - \rho)N^2 \geq p$ (so for the deployed FFT primes at $N = 2^{18}$), the error is $1 - 1/p$ at radius $1 - \rho - 1/(2 \log_2 n)$ on the Fermat fields $q = n + 1$ with $p \in \{17, 257, 65537\}$ at

rates $1/2$ and $1/4$, and infinitely many primes carry power-of-two subgroups of order $\Theta(\log p)$ with error at least $(\log_2 p)^{-6}$ at gap $1/N$ [1]. Concurrent near-capacity failure results over prime-field multiplicative subgroups [28, 29] apply related pressure through different mechanisms. We will recover and considerably extend the first of these from the slack calculus below. Two elementary results delimit the field-size dependence and force the shape of any true formulation.

Theorem 9.3 (One bad parameter per support). *For any linear code, any line, and every $S \subseteq D$: at most one z is bad with witness S , at any radius. Hence $\varepsilon_{\text{mca}}(C, \delta) \leq \#\{S : |S| \geq (1-\delta)n, C|_S \neq \mathbb{F}^S\}/q \leq 2^n/q$, and if $q \geq 2^n/\varepsilon^*$ then $\delta_C^*(\varepsilon^*) = 1$ exactly: over very large fields the no-slack principle holds trivially, for every line.*

Proof. Condition (ii) of Definition 9.2 does not involve z ; if it fails for S , no z has witness S . If it holds with $g|_S \in C|_S$ (so $f|_S \notin C|_S$): $u_z|_S \in C|_S$ would force $f|_S \in C|_S$. If $g|_S \notin C|_S$ and $z_1 \neq z_2$ both had $u_{z_i}|_S \in C|_S$: subtracting, $(z_1 - z_2)g|_S \in C|_S$, impossible. Supports with $C|_S = \mathbb{F}^S$ (for Reed–Solomon: all $|S| \leq k$, by interpolation) admit no witness since (ii) fails there. $\square$

Proposition 9.4 (Tangent floor). *For $C = \text{RS}[\mathbb{F}, D, k]$ and $\delta \in (0, 1 - \rho)$ with $m = \lfloor \delta n \rfloor \leq q$: $\varepsilon_{\text{mca}}(C, \delta) \geq m/q$.*

Proof. Fix $E \subseteq D$, $|E| = m$, codewords P_1, P_2 , and $f = P_1 + e_f$, $g = P_2 + e_g$ with errors supported on E , $e_g \equiv 1$ there and e_f injective there. For $x \in E$ put $z_x = -e_f(x)$ and $S_x = (D \setminus E) \cup \{x\}$: the error of u_{z_x} against the codeword $P_1 + z_x P_2$ vanishes at x and off E , so condition (i) holds; a codeword explaining g on S_x would agree with P_2 on $n - m > k$ points, hence equal P_2 , contradicting $g(x) \neq P_2(x)$: condition (ii). The m slopes z_x are distinct and bad. $\square$

Remark 9.5 (The phase transition, and the forced shape of the theory). For protocol-size fields $q = O(n)$ the floor alone gives $\varepsilon_{\text{mca}} \geq \Theta(\delta)$ at every radius, so no fixed- ε^* threshold statement can hold there, and even for $q = \text{poly}(n)$ the floor forces an additive n/q -scale correction into any positive conjecture; for $q \geq 2^n/\varepsilon^*$ the question trivializes. The live regime is $\text{poly}(n) \leq q \leq 2^{o(n)}$, the correct positive target is an error of the form $n^{1+o(1)}/q$, and — because of the failure results above — the radius must carry a slack reserve. This fixes the object of study: the bad-slope sets of lines at radius $1 - \rho - t/N$.

10 The slack calculus

10.1 The exact characterization

Definition 10.1 (The multi-symmetric image). For $1 \leq T \leq n - k$, $\delta_T = 1 - (k + T)/n$ and

$$\mathcal{B}_T(D, k) = \{(-1)^T e_T(S) : S \subseteq D, |S| = k + T, e_1(S) = \cdots = e_{T-1}(S) = 0\}.$$

For $T = 1$ this is $-(k + 1)^\wedge D$.

Theorem 10.2 (Exact slack characterization). *Let D be any evaluation set of size n in a field $\mathbb{F}$, $C = \text{RS}[\mathbb{F}, D, k]$, $1 \leq T \leq n - k$. For the line $u_z = x^{k+T} + z x^k$, the set of support-wise MCA-bad parameters at radius δ_T is exactly $\mathcal{B}_T(D, k)$: the line's error contribution is $|\mathcal{B}_T(D, k)|/q$, with error one iff $\mathcal{B}_T(D, k) = \mathbb{F}$.*

Proof. ($\subseteq$) If z is bad, u_z agrees with some R , $\deg R < k$, on a set S with $|S| \geq k + T$. The polynomial $P_z = X^{k+T} + zX^k - R$ is monic of degree $k + T$ with at least $k + T$ roots: exactly $k + T$, so $|S| = k + T$ and $P_z = L_S$. Comparing coefficients in degrees $k + T - j$ for $1 \leq j \leq T - 1$ (absent from both u_z and R) gives $e_j(S) = 0$; the degree- k coefficient gives $z = (-1)^T e_T(S)$.

($\supseteq$) If $z = (-1)^T e_T(S)$ with S as in [Definition 10.1](#), then $L_S = X^{k+T} + zX^k + R_S$ with $\deg R_S < k$, so u_z agrees with the codeword $-R_S$ on the $k + T$ points of S . Condition (ii): $g = x^k$ has degree exactly k , and two distinct polynomials of degree $\leq k$ agree on at most $k < k + T$ points, so no codeword explains g on S — (ii) holds for every witness set of more than k points, regardless of f . $\square$

The restricted-sumset description of the no-slack failures is therefore not an artifact of any proof method: it is the precise answer, and at every slack T the object to control is $\mathcal{B}_T$.

10.2 Quotients, and the collapse of symmetric witnesses

Lemma 10.3 (Quotient locator). *Let $D \leq \mathbb{F}^\times$ of order n , $N \mid n$, $a = n/N$, $Q = D^a$ the quotient subgroup of order N , $a \mid k$, $\ell = k/a + t \leq N$. For $A \subseteq Q$, $|A| = \ell$:*

$$L_A(X) := \prod_{b \in A} (X^a - b) = X^{k+ta} + \sum_{j=1}^t (-1)^j e_j(A) X^{k+(t-j)a} + R_A(X), \quad \deg R_A < k,$$

and L_A vanishes on $S_A = \{x \in D : x^a \in A\}$, of size $k + ta$. (Expand $\prod_b (Y - b)$ at $Y = X^a$: the displayed terms sit in degrees $a(\ell - j)$, and all remaining degrees are below $a(\ell - t) = k$.)

Theorem 10.4 (Slack- t lower bound). *With $V_t(Q, \ell) = \{z : \exists A \subseteq Q, |A| = \ell, e_1(A) = \dots = e_{t-1}(A) = 0, (-1)^t e_t(A) = z\}$: every $z \in V_t(Q, \ell)$ is bad for the line $f = x^{k+ta}$, $g = x^k$ at radius $1 - \rho - t/N$, so $\varepsilon_{\text{mca}}(\text{RS}[\mathbb{F}, D, k], 1 - \rho - t/N) \geq |V_t(Q, \ell)|/q$, with error one when $V_t = \mathbb{F}$. If $\text{char } \mathbb{F} > t$ the defining conditions are equivalent, by the Newton identities, to $p_1(A) = \dots = p_{t-1}(A) = 0$ and $p_t(A) = -tz$.*

Proof. [Lemma 10.3](#) gives $L_A = X^{k+ta} + zX^k + R_A$, so u_z agrees with $-R_A$ on S_A ; condition (ii) as in [Theorem 10.2](#). The fiber-union structure makes $e_j(S_A) = 0$ automatic for $a \nmid j$, so S_A is a $\mathcal{B}_{ta}$ -witness with exactly the displayed quotient conditions surviving. $\square$

Remark 10.5 (Strict subfamily). $\mathcal{B}_{ta}(D, k) \supseteq V_t(Q, \ell)$, and the containment can be strict: at $p = 17$, $n = 16$, $\rho = 1/2$, $T = 2$, the quotient family ($a = 2$, $N = 8$, $\ell = 5$) covers 16 of the 17 values and an asymmetric D -witness supplies the missing $z = 0$ (verified exhaustively; [appendix A](#)).

Proposition 10.6 (Coset collapse: symmetric slack is no slack). *Let $p > d$, $K \leq Q$ of order d , $d \mid \ell$, and let A be a union of ℓ/d distinct K -cosets. Then $p_j(A) = 0$ for all $d \nmid j$ — in particular $e_1(A) = \dots = e_{d-1}(A) = 0$ automatically — and $e_d(A) = (-1)^{d-1} \sum_i y_i$ with $y_1, \dots, y_{\ell/d}$ the distinct d -th powers of coset representatives. Hence the union-of-coset part of $V_d(Q, \ell)$ equals, up to an explicit nonzero scalar, the slack-one value set $m^\wedge(Q^d)$ at quotient order N/d , $m = \rho(N/d) + 1$: subgroup-symmetric witnesses produce nothing beyond the one-condition ladder.*

Proof. $\sum_{\kappa \in K} (x\kappa)^j = x^j \sum_{\kappa} \kappa^j$, which is dx^j if $d \mid j$ and 0 otherwise; sum over cosets, convert by Newton ($p > d$); distinct cosets biject with distinct d -th powers; and $\ell/d = \rho(N/d) + 1$ at $\ell = \rho N + d$. $\square$

New slack content therefore requires *asymmetric* witnesses with vanishing leading symmetric functions. The next two theorems prove that the two global models which decide no-slack questions — lifting to characteristic zero, and lifting to short digit ranges — cannot see them, even at $t = 2$.

Theorem 10.7 (Cyclotomic rigidity). *Let $N = 2^s$ and $A \subseteq \mu_N \subseteq \mathbb{Z}[\zeta_N]$ with $p_1(A) = 0$ in $\mathbb{Z}[\zeta_N]$. Then A is a union of antipodal pairs $\{\omega, -\omega\}$. Consequently every characteristic-zero witness for V_2 is of coset type, and the characteristic-zero slack-2 value set is exactly the ladder rung $2 \cdot m^\wedge(\mu_{N/2})$.*

Proof. In the power basis $1, \zeta, \dots, \zeta^{N/2-1}$, $\sum_{a \in A} a = \sum_j c_j \zeta^j$ with $c_j = \mathbf{1}[\zeta^j \in A] - \mathbf{1}[-\zeta^j \in A] \in \{-1, 0, 1\}$; vanishing forces $c \equiv 0$, i.e. each antipodal pair is contained in or disjoint from A . Then $p_2(A) = 2 \sum_{\text{pairs}} \zeta^{2j}$. $\square$

Theorem 10.8 (Digit rigidity on Fermat fields). *Let $p = 2^M + 1$ be a Fermat prime, $Q = \langle 2 \rangle = \{\pm 2^j : 0 \leq j < M\} \subseteq \mathbb{F}_p^\times$, and $A \subseteq Q$ with $\sum_{a \in A} a \equiv 0 \pmod{p}$. Then A is a union of antipodal pairs.*

Proof. The sum lifts to $s = \sum_j c_j 2^j$, $c_j \in \{-1, 0, 1\}$, with $|s| \leq 2^M - 1 < p$, so $s = 0$; uniqueness of binary representations forces $c \equiv 0$. $\square$

Remark 10.9 (Where the difficulty lives). Asymmetric elements of V_t require cancellations of size comparable to p — genuinely modular, i.e. subgroup equidistribution. [Section 12](#) pushes equidistribution as far as current exponential-sum technology reaches; [Section 14](#) develops the characteristic-zero theory at every slack, which governs the deepest level; [Section 18](#) proves where these methods stop.

10.3 Certified slack thresholds, and curve transfer

Lemma 10.10 (Dias da Silva–Hamidoune [8]; polynomial-method proofs in [9, 10]). *For every $A \subseteq \mathbb{F}_p$ with $|A| = m$ and every $1 \leq h \leq m$: $|h^\wedge A| \geq \min\{p, h(m-h) + 1\}$.*

Corollary 10.11 (Certified slack at the root barrier). *Let $0 < \rho \leq 1/2$ and let $N_{\text{full}}(p, n, \rho)$ be the smallest divisor $N \mid n$ with $\rho N \in \mathbb{Z}$, $\ell = \rho N + 1 \leq N$, and $\ell(N - \ell) + 1 \geq p$. For every divisor $N \geq N_{\text{full}}$ with $\rho N \in \mathbb{Z}$:*

$$\varepsilon_{\text{mca}}\left(\text{RS}[\mathbb{F}_p, D, \rho n], 1 - \rho - \frac{1}{N}\right) = 1,$$

hence error one on the interval $[1 - \rho - 1/N_{\text{full}}, 1]$, with $N_{\text{full}} \leq (1 + o(1))\sqrt{p/\rho(1-\rho)}$ when such divisors exist (automatic on smooth n of that size). ([Lemma 10.10](#) makes $V_1(Q, \ell) = -\ell^\wedge Q = \mathbb{F}_p$; [Theorem 10.4](#) at $t = 1$; monotonicity for the interval.)

Theorem 10.12 (Curve transfer). *Let $\varphi(Z) = \sum_{j=1}^c \varphi_j Z^j$ be nonconstant with $\varphi(0) = 0$, and consider the degree- c tuple $u_0 = x^{k+a}$, $u_j = \varphi_j x^k$. Every z with $\varphi(z) \in -\ell^\wedge Q$ is curve-bad at radius $1 - \rho - 1/N$, so $\varepsilon_{\text{mca}}^{(c)} \geq 1 - c(1 - |\ell^\wedge Q|/q)$ there: whenever $\ell^\wedge Q = \mathbb{F}$, every z is curve-bad and $\varepsilon_{\text{mca}}^{(c)} = 1$ on the whole interval, for every degree $c \geq 1$. Every slack floor proved for lines transfers to every curve class containing polynomial reparametrizations of lines.*

Proof. $u(z) = x^{k+a} + \varphi(z)x^k$; choose A with $\sum A = -\varphi(z)$ and apply Lemma 10.3 at $t = 1$. For condition (ii) one unexplained member suffices: any $u_j = \varphi_j x^k$ with $\varphi_j \neq 0$ is explained on at most $k < |S_A|$ points. Count preimages of φ ; monotonicity is verbatim for tuples. $\square$

11 The two-moment certificate, and fullness at the root barrier

The Dias da Silva–Hamidoune theorem is the $t = 1$ engine of the restricted-sum ladder, and its exterior-algebra proof invites an upgrade to the joint map $S \mapsto (e_1(S), e_2(S))$ on a multiplicative subgroup. The exact finite object keeps the subgroup relations and becomes a cyclic reduced alternant; and an elementary two-stage construction already gives *full* two-moment coverage at the square-root barrier, at every prime.

Let $Q \leq \mathbb{F}_p^\times$ have even order N . For $S \in \binom{Q}{\ell}$ set $\Psi_2(S) = (e_1(S), e_2(S))$; for $a \in \mathbb{F}_p$ let $\Pi_a(T) = 1 - (T - a)^{p-1}$ (the indicator of $T = a$); and put

$$\Delta(X) = \prod_{i < j} (X_j - X_i), \quad E_1 = \sum_i X_i, \quad E_2 = \sum_{i < j} X_i X_j, \quad F_{a,b} = \Delta \cdot \Pi_a(E_1) \Pi_b(E_2).$$

Theorem 11.1 (Exact reduced $t = 2$ certificate). *For every $(a, b) \in \mathbb{F}_p^2$:*

$$(a, b) \in \Psi_2\left(\binom{Q}{\ell}\right) \iff F_{a,b}(X) \not\equiv 0 \pmod{X_1^N - 1, \dots, X_\ell^N - 1}.$$

Proof. Since $p \nmid N$, $X^N - 1$ has the distinct root set Q , so $\mathbb{F}_p[X]/(X_i^N - 1)$ is the function algebra on Q^ℓ . On a tuple in Q^ℓ , Δ is nonzero exactly when the entries are distinct; on distinct tuples, $\Pi_a(E_1)\Pi_b(E_2)$ is nonzero exactly when the underlying set has $(e_1, e_2) = (a, b)$. $\square$

Equivalently, with $A_N = \mathbb{F}_p[u]/(u^N - 1)$, alternants identified with $\bigwedge^\ell A_N$, E_1, E_2 acting by multiplication by e_1, e_2 , and $\omega_\ell = 1 \wedge u \wedge \dots \wedge u^{\ell-1}$: the condition is $\Pi_a(E_1)\Pi_b(E_2)\omega_\ell \neq 0$ — a reduced cyclic exterior-algebra nonvanishing problem. Resolving it below the root barrier is open; at the barrier, no machinery is needed:

Theorem 11.2 (Two-moment fullness in the square-root range). *Let $Q \leq \mathbb{F}_p^\times$ have even order N and write $\ell = h + 2r$. If*

$$h(N - h) + 1 \geq p \quad \text{and} \quad r\left(\frac{N}{2} - h - r\right) + 1 \geq p,$$

then $\Psi_2\left(\binom{Q}{\ell}\right) = \mathbb{F}_p^2$. Consequently, for $\ell = \rho N + O(1)$, full two-moment coverage holds whenever

$$p \leq \left(\frac{\rho(1-\rho)}{5} - o(1)\right)N^2,$$

i.e. for all $N \geq (1 + o(1))\sqrt{5p/\rho(1-\rho)}$ — at every prime, with no thresholds.

Proof. Fix (a, b) . By [Lemma 10.10](#) on Q , the first hypothesis gives an h -set $B \subseteq Q$ with $e_1(B) = a$. Since N is even, $-1 \in Q$ and the squaring map on Q is two-to-one onto the index-two subgroup Q^2 of order $N/2$; let $A_B = Q^2 \setminus B^2$, of size at least $N/2 - h$. By [Lemma 10.10](#) on A_B , the second hypothesis gives an r -set $R \subseteq A_B$ with $\sum_{y \in R} y = e_2(B) - b$. For each $y \in R$ choose $x_y \in Q$ with $x_y^2 = y$; since $y \notin B^2$, neither x_y nor $-x_y$ lies in B , and pairs over distinct y are disjoint (their elements square to distinct values). Put $S = B \sqcup \bigsqcup_{y \in R} \{x_y, -x_y\}$, of size $h + 2r = \ell$. Each antipodal pair has $e_1 = 0$ and $e_2 = -x_y^2 = -y$; pairs of pairs and pair-versus- B cross terms cancel (e_2 of a disjoint union is $e_2 + e'_2 + e_1 e'_1$, and all the pair blocks have $e_1 = 0$). Hence

$$e_1(S) = e_1(B) = a, \quad e_2(S) = e_2(B) - \sum_{y \in R} y = b.$$

For the constant, set $h = \alpha N + O(1)$, $r = (\rho - \alpha)N/2 + O(1)$: the two hypotheses read $p \leq \alpha(1 - \alpha)N^2$ and $p \leq (\rho - \alpha)(1 - \rho - \alpha)N^2/4$, and the balance equation $4\alpha(1 - \alpha) = (\rho - \alpha)(1 - \rho - \alpha) = \rho(1 - \rho) - \alpha + \alpha^2$ simplifies to $5\alpha(1 - \alpha) = \rho(1 - \rho)$: the balanced value is exactly $\alpha(1 - \alpha) = \rho(1 - \rho)/5$. $\square$

Corollary 11.3 (Every slack-two monomial line fails at the root barrier, at every prime). *Let $a = n/N$, $a \mid k$, $\ell = \rho N + 2$, with the hypotheses of [Theorem 11.2](#) satisfiable at this ℓ (in particular $N \geq (1 + o(1))\sqrt{5p/\rho(1 - \rho)}$). Then for every base $c \in \mathbb{F}_p$, the line*

$$u_z = x^{k+2a} + c x^{k+a} + z x^k$$

has every slope z support-wise MCA-bad at radius $1 - \rho - 2/N$: error one on the whole slack-two monomial family, at every prime.

Proof. [Theorem 11.2](#) supplies $A \subseteq Q$, $|A| = \ell$, with $(-e_1(A), e_2(A)) = (c, z)$; [Lemma 10.3](#) turns A into an agreement set for u_z , and condition (ii) holds as always via $g = x^k$. (Parity of $\ell - h$ is absorbed in the $O(1)$'s.) $\square$

Remark 11.4 (Ledger). [Theorem 11.2](#) strengthens, at slack two, both ends of the unconditional ladder: [Theorem 13.1](#) covers only the diagonal base $c = 0$ (at constant 2 in place of $\sqrt{5}$), and [Theorem 12.2](#) reaches smaller N only with thresholds in p . It is not the entropy-scale result: the conjectural $t = 2$ target asks for coverage when $\log_2 p \lesssim \frac{1}{2} H_2(\rho)N$, far beyond $p \asymp N^2$; [Theorem 11.1](#) identifies the exact finite nonvanishing problem that going further requires. Verified: at $(p, N, h, r) = (31, 30, 2, 3)$, $\ell = 8$, the construction of [Theorem 11.2](#) built witnesses for all 961 targets $(a, b) \in \mathbb{F}_{31}^2$ with zero failures, and exhaustive enumeration of all $\binom{30}{8}$ subsets confirms Ψ_2 is onto ([appendix A](#)).

12 The failure ladder, I: the free pool

One lemma drives all the equidistribution-based failure results. For a subgroup $K \leq \mathbb{F}_p^\times$ of order d and $\vec{s} \in \mathbb{F}_p^t$ write $S_K(\vec{s}) = \sum_{\kappa \in K} e_p\left(\sum_{m \leq t} s_m \kappa^m\right)$.

Lemma 12.1 (Free pool). *Let $t \geq 1$, and suppose $K \leq Q \leq \mathbb{F}_p^\times$, $|K| = d$, $|Q| = N$, satisfy the equidistribution hypothesis*

$$\max_{\vec{s} \neq 0} |S_K(\vec{s})| \leq d p^{-\nu} \quad \text{for some } \nu > 0.$$

Put $C = \lceil (t - \log_p d)/\nu \rceil + 6$ and suppose $d \geq 4C^2$ and $p^{2\nu} \geq 2$. Then for every ℓ with $C \leq \ell \leq N - d$ and every $\vec{v} \in \mathbb{F}_p^t$ there is $A \subseteq Q$ with $|A| = \ell$ and $p_m(A) = v_m$ for all $m \leq t$: all power-sum fibers are nonempty. In particular $V_t(Q, \rho N + t) = \mathbb{F}_p$ and, by [Theorem 10.4](#), the MCA error at radius $1 - \rho - t/N$ is one.

Proof. Fix $\vec{v}$. Choose any $B_0 \subseteq Q \setminus K$ with $|B_0| = \ell - C$ (possible since $\ell - C \leq N - d$) and set $\vec{w} = \vec{v} - (p_1(B_0), \dots, p_t(B_0))$. For $j \geq 2$ let $T_j(\vec{w}) = \#\{(x_1, \dots, x_j) \in K^j : \sum_i x_i^m = w_m \ \forall m \leq t\}$. Orthogonality gives $T_j(\vec{w}) = p^{-t} \sum_{\vec{s}} S_K(\vec{s})^j e_p(-\langle \vec{s}, \vec{w} \rangle)$, and the t -dimensional Parseval identity reads $\sum_{\vec{s}} |S_K(\vec{s})|^2 = p^t \#\{(\kappa, \kappa') : \kappa^m = \kappa'^m \ \forall m \leq t\} = p^t d$, since the $m = 1$ equation already forces $\kappa = \kappa'$. Hence

$$\left| T_j(\vec{w}) - \frac{d^j}{p^t} \right| \leq \left(\max_{\vec{s} \neq 0} |S_K(\vec{s})| \right)^{j-2} \cdot p^{-t} \sum_{\vec{s} \neq 0} |S_K(\vec{s})|^2 \leq (dp^{-\nu})^{j-2} d.$$

For $j \in \{C - 2, C\}$ the main term dominates by a factor 2: $d^j/p^t \geq 2(dp^{-\nu})^{j-2}d$ is equivalent to $dp^{\nu(j-2)} \geq 2p^t$, and $j - 2 \geq C - 4 \geq (t - \log_p d)/\nu + 2$ gives $dp^{\nu(j-2)} \geq p^{t+2\nu} \geq 2p^t$. So $T_j(\vec{w}) \in [d^j/(2p^t), 2d^j/p^t]$ for these j . Tuples counted by $T_C(\vec{w})$ with two equal coordinates number at most $\binom{C}{2} \sum_{y \in K} T_{C-2}(\vec{w} - (2y, 2y^2, \dots, 2y^t)) \leq C^2 d^{C-1}/p^t$, which is at most $\frac{1}{2} \cdot d^C/(2p^t)$ once $d \geq 4C^2$. Some tuple therefore has C pairwise distinct coordinates in K with the prescribed power sums; adjoin B_0 . $\square$

Three exponential-sum inputs instantiate the hypothesis.

Theorem 12.2 (Conditional polynomial range from a free-pool input). *Let $D \leq \mathbb{F}_p^\times$ be smooth of order n , $C = \text{RS}[\mathbb{F}_p, D, \rho n]$, and fix a divisor scale $N \mid n$ with quotient subgroup Q of order N . Assume that Q contains a subgroup K of order d for which the exponential-sum hypothesis of [Lemma 12.1](#) holds for the relevant degree t , with a saving $\nu > 0$, and that $\ell = \rho N + t$ satisfies*

$$C(t, d, p, \nu) \leq \ell \leq N - d,$$

where $C(t, d, p, \nu) = \lceil (t - \log_p d)/\nu \rceil + 6$ is the pool size from [Lemma 12.1](#). Then

$$\varepsilon_{\text{mca}}(C, 1 - \rho - t/N) = 1.$$

In particular, the following are valid wherever the cited exponential-sum estimates supply the displayed hypothesis with the quoted savings.

(i) For every fixed $\delta > 0$ and fixed $t \geq 1$, the Bourgain–Glibichuk–Konyagin and Bourgain estimates [\[11, 12\]](#), in the required polynomial-phase form, imply error one for all sufficiently large p and all $N \geq p^\delta$.

(ii) For $t = 1$, Shkredov’s medium-range subgroup estimate [\[14\]](#), in the above form, gives explicit pool sizes in the range $N \geq 4p^{1/3+\varepsilon}$; sample pool sizes retained from the computation are $C(0.50) = 18$, $C(0.45) = 25$, $C(0.40) = 42$, $C(0.36) = 102$.

(iii) For fixed $t \geq 2$, the small-subgroup Weil bounds of Ostafe–Shparlinski–Voloch [13], in the required polynomial-phase form, give explicit pool sizes in the range $N \geq p^{3/7+\varepsilon}$; sample sizes retained from the computation are $C(2, 0.50) = 168$, $C(2, 0.55) = 99$, $C(3, 0.55) = 1563$, $C(3, 0.60) = 726$. The growing-slack statement is conditional on the same recursion remaining uniform up to

$$t \leq (1 + o(1)) \frac{\log \log p}{\log \log \log p}.$$

Proof. The theorem is exactly Lemma 12.1 followed by Theorem 10.4. The enumerated ranges are not independent algebraic arguments; they require the cited exponential-sum bounds in the precise uniform form stated in Lemma 12.1. This is why the result is recorded as conditional on that input rather than as a self-contained fixed-prime equidistribution theorem. $\square$

13 The failure ladder, II: descent, and the dichotomy

The pool method needs thresholds in p . At *dyadic* slack, an exact identity removes them entirely.

Theorem 13.1 (Dyadic descent: error one at every dyadic slack, DSH grade). *Let p be prime, $D \leq \mathbb{F}_p^\times$ have order n , let $N \mid n$, put $a = n/N$, and let $Q = D^a$ have order N . Assume $a \mid k$, $k/a = \rho N \in \mathbb{Z}$. Let t be a power of two with $t \mid N$, and assume*

$$\ell = \rho N + t \leq N, \quad t \mid \ell \quad \text{equivalently } t \mid \rho N.$$

For $B \subseteq Q^t$ with $|B| = \ell/t$, the orbit union

$$A = \{x \in Q : x^t \in B\}$$

satisfies

$$e_j(A) = 0 \quad (t \nmid j), \quad e_t(A) = -e_1(B)$$

(for even t ; for odd t the nonzero scalar is absorbed into the slope). Consequently

$$\{e_t(A) : |A| = \ell, e_1(A) = \cdots = e_{t-1}(A) = 0\} \supseteq -\left(\frac{\ell}{t}\right)^\wedge(Q^t).$$

If

$$\frac{\ell}{t} \left(\frac{N}{t} - \frac{\ell}{t} \right) + 1 \geq p,$$

then the monomial slack- t line has

$$\varepsilon_{\text{mca}}\left(\text{RS}[\mathbb{F}_p, D, k], 1 - \rho - \frac{t}{N}\right) = 1.$$

In particular, for fixed ρ and fixed dyadic t , this holds whenever

$$N \geq (t + o(1))\sqrt{p/\rho(1-\rho)},$$

subject to the divisibility assumptions above.

Proof. The t -power map $Q \rightarrow Q^t$ is exactly t -to-one and its fibers are the μ_t -orbits. Therefore

$$\prod_{a \in A} (X - a) = \prod_{y \in B} (X^t - y),$$

a polynomial in X^t . Hence all elementary symmetric coefficients whose indices are not divisible by t vanish. The coefficient of $X^{\ell-t}$ gives $e_t(A) = -e_1(B)$ for even t with the stated harmless sign convention. Dias da Silva–Hamidoune applied to the subgroup Q^t of order N/t gives full coverage of $\mathbb{F}_p$ when the displayed inequality holds. The conclusion follows from [Theorem 10.4](#). $\square$

Remark 13.2 (What descent adds). The radius $1 - \rho - t/N = 1 - \rho - 1/(N/t)$ coincides with the slack-one rung on the quotient Q^t ([Proposition 10.6](#) read in reverse), already reachable by [Corollary 10.11](#); the content is the *fiber-level* statement that every diagonal target $(0, \dots, 0, z)$ of the slack- t system is hit, by explicit witnesses, at DSH grade — the correct baseline against which [Theorem 12.2](#) must be measured (its gain is the non-diagonal targets and witness asymmetry, at the cost of thresholds in p), and the load-bearing input for the positive theorems' floors below. Verified at $p = 193$, $N = 64$, $t = 2$, $\ell = 34$: the identities, and coverage of all 193 residues by e_1 over the $\binom{32}{17}$ subsets of Q^2 ([appendix A](#)).

The characteristic-zero theory of [Section 14](#) will show ([Theorem 14.2](#)) that descent witnesses are the *only* characteristic-zero diagonal witnesses at dyadic slack, and that at non-dyadic slack none exist at all.

14 The rigidity theory at every slack

From here $p \equiv 1 \pmod{n}$, $n = 2^m$, and $\mathfrak{p} \subseteq \mathbb{Z}[\zeta_n]$ is a fixed degree-one prime over p , identifying $H \leq \mathbb{F}_p^\times$ with μ_n . For $h : \mu_n \rightarrow \mathbb{Z}$ write $\widehat{h}(j) = \sum_x h(x)x^j \in \mathbb{Z}[\zeta_n]$. For a slack level σ , $M_0 = 2^a$ is the least power of two exceeding σ and $N_0 = n/M_0$.

The engine is the antipodal identity, [Lemma 4.1](#), proved in [Section 4](#): for n' a power of two and $g : \mu_{n'} \rightarrow \mathbb{Z}$, the relation $\sum_y g(y)y = 0$ forces $g(y) = g(-y)$ for all y . The theorem below is the integer-valued form of [Theorem 4.2](#): the same dyadic induction applies to arbitrary $h : \mu_n \rightarrow \mathbb{Z}$, not only to indicator differences, and the optimality of the modulus is recorded.

Theorem 14.1 (Inverse quotient theorem, integer-valued form). *Let $2\sigma < n$ and $h : \mu_n \rightarrow \mathbb{Z}$ with $\widehat{h}(2^i) = 0$ for $0 \leq i < a$ (implied by $\widehat{h}(j) = 0$ for all $1 \leq j \leq \sigma$, since $2^{a-1} \leq \sigma$). Then h is invariant under μ_{M_0} . Conversely, every μ_{M_0} -invariant h has $\widehat{h}(j) = 0$ for every $M_0 \nmid j$. The modulus is optimal: μ_{M_0} -coset differences across distinct μ_{2M_0} -cosets witness every $\sigma \leq M_0 - 1$.*

Proof. Induction on the dyadic frequencies: $\widehat{h}(1) = \dots = \widehat{h}(2^i) = 0$ forces invariance under $\mu_{2^{i+1}}$. Base: [Lemma 4.1](#) at $n' = n$ gives $h(x) = h(-x)$. Step: h invariant under μ_{2^i} ; the map $x \mapsto x^{2^i}$ has fibers the μ_{2^i} -cosets, on which h is constant with common value h^* , and $0 = \widehat{h}(2^i) = 2^i \sum_{y \in \mu_{n'}} h^*(y)y$ with $n' = n/2^i \geq 4$. [Lemma 4.1](#) gives $h^*(y) = h^*(-y)$; choosing ω with $\omega^{2^i} = -1$ (present since $2^{i+1} \leq M_0 \mid n$), $h(x) = h^*(x^{2^i}) = h^*(-x^{2^i}) = h(\omega x)$, so h is invariant under $\langle \mu_{2^i}, \omega \rangle = \mu_{2^{i+1}}$. Converse: $\sum_{\kappa \in \mu_{M_0}} \kappa^j = 0$ for $M_0 \nmid j$. $\square$

[Theorem 10.7](#) is the case $\sigma = 1$: the antipodal mechanism propagates dyadically rather than stalling. Three consequences are used throughout. *Witness classification:* $A \subseteq \mu_n(\mathbb{C})$ with

$e_1(A) = \dots = e_{\sigma'}(A) = 0$ in $\mathbb{Z}[\zeta]$ is a union of $\mu_{M'_0}$ -orbits, M'_0 the least 2-power exceeding σ' (Newton converts e - to p -conditions in characteristic zero). *Class cap*: subsets with equal $\mathbb{Z}[\zeta]$ -prefixes $(e_1, \dots, e_\sigma)$ differ by μ_{M_0} -coset rearrangement, and the class of a set with a cosets inside and b outside has size exactly $\sum_u \binom{a}{u} \binom{b}{u} = \binom{a+b}{a} \leq 2^{N_0}$. *Norm bound*: a nonzero $\widehat{h}(j)$ has all archimedean conjugates of modulus at most $\|h\|_1$, so $0 < |\mathrm{N}_{\mathbb{Q}(\zeta)/\mathbb{Q}} \widehat{h}(j)| \leq \|h\|_1^{n/2}$ — the bridge from characteristic zero to $\mathbb{F}_p$: a mod- p coincidence that fails in $\mathbb{Z}[\zeta]$ (a *collision*) forces p to divide a bounded nonzero integer. The lattice identity $\{\widehat{h}(j) = 0, j \leq \sigma\} = \{\mu_{M_0}\text{-invariant}\}$ was verified by exact linear algebra at $n = 16$ and $n = 32$ across five slack levels ([appendix A](#)).

Theorem 14.2 (Witness dichotomy: non-dyadic slack is collision-only). *Let $t \geq 2$, $2(t-1) < n$, and $A \subseteq \mu_n(\mathbb{C})$ with $e_1(A) = \dots = e_{t-1}(A) = 0$ in $\mathbb{Z}[\zeta]$.*

- (i) *If t is a power of two: A is a union of μ_t -orbits — the descent witnesses of [Theorem 13.1](#) are the only characteristic-zero diagonal witnesses, and they realize exactly the values [Theorem 13.1](#) realizes.*
- (ii) *If t is not a power of two: $M'_0 > t$, the converse half of [Theorem 14.1](#) forces $p_t(A) = 0$, and Newton then $e_t(A) = 0$: characteristic-zero witnesses with $z \neq 0$ do not exist. Every mod- p witness with $z \neq 0$ is a collision: any unconditional non-dyadic slack claim is an instance of the collision problem — consistent with [Theorem 12.2\(iii\)](#) operating only at $p \leq N^{7/3}$, where collisions abound.*

Verified exhaustively at $n = 16$, $t = 3$, $\ell = 8$: all six characteristic-zero zero-prefix subsets are μ_4 -orbit unions, with $e_3 = 0$ ([appendix A](#)).

15 Positive theorems: the density regime

We now work at the deepest level: $N = n$, dyadic slack σ with $\sigma \mid k$, $s = k + \sigma$, and radius $1 - \rho - \sigma/n$. When one writes $\sigma \asymp Cn/\log_2 n$, the invariant quantity is

$$C_{\text{eff}} = \frac{\sigma \log_2 n}{n}, \quad \text{so that} \quad \frac{n}{\sigma} = \frac{\log_2 n}{C_{\text{eff}}}, \quad 2^{n/\sigma} = n^{1/C_{\text{eff}}}.$$

All asymptotic statements below are expressed in terms of C_{eff} ; if $\sigma = (1 + o(1))Cn/\log_2 n$, then $C_{\text{eff}} = C + o(1)$.

By [Theorem 10.2](#) read coefficientwise, the bad slopes of the line $u_z = U_{c^0} + zX^{s-j}$ through the base

$$U_{c^0} = X^s + \sum_{i < \sigma} (-1)^i c_i^0 X^{s-i}$$

in direction $j \in [1, \sigma]$ are exactly

$$\{(-1)^j(e_j(A) - c_j^0) : A \in \binom{H}{s}, e_i(A) = c_i^0 \text{ for all } i \in [1, \sigma] \setminus \{j\}\},$$

with $c_j^0 = 0$ at $j = \sigma$ or for the zero base. Condition (ii) of [Definition 9.2](#) is automatic on witness sets of more than k points. Write $B(p)$ for the bad-slope count of the canonical line ($c^0 = 0$, $j = \sigma$).

Theorem 15.1 (The canonical line, two-sided). (i) (Unconditional floor, every prime.)

$$B(p) \geq \min \left\{ p, \frac{s}{\sigma} \left(\frac{n}{\sigma} - \frac{s}{\sigma} \right) + 1 \right\} = \rho(1 - \rho) \left(\frac{\log_2 n}{C_{\text{eff}}} \right)^2 + O \left(\frac{\log n}{C_{\text{eff}}} \right),$$

realized by the descent witnesses of [Theorem 13.1](#) at $t = \sigma$, $N = n$.

(ii) (Density ceiling.) For every $X \geq n^2 \binom{n}{s}$ and at least half the primes $p \equiv 1 \pmod{n}$ in $(X, 2X]$ (Siegel–Walfisz [\[15\]](#); ineffective threshold),

$$B(p) \leq \underbrace{\min \left\{ p, \binom{n/\sigma}{s/\sigma} \right\}}_{\text{descent slopes}} + O_\rho(\log n).$$

Thus the canonical line's error contribution is at most

$$\frac{n^{1/C_{\text{eff}}} + O(\log n)}{p} \rightarrow 0$$

whenever p is in such a window and C_{eff} is bounded away from zero.

Proof. Part (i) is [Theorem 13.1](#) and [Lemma 10.10](#). For (ii), [Theorem 14.1](#) at level $\sigma - 1$ has modulus σ because σ is dyadic. Thus every characteristic-zero witness with $p_1(A) = \cdots = p_{\sigma-1}(A) = 0$ is a μ_σ -orbit union, hence a descent witness. There are at most $\binom{n/\sigma}{s/\sigma} \leq 2^{n/\sigma} = n^{1/C_{\text{eff}}}$ such descent slopes. Every other bad slope is carried by a collision single: a set A for which some nonzero cyclotomic integer $\widehat{1}_A(i)$, $i \leq \sigma - 1$, vanishes modulo $\mathfrak{p}$. Its norm is at most $n^{n/2}$, so it has at most $\kappa = \lceil n \log n / (2 \log X) \rceil$ prime divisors exceeding X . Summing over all s -subsets and dividing by the number $\gg X / (n \log X)$ of admissible primes gives the asserted $O_\rho(\log n)$ bound for at least half the primes. $\square$

Theorem 15.2 (All monomial bases, canonical direction). For every $X \geq n^2 \binom{n}{s}^2$ and at least half the primes $p \equiv 1 \pmod{n}$ in $(X, 2X]$, simultaneously for every base $c^0 \in \mathbb{F}_p^{\sigma-1}$, the canonical-direction line $U_{c^0} + zX^k$ has

$$\#\{\text{bad } z\} \leq |\Phi_{\sigma-1}^{-1}(c^0)| \leq 2^{n/\sigma} + O(\sqrt{\log n}) = n^{1/C_{\text{eff}}} + O(\sqrt{\log n}),$$

where $\Phi_{\sigma-1}(A) = (e_1(A), \dots, e_{\sigma-1}(A))$ on s -subsets.

Proof. Bad slopes refine the slab $\Phi_{\sigma-1}^{-1}(c^0)$ by the last coordinate e_σ , so their number is at most the slab size. In characteristic zero, any two sets in the same slab differ by a μ_σ -coset rearrangement; each class has size at most $P = 2^{n/\sigma}$. Let $C_p^{(\sigma-1)}$ count ordered pairs of s -subsets that have equal punctured prefixes modulo $\mathfrak{p}$ but not in $\mathbb{Z}[\zeta]$. For a slab of size x ,

$$x^2 \leq Px + C_p^{(\sigma-1)}, \quad \text{hence} \quad x \leq P + \sqrt{C_p^{(\sigma-1)}}.$$

For an ordered pair not equal in characteristic zero, choose the first nonzero prefix difference. Its norm is bounded by $(2s)^{n/2}$, so it is divisible by at most $\kappa \ll n \log n / \log X$ primes exceeding

X . Summed over $\binom{n}{s}^2$ ordered pairs and averaged over the admissible primes in $(X, 2X]$, the assumption $X \geq n^2 \binom{n}{s}^2$ gives $C_p^{(\sigma-1)} = O(\log n)$ for at least half the primes. This proves the displayed bound. $\square$

Remark 15.3 (Accounting the floor). The floor in [Theorem 15.1](#)(i) consists of quotient-periodic witnesses: quotient-codeword agreements, the natural candidates for what a final protocol definition of “mutual explanation” might absorb. Thus there are two numbers to track: the raw bad-slope count, which has a descent floor of order $(\log n/C_{\text{eff}})^2$ at every prime and, by [Proposition 19.1](#), a quotient-exact floor of order $n^{\beta(\rho)/C_{\text{eff}}(1-o(1))}$ above a quasipolynomial quotient threshold; and the aperiodic collision count, which is $O(\log n)$ for most primes in the density windows. Which number is operational depends on the protocol definition.

16 The direction-arithmetic of slack

Fix the zero base and vary the direction: bad slopes of $X^s + zX^{s-j}$ are carried by A with $e_i(A) = 0$ for $i \in [1, \sigma] \setminus \{j\}$ and $e_j(A) = \pm z$.

Lemma 16.1 (Ramification parity). *Let $B \subseteq \mu_N$, N a power of two, $|B| = s'$. Modulo the ramified prime $\lambda = (1 - \zeta_N)$ of $\mathbb{Z}[\zeta_N]$ [16], every root of unity is $\equiv 1$, so $e_i(B) \equiv \binom{s'}{i} \pmod{\lambda}$: $e_i(B) = 0$ forces $\binom{s'}{i}$ even. Moreover $\binom{s'}{i}$ is even for all $2 \leq i \leq \sigma'$ if and only if $s' \equiv 0$ or $1 \pmod{M'}$, M' the least power of two exceeding σ' .*

Proof. The congruence is immediate. Lucas, both directions: if $s' \equiv 0, 1 \pmod{M'}$, every $i \in [2, M' - 1]$ has a set bit in positions $1, \dots, \log_2 M' - 1$ where s' has none, so $\binom{s'}{i}$ is even. If $s' \equiv r \in [2, M' - 1]$: for $r \leq \sigma'$ take $i = r$; for $r > \sigma'$ the top bit of r sits at $M'/2$, and $i = r - M'/2$ (if ≥ 2) or $i = M'/2$ lies in $[2, \sigma']$ with bits inside r 's, making $\binom{s'}{i}$ odd. $\square$

Theorem 16.2 (Direction trichotomy). *Dyadic slack σ , zero base, direction $1 \leq j \leq \sigma$.*

- (i) $j = \sigma$ (canonical): the two-sided descent band of [Theorem 15.1](#).
- (ii) j non-dyadic: the constraints contain every dyadic frequency up to σ ; [Theorem 14.1](#) forces μ_{M_0} -invariance with $M_0 > \sigma \geq j$, hence $e_j(A) = 0$: no characteristic-zero slopes $z \neq 0$ exist; for most primes in the window of [Theorem 15.1](#), the bad-slope count is $O_\rho(\log n)$, all collision-borne.
- (iii) j dyadic, $j < \sigma$: the constraints at levels below j force μ_j -invariance, so $j \nmid s$ gives no witnesses at all; for $j \mid s$, the identity $\prod_{a \in A} (X - a) = \prod_{y \in B} (X^j - y)$ reduces the problem exactly to the direction-one problem at slack $\sigma' = \sigma/j$, size $s' = s/j$, on $\mu_{n/j}$: find B with $e_2(B) = \dots = e_{\sigma'}(B) = 0$ and $e_1(B) \neq 0$.
- (iv) Direction one: characteristic-zero witnesses require $s' \equiv 0, 1 \pmod{M'}$ ([Lemma 16.1](#)). Conversely they exist for $s' \equiv 1 \pmod{M'}$ at every σ' — take $B = \{y\} \sqcup (\mu_{M'}\text{-orbits})$, which

has $e_i(B) = e_i(\{y\}) = 0$ for $i \geq 2$ — and for $s' \equiv 0 \pmod{4}$ at $\sigma' = 2$: the quadruple $\{y, iy, \zeta_8 y, -\zeta_8 y\}$, whose six pairwise products $iy^2, \pm\zeta_8 y^2, \pm i\zeta_8 y^2, -iy^2$ cancel in pairs, padded by μ_4 -orbits. At $\sigma' = 2$ the classification is therefore complete: witnesses exist iff $s' \equiv 0, 1 \pmod{4}$. When witnesses exist they persist at every prime ($e_1(B) \neq 0$ in $\mathbb{Z}[\zeta]$ has finitely many prime divisors), and the scaling $A \mapsto uA$, $u \in H$, spreads e_j over at least n/j values: at least n/j bad slopes forever. When the parity obstructs: $O_\rho(\log n)$ bad slopes for most primes.

Proof. (ii) and the invariance step of (iii): [Theorem 14.1](#) applied to the dyadic frequencies present among the constraints (all of $1, \dots, j/2$ in case (iii); the full set in case (ii), since j is not among them), with Newton converting the e -conditions to $\mathbf{1}_A$ -conditions level by level. μ_j -invariance forces $j \mid |A| = s$. The reduction identity and the constructions are direct computations: e_2 of a disjoint union is $e_2(X) + e_2(Y) + e_1(X)e_1(Y)$, and $\mu_{M'}$ -orbits have $e_i = 0$ for $M' \nmid i$. Collision accounting as in [Theorem 15.1\(ii\)](#); persistence because the reduction of a characteristic-zero witness is a witness mod every $\mathfrak{p}$ not dividing $e_1(B)$. $\square$

The parity law leaves the residue $s' \equiv 0 \pmod{M'}$ at $\sigma' \geq 3$ undecided. The following exclusion, by a mechanism independent of ramification, settles its minimal case at every σ' .

Theorem 16.3 (Self-inversive exclusion). *Let $\sigma' \geq 2$ and let B be any set of $s' = \sigma' + 1$ points on the complex unit circle (in particular, roots of unity) with $e_2(B) = \dots = e_{\sigma'}(B) = 0$. Then $e_1(B) = 0$: direction-one witnesses of the minimal size $\sigma' + 1$ do not exist. Consequently, for $\sigma' = M' - 1$, where $\sigma' + 1$ is the minimal size in the residue class $s' \equiv 0 \pmod{M'}$, any witness in that class has at least $2M'$ elements.*

Proof. A monic polynomial with all roots on the unit circle is self-inversive: with $f(X) = \prod_{b \in B} (X - b)$ of degree s' and $f^*(X) := X^{s'} \bar{f}(1/X)$, unimodularity ($1/\bar{b} = b$) gives $f^* = \bar{q} f$ with $q = \prod b$, i.e. the coefficients satisfy $c_i = \bar{q} \bar{c}_{s'-i}$. Here $c_0 = 1$, $c_1 = -e_1$, and $c_2 = \dots = c_{\sigma'} = 0$; at $s' = \sigma' + 1$ the mirror of c_1 is $c_{\sigma'} = 0$, forcing $e_1 = 0$. $\square$

Remark 16.4 (The law, verified; the completed $\sigma' \leq 3$ picture). At $n = 16$, $\sigma = 2$, direction 1 ($M' = 4$): $s = 9 \equiv 1$: exactly $48 = 3 \cdot 16$ bad slopes at every prime tested from 257 to 524,353 — three witness classes times the H -orbit, as (iv) predicts; $s = 10 \equiv 2$ and $s = 11 \equiv 3$: zero at every large prime, with sporadic collision primes (257, 7489 at $s = 10$). At $\sigma = 4$, direction 3 (non-dyadic), $s = 12$: empty at every prime tested. Forty-eight forever against zero almost always, switched by $s \bmod 4$. At $\sigma' = 3$ ($M' = 4$): size 4 is excluded by [Theorem 16.3](#) at every N ; size 8 is characteristic-zero empty for $N \leq 32$ and size 12 for $N = 16$, proved by reduction modulo a single large prime — a characteristic-zero witness survives mod every prime, and exhaustive enumeration at $p = 65089$ finds none, while the witnesses visible mod 97 and 193 are collisions vanishing there ([appendix A](#)). We conjecture that at $\sigma' = 3$ witnesses exist iff $s' \equiv 1 \pmod{4}$: the parity-allowed residue 0 is empty, in sharp contrast to $\sigma' = 2$, where the ζ_8 -quadruple fills it. A $(1 - \zeta)$ -adic analysis at the next level yields the necessary condition $\sum_b j_b \equiv 0 \pmod{2}$ on exponents but no contradiction; the residue-zero classification at $\sigma' \geq 3$ beyond [Theorem 16.3](#) remains open.

Theorem 16.5 (Newton obstruction criterion for arbitrary bases). *Let $j \in [1, \sigma]$ be non-dyadic and fix a base vector $c^0 = (c_i)_{i \neq j}$. Let $P_m(e_1, \dots, e_m)$ be the Newton polynomial for the power sum p_m , and define, for a formal increment Δ in the j th elementary coefficient,*

$$\mathcal{N}_{m,j}(c^0; \Delta) = P_m(c_1, \dots, c_{j-1}, c_j + \Delta, c_{j+1}, \dots, c_m) - P_m(c_1, \dots, c_{j-1}, c_j, c_{j+1}, \dots, c_m),$$

where c_j is an auxiliary reference value. If two characteristic-zero sets A, B lie in the same punctured slab

$$e_i(A) = e_i(B) = c_i \quad (i \neq j, 1 \leq i \leq \sigma),$$

and if their increment $\Delta = e_j(A) - e_j(B)$ satisfies

$$\mathcal{N}_{2^r,j}(c^0; \Delta) = 0 \quad \text{for every dyadic } 2^r \leq \sigma,$$

then the rigidity theorem applies: $A - B$ is μ_{M_0} -invariant and consequently $\Delta = 0$. Therefore any attempt to prove a one-slope-per-slab statement for arbitrary bases must first prove that all dyadic Newton obstruction equations vanish. Without those extra equations, the rigidity argument does not apply.

Proof. The equations $\mathcal{N}_{2^r,j}(c^0; \Delta) = 0$ say exactly that the two sets have equal dyadic power sums up to σ . Applying [Theorem 14.1](#) to the difference of their indicator functions gives μ_{M_0} -invariance. Since $j < M_0$, the converse half of [Theorem 14.1](#) gives equality of the j th power sums. With all lower elementary symmetric functions already equal, Newton's identity at level j gives $j\Delta = 0$ in characteristic zero, hence $\Delta = 0$. $\square$

Example 16.6 (Why the earlier all-bases non-dyadic theorem was false). Take $\sigma = 4$ and $j = 3$. In a punctured slab, two sets have the same e_1, e_2, e_4 but may differ in e_3 by Δ . Newton's identities give

$$p_4 = e_1 p_3 - e_2 p_2 + e_3 p_1 - 4e_4, \quad p_3 = e_1 p_2 - e_2 p_1 + 3e_3,$$

and therefore

$$p_4(A) - p_4(B) = 4e_1 \Delta.$$

Thus the dyadic frequency p_4 is not fixed by the punctured slab unless $e_1 = 0$ or $\Delta = 0$. The zero-base non-dyadic theorem is valid because $e_1 = 0$; the arbitrary-base statement is open except under extra Newton-obstruction hypotheses.

17 The exact regime, and the closed-form count

Theorem 17.1 (Stable range: exactness for the proved strata). *Let $p \equiv 1 \pmod{n}$ with $p > (2s)^{n/2}$. Then, with no exceptional primes:*

(i) *the canonical line's bad-slope set is exactly the descent image*

$$\{\mp e_1(B) : B \subseteq H^\sigma, |B| = s/\sigma\};$$

- (ii) every canonical-direction line $U_{c^0} + zX^k$ has at most $2^{n/\sigma} = n^{1/C_{\text{eff}}}$ bad slopes;
- (iii) the zero-base direction trichotomy of [Theorem 16.2](#) holds with the “most primes” conclusions upgraded to “every prime”: parity- or rigidity-obstructed zero-base directions have no characteristic-zero slopes $z \neq 0$, while witness-bearing zero-base directions have exactly their characteristic-zero counts;
- (iv) arbitrary bases in non-canonical non-dyadic directions are not covered by this theorem, except under the additional Newton-obstruction conditions of [Theorem 16.5](#).

Proof. Any collision at a level up to σ requires $\mathfrak{p}$ to divide a nonzero cyclotomic integer $\widehat{h}(j_0)$ with all archimedean conjugates of modulus at most $\|h\|_1 \leq 2s$. Hence

$$p \leq |N \widehat{h}(j_0)| \leq (2s)^{n/2},$$

contrary to the stable-range hypothesis. Thus all modular coincidences are characteristic-zero coincidences. The canonical line and canonical-direction slabs are then governed by [Theorem 14.1](#); zero-base directions are governed by [Theorem 16.2](#). The arbitrary-base non-dyadic case is excluded from the conclusion because equality in a punctured slab does not force equality at the dyadic power sums; Example [16.6](#) gives the basic obstruction. $\square$

In the stable range the canonical count is not merely pinned: it has a closed form.

Theorem 17.2 (Exact slope count). *Let $N' = n/\sigma$, $n_1 = N'/2$, $\ell' = \rho N' + 1$, and let*

$$A(N', \ell') = \sum_{\substack{u \geq 0, \ t = \ell' - 2u \geq 0 \\ u \leq n_1 - t}} \binom{n_1}{t} 2^t$$

be the number of antipodal-rearrangement classes of ℓ' -subsets of $\mu_{N'}$; the class invariant is the signed set of singleton antipodal pairs together with the number of full pairs. For every prime in the stable range, the canonical line’s bad-slope count is exactly

$$B(p) = A(N', \ell'), \quad \text{and at } \rho = \tfrac{1}{2}: \quad B(p) = \frac{3^{n/(2\sigma)} - 1}{2}.$$

Consequently, at safe slack,

$$B(p) = n^{\beta(\rho)/C_{\text{eff}} + o(1)},$$

where

$$\beta(\rho) = \frac{1}{2} \max_{0 \leq \theta \leq 2 \min(\rho, 1-\rho)} (H_2(\theta) + \theta) = \begin{cases} \frac{1}{2}(H_2(2\rho) + 2\rho), & 0 < \rho < 1/3, \\ \frac{1}{2} \log_2 3, & 1/3 \leq \rho \leq 2/3, \\ \frac{1}{2}(H_2(2(1-\rho)) + 2(1-\rho)), & 2/3 < \rho < 1. \end{cases}$$

At the deployed rate $\rho = 1/2$, this is $n^{(\log_2 3)/(2C_{\text{eff}}) + o(1)}$; the lower deployed rates use the small- ρ branch above.

Proof. By [Theorem 17.1\(i\)](#), $B(p)$ is the number of distinct values of $e_1(B)$ for $B \in \binom{\mu_{N'}}{\ell'}$. Two such subsets have equal e_1 in $\mathbb{Z}[\zeta_{N'}]$ iff their difference h satisfies $\widehat{h}(1) = 0$, which by [Theorem 10.7](#) is exactly antipodal-pair rearrangement. Choose t singleton antipodal pairs and their signs, and choose u full pairs; feasibility is $t + 2u = \ell'$ and $u \leq n_1 - t$, giving the displayed sum. The stable range prevents distinct characteristic-zero values from colliding modulo $\mathfrak{p}$. At $\rho = 1/2$ the parity-restricted trinomial sum equals $\frac{1}{2}(3^{n_1} - 1)$. The asymptotic formula is the entropy maximization of $\binom{n_1}{t} 2^t$ under the feasibility cap $t \leq \min(\ell', N' - \ell')$ built into the sum: the constraint $t \leq \ell'$ binds for $\rho \leq 1/2$ and $t \leq N' - \ell'$, from $u \leq n_1 - t$, binds for $\rho > 1/2$. $\square$

Remark 17.3 (Where the stable threshold sits). $(2s)^{n/2}$ lies above the density windows of [Theorems 15.1](#) and [15.2](#). Those theorems can be viewed as averaged-prime descents of the exact regime. Empirically, collapse may occur far earlier, but the stable threshold is the price of every-prime certainty.

18 The boundary, and the conjecture

18.1 The exact normal form

The positive theorems above concern monomial lines. The general question is posed exactly by the following coordinates. Throughout, $s_\delta = \lceil (1 - \delta)n \rceil$ and $r = n - k$.

Definition 18.1 (Residue-line datum and packing number). Fix $1 \leq t \leq r$. A *degree- t residue-line datum* over $(\mathbb{F}, D, k)$ is a denominator $E \in \mathbb{F}[X]$ of degree t nonzero on D , a numerator B with $\deg B < t$, and an anchor word $w : D \rightarrow \mathbb{F}$. A *witness* for slope z at radius δ is (Q_z, S_z) with $\deg Q_z < k + t$, $|S_z| \geq s_\delta$, $Q_z \equiv zB \pmod{E}$, and $Q_z = w$ on S_z ; it is *noncontained* if no $A, G \in \mathbb{F}_{<k}[X]$ satisfy $A = w/E$ and $G = -B/E$ on S_z . Let $\Lambda_{t,\delta}^{\text{NC}}(D, k)$ be the maximum, over data, of the number of slopes admitting a noncontained witness.

Lemma 18.2 (Denominator representation). *For a word $b : D \rightarrow \mathbb{F}$ and E of degree t nonzero on D : a representation $b = R - B/E$ on D with $\deg R < k$, $\deg B < t$ exists iff the word Eb extends to a polynomial of degree $< k + t$ on D ; for $t = r$ this always holds, by interpolation. (Euclidean division of the extension by E , both ways.)*

Theorem 18.3 (Exact normal form). *If some polynomial of degree r is nonzero on D (automatic for multiplicative domains, via X^r), then for every δ :*

$$\varepsilon_{\text{mca}}(C, \delta) = \frac{1}{q} \max_{1 \leq t \leq r} \Lambda_{t,\delta}^{\text{NC}}(D, k).$$

Proof. ($\leq$) Fix a line $u_z = f + zg$ with bad-slope set Γ and witnesses $u_z = P_z$ on S_z . By [Lemma 18.2](#) at $t = r$ write $g = R - B/E$ on D ; set $w = E \cdot f$ and $Q_z = E(P_z - zR) + zB$. Then $\deg Q_z < k + t$, $Q_z \equiv zB \pmod{E}$, and on S_z : $Q_z = E(f + zg - zR) + zB = Ef - zB + zB = w$. Containment would yield degree- $< k$ explanations of both f and g on S_z , contradicting badness (ii); all of Γ uses one datum. ($\geq$) Given a datum and noncontained-witnessed slopes, the line $f = w/E$, $g = -B/E$ has, via $P_z = (Q_z - zB)/E$ (degree $< k$ by count), $P_z = f + zg$ on S_z : condition (i); noncontainedness is verbatim condition (ii). $\square$

Remark 18.4 (Strata). The monomial slack line $x^{k+T} + zx^k$ on a subgroup is the datum $(E, B, w) = (X^r, -1, x^T)$ — maximal denominator, simplest numerator — and its noncontained slope set is exactly $\mathcal{B}_T(D, k)$ (Theorem 10.2): every theorem of Sections 15 to 17 is a statement about this stratum and its quotient refinements. Theorem 9.3 is the trivial packing bound ($\Lambda^{\text{NC}} \leq 2^n$), and the tangent floor (Proposition 9.4) is a residue-line family with $\Lambda^{\text{NC}} \geq \lfloor \delta n \rfloor$, which is why Conjecture 18.8 below carries the $n^{1+o(1)}/q$ correction — the floor saturates it, so the conjectured positive half is sharp if true.

Proposition 18.5 (Denominator closure as a coordinate normal form). *In the coordinate formulation of line obstructions, assuming $|S_z| > k$ for active rows, every obstruction class is either zero, tangent (the direction word is explained on the active coordinates but the anchor is not), or a residue-line datum as in Definition 18.1. Conversely, every residue-line datum gives such a class. This proposition is a normal-form statement only; it does not bound the packing number Λ^{NC} .*

Proof. If the direction word b is explained by a degree- $< k$ polynomial on the active set, subtract that global completion; the remaining class is zero or tangent according as the anchor is explained. Otherwise Lemma 18.2 writes $b = R - B/E$. Subtracting the global direction R reduces to $b = -B/E$ with $B \neq 0$, and multiplying the row relation by E gives $EP_z + zB = w$ on S_z , exactly the residue-line form. The converse is obtained by reversing these formulas. $\square$

18.2 The obstruction, the open core, and the conjecture

Proposition 18.6 (No-anchor obstruction). *Every pair (S, T) of s -subsets of D is simultaneously an agreement pattern of some received word $(U$ with interpolant $L_{S \cup T} \cdot G$ has both explaining codewords zero). Hence no word-free pair condition exists, lines with arbitrary f, g admit no characteristic-zero anchor, and methods that certify statements by averaging over primes — bounding, for each characteristic-zero object, the number of primes at which it degenerates, as in every density theorem of this paper — are structurally incapable of bounding ε_{mca} over all lines, the words being quantified after the prime. The all-lines positive half rests entirely on fixed-prime technology.*

Problem 18.7 (Per-fiber collision problem). *For $\sigma \geq Cn/\log_2 n$ and all large $p \equiv 1 \pmod{n}$ with $n^{c_0} \leq p \leq 2^{o(n)}$: every fiber of the prefix map $\Phi_\sigma(A) = (e_1(A), \dots, e_\sigma(A))$ on s -subsets of H contains at most $n^{O(1)}$ ordered pairs that are prefix-equal mod $\mathfrak{p}$ but not in $\mathbb{Z}[\zeta]$. This single divisibility statement implies the monomial-line positive half throughout the window, beyond the density and stable regimes proved above; the natural attacks — prime averaging (blocked at fixed p by construction), the polynomial method (whose Nullstellensatz coefficients on subgroup grids become complete character sums), subgroup exponential sums (whose known ranges end at $N \geq p^{3/7+\varepsilon}$), and anticoncentration technology (whose forward inequalities reduce to the same even-moment counts) — all terminate at it.*

Conjecture 18.8 (Slack MCA threshold, floor- and quotient-corrected). *Let $H_n \leq \mathbb{F}_{q_n}^\times$ be smooth of order n , $k_n = \rho n$, $\text{poly}(n) \leq q_n \leq 2^{o(n)}$, and let $q_{D,n}$ be the generated-field size. Suppose η_n clears the corrected list-side scales,*

$$\eta_n \geq (1 + \varepsilon)\tau^*(\rho, q_{D,n})$$

for some fixed $\varepsilon > 0$, and

$$\mathcal{Q}_{H_n}(\eta_n) \leq B_Q \log_2 n + O(\log \log n)$$

for some fixed profile exponent B_Q . With Λ^{aper} the quotient-separated packing number of [Remark 18.9](#), the positive MCA prediction is

$$\max_t \Lambda_{t, 1-\rho-\eta_n}^{\text{aper}}(H_n, k_n) \leq n^{1+o(1)},$$

and, after the tangent-floor and generated-field corrections, the raw error obeys the two-term bound

$$\varepsilon_{\text{mca}}(C_n, 1 - \rho - \eta_n) \leq \frac{n^{1+o(1)} + 2^{(\beta(\rho)/H_2(\rho))\mathcal{Q}_{H_n}(\eta_n)(1+o(1))}}{q_{D,n}} \leq \frac{n^{1+o(1)} + n^{(\beta(\rho)/H_2(\rho))B_Q+o(1)}}{q_{D,n}}.$$

In particular $\varepsilon_{\text{mca}}(C_n, 1 - \rho - \eta_n) \leq n^{1+o(1)}/q_{D,n}$ whenever the profile obeys the stronger bound $\mathcal{Q}_{H_n}(\eta_n) \leq (H_2(\rho)/\beta(\rho)) \log_2 n + o(\log n)$. The quotient term is not removable: [Theorem 19.4](#) gives the exact monomial-line lower bound $A(N^*, \rho N^* + 1)/q$ above the quotient norm threshold, and the profile-asymptotic form $2^{(\beta/H_2)\mathcal{Q}_H(\eta)(1-o(1))}/q$ when $N^* \rightarrow \infty$. Thus, in that asymptotic quotient range, a target raw bound n^A/q requires $B_Q \leq (H_2(\rho)/\beta(\rho))A$, equivalently $\mathcal{Q}_H(\eta) \lesssim (H_2(\rho)/\beta(\rho))A \log_2 n$. The failure half remains a separate conditional prediction: below the generated-field entropy gap, MCA error should be $1 - o(1)$ only in strata where the relevant power-sum or residue-line map is sufficiently equidistributed to hit almost every slope. Coefficient pigeonhole and quotient-core mechanisms prove list-side obstructions and motivate this failure prediction, but they do not by themselves force many bad slopes of an arbitrary line.

Remark 18.9 (Aperiodic packing normalization). The phrase “after quotient-periodic denominator families are separated” in [Conjecture 18.8](#) and [Conjecture 27.2](#) can be made precise. Call a residue line in the normal form of [Theorem 18.3](#) *quotient-periodic* if its denominator is a pull-back through a proper quotient $x \mapsto x^M$ with $M \mid \gcd(n, k)$, $M > 1$, and let $\Lambda_{t,\delta}^{\text{aper}}(H, k)$ be the maximum size of a noncontained packing containing no quotient-periodic line. The conjectures assert $\max_t \Lambda_{t, 1-\rho-\eta_n}^{\text{aper}}(H_n, k_n) \leq n^{1+o(1)}$, while the removed quotient-periodic lines are predicted to contribute at most the displayed quotient-profile term, matching the lower bound of [Theorem 19.4](#).

19 The quotient-exact floor, and the necessity of the quotient profile

The descent floor of [Theorem 15.1\(i\)](#) counts canonical-line bad slopes produced on a quotient. At quotient order $N' = n/\sigma$ and quotient agreement size $\ell' = \rho N' + 1$, the Dias da Silva–Hamidoune

lower bound gives only the no-threshold floor $\rho(1 - \rho)(N')^2 + O(N')$. The closed-form quotient count $A(N', \ell') = 2^{\beta(\rho)N'(1-o(1))}$ from [Theorem 17.2](#) is exponentially larger in N' , but the proof of exactness there used the full-domain stable range $p > (2s)^{n/2}$. The point of this section is that, for the floor alone, only quotient-level collisions matter. The relevant norm bound is therefore $(2\ell')^{N'/2}$, which is quasipolynomial at safe slack, not exponential in $n \log n$.

Throughout this section

$$\beta(\rho) = \frac{1}{2} \max_{0 \leq \theta \leq 2 \min(\rho, 1-\rho)} (\mathbf{H}_2(\theta) + \theta), \quad A(N', \rho N' + 1) = 2^{\beta(\rho)N'(1-o(1))},$$

as in [Theorem 17.2](#). Thus $\beta(\rho) = \frac{1}{2}(\mathbf{H}_2(2\rho) + 2\rho)$ for $\rho < 1/3$, $\beta(\rho) = \frac{1}{2} \log_2 3$ for $1/3 \leq \rho \leq 2/3$, and $\beta(\rho) = \frac{1}{2}(\mathbf{H}_2(2(1-\rho)) + 2(1-\rho))$ for $\rho > 2/3$; at the deployed rates $\rho \leq 1/2$ this agrees with the one-sided maximization over $\theta \leq 2\rho$.

19.1 The quotient-exact floor

Proposition 19.1 (Quotient-exact floor). *Let $D \leq \mathbb{F}_p^\times$ be smooth of order n , let $C = \text{RS}[\mathbb{F}_p, D, k]$ with $k = \rho n$, and let $N' \mid n$ satisfy $\rho N' \in \mathbb{Z}$ and $\ell' = \rho N' + 1 \leq N'$. Put $\sigma = n/N'$ and $Q = D^\sigma$, so Q has order N' . If*

$$p > (2\ell')^{N'/2},$$

then the canonical line $x^{k+\sigma} + zx^k$ at radius $1 - \rho - 1/N'$ has at least

$$\left| \{-e_1(B) \bmod \mathfrak{p} : B \in \binom{Q}{\ell'}\} \right| = A(N', \ell')$$

bad slopes, for every prime $p \equiv 1 \pmod{n}$ above the displayed threshold, where $\mathfrak{p} \subset \mathbb{Z}[\zeta_{N'}]$ is any degree-one prime above p ; if $N' \rightarrow \infty$, then $A(N', \ell') = 2^{\beta(\rho)N'(1-o(1))}$ by [Theorem 17.2](#). Equivalently, above the quotient norm threshold, the descent part of the canonical bad-slope set is exactly the reduction of its characteristic-zero value set. The threshold is

$$2^{(N'/2) \log_2(2\ell')},$$

for example $2^{33.4}$ at $(N', \rho) = (16, 1/2)$ and $2^{81.4}$ at $(32, 1/2)$.

Proof. By [Lemma 10.3](#) at $t = 1$ and [Theorem 10.4](#), every value $z = -e_1(B)$ with $B \in \binom{Q}{\ell'}$ is support-wise bad for $x^{k+\sigma} + zx^k$ at radius $1 - \rho - 1/N'$. It remains only to count the image of e_1 modulo p .

Fix a degree-one prime $\mathfrak{p} \subset \mathbb{Z}[\zeta_{N'}]$ over p , identifying Q with $\mu_{N'}$. If $B \neq B'$ have $e_1(B) \neq e_1(B')$ in $\mathbb{Z}[\zeta_{N'}]$, then $h = \mathbf{1}_B - \mathbf{1}_{B'}$ has $\widehat{h}(1) \neq 0$. Every archimedean conjugate of $\widehat{h}(1)$ has modulus at most $\|h\|_1 \leq 2\ell'$, so

$$0 < \left| \mathbf{N}_{\mathbb{Q}(\zeta_{N'})/\mathbb{Q}} \widehat{h}(1) \right| \leq (2\ell')^{N'/2} < p.$$

Thus $\mathfrak{p} \nmid \widehat{h}(1)$: distinct characteristic-zero values remain distinct modulo $\mathfrak{p}$. By [Theorem 10.7](#) and the class enumeration in [Theorem 17.2](#), the number of distinct characteristic-zero values is exactly the antipodal-class count $A(N', \ell')$. $\square$

Remark 19.2 (The threshold is real). Below the quotient norm threshold, quotient-level collisions can reduce the image below $A(N', \ell')$. At $(N', \ell') = (16, 9)$ the image equals $3280 = A(16, 9)$ for every tested prime $p \geq 67,057$, but drops to 3264 at $p = 58,337$ and 3232 at $p = 66,593$, and fluctuates in $[2576, 3280]$ throughout $p \in [5281, 66593]$. Every observed deficit is a multiple of 16, i.e. an H -orbit of collided classes, as forced by the dilation invariance of [Lemma 26.3](#); see V13 in [appendix A](#). Thus an every-prime floor of the exact quotient size genuinely needs a norm threshold. Below it, the appropriate unconditional statement is a density statement over primes, as in [Theorem 15.1\(ii\)](#).

Corollary 19.3 (Canonical line at safe slack, two-sided at every large quotient prime). *In the setting of [Theorem 15.1](#), put $N' = n/\sigma = (\log_2 n)/C_{\text{eff}}$ and $\ell' = \rho N' + 1$. For every prime $p \equiv 1 \pmod{n}$ with*

$$p > (2\ell')^{N'/2} = n^{O_{C_{\text{eff}}}(\log \log n)},$$

one has

$$B(p) \geq A(N', \ell') = n^{\beta(\rho)/C_{\text{eff}}(1-o(1))}.$$

Moreover, for at least half the primes in the windows of [Theorem 15.1\(ii\)](#), the descent term is bounded by the same class count and hence

$$B(p) \leq A(N', \ell') + O_\rho(\log n).$$

Where the two prime ranges overlap, the canonical-line count at safe slack is pinned at exponent $\beta(\rho)/C_{\text{eff}}$.

19.2 Necessity of the quotient profile for MCA

Theorem 19.4 (The quotient profile is necessary for MCA at large primes). *Let $H_n \leq \mathbb{F}_p^\times$ be smooth of order n , let $C = \text{RS}[\mathbb{F}_p, H_n, \rho n]$, and let $\eta \in (0, 1 - \rho)$. Assume the maximum defining $\mathcal{Q}_{H_n}(\eta)$ in [Definition 3.8](#) is nonempty, let M^* be a divisor attaining it, and put $N^* = n/M^*$. Suppose*

$$p > (2(\rho N^* + 1))^{N^*/2}.$$

Then the exact quotient-floor lower bound is

$$\varepsilon_{\text{mca}}(C, 1 - \rho - \eta) \geq \frac{A(N^*, \rho N^* + 1)}{p}.$$

If moreover $N^ \rightarrow \infty$ along the sequence, then*

$$A(N^*, \rho N^* + 1) = 2^{(\beta(\rho)/H_2(\rho))\mathcal{Q}_{H_n}(\eta)(1-o(1))}.$$

Consequently, along sequences with $N^ \rightarrow \infty$, if $\varepsilon_{\text{mca}}(C, 1 - \rho - \eta) \leq n^A/p$ at such a prime, then*

$$\mathcal{Q}_{H_n}(\eta) \leq \frac{H_2(\rho)}{\beta(\rho)} A \log_2 n (1 + o(1)).$$

For bounded N^* the exact bound stands while the profile is $O(1)$, so the consequence is vacuous there. At the deployed rates $1/2, 1/4, 1/8, 1/16$, the constants $H_2(\rho)/\beta(\rho)$ are approximately 1.262, 1.082, 1.024, and 1.009.

Proof. A divisor qualifying for [Definition 3.8](#) satisfies $M^* > \eta n$, equivalently $1/N^* > \eta$. The radius $1 - \rho - 1/N^*$ is therefore below $1 - \rho - \eta$, and bad slopes at the former radius remain bad at the latter by monotonicity in [Definition 9.2](#). Applying [Proposition 19.1](#) at quotient order N^* gives at least $A(N^*, \rho N^* + 1)$ bad slopes, which is the exact bound. If $N^* \rightarrow \infty$, then

$$\mathcal{Q}_{H_n}(\eta) = \log_2 \left(\frac{N^* - 1}{\rho N^*} \right) = H_2(\rho)N^*(1 - o(1)), \quad \log_2 A(N^*, \rho N^* + 1) = \beta(\rho)N^*(1 - o(1)),$$

by Stirling and [Theorem 17.2](#), which gives the asymptotic identity and the displayed necessary condition. $\square$

Corollary 19.5 (Every-prime polynomial lower bounds at logarithmic gap). *In the setting of [Theorem 19.4](#), with $C = \text{RS}[\mathbb{F}_p, H_n, \rho n]$, let $N' = N'(n)$ be a qualifying quotient order with $N' = (1 - o(1))c \log_2 n$ and $1/N' > 1/(c \log_2 n)$. For every prime $p \equiv 1 \pmod{n}$ with*

$$\log_2 p \geq \frac{c}{2}(1 + o(1)) \log_2 n \cdot \log_2(c \log_2 n),$$

$$\varepsilon_{\text{mca}} \left(C, 1 - \rho - \frac{1}{c \log_2 n} \right) \geq \frac{n^{\beta(\rho)c(1-o(1))}}{p}.$$

Thus, above the quotient norm threshold, a gap of order $1/\log n$ carries polynomially many quotient-periodic bad slopes of degree $\beta(\rho)c(1 - o(1))$ — any prescribed polynomial degree once the constant c is large. Superpolynomial counts require $c \rightarrow \infty$, i.e. gaps of order $o(1/\log n)$.

Remark 19.6 (Margin against the conjectured positive half). [Theorem 19.4](#) gives the strongest value-counting attack available from the quotient rigidity theory. If $\mathcal{Q}_{H_n}(\eta) \leq B_Q \log_2 n$, the attack produces $n^{(\beta/H_2)B_Q(1-o(1))}$ bad slopes, while a target upper bound of the form n^{B_Q}/p would still survive by the margin

$$H_2(\rho) - \beta(\rho) = \begin{cases} \frac{1}{2}(2H_2(\rho) - H_2(2\rho) - 2\rho), & 0 < \rho < 1/3, \\ H_2(\rho) - \frac{1}{2} \log_2 3, & 1/3 \leq \rho \leq 2/3, \\ \frac{1}{2}(2H_2(\rho) - H_2(2(1-\rho)) - 2(1-\rho)), & 2/3 < \rho < 1, \end{cases}$$

which is positive for $0 < \rho < 1$: the outer branches equal $f(\min(\rho, 1 - \rho))$ for $f(\delta) = H_2(\delta) - \frac{1}{2} H_2(2\delta) - \delta$, which has $f(0) = 0$ and $f'(\delta) = \log_2 \frac{1-\delta}{1-2\delta} \geq 0$, while $H_2(\rho) \geq H_2(1/3) > \frac{1}{2} \log_2 3$ on the middle branch. The margin is $\rho^2/(2 \ln 2)(1 + O(\rho))$ as $\rho \rightarrow 0$ and $(1 - \rho)^2/(2 \ln 2)(1 + O(1 - \rho))$ as $\rho \rightarrow 1$. Numerically this margin is approximately 0.2075, 0.0613, 0.0129, 0.0030 at the deployed rates $1/2, 1/4, 1/8, 1/16$. Any refutation of the positive half beyond this quotient attack must therefore produce, at some quotient scale N' , at least $2^{(H_2(\rho) - \beta(\rho))N'(1-o(1))}$ additional bad slopes outside the characteristic-zero quotient image: collision-borne or non-quotient slopes of the kind isolated in [Problem 18.7](#). The empirical frontier in V13–V14 is consistent with this division.

19.3 A concrete cap for challenge-size prime fields

Proposition 19.7 (Grand-challenge cap below roughly 150–162 bits). *Let $\rho \in \{1/2, 1/4, 1/8, 1/16\}$, let $D \leq \mathbb{F}_p^\times$ be smooth of order n with $64 \mid n$, and let $C = \text{RS}[\mathbb{F}_p, D, \rho n]$. Let $R(\rho)$ be the total reach in the table below. For every prime $p \leq 2^{R(\rho)}$,*

$$\varepsilon_{\text{mca}} \left(C, 1 - \rho - \frac{1}{64} \right) > 2^{-128}, \quad \delta_C^*(2^{-128}) < 1 - \rho - \frac{1}{64}.$$

In particular, the conclusion holds uniformly for all four displayed rates for every $p \leq 2^{150}$.

ρ	best N'	gap $1/N'$	DSH reach (bits)	norm floor $A(N', \ell')$	$R(\rho)$ (bits)
1/2	32	1/32	136.0	$21,523,360 = \frac{1}{2}(3^{16} - 1)$	152.4
1/4	32	1/32	135.7	7,465,888	150.8
1/8	64	1/64	137.0	14,798,298,944	161.8
1/16	64	1/64	136.2	6,483,776	150.6

Proof. At quotient order N' with $\ell' = \rho N' + 1$, the slack-one value set contains $-\ell'^\wedge Q$. By [Lemma 10.10](#), this set has size at least $\min\{p, \ell'(N' - \ell') + 1\}$ at every prime. Therefore

$$\varepsilon_{\text{mca}} \geq \frac{\min\{p, \ell'(N' - \ell') + 1\}}{p} \geq 2^{-128}$$

whenever $\log_2 p \leq 128 + \log_2(\ell'(N' - \ell') + 1)$; this is the DSH reach. Above the quotient norm threshold, [Proposition 19.1](#) gives

$$\varepsilon_{\text{mca}} \geq \frac{A(N', \ell')}{p} \geq 2^{-128}$$

for $(2\ell')^{N'/2} < p \leq 2^{128} A(N', \ell')$; this is the norm reach. In each row, the norm threshold lies below the DSH reach, so the two bands cover every prime up to $R(\rho)$. For $N' = 32$, bad slopes at radius $1 - \rho - 1/32$ remain bad at $1 - \rho - 1/64$ by monotonicity; for $N' = 64$ the radii are equal. $\square$

Remark 19.8 (The remaining 153–256 bit band). The norm-certificate mechanism is structurally capped near this range: its reach is $128 + \beta(\rho)N'$ bits, while its threshold is $(N'/2)\log_2(2\ell')$ bits. For prime fields below the per-rate reach in [Proposition 19.7](#), the challenge threshold is pinned below $1 - \rho - 1/64$. Between that reach and 2^{256} , whether the diagonal value sets survive reduction is a fixed-prime collision question of the kind posed in [Problem 18.7](#). The heuristic suggested by V13–V14 is that the value sets stay near $\min\{p, A\}$ until the counting limit, but proving such a statement would require equidistribution for power-of-two subgroups of order $O(\log p)$, far below currently available exponential-sum ranges. [Theorem 22.3](#) closes this band at a different strength grade: conditional on the imported list-to-agreement conversion [25, 24], every field of size at most 2^{256} admits, for challenge codes with $n \leq q/2$, MCA error at least 2^{-42} already at gap 2^{-11} ([Corollary 22.4](#)), so the remaining fixed-prime collision question governs only the error-one-grade strengthening. The sharpened form of this cap — gap 2^{-9} , error 2^{-86} , no field-size condition, extension fields included — is the main theorem of [2].

20 Solved high-agreement boundary layer

Status: proved finite Reed–Solomon/MDS boundary theorem. The quotient and entropy mechanisms in the preceding sections control the near-capacity side of the problem. There is also a complementary boundary layer, at very high agreement and small Hamming radius, where the staircase is exactly computable and no smoothness or quotient structure is involved. This section promotes the stable high-agreement material from the experimental notes into the main slack theory.

Let

$$C = \text{RS}[F, D, k], \quad |D| = n,$$

and write

$$a = n - r, \quad R = n - k,$$

so that r is the integer Hamming radius and a is the agreement threshold. For a line ledger, let $\text{LD}_{\text{sw}}(C, a)$ be the largest number of finite slopes on a received affine line that are support-wise noncontained at agreement at least a . Let $\text{LD}_{\text{ca}}(C, a)$ be the analogous no-loss correlated-agreement numerator, and let $\text{LD}_{\text{sw,proj}}(C, a)$ be the projective-slope version where the parameter space is $\mathbb{P}^1(F)$.

Theorem 20.1 (high-agreement tangent line staircase). *Let $C = \text{RS}[F, D, k]$ with $|D| = n$, and let $a = n - r$. If*

$$3a - 2n \geq k, \quad \text{equivalently} \quad r \leq \frac{n - k}{3},$$

then

$$\text{LD}_{\text{sw}}(C, a) = \text{LD}_{\text{ca}}(C, a) = \text{LD}_{\text{sw,proj}}(C, a) = r + 1.$$

Proof. For the lower bound, choose a set $A \subset D$ of size $a - 1$ and let $T = D \setminus A$, so $|T| = r + 1$. The displayed range implies $2a - n - 1 \geq k$: indeed, with $R = n - k$ and $r = n - a$, the hypothesis is $r \leq \lfloor R/3 \rfloor$, hence $2a - n - 1 - k = R - 2r - 1 \geq 0$. Choose distinct finite parameters $\theta(t) \in F$, $t \in T$, and define words

$$f|_A = g|_A = 0, \quad f(t) = -\theta(t), \quad g(t) = 1 \quad (t \in T).$$

For $z = \theta(t)$, the word $f + zg$ is zero on $A \cup \{t\}$, hence is explained by the zero codeword on a positions. It is support-wise noncontained: if g were explained on $A \cup \{t\}$ by a degree- $< k$ polynomial, it would vanish on $|A| = a - 1 \geq k$ points and hence be zero, but $g(t) = 1$. In the displayed range the same example is globally far from C^2 at agreement a , so the same slopes are no-loss CA-bad.

For the upper bound, take two bad affine slopes with witnesses on supports S_1, S_2 . Their intersection has size at least $2a - n \geq k$, so the two explaining codewords recover a common code-line (c_f, c_g) on that intersection by solving a two-by-two linear system in the slopes. Let S_0 be the maximal support on which $f = c_f$ and $g = c_g$. Since $|S_0| \geq 2a - n$, every further bad slope agrees with its recovered codeword on at least

$$a + |S_0| - n \geq 3a - 2n \geq k$$

points of S_0 , and therefore must be explained by the same common code line. The remaining agreement must come from residual zeros outside S_0 . No coordinate outside S_0 can be a residual

zero for two different finite slopes, and each noncontained witness needs at least $\max\{1, a - |S_0|\}$ such zeros. The residual charging bound gives at most $r + 1$ bad affine slopes.

For projective slopes, if more than $r + 1$ projective points were bad, either choose a nonbad point as infinity and reduce all bad points to finite affine slopes, or, if every projective point is bad, choose any point as infinity. The remaining finite bad slopes would contradict the affine upper bound. Thus the projective numerator is also $r + 1$. Since CA-bad slopes are support-wise MCA-bad and the lower construction is CA-bad in the displayed range, the CA numerator is exactly $r + 1$ as well. $\square$

Theorem 20.2 (interleaved uniqueness in the high-agreement layer). *Let $C \subseteq F^D$ be an MDS linear code of dimension k and length n . For every interleaving arity $\mu \geq 1$, if $a = n - r$ and*

$$2a - n \geq k, \quad \text{equivalently} \quad r \leq \frac{n - k}{2},$$

then every μ -row received word has at most one interleaved codeword tuple agreeing on a common positions. Thus the maximum interleaved list numerator is exactly 1.

Proof. If two distinct μ -row codeword tuples agree with the same received word on supports of size at least a , then their common supports intersect in at least $2a - n \geq k$ positions. On that intersection every row of the two tuples is equal. Since each row belongs to an MDS code of dimension k , equality on k positions forces equality of that row everywhere. Hence the tuples are equal. A received word that is itself a codeword tuple gives one element, so the maximum is exactly 1. $\square$

For finite-parameter curves the high-agreement upper bound is just as clean, but one must state the lower bound with the correct splitting hypothesis. This avoids confusing a theorem about finite power-curve models with arbitrary protocol curve samplers.

Definition 20.3 (split moving-root set). Let $R \subseteq D$ have size $r + 1$. A polynomial $\phi \in F[T]$ of degree at most d is *d-split over R* if

$$\#\{\gamma \in F : \phi(\gamma) \in R\} = d(r + 1).$$

Theorem 20.4 (degree- d curve envelope and split exactness). *Let $C = \text{RS}[F, D, k]$, let $a = n - r$, and fix $d \geq 1$. In the range*

$$(d + 2)a - (d + 1)n \geq k, \quad \text{equivalently} \quad r \leq \frac{n - k}{d + 2},$$

every finite power curve

$$W_\gamma = f_0 + \gamma f_1 + \cdots + \gamma^d f_d, \quad \gamma \in F,$$

has at most $d(r + 1)$ support-wise curve-MCA-bad parameters at agreement a ; the same upper bound applies to no-loss curve-CA. Conversely, if there is a set $R \subseteq D$ of size $r + 1$ and a degree- $\leq d$ polynomial ϕ that is d -split over R , then the support-wise upper bound is attained:

$$\text{CurveLD}_{\text{sw}}^{(d)}(C, a) = d(r + 1).$$

Moreover, since $r \leq \lfloor (n-k)/(d+2) \rfloor \leq \lfloor (n-k)/3 \rfloor$, the same construction is no-loss curve-CA-bad, so

$$\text{CurveLD}_{\text{ca}}^{(d)}(C, a) = d(r+1).$$

Without the splitting hypothesis, the unconditional high-agreement statement is the envelope

$$r+1 \leq \text{CurveLD}_{\text{sw}}^{(d)}(C, a) \leq d(r+1).$$

The lower bound in this envelope is the ordinary moving-root line witness, viewed as a degree- d power curve by setting the higher curve coefficients to zero.

Proof. The upper bound is a residual-root count. Choose $d+1$ bad parameters and intersect their supports. The intersection has size at least

$$(d+1)a - dn.$$

Interpolating the $d+1$ explaining codewords in the parameter γ recovers a common code-valued curve on that intersection. For any further bad parameter, the witness intersects the common support in at least

$$a + ((d+1)a - dn) - n = (d+2)a - (d+1)n \geq k$$

positions, so MDS rigidity forces the explaining codeword to be the value of the same recovered code curve. The remaining agreement must therefore come from residual coordinate polynomials in γ of degree at most d . No residual coordinate carries the zero polynomial after maximalizing the common support, so each coordinate contributes at most d parameter roots. The same residual budget as in the line case gives at most $d(r+1)$ bad parameters.

For the lower bound under the split hypothesis, choose $A = D \setminus R$ and put $B_A(X) = \prod_{x \in A} (X - x)$. Define

$$W_\gamma = ((X - \phi(\gamma))B_A(X))_{\geq k}.$$

Expanding $\phi(\gamma)$ gives a power curve of degree at most d . Whenever $\phi(\gamma) = t \in R$, the word W_γ is explained on $A \cup \{t\}$, and the tangent-star noncontainment argument applies. In the same high-agreement range the coefficient tuple is globally far from C^{d+1} , yielding no-loss curve-CA badness as well. $\square$

Corollary 20.5 (multiplicative subgroup split criterion). *Let $D = H \subset F^\times$ be a multiplicative subgroup of order n . Suppose $\text{char } F \nmid d$, $d \mid (|F| - 1)$, and*

$$r+1 \leq \frac{n}{\gcd(d, n)}.$$

Then $\phi(\gamma) = \gamma^d$ is d -split over a suitable residual set $R \subset H$, and the exact degree- d curve formula of Theorem 20.4 holds in the displayed high-agreement range.

Proof. The image H^d has size $n/\gcd(d, n)$. Choose $R \subseteq H^d$ of size $r+1$. Since $d \mid (|F| - 1)$ and $\text{char } F \nmid d$, every nonzero element of R has exactly d preimages under $\gamma \mapsto \gamma^d$. $\square$

Corollary 20.6 (finite-grid tangent compiler). *Fix a target $2^{-\lambda}$ and a sampler denominator Q . Define*

$$B(Q, \lambda) = \left\lfloor \frac{Q}{2^\lambda} \right\rfloor.$$

Inside the exact line range $r \leq \lfloor (n - k)/3 \rfloor$, a single line, CA, or projective-line term is safe at target $2^{-\lambda}$ exactly when

$$r + 1 \leq B(Q, \lambda).$$

If an interleaved-list term over the same denominator is also consumed, the exact high-agreement ledger is safe exactly when

$$r + 2 \leq B(Q, \lambda).$$

For a degree- d curve term in its split-exact range, the corresponding conditions are

$$d(r + 1) \leq B(Q, \lambda), \quad d(r + 1) + 1 \leq B(Q, \lambda)$$

without and with one same-denominator list term, respectively. With different denominators, the numerators must not be merged; one must check the rational sum term by term.

Remark 20.7 (effect on the frontier). Theorems 20.1–20.4 solve the small-radius/high-agreement boundary layer. They do not address the near-capacity region $a = k + \sigma$, where $r = n - k - \sigma$ is generally much larger than $(n - k)/3$. The remaining positive frontier is therefore unchanged in substance but sharper in scope: quotient-profile floors must be budgeted, and after quotient/tangent structures are removed one needs finite-field local-limit upper bounds for arbitrary locator fibers and all-line residue packing.

21 Summary, and open problems

The main open problem is the fixed-prime, per-fiber collision problem Problem 18.7. It is the missing ingredient behind a positive all-lines MCA theorem in the polynomial-field window. Other open directions are: the reduced cyclic exterior-algebra nonvanishing problem of Theorem 11.1 below the square-root barrier; sufficiency and sharp constants for the quotient-profile condition beyond the lower bound of Theorem 19.4; the residue-zero classification at $\sigma' \geq 3$ beyond Theorem 16.3; arbitrary bases in non-canonical directions, where the Newton obstruction equations replace the earlier false rigidity shortcut; and the subpolynomial- N failure half, where only density-over-primes statements currently look accessible. Nothing here changes systems operating purely in the Johnson regime. The list-side open problems, and the final two-scale conjectures unifying both sides, are collected in Part IV; Part III first develops transfers, caps, and repairs, several of which reshape these problems directly.

Table 1: Corrected status of slack MCA for $\text{RS}[\mathbb{F}_p, D, \rho n]$ on smooth domains, radius $1 - \rho - t/N$.

regime	behavior	status/source
t dyadic, DSH inequality holds	error 1, every prime	Theorem 13.1
$t = 1$, DSH inequality holds	error 1, every prime	Corollary 10.11
$t = 2$, two-moment inequalities hold	error 1 for the slack-two monomial family: all middle bases c , all slopes z	Theorem 11.2 and Corollary 11.3
free-pool exponential-sum input	error 1	conditional, Theorem 12.2
t non-dyadic, characteristic zero	no nonzero diagonal slopes	Theorem 14.2
zero-base non-dyadic directions	collision-only; most-prime bounds	Theorem 16.2
canonical line, safe slack	no-threshold floor $\asymp (\log n / C_{\text{eff}})^2$; most-prime ceiling $n^{1/C_{\text{eff}}} + O(\log n)$	Theorem 15.1
canonical quotient floor	$\#\text{bad} \geq A(N', \ell') = n^{\beta(\rho)/C_{\text{eff}}(1-o(1))}$ for every $p > (2\ell')^{N'/2}$	Proposition 19.1 and Corollary 19.3
all bases, canonical direction	$\#\text{bad} \leq n^{1/C_{\text{eff}}} + O(\sqrt{\log n})$ for most p	Theorem 15.2
arbitrary base, non-dyadic direction	needs Newton obstruction equations; the former all-bases theorem is not claimed	Theorem 16.5 , Example 16.6
stable range $p > (2s)^{n/2}$	exact for canonical and zero-base strata	Theorem 17.1
canonical stable count	$\frac{1}{2}(3^{n/(2\sigma)} - 1)$ at $\rho = 1/2$	Theorem 17.2
quotient profile, large primes	$\varepsilon_{\text{mca}} \geq A(N^*, \rho N^* + 1)/p$ exactly; $= 2^{(\beta/H_2)\mathcal{Q}_H(\eta)(1-o(1))}/p$ as $N^* \rightarrow \infty$	Theorem 19.4
deployed rates, challenge threshold	$\delta_C^*(2^{-128}) < 1 - \rho - 1/64$ below the per-rate prime-field reach	Proposition 19.7
$q \geq 2^n/\varepsilon^*$	no-slack principle trivially safe	Theorem 9.3
any q , any $\delta < 1 - \rho$	tangent floor $\geq \lfloor \delta n \rfloor / q$	Proposition 9.4
arbitrary lines	reduced to residue-line packing	Theorem 18.3 , Problem 18.7

Part III

Transfers, caps, and repairs

22 A universal field-size cap via list-to-agreement conversion

The imported Crites–Stewart conversion states that correlated agreement with sufficiently small error forces small list sizes [25]; the survey [24] records this as the basic conversion between the two grand challenges. Composing the contrapositive with the quotient-core lists of [Theorem 3.5](#) conditionally caps the achievable mutual-correlated-agreement threshold over *every* field of challenge size, prime or extension, with no value-set counting at all. Throughout, $\varepsilon_{\text{ca}}(C, \delta)$ denotes the correlated-agreement error in the normalization of [Definition 9.2](#): the maximal density of parameters z for which $f + zg$ is δ -close to C while no common agreement set of density $1 - \delta$ explains f and g simultaneously. Every such z is in particular support-wise MCA-bad, so

$$\varepsilon_{\text{mca}}(C, \delta) \geq \varepsilon_{\text{ca}}(C, \delta). \quad (1)$$

Theorem 22.1 (Imported: list-to-agreement conversion; [25], Theorem 2, as restated in [24]). *Let*

$C = \text{RS}[\mathbb{F}_q, D, k]$ with $|D| = n$, and let $\delta \in (0, 1 - \rho]$. If

$$\varepsilon_{\text{ca}}(C, \delta) \leq \eta \left(\frac{1}{k} - \frac{n}{kq} \right) \quad \text{for some } \eta \in (0, 1),$$

then the δ -list size of the relevant augmented code $C^+ \supseteq C$ satisfies $\text{List}(C^+, \delta) \leq \lceil q \varepsilon_{\text{ca}}(C, \delta) / (1 - \eta) \rceil$. We use only the consequence $\text{List}(C, \delta) \leq \text{List}(C^+, \delta)$; see [Remark 22.7](#) for the bookkeeping this import requires.

Lemma 22.2 (Contrapositive at $\eta = 1/2$). *Under the import, if $\text{List}(C, \delta) > q/k + 1$, then $\varepsilon_{\text{ca}}(C, \delta) > \frac{1}{2k} \left(1 - \frac{n}{q} \right)$.*

Proof. If instead $\varepsilon_{\text{ca}} \leq \frac{1}{2k} \left(1 - \frac{n}{q} \right) = \frac{1}{2} \left(\frac{1}{k} - \frac{n}{kq} \right)$, then [Theorem 22.1](#) at $\eta = \frac{1}{2}$ gives $\text{List}(C, \delta) \leq \lceil 2q \varepsilon_{\text{ca}} \rceil \leq \lceil \frac{q}{k} \left(1 - \frac{n}{q} \right) \rceil \leq q/k + 1$. $\square$

Theorem 22.3 (Universal cap, conditional on the import). *Let q be any prime power, let $H \leq \mathbb{F}_q^\times$ be smooth of order $n \leq q/2$, and put $C = \text{RS}[\mathbb{F}_q, H, k]$, where $k = \rho n$ and $\rho = 2^{-a}$. Let $N = 2^c$ with $a \leq c$, assume $N \mid n$, and put $M := n/N \geq 2$. Suppose*

$$\binom{N-1}{\rho N} > \frac{q}{k} + 1.$$

Then, conditional on [Theorem 22.1](#),

$$\varepsilon_{\text{mca}} \left(C, 1 - \rho - \frac{M-1}{n} \right) > \frac{1}{2k} \left(1 - \frac{n}{q} \right) \geq \frac{1}{4k}, \quad \text{hence} \quad \delta_C^* \left(\frac{1}{4k} \right) < 1 - \rho - \frac{1}{2N}.$$

Proof. $M \mid k$ and $k/M = \rho N \leq N - 1$ hold by the divisibility hypotheses, so [Theorem 3.5](#) at quotient order N with $\sigma = M - 1$ gives $\text{List}(C, 1 - \rho - \frac{M-1}{n}) \geq \binom{N-1}{\rho N} > q/k + 1$. [Lemma 22.2](#) converts this into $\varepsilon_{\text{ca}} > \frac{1}{2k} \left(1 - \frac{n}{q} \right)$ at the same radius, and [equation \(1\)](#) transfers it to ε_{mca} ; $n \leq q/2$ gives the clean constant. Finally $\frac{M-1}{n} = \frac{1}{N} \left(1 - \frac{1}{M} \right) \geq \frac{1}{2N}$ and $\varepsilon_{\text{mca}}(C, \cdot)$ is nondecreasing. $\square$

Corollary 22.4 (Two-tier answer at challenge parameters). *Let $q \leq 2^{256}$ be any prime power, $\rho \in \{\frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \frac{1}{16}\}$, and let $H \leq \mathbb{F}_q^\times$ be smooth of order n with $2^{11} \mid n$, $n \leq q/2$, and $k = \rho n \leq 2^{40}$, and put $C = \text{RS}[\mathbb{F}_q, H, k]$. Then, conditional on [Theorem 22.1](#),*

$$\delta_C^*(2^{-128}) \leq \delta_C^*(2^{-42}) < 1 - \rho - 2^{-11}.$$

Combined with [Proposition 19.7](#): below the per-rate reach of roughly $2^{150} - 2^{162}$ bits the challenge threshold is pinned below $1 - \rho - \frac{1}{64}$ unconditionally, over every field up to 256 bits it is pinned below $1 - \rho - 2^{-11}$ conditionally on the import, and below $1 - \rho - 2^{-9}$ (2^{-10} at rate $1/16$) by the sharpened cap of [\[2\]](#).

Proof. Take $N = 2^{10}$ in [Theorem 22.3](#); the hypothesis $2^{11} \mid n$ ensures $M = n/N \geq 2$. Since $k \geq 1$ and $q \leq 2^{256}$, we have $\log_2(q/k + 1) \leq 257$, while $\log_2 \binom{N-1}{\rho N}$ equals 1017.7, 825.2, 551.7, and 341.0 at the four rates, so the binomial hypothesis holds comfortably. The constraint $k \leq 2^{40}$ enters only at the end: the cap gives $\varepsilon_{\text{mca}} \geq \frac{1}{4k} \geq 2^{-42} > 2^{-128}$ at gap at least 2^{-11} . $\square$

Remark 22.5 (Relation to the companion cap paper). The composition above is extracted, sharpened, and made the main theorem of [2]: running the pigeonhole on slack-two locator fibers over the field of definition, rather than on quotient cores, improves the gap to 2^{-9} (2^{-10} at rate $1/16$) under $2^{10} \mid n$ (resp. $2^{11} \mid n$), yields error $\frac{1}{2k}(1 - n/q) \geq 2^{-86}$ with no condition relating n and q , and covers extension fields $\mathbb{F} \supseteq B \supseteq H$ directly through the subfield pigeonhole. In its v5 form, that companion cap is self-contained for the main MCA statement. [Theorem 22.3](#) is retained here because the quotient-core mechanism is the one native to this paper’s framework and the proof is self-contained given Part I plus the imported conversion; for citation purposes, the cap of [2] supersedes [Corollary 22.4](#).

Remark 22.6 (The pinch on the list–agreement equivalence). The survey [24] discusses the hoped-for equivalence “lists small $\Leftrightarrow$ correlated agreement small” in the regime $n \ll q$. On smooth domains it is now pinched from both sides: [Theorem 22.1](#) converts small agreement error into small lists, while the dyadic pigeonhole of [Theorem 3.5](#) produces lists of size $2^{\Theta(N)}$ at gap $\Theta(1/N)$ for every divisor scale, which by the same conversion *forces* agreement failure. Any positive MCA theory above the corrected reserve must therefore live on domains where the divisor pigeonhole collapses — which is precisely the structural definition of non-smoothness this paper has been quantifying.

Remark 22.7 (Due diligence on the import). [Theorem 22.1](#) is used as a black box in the older quotient-core route retained in this section, and three items must be checked against [25] before that route is cited as unconditional: (i) the exact admissible range of δ in their Theorem 2, and the definition of the augmented code C^+ together with the monotonicity $\text{List}(C, \delta) \leq \text{List}(C^+, \delta)$; (ii) the normalization of ε_{ca} (density over z , including extension-field sampling of z in the sense of the proximity-gaps literature); (iii) the constants in the displayed implication. Should any constant need repair, the slacked conversion of Ben-Sasson–Carmon–Haböck–Kopparty–Saraf [26], recorded as Theorem 5.2 of [24], yields the same corollary with 2^{-42} replaced by a slightly smaller explicit constant and 2^{-11} by a slightly smaller gap. This import audit is no longer a caveat on the main companion cap: Paper D v5 proves the main MCA universal cap directly.

23 Subfield confinement over extension fields

Deployed instantiations sample the slope from an extension field $\mathbb{F} = B^d$ of the base field B containing the domain. The following confinement theorem shows the entire prime-field witness theory deflates by a factor $|B|/|\mathbb{F}|$ in that setting, while the list-side obstructions persist — so the two grand challenges genuinely decouple over extensions. A two-radius refinement of [Theorem 23.1](#), covering the proximity-loss form of correlated agreement, is proved in [2, Lem. “subfield confinement”].

Theorem 23.1 (Subfield confinement). *Let $B \leq \mathbb{F}$ be a subfield, $D \subseteq B$, $C = \text{RS}[\mathbb{F}, D, k]$, and let $f, g \in B^D$ be B -valued words. Then every support-wise MCA-bad parameter of the line $f + zg$, at every radius, lies in B .*

Proof. Fix a B -basis $e_1 = 1, e_2, \dots, e_d$ of $\mathbb{F}$ and the B -linear coordinate projections $\pi_i : \mathbb{F} \rightarrow B$. For $P \in \mathbb{F}[X]_{<k}$ and $x \in D \subseteq B$ we have $\pi_i(P(x)) = (\pi_i P)(x)$, where $\pi_i P$ applies π_i to coefficients;

each $\pi_i P$ is a codeword. Let $z \notin B$, so $z = \sum_i z_i e_i$ with $z_{i_0} \neq 0$ for some $i_0 \geq 2$, and let S be any set on which $f + zg$ agrees with a codeword P . Comparing coordinates on S : the i_0 -th coordinate of $(f + zg)(x) = f(x) + zg(x)$ is $z_{i_0} g(x)$, so $g|_S = (z_{i_0}^{-1} \pi_{i_0} P)|_S$, and the first coordinate is $f(x) + z_1 g(x)$, so $f|_S = (\pi_1 P - z_1 z_{i_0}^{-1} \pi_{i_0} P)|_S$. Both explaining polynomials have degree $< k$, so f and g are simultaneously explained on S itself, and z fails condition (ii) of [Definition 9.2](#) for every candidate S . $\square$

Corollary 23.2 (Deflation of the prime-field witness theory). *In the sextic-extension instantiation discussed in [24] ($\mathbb{F} = B^6$ over a 31-bit base field, $D \subseteq B$), every line with B -valued f, g — in particular every monomial line and every witness line constructed in Parts I–II — contributes MCA error at most $|B|/|\mathbb{F}| = p^{-5} < 2^{-154} < 2^{-128}$.*

Corollary 23.3 (Decoupling). *List obstructions do not deflate: $C_B := \text{RS}[B, D, k] \subseteq C$ and a B -valued received word has C -list at least its C_B -list, so the quotient-core lists of [Theorem 3.5](#) persist verbatim over $\mathbb{F}$. Consequently, at deployed extension parameters, lists stay superpolynomial while every MCA attack below 2^{-128} must use genuinely $\mathbb{F}$ -valued line data — and such data exists nonconstructively, since [Theorem 22.3](#) applies to $q = |\mathbb{F}| \leq 2^{256}$ directly; sharper, the deployed-parameter corollary of [2] certifies $\varepsilon_{\text{ca}} > 2^{-22}$ at gap 2^{-7} on the sextic itself.*

Remark 23.4 (Where the explicit $\mathbb{F}$ -valued attack should live). By [Theorem 23.1](#) the residue-line normal form of [Theorem 18.3](#) is the natural source: denominators $E \in \mathbb{F}[X]$ not defined over B evade confinement, and the quotient locator calculus applies to them unchanged. Making the $\mathbb{F}$ -valued bad lines explicit over the deployed sextic extension is the concrete open target: by [2] such lines exist nonconstructively already at gap 2^{-7} with bad-slope density $> 2^{-22}$, and exhibiting them is posed there as the explicit-lines problem.

24 Chebyshev domains and circle codes

The locator mechanism uses only two facts about $x \mapsto x^M$: the domain is a union of complete fibers, and fiber locators are value-translates of a single degree- M polynomial. Both transfer to Dickson–Chebyshev quotient chains, hence to the x -coordinate domains of circle groups used by circle STARKs [33], and in principle to the isogeny chains of ECFFT [34].

Lemma 24.1 (Dickson quotient chain). *Let p be odd, $N \mid p+1$ even, and $X_N := \{\zeta + \zeta^{-1} : \zeta \in \mu_N(\mathbb{F}_{p^2})\} \subseteq \mathbb{F}_p$, of size $N/2 + 1$. For $m \mid N$ let D_m be the degree- m monic Dickson polynomial, $D_m(\zeta + \zeta^{-1}) = \zeta^m + \zeta^{-m}$. Then $D_m(X_N) = X_{N/m}$, and for every $w \in X_{N/m} \setminus \{\pm 2\}$ the fiber $D_m^{-1}(w) \cap X_N$ has exactly m elements and locator polynomial exactly $D_m(X) - w$.*

Proof. $D_m(\zeta + \zeta^{-1}) = \xi + \xi^{-1}$ iff $\zeta^m \in \{\xi, \xi^{-1}\}$; for $\xi \neq \xi^{-1}$ this gives $2m$ values of $\zeta \in \mu_N$ pairing into m distinct values of $\zeta + \zeta^{-1}$, all roots of the degree- m polynomial $D_m(X) - w$, which is therefore their monic locator. $\square$

Theorem 24.2 (Quotient cores for Chebyshev domains). *Let $C = \text{RS}[\mathbb{F}_p, X_N, k]$ with $m \mid k$, $\kappa = k/m$, $\kappa \leq \frac{N}{2m} - 3$, and $1 \leq \sigma < m$. Then*

$$\text{List} \left(C, 1 - \frac{k + \sigma}{N/2 + 1} \right) \geq \binom{N/(2m) - 2}{\kappa}.$$

Proof. Fix $w_0 \in X_{N/m} \setminus \{\pm 2\}$ with fiber F_0 , choose $T \subset F_0$ with $|T| = \sigma$, and set $Y := D_m(X)^\kappa L_T(X)$ with $L_T = \prod_{t \in T} (X - t)$. For each $A \subseteq X_{N/m} \setminus \{\pm 2, w_0\}$ of size κ , the union of fibers over A has locator $L_A = \prod_{w \in A} (D_m(X) - w)$ by Lemma 24.1, and as a polynomial in D_m ,

$$D_m^\kappa - L_A = e_1(A)D_m^{\kappa-1} - e_2(A)D_m^{\kappa-2} + \dots,$$

of degree at most $m(\kappa - 1) = k - m$. Thus $P_A := L_T \cdot (D_m^\kappa - L_A)$ has degree at most $\sigma + k - m < k$, hence is a codeword, and $Y - P_A = L_T L_A$ vanishes on $T \cup D_m^{-1}(A)$, an agreement set of size $\sigma + m\kappa = k + \sigma$. Distinct A give distinct P_A . (The pullback form of Y is essential: unlike $X^M - \alpha$, the polynomial D_m has internal terms, so $X^k - L_A$ would only have degree $k - 2$.) $\square$

Proposition 24.3 (Slack-one canonical stratum on a Chebyshev quotient). *With the notation of Theorem 24.2, consider the line with $f = D_m^{\kappa+1}$ and $g = D_m^\kappa$ as words on X_N , and let $X' = X_{N/m} \setminus \{\pm 2\}$, $N' = |X'| = \frac{N}{2m} - 1$. For every $A \subseteq X'$ with $|A| = \kappa + 1 \leq N'$, the parameter $z = -e_1(A)$ is support-wise MCA-bad at radius $1 - \frac{k+m}{N/2+1}$. Consequently, by Dias da Silva–Hamidoune [8],*

$$\varepsilon_{\text{mca}} \left(C, 1 - \frac{k + m}{N/2 + 1} \right) \geq \frac{\min\{p, (\kappa + 1)(N' - \kappa - 1) + 1\}}{p}.$$

Proof. Let $S = D_m^{-1}(A) \cap X_N$, of size $m(\kappa + 1) = k + m$. In the variable $u = D_m$, the polynomial $u^{\kappa+1} - e_1(A)u^\kappa - \prod_{w \in A} (u - w)$ has degree at most $\kappa - 1$, so $P := D_m^{\kappa+1} - e_1(A)D_m^\kappa - L_A$ is a codeword, and $f + zg - P = L_A$ vanishes on S : the line is S -close. But $g = D_m^\kappa$ cannot be explained on S : for any codeword c_g , the difference $D_m^\kappa - c_g$ is monic of degree exactly k and would need $k + m > k$ roots. So condition (ii) of Definition 9.2 holds at S , and $z = -e_1(A)$ is bad. The count is the restricted sumset bound applied to the arbitrary set $X' \subset \mathbb{F}_p$. $\square$

Corollary 24.4 (Failure on circle x -coordinate domains at Mersenne-31 parameters). *Let $p = 2^{31} - 1$, so the circle group has order $p + 1 = 2^{31}$ and every X_N with $N \mid 2^{31}$ is available. For $N \geq 2^{19}$, $m = N/2^{18}$, and $k = \rho(N/2 + 1) + O(1)$ at $\rho = 1/2$ with $m \mid k$:*

$$(\kappa + 1)(N' - \kappa - 1) + 1 \approx \frac{1}{4}(2^{17})^2 = 2^{32} > p,$$

so the slack-one canonical stratum already has error one at radius approximately $1 - \rho - 2^{-17}$. The up-to-capacity hope thus fails for the Reed–Solomon constituents on circle-group x -coordinate domains exactly as on multiplicative smooth domains.

Remark 24.5 (Circle codes proper, and the general isogeny transfer). Circle codes decompose as $f(x, y) = f_0(x) + yf_1(x)$ with f_0, f_1 supported on x -coordinate domains [33]; the stratum above lives in the even part, and carrying the full bivariate MCA bookkeeping through this decomposition is routine but deferred. More generally, the proofs above used only that the domain is a union of complete fibers of a chain of low-degree maps whose fiber locators are value-translates of one polynomial: multiplicative powers, Dickson–Chebyshev maps, linearized polynomials, and the 2-isogeny chains of ECFFT [34] all qualify. This gives structural content to the degenerate-code characterization requested in [24]: degeneracy is a large cover structure compatible with the degree filtration.

25 Torsion-coset repairs: $\{2, 3\}$ -smooth rigidity and arbitrary-base templates

We isolate the external input in exactly the form used, matching the pattern of [Theorems 22.1](#) and [25.4](#).

Theorem 25.1 (Imported: vanishing sums over two primes; [38, 5, 6]). *Let n have prime support contained in $\{2, 3\}$. (i) The $\mathbb{Z}$ -module of solutions of $\sum_{y \in \mu_n} h(y)y = 0$ is generated by the rotations of the prime relations R_p for the primes $p \mid n$ [38]. (ii) Every minimal vanishing sum of n -th roots of unity with nonnegative integer coefficients is a rotation of R_p for some prime $p \mid n$ [5, 6].*

Theorem 25.2 ($\{2, 3\}$ -smooth slack-one rigidity, conditional on the import). *Let $n = 2^a 3^b$ with $a \geq 1$, and when $b \geq 1$ let $\omega \in \mu_n$ have order 3. A function $h : \mu_n \rightarrow \mathbb{Z}$ satisfies $\sum_y h(y)y = 0$ if and only if h is a $\mathbb{Z}$ -combination of the pair functions $\delta_\zeta + \delta_{-\zeta}$ and, when $b \geq 1$, the triangle functions $\delta_\zeta + \delta_{\zeta\omega} + \delta_{\zeta\omega^2}$. Moreover, for sets $S, T \subseteq \mu_n$: $e_1(S) = e_1(T)$ if and only if the multiset $S \sqcup (-T)$ is an $\mathbb{N}$ -linear combination of rotated pair relations $\{\zeta, -\zeta\}$ and, when $b \geq 1$, rotated triangle relations $\{\zeta, \zeta\omega, \zeta\omega^2\}$.*

Proof. The displayed generators visibly lie in the kernel of $h \mapsto \sum_y h(y)y$; conversely the kernel is generated by them, by part (i) of [Theorem 25.1](#), since for $n = 2^a 3^b$ the rotated prime relations are exactly the pair functions and, when $b \geq 1$, the triangle functions. For the set version, suppose $e_1(S) = e_1(T)$. Since n is even, $-T \subseteq \mu_n$, and $\sum_{y \in S} y + \sum_{y \in T} (-y) = 0$ exhibits the multiset $S \sqcup (-T)$ — multiplicities at most two, attained exactly when $y \in S$ and $-y \in T$ — as a vanishing sum of n -th roots with nonnegative coefficients. Every such sum decomposes as a sum of minimal vanishing sums with nonnegative coefficients, by induction on the total multiplicity: any nonempty vanishing sum contains a vanishing subsum of minimal cardinality, which is itself minimal, and removing it leaves a shorter vanishing sum. By part (ii) of [Theorem 25.1](#), each minimal summand is a rotation of R_p for a prime $p \mid n$: a pair, or a triangle when $b \geq 1$. This is the asserted $\mathbb{N}$ -combination. The converse is immediate, since every rotated pair and triangle sums to zero. At $b = 0$ only the pair relations occur, recovering the antipodal classification underlying [Theorem 4.2](#). $\square$

Remark 25.3 (Mixed-radix counts and domains). [Theorem 25.2](#) supplies the characteristic-zero relation-module and set-decomposition input for the $\{2, 3\}$ -smooth (mixed-radix FFT) extension of the slack-one theory. The class invariant should consist of a signed pair profile together with a

triangle profile, and the analogue of $A(N', \ell')$ should be a two-parameter sum whose exponent is a constrained entropy maximization over pair and triangle densities. The norm sieve itself transfers unchanged once the characteristic-zero classes are fixed, because the archimedean bounds never used 2-smoothness. The exact class enumeration and the corresponding two-parameter analogue of $\beta(\rho)$ are left as future combinatorics rather than claimed here; in particular, no canonical positive move normal form for set-to-set transformations is asserted.

We isolate the external input in exactly the form used.

Theorem 25.4 (Imported: degree-only torsion-coset bound; [35], with [7, 36, 37]). *For all $s, d \geq 1$ there is an explicit constant $C_{tc}(s, d)$ such that for every subvariety $V \subseteq \mathbb{G}_m^s$ cut out by polynomials of degree at most d , the torsion points of V lie on at most $C_{tc}(s, d)$ maximal torsion cosets contained in V , independently of the coefficients of the defining polynomials.*

As with Theorem 22.1, this statement is imported as a black box in exactly the form used, and Theorem 25.5 below is conditional on it in the same sense.

Theorem 25.5 (Torsion-coset template structure for arbitrary-base slabs, conditional on the import). *Fix $s \geq 1$, $\sigma \geq 1$. There is a constant $C(\sigma, s)$, independent of n and of the pinned values, with the following property. For any $c_1, \dots, c_\sigma \in \overline{\mathbb{Q}}$, let $V_c = \{x \in \mathbb{G}_m^s : p_i(x) = c_i, 1 \leq i \leq \sigma\}$ be the slab variety in the power sums. Then the torsion points of V_c lie on at most $C(\sigma, s)$ maximal torsion cosets $u\mathcal{H} \subseteq V_c$. Along each positive-dimensional coset all pinned power sums are constant (a padding family); the zero-dimensional cosets are at most $C(\sigma, s)$ isolated templates. In particular, for every n and every base, the characteristic-zero witness set of a σ -slab on s letters decomposes into at most $C(\sigma, s)$ classes.*

Proof. Witnesses are torsion points of $\mathbb{G}_m^s$ on V_c . Finiteness of the maximal torsion cosets is Laurent’s theorem [7]; the uniform count is Theorem 25.4 with $d = \sigma$, since the defining degrees $1, \dots, \sigma$ depend on neither c nor n . This gives $C(\sigma, s) = C_{tc}(s, \sigma)$. Constancy of the p_i along a coset contained in V_c holds by definition of V_c . $\square$

Corollary 25.6 (Finiteness of the residue-zero classification). *For each fixed (σ', s') the variety $\{e_2 = \dots = e_{\sigma'} = 0\}$ has fixed coefficients, so Theorem 25.5 reduces its torsion points to a finite, effectively computable list of coset families, uniformly in n . Positive-dimensional cosets may still contain padding families of unbounded order; once those families are enumerated, the remaining isolated templates form a finite set. Thus the residue-zero classification at $\sigma' \geq 3$ left open after Theorem 16.3 becomes a finite template computation for each fixed size.*

Remark 25.7 (The principled replacement for the retracted all-bases theorem). Theorem 25.5 is the structure that the retracted “all bases, all non-dyadic directions” claim was missing: every characteristic-zero witness of an arbitrary-base slab is “one of at most $C(\sigma, s)$ templates, padded by vanishing-sum families,” and the Newton-obstruction criterion of Theorem 16.5 determines which templates survive a given direction. Because the bases themselves are n -dependent torsion data,

isolated templates can have unbounded order, so the finite-field transfer is per-class: the stable-range and density norm sieves of Part II apply to each of the boundedly many classes separately. Turning this into a clean every-prime arbitrary-base theorem would require carrying out that finite-template verification and the corresponding per-class norm-sieve bookkeeping; it is no longer a single rigidity shortcut.

26 Additive domains, dilation invariance, and further transfers

Proposition 26.1 (Additive confinement of the canonical locator stratum). *Let $q = 2^e$, let $V \leq \mathbb{F}_q$ be an $\mathbb{F}_2$ -subspace, $D = a + V$ with $|D| = n$, and $C = \text{RS}[\mathbb{F}_q, D, k]$. The slack-one locator stratum of the canonical line — the bad slopes $\{e_1(A) : A \in \binom{D}{k+1}\}$ produced by full agreement sets — is contained in the single additive coset $(k+1)a + V$, hence has size at most n and density at most n/q .*

Proof. $e_1(A) = \sum_i (a + v_i) = (k+1)a + \sum_i v_i \in (k+1)a + V$, since V is closed under addition. $\square$

Remark 26.2. Restricted sums inside an $\mathbb{F}_2$ -affine space cannot expand, so the smooth-domain catastrophe is a multiplicative phenomenon: on binary-tower domains the canonical locator stratum is confined to the tangent-floor scale n/q , versus error one multiplicatively. This is a statement about the locator stratum only; other failure mechanisms in characteristic two are not excluded, so the separation is phrased stratum-wise.

Lemma 26.3 (Dilation invariance of restricted subgroup sums). *Let $Q \leq \mathbb{F}_p^\times$ and $1 \leq \ell \leq |Q|$. Then $u \cdot (\ell^\wedge Q) = \ell^\wedge Q$ for every $u \in Q$: the restricted sumset is a union of multiplicative Q -cosets, together possibly with $\{0\}$.*

Proof. $u \cdot \sum_{x \in A} x = \sum_{x \in uA} x$ and $A \mapsto uA$ permutes $\binom{Q}{\ell}$. $\square$

Remark 26.4 (Why the V13 deficits are multiples of 16). At $(N', \ell') = (16, 9)$ the image of e_1 is a union of order-16 cosets (plus possibly 0) by [Lemma 26.3](#), so every deficit from the characteristic-zero plateau is a multiple of 16 — exactly the pattern recorded in item V13 of [appendix A](#). The same invariance shows that energy-increment arguments must work coset-wise, which is one structural reason unconditional technology below $N \sim \sqrt{p}$ cannot currently beat the Dias da Silva–Hamidoune bound: few-fold growth estimates for subgroup sumsets [\[14\]](#) act on the coset count, which dilation already fixes.

Remark 26.5 (Polynomial generators, folding, and interleaving). Three transfer notes toward the questions of [\[24\]](#). (i) *Polynomial generators*: by [Corollary 11.3](#), any generator family whose bad-parameter event contains the slack-two monomial stratum inherits error one at every prime below the root barrier, so on smooth domains in that range polynomial generators [\[31\]](#) cannot improve the threshold beyond what the distance-preservation reductions already give. (ii) *Folded Reed–Solomon*: running the locator through an s -fold packet $\prod_{b \in A} \prod_{0 \leq i < s} (X^a - b \omega^{ai})$ leaves the agreement structure intact, contributes the slope through $e_1(A)$ scaled by $\sum_i \omega^{-ai}$, and pins the intermediate coefficients — exactly a slab condition at $\sigma = s - 1$. The fullness theory of Part II then predicts that small folding parameters cannot escape the two-moment barrier, while admissible

large- s folding [30] acts precisely as an anti-stabilizer device; making the $s \in \{2, 3\}$ case unconditional is a concrete short-term target. (iii) *Interleaving*: lower bounds transfer to interleaved codes by equal components, and the residue-line normal form is the natural tool for deciding whether the noncontained packing number of $C^{\equiv s}$ exceeds that of C at slack radii, extending [32] beyond unique decoding.

Part IV

Integrated frontier and synthesis

27 Integrated frontier and final conjectural asymptotics

The list and MCA sides now have the same two-scale shape. Let $s = k + \sigma$ and assume the generated-field condition unless a field-of-definition parameter is explicitly inserted. The proved lower-bound side says that polynomial behavior cannot begin before

$$\sigma \gtrsim \max \left\{ \frac{\log_2 \binom{n}{k+\sigma}}{\log_2 q}, \frac{n}{\log n} \right\}.$$

The first term is the ambient entropy reserve and the second is the quotient-core reserve. For a prescribed polynomial list exponent B , the quotient reserve has the sharper form

$$\sigma \geq \frac{(H_2(\rho) + o(1))n}{B \log_2 n}$$

for the explicit quotient-core family to be at most n^B .

On the MCA side, the same two-scale shape is now matched by proved statements in Part II: the quotient-exact floor (Proposition 19.1) shows the closed-form quotient count survives at every prime above the quotient norm threshold, the necessity theorem (Theorem 19.4) shows the quotient profile of Definition 3.8 is genuinely necessary for polynomial MCA slope bounds at large primes, and the deployed-rate challenge cap (Proposition 19.7) translates the floor into concrete prime-field parameter statements.

The proved positive side currently reaches two important but not final regimes. In characteristic zero, prefix fibers are exactly quotient-periodic and therefore polynomial at $\sigma \geq Cn/\log n$. In prime fields $p \equiv 1 \pmod{n}$, the Galois-amplified no-collision theorem transfers this to finite fields once

$$p > \exp \left(C \frac{n \log n}{\sigma} \right).$$

At the corrected reserve $\sigma \asymp n/\log n$, this is the quasi-polynomial split-prime range $p > \exp(O((\log n)^2))$.

The expected final list theorem is the following local-limit statement, which strengthens the monomial-prefix conjecture to arbitrary received words.

Conjecture 27.1 (Final locator local limit). *For every fixed $\rho \in (0, 1)$ and every $B, \varepsilon > 0$, there is $C_{\rho, B, \varepsilon}$ such that for every generated-field smooth domain $H \leq \mathbb{F}_q^\times$ of order $n = 2^m$ (so $q = q_D$), every $k = \rho n + O(1)$, and every $\deg U < n$,*

$$|\mathcal{F}_U(k + \sigma)| \leq n^B$$

whenever

$$\sigma \geq C_{\rho, B, \varepsilon} \frac{n}{\log n} \quad \text{and} \quad \sigma \log_2 q \geq (1 + \varepsilon) \log_2 \binom{n}{k + \sigma}.$$

Equivalently, after quotient-periodic classes are removed, the map $S \mapsto U \bmod L_S$ behaves like a random codimension- σ hash on $(k + \sigma)$ -subsets.

The expected final MCA theorem is the residue-line analogue.

Conjecture 27.2 (Final floor- and quotient-corrected MCA asymptotic). *For every fixed $\rho \in (0, 1)$ and $\varepsilon > 0$, there is $C_0 = C_0(\rho, \varepsilon)$ such that the following holds for every fixed reserve constant $C \geq C_0$. Let $H_n \leq \mathbb{F}_{q_n}^\times$ be generated-field smooth domains of order n (so $q_n = q_{D,n}$), with $k_n = \rho n + O(1)$ and $\text{poly}(n) \leq q_n \leq 2^{o(n)}$. Assume the corrected reserve*

$$\sigma_n \geq C \frac{n}{\log_2 n}, \quad \sigma_n \log_2 q_n \geq (1 + \varepsilon) \log_2 \binom{n}{k_n + \sigma_n},$$

and put $\eta_n = \sigma_n/n$ and $C_n = \text{RS}[\mathbb{F}_{q_n}, H_n, k_n]$. Then, with the unavoidable tangent floor and the separated quotient-periodic floor included on the right,

$$\varepsilon_{\text{mca}}(C_n, 1 - \rho - \eta_n) \leq \frac{n^{1+o(1)} + 2^{(\beta(\rho)/H_2(\rho))\mathcal{Q}_{H_n}(\eta_n)(1+o(1))}}{q_n}.$$

In packing form ([Remark 18.9](#)), the conjectural aperiodic noncontained residue-line packing has size $n^{1+o(1)}$, while the separated quotient-periodic contribution is bounded at the displayed quotient-profile scale. In particular, if

$$\mathcal{Q}_{H_n}(\eta_n) \leq \frac{H_2(\rho)}{\beta(\rho)} \log_2 n + o(\log n)$$

— automatic whenever the reserve constant satisfies $C > \beta(\rho)$, since every qualifying divisor has $n/M < 1/\eta_n \leq (\log_2 n)/C$ and hence $\mathcal{Q}_{H_n}(\eta_n) \leq (H_2(\rho)/C) \log_2 n (1 + o(1))$ — then the raw prediction takes the clean form

$$\varepsilon_{\text{mca}}(C_n, 1 - \rho - \eta_n) \leq \frac{n^{1+o(1)}}{q_n}.$$

This is the strongest raw bound compatible with [Theorem 19.4](#).

These conjectures are deliberately stated as finite combinatorial-algebraic assertions. They can be attacked through BCH coset enumerators [\[21\]](#), subgroup Hayes-equivalence estimates [\[19, 20\]](#), finite-field harmonic analysis on the prefix map and subgroup subset-sum asymptotics [\[17, 18\]](#), or inverse Littlewood–Offord style template counting. The results proved in this manuscript remove the characteristic-zero ambiguity and reduce the remaining problem to finite-field local limits.

28 Conclusion

The final form of the program is two-scale and floor-corrected. The list-side negative mechanisms are settled by entropy pigeonholes and quotient cores. The characteristic-zero structure is settled by the inverse quotient theorem. The quasi-polynomial split-prime monomial-prefix positive theorem follows from Galois-amplified collision divisibility. The MCA canonical-line theory is exact at every slack, and the all-line theory reduces to residue-line packing — under the corrected bookkeeping of Part II: the free-pool ladder is conditional on its exponential-sum input, the former all-bases non-dyadic claim is replaced by the Newton-obstruction criterion, and the quotient-exact floor and deployed-rate challenge cap are the new unconditional positive-side anchors. Part III extends these mechanisms beyond the smooth multiplicative prime-field setting — subfield confinement, additive confinement, Chebyshev and circle x -coordinate transfer, mixed-radix relation modules, and torsion-coset templates — and, conditional on one imported conversion theorem, caps the challenge threshold over every field of challenge size.

What is not proved here is equally important: polynomial-field arbitrary-word list decoding and all-line MCA above the corrected reserve require new finite-field local-limit estimates. The paper isolates those estimates as Conjectures 27.1 and 27.2; proving them would close the corrected entropy-gap program for smooth-domain Reed–Solomon codes.

A Verification record

The following computational checks are retained as a verification record. They should be accompanied by reproducible scripts before submission; items depending on the conditional free-pool theorem are evidence for the stated constants, not independent proofs of the exponential-sum input.

- V1. *Exact characterization; strict subfamily.* Theorem 10.2 confirmed exhaustively at small parameters. At $p = 17$, $n = 16$, $\rho = 1/2$, $T = 2$: the quotient family covers 16 of 17 values; an asymmetric witness supplies $z = 0$.
- V2. *Two-moment fullness.* $(p, N, h, r) = (31, 30, 2, 3)$, $\ell = 8$: the construction of Theorem 11.2 built witnesses for all 961 targets $(a, b) \in \mathbb{F}_{31}^2$ with zero failures, and exhaustive enumeration of all $\binom{30}{8} = 5,852,925$ subsets confirms Ψ_2 is onto $\mathbb{F}_{31}^2$.
- V3. *Rigidity lattices.* $\{\widehat{h}(j) = 0, j \leq \sigma\} = \{\mu_{M_0}\text{-invariant}\}$ as $\mathbb{Z}$ -lattices at $n = 16$ ($\sigma = 2, 3$) and $n = 32$ ($\sigma = 3, 4, 5, 7, 8$; $M_0 = 4, 8, 8, 8, 16$), with optimality witnesses.
- V4. *Descent.* $p = 193$, $N = 64$, $t = 2$, $\ell = 34$: the identities $e_1(A) = e_3(A) = 0$, $e_2(A) = -e_1(B)$, and coverage of all 193 residues by e_1 over the $\binom{32}{17} = 565,722,720$ subsets of Q^2 .
- V5. *Dichotomy.* $n = 16$, $t = 3$, $\ell = 8$: all six characteristic-zero zero-prefix subsets are μ_4 -orbit unions, all with $e_3 = 0$.
- V6. *Canonical line ladder* ($n = 16$, $k = 8$, $\sigma = 2$, $s = 10$; descent set has 40 elements at large p):

p	17	97	257	769	1153	2129	4001	7489	15217	65089	524353
bad slopes	17	56	40	40	40	40	40	40	40	40	40
excess over descent	1	16	0	0	0	0	0	0	0	0	0

V7. *All bases, canonical direction.* Maximum slope count over all nonzero canonical-direction bases collapses $56 \rightarrow 14$, stable from $p = 7489$, below the class cap $\binom{8}{5} = 56$.

V8. *Direction law* (direction 1, $\sigma = 2$; bad-slope counts over the prime ladder of V5):

s	17	97	257	769	1153	2129	4001	7489	15217	65089	524353
9 ($\equiv 1 \pmod{4}$)	—	—	48	48	48	48	48	48	48	48	48
10 ($\equiv 2$)	17	16	64	0	0	0	0	16	0	0	0
11 ($\equiv 3$)	—	—	0	0	0	0	0	0	0	0	0

Explicit witnesses $e_2(\{y\} \sqcup \text{two } \mu_4\text{-orbits}) = 0$ and $e_2(\{1, i, \zeta_8, -\zeta_8\}) = 0$, verified mod 257, 1153, 4001 and structurally over $\mathbb{Z}[\zeta]$.

V9. $\sigma' = 3$, *residue zero*. Size 4: empty in exact $\mathbb{Z}[\zeta_N]$ arithmetic for $N \in \{16, 32\}$ (and for all N by [Theorem 16.3](#)). Sizes 8 and 12 at $N = 16$: empty in exact arithmetic. Size 8 at $N = 32$: empty modulo $p = 65089$, hence in characteristic zero; the 200 resp. 72 witnesses visible mod 97 and 193 are collisions, all vanishing at 65089.

V10. *Non-dyadic slabs*. $\sigma = 4$, $j = 3$, $n = 16$, $s = 12$: all 1820 nonempty punctured slabs carry exactly one e_3 -value, at $p = 65089$ and 524,353. This is numerical evidence; the general all-bases theorem is not claimed because of the Newton obstruction in [Example 16.6](#).

V11. *Stable range and exact counts*. $(n, s, \sigma) = (16, 9, 2)$: collision count 0 and fibers equal to their rearrangement classes exactly, from $p = 769$, re-verified by direct enumeration at $p = 15217$, 65089, 524,353. Exact counts: $A(8, 5) = 40 = \frac{1}{2}(3^4 - 1)$, $A(8, 3) = 40$, $A(4, 3) = 4 = \frac{1}{2}(3^2 - 1)$, each matching the enumerated bad-slope count at $p = 65089$ and 524,353, for $(n, \rho, \sigma) = (16, \frac{1}{2}, 2)$, $(16, \frac{1}{4}, 2)$, $(16, \frac{1}{2}, 4)$ respectively.

V12. *Pool constants*. The exponent recursion $\eta_r(\varepsilon)$ of [\[13\]](#) and all pool constants used in the conditional [Theorem 12.2](#) were recomputed from the published recursion.

V13. *Slack-one quotient channel*. For $(N', \ell', \rho) = (16, 9, 1/2)$, the image of e_1 on $\binom{Q}{9}$ was enumerated over 186 primes $p \equiv 1 \pmod{16}$, $17 \leq p \leq 524,353$. The image is p or $p - 1$ for every $p \leq 1153$ and sporadically to 1697; then it enters a collision band; it first reaches the characteristic-zero plateau $3280 = A(16, 9) = \frac{1}{2}(3^8 - 1)$ at $p = 5281$, equals 3280 for every tested $p \geq 67,057$, and has sporadic sub-plateau values such as 3264 at $p = 58,337$ and 3232 at $p = 66,593$. Every observed deficit is a multiple of 16.

V14. *Slack-two frontier*. For $(N, \ell, \rho) = (32, 18, 1/2)$, a meet-in-the-middle enumeration of Ψ_2 on $\binom{Q}{18}$ over $\mathbb{F}_p^2$ was used for primes $p \equiv 1 \pmod{32}$. The map is onto $\mathbb{F}_p^2$ for every tested prime $p \leq 3617$; the first tested failure occurs at $p = 4001$ with 736 missing targets out of 4001^2 . Missing targets occur in scaling orbits $(a, b) \mapsto (ua, u^2b)$ and are strongly enriched on the rigidity diagonal $e_1 = 0$. A direct run of the companion meet-in-the-middle script confirmed onto-ness at $p = 97$ in this session.

V15. *Independent replications*. The two-moment fullness check at $(p, N, \ell) = (31, 30, 8)$, the six characteristic-zero zero-prefix subsets at $(n, t, \ell) = (16, 3, 8)$, and the canonical ladder 17, 56, 40, 40, $\dots$ at $(n, \sigma) = (16, 2)$ were re-derived from independent code and match the earlier verification record.

References

- [1] P. Chojek, *Capacity-edge obstructions to Reed–Solomon mutual correlated agreement over smooth multiplicative domains*, companion manuscript, June 2026.
- [2] P. Chojek, *A universal field-size cap for mutual correlated agreement on smooth Reed–Solomon domains*, companion manuscript, June 2026.
- [3] E. Ben-Sasson, I. Bentov, Y. Horesh, and M. Riabzev, *Fast Reed–Solomon interactive oracle proofs of proximity*, ICALP 2018.
- [4] G. Arnon, A. Chiesa, G. Fenzi, and E. Yogev, *WHIR: Reed–Solomon proximity testing with super-fast verification*, 2024.
- [5] J. H. Conway and A. J. Jones, *Trigonometric Diophantine equations (On vanishing sums of roots of unity)*, Acta Arith. **30** (1976), 229–240.
- [6] T. Y. Lam and K. H. Leung, *On vanishing sums of roots of unity*, J. Algebra **224** (2000), 91–109.
- [7] M. Laurent, *Equations diophantiennes exponentielles*, Invent. Math. **78** (1984), 299–327.
- [8] J. A. Dias da Silva and Y. O. Hamidoune, *Cyclic spaces for Grassmann derivatives and additive theory*, Bull. London Math. Soc. **26** (1994), 140–146.
- [9] N. Alon, M. B. Nathanson, and I. Ruzsa, *The polynomial method and restricted sums of congruence classes*, J. Number Theory **56** (1996), 404–417.
- [10] N. Alon, *Combinatorial Nullstellensatz*, Combin. Probab. Comput. **8** (1999), 7–29.
- [11] J. Bourgain, A. A. Glibichuk, and S. V. Konyagin, *Estimates for the number of sums and products and for exponential sums in fields of prime order*, J. London Math. Soc. (2) **73** (2006), 380–398.
- [12] J. Bourgain, *Mordell’s exponential sum estimate revisited*, J. Amer. Math. Soc. **18** (2005), 477–499.
- [13] A. Ostafe, I. E. Shparlinski, and J. F. Voloch, *Weil sums over small subgroups*, Math. Proc. Cambridge Philos. Soc. **176** (2024), 39–53.
- [14] I. D. Shkredov, *On exponential sums over multiplicative subgroups of medium size*, arXiv:1311.5726.
- [15] H. Davenport, *Multiplicative Number Theory*, 3rd ed., GTM 74, Springer, 2000. (Siegel–Walfisz, §22.)

- [16] L. C. Washington, *Introduction to Cyclotomic Fields*, 2nd ed., GTM 83, Springer, 1997. (Ramification in $\mathbb{Q}(\zeta_{2^a})$, Ch. 2.)
- [17] J. Li and D. Wan, *On the subset sum problem over finite fields*, Finite Fields Appl. **14** (2008), 911–929.
- [18] G. Zhu and D. Wan, *An asymptotic formula for counting subset sums over subgroups of finite fields*, arXiv:1101.0289, 2010.
- [19] S. Gao, *Polynomials with prescribed leading coefficients and a prescribed number of linear factors over finite fields*, arXiv:2105.12845, 2021.
- [20] S. Gao, *Random polynomials in Hayes equivalence classes and Reed–Solomon distance distributions*, arXiv:2203.06729, 2022.
- [21] A. Davydov, S. Marcugini, and F. Pambianco, *On coset weight distributions of MDS codes*, arXiv:2101.12722, 2021.
- [22] J. Brakensiek, S. Gopi, and V. Makam, *Generic Reed–Solomon codes achieve list-decoding capacity*, STOC 2023, 1488–1501.
- [23] Z. Guo and Z. Zhang, *Randomly punctured Reed–Solomon codes achieve the list decoding capacity over polynomial-size alphabets*, FOCS 2023, 164–176.
- [24] G. Arnon, D. Boneh, G. Fenzi, *Open problems in list decoding and correlated agreement*, Cryptology ePrint Archive, Paper 2026/680, 2026.
- [25] E. Crites, A. Stewart, *On Reed–Solomon proximity gaps conjectures*, Cryptology ePrint Archive, Paper 2025/2046, 2025.
- [26] E. Ben-Sasson, D. Carmon, U. Haböck, S. Kopparty, S. Saraf, *On proximity gaps for Reed–Solomon codes*, Cryptology ePrint Archive, Paper 2025/2055, 2025.
- [27] E. Ben-Sasson, D. Carmon, Y. Ishai, S. Kopparty, S. Saraf, *Proximity gaps for Reed–Solomon codes*, FOCS 2020; J. ACM **70** (2023), no. 5.
- [28] D. Krachun, S. Kazanin, U. Haböck, *Failure of proximity gaps close to capacity*, Cryptology ePrint Archive, Paper 2026/782, 2026.
- [29] A. Kambiré, *Proximity gaps conjecture fails near capacity over prime fields*, arXiv:2604.09724, 2026.
- [30] R. Goyal, V. Guruswami, *Optimal proximity gaps for subspace-design codes and (random) Reed–Solomon codes*, Cryptology ePrint Archive, Paper 2025/2054, 2025.

- [31] S. Bordage, A. Chiesa, Z. Guan, I. Manzur, *All polynomial generators preserve distance with mutual correlated agreement*, Cryptology ePrint Archive, Paper 2025/2051, 2025.
- [32] B. E. Diamond, A. Gruen, *Proximity gaps in interleaved codes*, IACR Communications in Cryptology 1(4), 2025.
- [33] U. Haböck, D. Levit, S. Papini, *Circle STARKs*, Cryptology ePrint Archive, Paper 2024/278, 2024.
- [34] E. Ben-Sasson, D. Carmon, S. Kopparty, D. Levit, *Elliptic curve fast Fourier transform (ECFFT) Part I: fast polynomial algorithms over all finite fields*, arXiv:2107.08473, 2021.
- [35] I. Aliev, C. J. Smyth, *Solving algebraic equations in roots of unity*, Forum Math. 24 (2012), 641–665.
- [36] J.-H. Evertse, H. P. Schlickewei, W. M. Schmidt, *Linear equations in variables which lie in a multiplicative group*, Ann. of Math. 155 (2002), 807–836.
- [37] F. Beukers, C. J. Smyth, *Cyclotomic points on curves*, in: Number Theory for the Millennium I, A K Peters, 2002, 67–85.
- [38] I. J. Schoenberg, *A note on the cyclotomic polynomial*, Mathematika 11 (1964), 131–136.