FRI Soundness Above the Johnson Bound via Threshold Halving

Raullen Chai*[1] and Xinxin Fan†[1]

[1]IoTeX Network

April 29, 2026

Abstract

We prove the first unconditional soundness theorem above the Johnson bound for FRI [5], STIR [8], and WHIR [9] — the proximity-testing protocols underlying every deployed STARK, zkVM, and FRI-based system on Ethereum's roadmap. For RS$[F, L, k]$ with $k = 2^m$ and L admitting a fixed-point-free involution (standard for deployed FRI, in either characteristic), for every $\delta \in (\delta_J,\, 1-\rho)$:

$$\varepsilon_{\mathrm{FRI}} \;\le\; \frac{nR}{|F|} \;+\; \left(1 - \tfrac{\delta}{2}\right)^q.$$

Three results: **(A)** the bound above, via threshold halving from RVW13 [28]; the protocol, prover messages, and verifier checks are unchanged — only the query parameter is recalibrated. The argument enters the unique-decoding regime after one round, where BCIKS [5] locks the distance, making it immune to any open-zone counterexample. **(B)** The $\sim$2$\times$ query overhead is optimal within the correlated-agreement (CA) framework: $\varepsilon_{\mathrm{ca}}(C, \delta, \delta) = \binom{n}{w}/|F|$ (Theorem 7, tight by Proposition 9), exponential in code length, vacuous at FRI scale. **(C)** First p-dependent list-size bounds: $\mathbb{E}[M] = \binom{n}{w}/p^c$ at all codimension excesses $c \ge 1$ (Corollary 28); sub-Poisson tail at c=2 (Theorem 39); phase-diagram conjecture at $c \ge 3$ (Conjecture 41, the deployment regime).

Impact for Ethereum. Every deployed and proposed FRI-based system — SP1, RISC Zero, Plonky3, the $\sim$30 zkVMs on EthProofs, Stwo (Mersenne31 / circle FRI), the planned post-quantum signature aggregation layer — now has a positive proven soundness floor in the open zone above Johnson, replacing a conjecture whose up-to-capacity form was disproved in late 2025 [3, 10]. The cost is a factor $\sim$2 in queries; the bound holds in characteristic 2 (Appendix B), on the unit circle (Appendix C), and compiles to non-interactive knowledge soundness via Fiat–Shamir and DEEP (Appendix D).[1]

Keywords: Reed–Solomon codes, FRI protocol, proximity gap, correlated agreement, Johnson bound, STARK soundness, interactive oracle proofs.

MSC 2020: 94B35 (primary), 11T71, 68Q10.

*`raullen@iotex.io`

†`xinxin@iotex.io`

[1]Half-threshold CA (Theorem 5) is formally verified in Lean 4 with Mathlib (zero `sorry`). Lean proofs, Python verification scripts, and reproduction commands are in the companion repository [36].

1 Introduction

1.1 The problem

Why this matters for Ethereum. Ethereum's medium-term scaling and post-quantum roadmap rests on succinct proofs: zkEVMs, zkVMs, and a planned post-quantum signature-aggregation layer. EthProofs[2] alone tracks ∼30 such systems in active deployment or pre-deployment. Almost all of them are STARK-based, and almost all STARKs are built on a single low-level cryptographic component: the FRI proximity test (or its successors STIR and WHIR). FRI's soundness is what stops a malicious prover from convincing the chain — and through it, every smart contract and bridge — of a statement that is in fact false. To meet performance budgets (proof size, prover memory, verifier gas), every deployed system runs FRI in the so-called *open zone above the Johnson bound*: the parameter regime where, until now, the soundness analysis depended on an *unproved* conjecture (the Ben-Sasson–Carmon–Ishai–Kopparty–Saraf "correlated agreement above Johnson" conjecture [2]). In late 2025 the strongest form of that conjecture — correlated agreement up to capacity — was disproved [3, 10]. This left every above-Johnson soundness argument used in production technically *unproven*: a positive security number quoted in audit reports, but with no theorem behind it in the deployed regime.

This paper closes that gap. We give an unconditional soundness theorem in the full open zone, with no change to any deployed protocol — only the verifier's query parameter is recalibrated. Concretely, the system designer pays approximately one extra query round (§1.6) and, in return, gets a proven soundness bound that is *immune* to any future open-zone counterexample, including any sharper version of the Crites–Stewart construction. Operators of SP1, RISC Zero, Plonky3, Stwo, and the rest of the EthProofs slate can replace the conjectural baseline they currently assume with a theorem; this paper supplies the theorem and quantifies the recalibration cost (Table 1, §1.6, §9).

Formal statement of the prize and the prior gap. The Ethereum Foundation Proximity Prize[3] [1] targets the Reed–Solomon proximity gap conjectures of [2] that control the soundness of FRI [5], STIR [8], and WHIR [9]. BCHKS [4] proved zero-loss proximity gaps for $\delta \le \delta_J$ with $O(n)/|F|$ per-round error (equivalently $O(Rn)/|F|$ over R rounds, as in Table 1). Goyal–Guruswami [6] achieved optimal proximity gaps up to capacity for subspace-design codes and for *random* plain RS codes; Jeronimo–Liu–Rajpal [7] do the same for *folded* RS codes. For *deployed* plain RS codes above Johnson — which use a *fixed* multiplicative-coset evaluation domain L and are therefore neither random nor folded — no positive result existed.

The deployment baseline is the conjecture $\varepsilon_{\mathrm{ca}}(C, \delta, \delta) = O(1)/|F|$ for all $\delta < 1 - \rho$ [2], disproved at capacity by Crites–Stewart [3] (ePrint 2025/2046, November 2025) and independently by Kambiré [10] (arXiv 2604.09724, 2026). For the intermediate zone ($\delta_J < \delta < 1 - \rho$): the bound $\varepsilon_{\mathrm{ca}}(C, \delta, \delta) \le \binom{n}{w}/|F|$ (Theorem 7, tight by Proposition 9) is formally $O(1)/|F|$ but with constant exponential in the code length, making it vacuous at FRI scale. Without the contribution of this paper, every above-Johnson FRI deployment on Ethereum has zero proven soundness in the open zone.

[2] `https://ethproofs.org`

[3] `https://proximityprize.org`

1.2 What we deliver

For FRI (Theorem 18; Theorem 61 for char = 2) on plain Reed–Solomon codes over any finite field, with $k = 2^m$ and domain L admitting a fixed-point-free involution, for every $\delta_J < \delta < 1 - \rho$:

$$\boxed{\varepsilon_{\mathrm{FRI}} \;\le\; \frac{nR}{|F|} \;+\; \Big(1 - \frac{\delta}{2}\Big)^q}$$

A *theorem*, not a conjecture. No change to the FRI protocol structure, prover messages, or verifier checks — only the soundness analysis and the query parameter q it calibrates. Increasing q is still a real implementation-parameter change: it increases Merkle openings, consistency checks, proof bytes, and wrapper-verifier work. The same bound holds in characteristic 2 via additive folding (Appendix B), and is extended via a batch-CA theorem to STIR (Theorem 46) and to WHIR via constrained RS subcodes (Corollary 48); deployment-grade non-interactive knowledge soundness is established via round-by-round soundness and Fiat–Shamir compilation (Appendix D).

The key idea is *threshold halving*: use the CA bound at conclusion threshold $\delta/2$ instead of δ (Theorem 5). The algebraic content is a single inequality $\delta/2 < (1-\rho)/2$ and a 10-line proof. But this simple step has a profound consequence: it places the protocol inside the unique-decoding regime after round 1, where BCIKS [5] supplies zero-loss distance preservation — making the soundness argument *immune to any counterexample construction in the open zone*. Prior results [5, 4] proved soundness only at or below the Johnson bound, where zero-loss CA ($\delta_{\mathrm{hd}} = \delta$) holds. Extending to the open zone above Johnson was blocked by the equal-threshold CA conjecture [2], which requires $M = 0$ on all RS-compatible flats — now known to be false in its strongest form (Proposition 33). Threshold halving sidesteps the conjecture entirely: by accepting a $\sim 2\times$ query overhead (optimal within the CA framework, Remark 12), it yields unconditional soundness in the full open zone.

1.3 Coverage of Ethereum's deployed and proposed stack

The systems covered, all operating above the Johnson bound:

- **Odd characteristic (deployed).** SP1 ($\mathrm{BabyBear}_4$, rate 1/2), RISC Zero (BabyBear, rate 1/4), Plonky3 (BabyBear / Koala-Bear, Polygon zkEVM), and the ∼30 zkVMs on EthProofs[4] that use multiplicative FRI.
- **Circle FRI (covered).** Stwo (Starkware, Mersenne31) uses circle FRI [20] on the unit circle $\{(x, y) \in \mathbb{F}_p^2 : x^2 + y^2 = 1\}$, whose folding structure differs from the multiplicative and additive cases (the involution is $(x, y) \mapsto (x, -y)$, not negation on a subgroup). Appendix C proves the required coupling lemma for circle codes and establishes FRI soundness (Theorem 69) with the same $nR/|F| + (1-\delta/2)^q$ bound.
- **Characteristic 2 (proposed).** Systems using additive FRI on $\mathbb{F}_2$-linear subspaces (binary tower fields $\mathbb{F}_{2^{128}}$). Our Theorem 61 covers the basic additive fold $s(x) = x^2 + \beta x$. Binius [19] uses a more elaborate block-decomposition/Gao-Mateer recursion that is compatible with — but distinct from — this basic fold.

Protocol coverage:

- **FRI (proved).** Theorem 18 (char ≠ 2) and Theorem 61 (char = 2).
- **STIR (fully proved).** Theorem 46 proves STIR soundness above Johnson via our batch-CA bound (Theorem 44).
- **WHIR (fully proved).** Corollary 48 proves our CA bound holds for any *linear* subcode of RS_k; end-to-end WHIR soundness with sumcheck composition is established in Theorem 52.

[4] `https://ethproofs.org`

The equal-threshold bound $\binom{n}{w}/|F|$ is vacuous at FRI scale (§1.7); our half-threshold bound is optimal within the CA framework.

1.4 Technical contributions

(i) **FRI soundness above Johnson** (Theorem 18; Theorem 61 for char 2): the bound $nR/|F| + (1-\delta/2)^q$. The commit-phase term $O(nR/|F|)$ matches BCHKS up to constants; the contribution is that this is *proved* rather than conjectured in the open zone $(\delta_J, 1-\rho)$.

(ii) $O(1)$ **per-round proximity gap** (Theorem 14): At most 1 value of α makes the FRI fold $(\delta/2)$-close to $\mathrm{RS}_{k/2}$. Holds for all even k, any field.

(iii) **List-size landscape and OP2** (§7): exact second moment $\mathbb{E}[M] = \binom{n}{w}/p$, $\mathrm{Var}[M] = \mathbb{E}[M](1 - 1/p)$ (Theorem 23), via the error-locator normal vector identity (Lemma 21) — the first p-dependent moment bound on the list size. This extends to $c \geq 2$ (Theorem 26): the $2c$ error-locator normals from two distinct subsets are linearly independent whenever $|E_1 \cap E_2| \leq w-c$, giving Poisson dispersion $\mathrm{Var} \approx \mathbb{E}$ at all codimension excesses (Corollary 28). At the FRI-relevant codimension $c = 2$, a Möbius identity on $(w-1)$-cores (Lemma 37) yields the universal worst-case bound $M_{\mathrm{true}} \leq \min(p, 2\binom{n}{w-1})$ (Theorem 38) and a sub-Poisson Bernstein tail on bin sizes in expectation (Theorem 39, via the Möbius-induced negative correlation); the empirical worst case grows as $\sim 0.66 \cdot 1.36^n$ through $n = 24$ (Remark 40). The codimension picture is summarized as a phase diagram (§7.6): $c = 1$ saturating, $c = 2$ exponential, and $c \geq 3$ conjecturally linear $M_{\mathrm{true}} \leq \lfloor (2D-1)/c \rfloor$ (Conjecture 41, empirically verified through $n = 40$); this places deployment FRI ($c = \Theta(n)$) in a qualitatively different regime from $c = 2$. OP2 obstruction: $M_{\max} \geq 2$ at FRI parameters for all p (Proposition 33).

(iv) **Tightness of the threshold ratio** (Remark 12): The 1:2 ratio ($\delta_{\mathrm{hd}} = \delta/2$) is tight: two bad γ's give joint distance $\leq 2\delta_{\mathrm{hd}}$, and $2\delta_{\mathrm{hd}} \leq \delta$ iff $\delta_{\mathrm{hd}} \leq \delta/2$. The equal-threshold CA bound $\varepsilon_{\mathrm{ca}}(C, \delta, \delta) = \binom{n}{w}/|F|$ is $O(1)/|F|$ but with constant exponential in n, making the $2\times$ query overhead intrinsic at deployment scale.

(v) **STIR/WHIR extension** (Appendix A): Batch CA theorem (Theorem 44) via conditioning and union bound; STIR soundness above Johnson (Theorem 46); WHIR soundness via sumcheck composition and half-threshold CA (Theorem 52).

(vi) **Characteristic-2 extension** (Appendix B): The multiplicative $\pm\sqrt{y}$ decomposition is replaced by additive subspace-polynomial folding $s(x) = x^2 + \beta x$, yielding FRI soundness (Theorem 61) with the same bound over binary tower fields.

(vii) **Circle FRI extension** (Appendix C): A coupling lemma for circle RS codes (Lemma 67) establishes FRI soundness on the unit circle of $\mathbb{F}_{p^2}$ (Theorem 69), covering Stwo's Mersenne31 deployment with the same $NR/|F| + (1-\delta/2)^q$ bound.

(viii) **Fiat–Shamir compilation and knowledge soundness** (Appendix D): Round-by-round soundness (Theorem 72), non-interactive FRI soundness via BCS compilation (Theorem 73), and DEEP-FRI polynomial commitment knowledge soundness (Theorem 75, Corollary 76).

Result	Regime	FRI commit error	Loss	Ref.
BCIKS [5]	$\delta < \delta_J$	$O(Rn^2)/\lvert F\rvert$	0	FOCS'20
BCHKS [4]	$\delta \le \delta_J$	$O(Rn)/\lvert F\rvert$	0	2025
Goyal–Gur. [6]	$\delta < 1-\rho$	$O(R)/\lvert F\rvert$	0	2025
		(random plain RS; subspace-design codes)		
This paper	$\delta_J < \delta < 1-\rho$	$\mathbf{O(nR)/\lvert F\rvert}$	$\delta/2$	**new**

Table 1: FRI commit-phase soundness for **plain fixed-L** RS (the deployment regime). Goyal–Guruswami applies to random plain RS and subspace-design codes, not to the fixed-L setting in which every deployed multiplicative-FRI system operates (Stwo's circle FRI is covered in Appendix C).

Field ($\lvert F\rvert$)	Commit	Query (q=128, δ=0.4)		Proved total
	(both)	Conj. $(1-\delta)^q$	**Proved** $(1-\delta/2)^q$	
BabyBear base (2^{31})	7	*94 (unproved)*	41	**7**
BabyBear_4 (2^{124})	100	*94 (unproved)*	41	**41**
Goldilocks (2^{64})	40	*94 (unproved)*	41	**40**

Table 2: Interactive FRI security (bits), $n = 2^{20}$, $\rho = 1/2$, $R = 20$. The commit-phase term matches BCHKS [4]. Prior to this work, the 94-bit query-phase figure was *conjectured*; we *prove* 41 bits at the same $q = 128$. Proved total = min(commit, proved query): on BabyBear_4 the query phase is the bottleneck and our proved total is 41 bits; on Goldilocks and base-field BabyBear the commit phase dominates. Increasing to $q \approx 398$ achieves 128-bit *interactive* query-phase security; BabyBear_4 remains capped at $\sim$100 total interactive bits by the commit term. Table 3 accounts separately for Fiat–Shamir's non-interactive Q-loss. Circle FRI (Stwo/Mersenne31) achieves the same bound (Theorem 69).

1.5 Comparison

1.6 Concrete deployment cost

Against what is actually proved today, our bound provides positive proven security at every δ in the open zone where BCHKS provides none: above the Johnson bound, BCHKS gives zero proven soundness, whereas we give $\varepsilon_{\text{FRI}} \le nR/|F| + (1-\delta/2)^q$. There is no proved-vs-proved regression; there is only added coverage.

Against the *conjectured* zero-loss CA baseline (which has $\varepsilon_{\text{ca}} = \binom{n}{w}/|F|$ above Johnson, vacuous at deployment scale — see Remark 12 and §8, item 1 — and independently disproved at capacity [3, 10]), the half-threshold exacts a precise toll on query complexity. To achieve λ-bit security from the query term alone:

$$q_{\text{zero-loss}} = \frac{\lambda}{\log_2(1/(1-\delta))}, \qquad q_{\text{ours}} = \frac{\lambda}{\log_2(1/(1-\delta/2))}.$$

The overhead ratio $q_{\text{ours}}/q_{\text{zero-loss}} = \log_2(1/(1-\delta))/\log_2(1/(1-\delta/2))$:

δ	Regime	$q_{\text{ours}}/q_{\text{zero-loss}}$	Status of zero-loss
0.293	Johnson ($\rho = 1/2$)	2.2×	$\binom{n}{w}/\lvert F\rvert$ (Remark 12)
0.40	Intermediate	2.3×	$\binom{n}{w}/\lvert F\rvert$, vacuous at FRI scale
0.45	Near capacity	2.4×	Disproved at capacity [3, 10]

The overhead is paid only in queries, and only against a baseline that is itself unproven. The commit-phase term $O(nR)/|F|$ matches BCHKS up to constants; we do not claim a Pareto improvement over BCHKS in the Johnson regime. The contribution is coverage of the open zone, where BCHKS provides no bound. Section 9 quantifies both terms in deployment-grade parameters.

1.7 The correlated agreement framework and its limits

Every known FRI soundness proof — BBHR [18], BCIKS [5], BCHKS [4], Goyal–Guruswami [6], and the present paper — follows the same two-step strategy:

1. **Correlated agreement (CA) lemma**: if (f_1, f_2) has joint distance $> \delta$ from $C \times C$, then for random γ, the fold $f_1 + \gamma f_2$ stays far from C with high probability.
2. **Coupling**: the FRI decomposition $f \mapsto (f_{\text{even}}, f_{\text{odd}})$ maps distance from RS_k to joint distance from $\mathrm{RS}_{k/2} \times \mathrm{RS}_{k/2}$, linking the CA lemma to the protocol.

We call this the **CA framework**. The query-phase soundness of FRI is controlled by the *conclusion threshold* δ_{hd} of the CA lemma: each query catches a far fold with probability δ_{hd}, giving soundness $(1 - \delta_{\text{hd}})^q$ after q queries. The ratio $\delta_{\text{hd}}/\delta$ determines the *query overhead*: at ratio 1:1 (zero loss), queries are optimal; at ratio 1:r, the overhead is $\sim r\times$.

The CA framework is settled. The $\delta_{\text{hd}} = \delta/2$ bound (ratio 1:2) was established by Rothblum–Vadhan–Wigderson [28] in 2013 for arbitrary linear codes (Theorem 5). At equal threshold $\delta_{\text{hd}} = \delta$: $\varepsilon_{\text{ca}} = \binom{n}{w}/|F|$ with constant exponential in the code length (Remark 12), vacuous at FRI scale. The 2× query overhead is intrinsic — every CA-based FRI proof at current code lengths must pay this cost.

What lies beyond? Reducing the overhead below 2× requires a proof that does not factor through a CA lemma. No such argument exists in the literature. We discuss concrete candidates — character sums on multiplicative subgroups, cross-correlation techniques — in Section 10.

1.8 Scope and limitations

To prevent misreading: this paper (i) proves an unconditional FRI soundness theorem above Johnson with a $\sim$2× query overhead that is optimal within the CA framework (§1.7); (ii) is *not* a protocol change — FRI/STIR/WHIR as deployed today are unchanged; (iii) is *not* free — the $\sim$2× overhead vs. the zero-loss baseline is the price of an unconditional proof.

Structural assumptions. The $k = 2^m$ condition is necessary: for odd $k \geq 3$ the monomial x^k gives a counterexample with $\Delta(x^k, \mathrm{RS}_k) = 1 - \rho$ yet the FRI fold is always a codeword (Remark 4). Standard FRI uses $k = 2^m$ already. The second-moment bound (Theorem 23) holds for all k but at codimension $c = 1$; it is not on the critical path to the main FRI theorem. There is *no* characteristic restriction: the main body handles $\mathrm{char}(F) \neq 2$ via multiplicative folding; Appendix B extends to $\mathrm{char}(F) = 2$ via additive folding, with the identical $nR/|F| + (1 - \delta/2)^q$ bound. **There is also no rate restriction:** the main theorem holds for every rate $\rho = k/n$ with $k = 2^m$ (the deployment tables in §1.6 use $\rho = 1/2$ as the BabyBear/Goldilocks reference; rates such as $\rho = 1/4$ or $\rho = 1/8$, common in zkVM deployments for soundness amplification, satisfy the same bound with $R = m$ rounds and any $\delta \in (\delta_J(\rho), 1-\rho)$).

1.9 Relation to ABF open problems

The EF Proximity Prize is tied to the ABF open problems [2]: OP1 (prove/disprove zero-loss CA above Johnson), OP2 (proximity gap with sub-$O(n)$ loss above Johnson), OP3 (improved CA constants).

The table below is the claim map used in the rest of the paper. It separates theorem-level results from partial progress and from the original zero-loss targets that remain open.

Prize target	Paper result	Status	Reference	Remaining gap
OP1: zero-loss / equal-threshold CA above Johnson	Equal-threshold CA satisfies $\varepsilon_{\mathrm{ca}}(C,\delta,\delta) \le \binom{n}{w}/\lvert F\rvert$, tight for large fields; the constant is exponential in n and vacuous at FRI scale	Resolved within the CA framework	Thm. 7; Prop. 9; Rem. 12	A non-CA proof may still beat the $2\times$ overhead
FRI soundness above Johnson for deployed plain RS	Unconditional bound $nR/\lvert F\rvert + (1-\delta/2)^q$ for multiplicative, additive, and circle FRI; no protocol change	Solved as theorem	Thms. 18, 61, 69	Query overhead remains relative to the conjectured zero-loss baseline
STIR / WHIR above Johnson	Same half-threshold safety net via batch CA, constrained-code composition, and proved unique-decoding MCA after the first fold	Solved as theorem for stated models	Thms. 46, 52	Full SNARK composition remains protocol-specific
OP2: sub-$O(n)$ / list-size route to proximity gaps	First p-dependent moment bounds; sub-Poisson tail at $c = 2$ in expectation; codimension phase diagram with linear-n conjecture at $c \ge 3$ (deployment regime); explicit obstruction to worst-case $M = 0$	Partially addressed; strongest $M = 0$ form refuted; phase diagram conjectural at $c \ge 3$	Thm. 23; Cor. 28; Thm. 39; Conj. 41; Prop. 33	Worst-case $c = 2$ extension; rigorous proof of Conj. 41 at $c \ge 3$
OP3: improved CA constants	Half-threshold CA constant is $1/\lvert F\rvert$ and commit phase matches $O(nR)/\lvert F\rvert$ up to constants	Partially addressed	Thm. 5; Thm. 18	Does not prove zero-loss CA
Original up-to-capacity MCA / zero-loss proximity gap	Not proved; capacity-level variants are now known false in related settings	Not solved here	Crites–Stewart [3]; Kambiré [10]; §8	Requires different techniques or revised conjectures

Interpretation. The paper does not claim the original zero-loss proximity gap. Its prize-facing contribution is narrower and unconditional: deployed FRI/STIR/WHIR get soundness above Johnson with a half-threshold CA safety net, while the equal-threshold CA route is shown to be vacuous at deployment scale. The remaining route to better-than-$2\times$ query overhead must therefore avoid the CA framework or prove a strictly stronger non-CA proximity statement.

1.10 Guide to the proof

The paper has two logically independent tracks: the **mainline** (§3–§6), which proves the FRI soundness theorem in five steps, and the **structural track** (§7), which establishes the first p-dependent list-size bounds and resolves Open Problem 2 (negatively). The mainline is entirely self-contained; the structural track is not on the critical path to any soundness theorem.

Mainline: FRI soundness in five steps.

Step 1. Half-threshold CA (Theorem 5, §3). For any linear code C: if (f_1, f_2) has joint distance $> \delta$ from $C \times C$, at most one $\gamma \in F$ makes $f_1 + \gamma f_2$ within distance $\delta/2$ of C. *Input*: linearity of C only (RVW13 [28]).

Step 2. Even/odd coupling (Lemma 13, §5). For even k: $\Delta(f, \mathrm{RS}_k) \le \Delta_{\mathrm{joint}}((f_{\mathrm{even}}, f_{\mathrm{odd}}), \mathrm{RS}^{=2}_{k/2})$. *Input*: the isomorphism $\mathrm{RS}_k \xrightarrow{\sim} \mathrm{RS}^{=2}_{k/2}$ (requires k even; Remark 4).

Step 3. Proximity gap (Theorem 14, §5). At most 1 value of α makes the FRI fold $(\delta/2)$-close to $\mathrm{RS}_{k/2}$. *Input*: Steps 1 + 2 composed.

Step 4. Distance locking (Remark 19, §6). After round 1, the effective distance is $\delta/2 < (1-\rho)/2$ (unique decoding radius). The BCIKS proximity gap [5] (Theorem 1.2, Eq. (1.1)) at threshold $\gamma = \delta/2 < (1-\rho)/2$ gives $\le n$ bad α's per round; distance stays at $\delta/2$ for all subsequent rounds. *Input*: the structural observation $\delta/2 < (1-\rho)/2$ (always true for $\delta < 1-\rho$) + BCIKS [5] (external, proved).

Step 5. FRI soundness (Theorem 18, §6). Strategy A (honest fold): multilinear Schwartz–Zippel gives $\le R/|F|$. Strategy B (deviate): Steps 3–4 bound bad α's to ≤ 1 (round 1) and $\le n$ (rounds ≥ 2) per round and the consistency check catches a $(\delta/2)$-far deviation with probability $\ge 1-(1-\delta/2)^q$. Union: $nR/|F| + (1-\delta/2)^q$. *Input*: Steps 3 + 4 + Schwartz–Zippel.

The mainline uses *no* result from §7. The entire proof rests on three external ingredients: linearity of RS_k (trivial), the BCIKS proximity gap at $\delta \le \delta_J$ (proved [5]), and the Schwartz–Zippel lemma.

Extensions (from the same CA bound). Theorem 5 is characteristic-independent, so three extensions branch directly from Step 1 (see Figure 1):

- *STIR*: Batch CA (Theorem 44) → STIR soundness (Theorem 46), Appendix A.
- *WHIR*: Sumcheck composition + half-threshold CA (Theorem 52), Appendix A.
- *Char* 2: Additive coupling (Lemma 59) → FRI soundness in char 2 (Theorem 61), Appendix B.
- *Circle FRI*: Circle coupling (Lemma 67) → FRI soundness on unit circle (Theorem 69), Appendix C.

Structural track (§7): independent of the soundness proof. Section 7 studies the per-word list size M (how many codewords lie within distance w of a received word), aiming at Open Problem 2 (§8). Four results:

1. *First p-dependent moment bound.* $\mathbb{E}[M] = \binom{n}{w}/p^c$ for all codimension excess $c \ge 1$ (Corollary 28), via the error-locator normal vector identity (Lemma 25) and pairwise independence for non-overlapping supports (Theorem 26). At FRI-relevant codimension ($c = \Theta(n)$), $\mathbb{E}[M]$ is exponentially small.
2. *Worst-case landscape by codimension.* At $c = 2$, a Möbius identity (Lemma 37) yields a universal $M_{\mathrm{true}} \le \min(p, 2\binom{n}{w-1})$ (Theorem 38) and a Bernstein-style sub-Poisson tail on bin sizes in expectation (Theorem 39); the empirical worst-case grows as $\sim 0.66 \cdot 1.36^n$ through $n = 24$ (Remark 40). At $c \ge 3$, an Open-Set Rank Lemma (Conjecture 41) predicts a linear bound $M_{\mathrm{true}} \le \lfloor (2D-1)/c \rfloor$, empirically verified through $n = 40$ (§7.6). The deployment regime $c = \Theta(n)$ falls in this third phase, where the conjecture predicts $M_{\mathrm{true}} = O(1)$ at the Johnson radius.
3. *OP2 obstruction.* $M_{\max} \ge 2$ at FRI parameters ($n \ge 32$, rate 1/2) for all p (Proposition 33). The strongest $M = 0$ conjecture is false.
4. *Practical $M = 0$.* At deployment parameters ($p \approx 2^{31}$, $c = \Theta(n)$): $\mathbb{E}[M] = \binom{n}{w}/p^c \ll 1$ (Corollary 28); $M = 0$ in practice even though worst-case $M = 0$ fails.

Additional structural results (fiber bounds, Bézout bounds, codimension bounds, incidence bounds, irreducibility of the remainder hypersurface) are developed in a companion note [22].

Reader's guide. A reader interested only in the FRI soundness theorem may read §2–§6 and skip §7 entirely.

2 Definitions

Let F be a finite field and $L \subset F$ with $|L| = n$. The FRI folding requires a *fixed-point-free involution* on L: a bijection $\iota : L \to L$ with $\iota^2 = \mathrm{id}$ and $\iota(x) \neq x$ for all $x \in L$. In the main body (§3–§6), we use $\iota(x) = -x$ (i.e., $\mathrm{char}(F) \neq 2$ and L closed under negation). Appendix B uses $\iota(x) = x + \beta$ (i.e., $\mathrm{char}(F) = 2$ and L an $\mathbb{F}_2$-linear subspace containing $\beta \neq 0$). $\mathrm{RS}_k = \mathrm{RS}[F, L, k]$ is the Reed–Solomon code of evaluations of polynomials of degree $<k$ on L, with rate $\rho = k/n$. We write $\delta_J = 1 - \sqrt{\rho}$ for the *Johnson radius* (the classical relative list-decoding radius with polynomial list size [13, 14]) and use this same symbol throughout; the phrase "Johnson bound" is used interchangeably. Theorems in this paper operate in the *open zone* $\delta_J < \delta < 1 - \rho$ (strict on both ends); at the boundary $\delta = \delta_J$ the proved zero-loss CA of BCHKS [4] applies and our reduced-threshold statement is strictly weaker.

Remark 1 (Base vs. extension fields)**.** *The evaluation domain L typically sits inside a prime subfield or a small extension (e.g., $L \subset \mathbb{F}_p$ with $|L| = n \mid (p-1)$). The soundness bounds in this paper depend on $|F|$, not on the subfield containing L. In deployed systems, FRI randomness α_i is sampled from an extension field $\mathbb{F}_{p^d}$ to achieve sufficient Schwartz–Zippel probability: $|F| = p^d$. For example, BabyBear ($p = 2^{31} - 1$) with degree-4 extension gives $|F| = p^4 \approx 2^{124}$, yielding $\log_2(|F|/(nR)) \approx 100$ bits of commit-phase security (at $n = 2^{20}$, $R = 20$). The $|F|$ in Theorems 5, 18, and 46 refers to the field from which random challenges are drawn, which is the extension field in practice.*

Definition 2 (Correlated Agreement [2])**.** $\varepsilon_{\mathrm{ca}}(C, \delta_{\mathrm{hd}}, \delta_{\mathrm{nt}}) = \max_{f_1, f_2 : L \to F} \Pr_{\gamma \leftarrow F}[\Delta(f_1 + \gamma f_2, C) \leq \delta_{\mathrm{hd}}$ *and* $\Delta((f_1, f_2), C^2) > \delta_{\mathrm{nt}}]$*, where γ is drawn uniformly at random from F, the max is over all functions $f_1, f_2 : L \to F$ (not necessarily codewords), and $\Delta((f_1, f_2), C^2) = \min_{g_1, g_2 \in C} \frac{1}{|D|} |\{x \in D : f_1(x) \neq g_1(x)$ or $f_2(x) \neq g_2(x)\}|$ is the joint Hamming distance to the product code $C^2 = C \times C$ [5], where D is the common domain of f_1, f_2 (always the domain of C; when the CA bound is applied to $(f_{\mathrm{even}}, f_{\mathrm{odd}})$ on L', $D = L'$ and $|D| = n/2$). We call δ_{hd} the conclusion threshold (distance of the random combination to C) and δ_{nt} the premise threshold (joint distance from C^2). Standard zero-loss CA [2] corresponds to $\delta_{\mathrm{hd}} = \delta_{\mathrm{nt}} = \delta$.*

Definition 3 (FRI Folding)**.** *$L' = \{x^2 : x \in L\}$ with $|L'| = n/2$. For each $y \in L'$, fix a square root $\sqrt{y} \in L$ (the expressions below are invariant under the choice $\sqrt{y} \mapsto -\sqrt{y}$). For $f : L \to F$: $f_{\mathrm{even}}(y) = (f(\sqrt{y}) + f(-\sqrt{y}))/2$, $f_{\mathrm{odd}}(y) = (f(\sqrt{y}) - f(-\sqrt{y}))/(2\sqrt{y})$. FRI fold: $f'(y) = f_{\mathrm{even}}(y) + \alpha \cdot f_{\mathrm{odd}}(y)$.*

Remark 4 (Even/odd isomorphism for even k)**.** *For k even: $g \mapsto (g_{\mathrm{even}}, g_{\mathrm{odd}})$ is a linear isomorphism $\mathrm{RS}_k \xrightarrow{\sim} \mathrm{RS}^{=2}_{k/2}$ (both have dimension k; the map permutes the k coefficients). Consequently: $f \notin \mathrm{RS}_k$ iff $(f_{\mathrm{even}}, f_{\mathrm{odd}}) \notin \mathrm{RS}^{=2}_{k/2}$.*

For k odd: $\mathrm{RS}^{=2}_{\lceil k/2 \rceil}$ has dimension $k+1 > k$, and the isomorphism fails. The monomial x^k (for $k \geq 3$ odd) satisfies $\Delta(x^k, \mathrm{RS}_k) = 1 - \rho$ (since $x^k - g$ has at most k roots for any g of degree $<k$), yet $f_{\mathrm{even}} = 0$ and $f_{\mathrm{odd}}(y) = y^{(k-1)/2} \in \mathrm{RS}_{\lceil k/2 \rceil}$ (since $(-1)^k = -1$), so the fold $f' = \alpha \cdot y^{(k-1)/2} \in \mathrm{RS}_{\lceil k/2 \rceil}$ for all α and FRI always accepts. The degenerate case $k = 1$ is vacuous: $\mathrm{RS}_1 = \{$constants$\}$ and FRI is not invoked ($R = 0$). This is why we require $k = 2^m$ with $m \geq 1$.

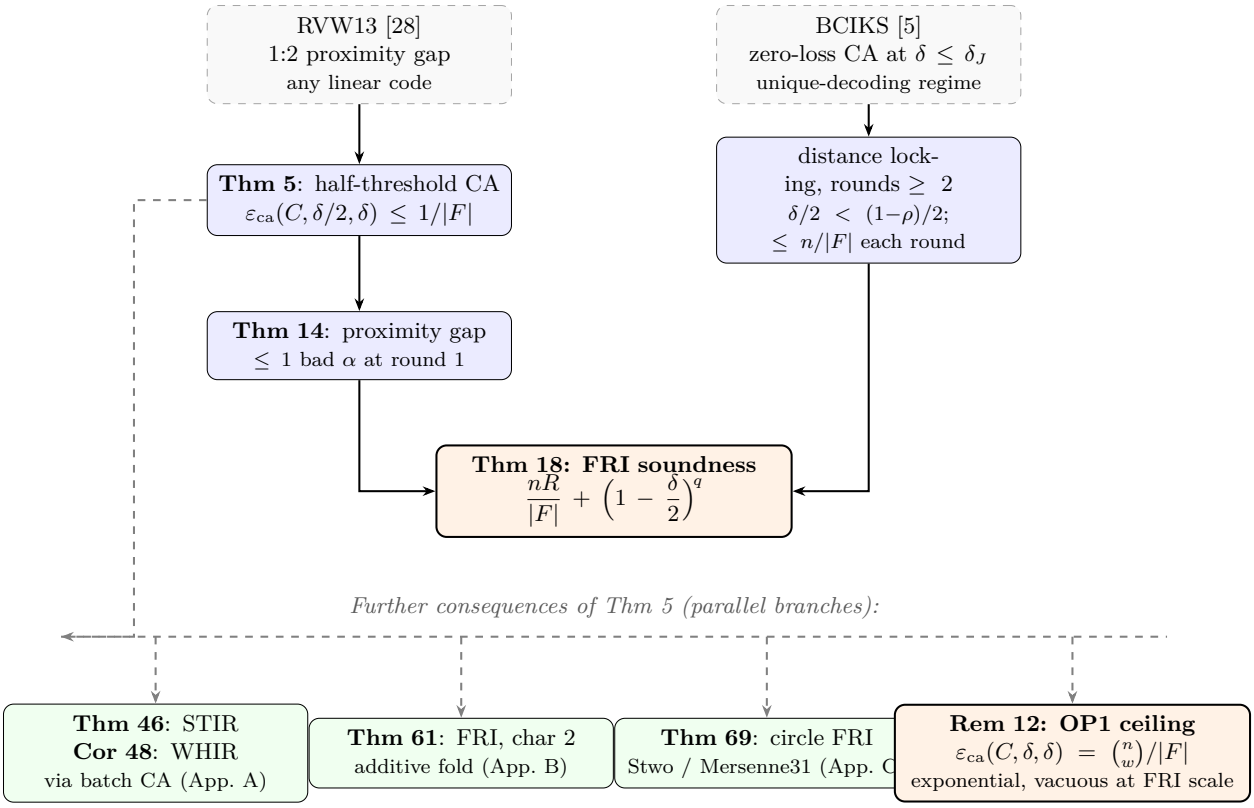

Figure 1: Proof map. **Top half** (mainline): RVW13 yields the half-threshold CA bound (Thm 5), which composes with the FRI even/odd coupling to give the round-1 proximity gap (Thm 14); BCIKS supplies the round-≥ 2 distance lock (the structural observation $\delta/2 < (1-\rho)/2$); the two combine into the FRI soundness theorem. **Bottom half** (parallel branches from the same CA bound): STIR/WHIR via batch CA; characteristic-2 FRI via additive coupling; circle FRI via the Stwo coupling; equal-threshold CA (the OP1 ceiling) showing the $\binom{n}{w}/|F|$ bound is vacuous at deployment scale, so the $2\times$ overhead is intrinsic to any CA-based proof. The list-size structural track (§7) is independent of the soundness proof and is not shown.

3 Half-Threshold CA Bound

Theorem 5 (RVW13 [28]; BCHKS [4] Table 1 row 1)**.** *For any linear code C over F with $|L| = n$: $\varepsilon_{\text{ca}}(C, \delta/2, \delta) \le 1/|F|$.*

Equivalently: for any $f_1, f_2 : L \to F$ with $\Delta_{\text{joint}}((f_1, f_2), C \times C) > \delta$, at most one $\gamma \in F$ satisfies $\Delta(f_1 + \gamma f_2, C) \le \delta/2$.

Proof (due to [28])**.** Two bad γ's give two linear equations in (f_1, f_2) on overlapping agreement sets. Solving for both components simultaneously yields a pair $(g_1, g_2) \in C \times C$ matching (f_1, f_2) on $\ge (1-\delta)n$ points, so joint distance $\le \delta$ — contradicting the premise.

Proof. Fix f_1, f_2 with $\Delta_{\text{joint}} > \delta$. Call γ "bad" if $\exists\, h_\gamma \in C$ with $\Delta(f_1 + \gamma f_2, h_\gamma) \le \delta/2$; let $S_\gamma = \{x \in L : f_1(x) + \gamma f_2(x) = h_\gamma(x)\}$, so $|S_\gamma| \ge (1 - \delta/2)n$.

Suppose $\gamma_1 \ne \gamma_2$ are both bad. By inclusion-exclusion:

$$|S_{\gamma_1} \cap S_{\gamma_2}| \;\ge\; 2(1 - \delta/2)n - n \;=\; (1 - \delta)n.$$

On $S_{\gamma_1} \cap S_{\gamma_2}$, the system $f_1 + \gamma_i f_2 = h_i$ $(i = 1, 2)$ has the unique solution

$$g_1 := \frac{\gamma_2 h_1 - \gamma_1 h_2}{\gamma_2 - \gamma_1} \in C, \qquad g_2 := \frac{h_1 - h_2}{\gamma_1 - \gamma_2} \in C$$

(using linearity of C). So $(f_1, f_2) = (g_1, g_2)$ on $S_{\gamma_1} \cap S_{\gamma_2}$, giving:

$$\Delta_{\text{joint}}\big((f_1, f_2),\, C \times C\big) \;\le\; \frac{|L \setminus (S_{\gamma_1} \cap S_{\gamma_2})|}{n} \;\le\; \delta.$$

This contradicts $\Delta_{\text{joint}} > \delta$. Hence at most 1 bad γ. □

Remark 6 (Attribution and scope)**.** *Theorem 5 is the classical* 1:2 *proximity-gap inequality of Rothblum–Vadhan–Wigderson [28] (STOC 2013), formalized in BCHKS [4] Table 1 as the row [RVW13]: γ arbitrary, $a \ge 2$, $\varepsilon^* = \gamma$. The proof uses only linearity of C (closure under F-linear combinations): no polynomial structure, no multiplicative structure on L, no restriction on k or δ. It applies to any linear code over any finite field.*

What is new in this paper is not the CA bound, but (a) its application above the Johnson bound to obtain an unconditional FRI soundness theorem in the open zone $(\delta_J, 1-\rho)$, where no such theorem was previously available; (b) the analytic equal-threshold CA bound $\varepsilon_{\text{ca}} = \binom{n}{w}/|F|$ (Remark 12); and (c) the tightness argument showing the 1:2 *ratio is optimal at deployment scale (Remark 12). The structural observation enabling (a) is that $\delta/2 < (1-\rho)/2$ always, so after one round the protocol enters the unique-decoding regime where BCIKS [5] supplies the round-by-round gap.*

The prior proximity-gap results of AHIV17 [29] ($\delta < \delta_V/4$), RZ18 [30] ($\delta < \delta_V/3$), and BKS18 [31] (δ above unique decoding) each proved successively stronger gap statements. The BKS18 FRI soundness proof [31] (Theorem 16, the one used by BCIKS [5]) splits on $\Delta(f_2, C)$ and bounds bad α's in each case separately (Case 1/2), yielding a 1:3 ratio ($\delta_{\text{hd}} = \delta/3$), tight for that framework because the case bookkeeping gives $2\delta_{\text{hd}} + \delta_{\text{hd}} = 3\delta_{\text{hd}} \le \delta$ only when $\delta_{\text{hd}} \le \delta/3$. The direct proof above bypasses Case 1/2 entirely: it solves for both f_1 and f_2 simultaneously from two bad γ's, obtaining the joint distance bound $\le 2\delta_{\text{hd}}$ in one step. This is $\le \delta$ when $\delta_{\text{hd}} \le \delta/2$, improving the ratio from 1:3 to 1:2.

We state $nR/|F|$ (not $R/|F|$) in Theorem 18 to maintain a unified presentation with STIR (Theorem 46), where the per-round count cannot be tightened below n in general.

4 Equal-Threshold CA: Exact Bound and Tightness of the 1:2 Ratio

This section establishes the exact form of the equal-threshold CA bound (Theorem 7) and shows that it is tight up to lower-order terms (Proposition 9). The conclusion (Remark 12) is that the $\binom{n}{w}$ constant is exponential in the code length, vacuous at FRI scale; this is the formal statement behind the "OP1 obstruction" framing of §1.9.

4.1 Upper bound at equal threshold

Theorem 7 (Equal-threshold CA upper bound)**.** *Let $C = \text{RS}[F, L, k]$ with $|L| = n$, $w = n - k - 1$, $\delta = w/n$. For any $f_1, f_2 : L \to F$ satisfying $\Delta_{\text{joint}}((f_1, f_2), C \times C) > \delta$:*

$$\big|\{\gamma \in F : \Delta(f_1 + \gamma f_2, C) \le \delta\}\big| \;\le\; \binom{n}{w}.$$

Consequently $\varepsilon_{\text{ca}}(C, \delta, \delta) \le \binom{n}{w}/|F|$.

Proof. Let $\Gamma = \{\gamma \in F : \Delta(f_1 + \gamma f_2, C) \le \delta\}$. For each $\gamma \in \Gamma$ there exists a witness codeword $c_\gamma \in C$ with $d(f_1 + \gamma f_2,\, c_\gamma) \le \delta n = w$, so the agreement set has size $\ge n - w = k + 1$. Choose any $A_\gamma \subseteq L$ with $|A_\gamma| = k + 1$ on which $f_1 + \gamma f_2 = c_\gamma$, and define the assignment

$$\varphi : \Gamma \to \binom{L}{k+1}, \qquad \varphi(\gamma) = A_\gamma.$$

Claim: φ is injective. Suppose $\gamma \ne \gamma'$ in Γ with $\varphi(\gamma) = \varphi(\gamma') = A$. On A (size $k + 1$):

$$f_1 + \gamma f_2 = c_\gamma, \qquad f_1 + \gamma' f_2 = c_{\gamma'}.$$

Subtracting and dividing by $\gamma - \gamma' \ne 0$ (allowed since C is closed under F-linear combinations):

$$f_2\big|_A \;=\; \tfrac{c_\gamma - c_{\gamma'}}{\gamma - \gamma'}\big|_A, \qquad g_2 := \tfrac{c_\gamma - c_{\gamma'}}{\gamma - \gamma'} \in C.$$

Likewise $f_1|_A = c_\gamma|_A - \gamma g_2|_A$, so setting $g_1 := c_\gamma - \gamma g_2 \in C$ gives $(f_1, f_2) = (g_1, g_2)$ on A. Hence the joint disagreement is at most $|L \setminus A| = n - (k + 1) = w$, so

$$\Delta_{\text{joint}}\big((f_1, f_2),\, C \times C\big) \;\le\; \frac{w}{n} \;=\; \delta,$$

contradicting the hypothesis $\Delta_{\text{joint}} > \delta$. Therefore φ is injective and $|\Gamma| \le \binom{n}{k+1} = \binom{n}{w}$. □

Remark 8 (Scope of Theorem 7)**.** *The bound holds (i) for all (f_1, f_2) with $\Delta_{\text{joint}} > \delta$ — not just generic ones; (ii) over any field F, including small fields where $|F| \le \binom{n}{w}$ (in which case the bound $\binom{n}{w}/|F| \ge 1$ is trivial); (iii) at the specific weight $w = n - k - 1$ (the largest w with $\delta n < n - k$, i.e., one below the covering radius). For $w' \le w$ (smaller decoding ball, same code): the same $(k+1)$-subset injection yields $|\Gamma| \le \binom{n}{k+1} = \binom{n}{w}$, since the agreement set still has size $\ge n - w' \ge k + 1$. Tighter bounds at $w' < w$ require a different argument (e.g., injection into $(n-w')$-subsets together with intersection constraints).*

4.2 Tightness

Proposition 9 (Equal-threshold CA lower bound, large field)**.** *With notation as in Theorem 7, for every field F with $|F| > \binom{n}{w}^2$ there exist $f_1, f_2 : L \to F$ with $\Delta_{\text{joint}}((f_1, f_2),\, C \times C) > \delta$ and*

$$\big|\{\gamma \in F : \Delta(f_1 + \gamma f_2,\, C) \le \delta\}\big| \;=\; \binom{n}{w}.$$

Proof. For each $(k + 1)$-subset $A \subset L$, encode A-consistency of $f_1 + \gamma f_2$ as a single linear equation in γ. Fix A and let $\pi_A : F^L \to F$ be the linear functional defined by: $\pi_A(f) :=$ (value of the unique degree-$<k$ Lagrange interpolant of $f|_{A_0}$ at the remaining point of A) $- f$ at that point, where A_0 is any fixed k-subset of A (the choice of A_0 does not affect the kernel). Then f agrees with some codeword of C on all of A iff $\pi_A(f) = 0$. Set

$$a_A := \pi_A(f_1) \in F, \qquad b_A := \pi_A(f_2) \in F.$$

By linearity of π_A: $\pi_A(f_1 + \gamma f_2) = a_A + \gamma \cdot b_A$. Hence $f_1 + \gamma f_2$ is A-consistent iff

$$a_A + \gamma \cdot b_A \;=\; 0. \tag{1}$$

For $b_A \ne 0$, equation (1) has the unique solution $\gamma_A = -a_A / b_A \in F$. As A ranges over the $\binom{n}{k+1} = \binom{n}{w}$ choices, the values $\{\gamma_A\}_A$ exhaust all candidate decodable γ's by Theorem 7's injection (when each agreement set is exactly $(k+1)$-sized, the injection becomes a bijection onto its image).

Construction. Fix $f_2 \in F^L$ such that $b_A(f_2) \neq 0$ for every $A \in \binom{L}{k+1}$; the bad set $\{f : \pi_A(f) = 0\}$ is a single proper hyperplane in F^L for each A, contributing $|F|^{n-1}/|F|^n = 1/|F|$ measure, so the union of $\binom{n}{w}$ such hyperplanes leaves at least $|F|^n(1 - \binom{n}{w}/|F|) > 0$ valid f_2 when $|F| > \binom{n}{w}$. Pick $f_1 \in F^L$ uniformly at random. For $A \neq A'$: the equation $\gamma_A = \gamma_{A'}$ becomes

$$a_A b_{A'} - a_{A'} b_A \;=\; 0 \iff b_{A'} \cdot \pi_A(f_1) - b_A \cdot \pi_{A'}(f_1) \;=\; 0,$$

a single linear constraint on f_1. Since π_A and $\pi_{A'}$ are non-proportional functionals (their kernels differ — they encode interpolation residuals on different $(k+1)$-subsets), the linear combination $b_{A'}\pi_A - b_A\pi_{A'}$ is a nonzero functional on F^L, hence vanishes on a hyperplane of measure $1/|F|$. By union bound over $\binom{\binom{n}{w}}{2} \leq \binom{n}{w}^2/2$ pairs:

$$\Pr_{f_1}[\text{some } \gamma_A = \gamma_{A'}] \;\leq\; \frac{\binom{n}{w}^2}{2\,|F|} \;<\; 1$$

provided $|F| > \binom{n}{w}^2/2$. The hypothesis $|F| > \binom{n}{w}^2$ suffices, giving an f_1 with all $\binom{n}{w}$ values γ_A pairwise distinct in F. By construction each γ_A is decodable (witnessed by codeword agreement on the $(k+1)$-subset A), so $|\Gamma| \geq \binom{n}{w}$. Combined with the upper bound (Theorem 7), $|\Gamma| = \binom{n}{w}$.

Joint-distance check. We must verify $\Delta_{\text{joint}}((f_1, f_2), C \times C) > \delta$. If $\Delta_{\text{joint}} \leq \delta$, then there exists a codeword pair $(g_1, g_2) \in C \times C$ with $|\{x : (f_1(x), f_2(x)) \neq (g_1(x), g_2(x))\}| \leq \delta n$. For *any* $\gamma \in F$: $f_1 + \gamma f_2$ agrees with $g_1 + \gamma g_2 \in C$ on $\geq (1 - \delta)n$ points, so every $\gamma \in F$ would be decodable, giving $|\Gamma| = |F|$. But the construction gives $|\Gamma| = \binom{n}{w} < |F|$ (since $|F| > \binom{n}{w}^2 \geq \binom{n}{w}$). Therefore $\Delta_{\text{joint}} > \delta$. □

Corollary 10 (Asymptotic equality)**.** *For $|F|$ large enough (specifically $|F| > \binom{n}{w}^2$):*

$$\varepsilon_{\text{ca}}(C, \delta, \delta) \;=\; \frac{\binom{n}{w}}{|F|}.$$

4.3 Computational verification

The exact bound is verified for small parameters where exhaustive enumeration is feasible.

Remark 11 (Empirical confirmation)**.** *For* RS$[6, 3]$ *over $\mathbb{F}_q$, $q \in \{7, 13, 19, \ldots, 163\}$: the empirical maximum (over all tested (f_1, f_2)) of the number of δ-close γ's equals* $\min(q, \binom{6}{2}) = \min(q, 15)$ *exactly. For* RS$[8, 4]$ *over $q \in \{17, \ldots, 281\}$:* $\min(q, \binom{8}{3}) = \min(q, 56)$*, also exactly. See Note [23]; scripts* `op1_scaling.py` *and* `op1_scaling_n8.py` *in [36].*

4.4 Why the 1:2 ratio is optimal at deployment scale

Remark 12 (Why the 1:2 ratio is tight)**.** *Algebraic ceiling. The proof of Theorem 5 shows: two bad γ's at conclusion threshold δ_{hd} produce a joint pair $(g_1, g_2) \in C \times C$ matching (f_1, f_2) on $\geq (1 - 2\delta_{\text{hd}})n$ points, giving $\Delta_{\text{joint}} \leq 2\delta_{\text{hd}}$. For this to contradict $\Delta_{\text{joint}} > \delta$: need $2\delta_{\text{hd}} \leq \delta$, i.e., $\delta_{\text{hd}} \leq \delta/2$. The bound $\delta_{\text{hd}} = \delta/2$ is therefore optimal for any argument that extracts a joint codeword pair from two bad random-linear-combination values.*

Why pushing to equal threshold is vacuous. By Corollary 10, the equal-threshold bound is exactly $\binom{n}{w}/|F|$ for sufficiently large fields. The constant $\binom{n}{w}$ grows as $2^{nH(\delta)}$ (Stirling). At FRI-relevant code lengths ($n \geq 2^{10}$, $|F| \sim 2^{31}$): $\binom{n}{w} \gg |F|$, so $\binom{n}{w}/|F| \gg 1$ and the equal-threshold bound

is vacuous. The 2× query overhead is intrinsic at deployment scale, even though $\varepsilon_{\mathrm{ca}}(C,\delta,\delta)$ is technically $O(1)/|F|$ for any fixed n.

Half-threshold is tight. At $\delta_{\mathrm{hd}} = \delta/2$, $\max_\gamma \le 1$ in all tested cases [36], matching the ≤ 1 bound of Theorem 5. The constant 2 in the half-threshold bound is independent of n, in sharp contrast to the $\binom{n}{w}$ constant at equal threshold.

Comparison with the BKS18 1:3 ratio. The BKS18 [31] Case 1/2 approach (split on $\Delta(f_2, C)$, bound each case separately) achieves ratio 1:3 with zero proximity loss at $\delta_{\mathrm{hd}} < \delta/3$. The RVW13 [28] 2×2-solve argument of Theorem 5 trades zero loss for a better ratio (1:2 with loss δ_{hd}); the 1:2 ratio has been folklore in the proximity-gap community since 2013.

5 FRI Coupling and Proximity Gap

Lemma 13 (FRI Coupling, even k)**.** *For even k and any $f : L \to F$: $\Delta(f, \mathrm{RS}_k) \le \Delta((f_{\mathrm{even}}, f_{\mathrm{odd}}), \mathrm{RS}^{=2}_{k/2})$.*

Proof. By Remark 4: $g \mapsto (g_{\mathrm{even}}, g_{\mathrm{odd}})$ is an isomorphism $\mathrm{RS}_k \to \mathrm{RS}^{=2}_{k/2}$. For any $g \in \mathrm{RS}_k$ and $y \in L'$ let $c_y := |\{x \in \{\sqrt{y}, -\sqrt{y}\} : f(x) \ne g(x)\}| \in \{0,1,2\}$ denote the number of base-level errors in the fibre over y. If $f_{\mathrm{even}}(y) = g_{\mathrm{even}}(y)$ and $f_{\mathrm{odd}}(y) = g_{\mathrm{odd}}(y)$, then $f(\pm\sqrt{y}) = g(\pm\sqrt{y})$, so every joint-error position y contributes $c_y \ge 1$ error on L and at most 2. Let $E = |\{y \in L' : c_y \ge 1\}|$ denote the number of joint-error positions. Then $\Delta(f,g) = \sum_y c_y/n \le 2E/n = E/|L'| = \Delta_{\mathrm{joint}}$, where the cancellation uses $n = 2|L'|$. Taking $\min_g$ (using the isomorphism, the min over $\mathrm{RS}^{=2}_{k/2}$ equals the min over RS_k): $\Delta(f, \mathrm{RS}_k) \le \Delta_{\mathrm{joint}}$. □

Theorem 14 (Half-Threshold Proximity Gap)**.** *For even k, $\delta_J < \delta < 1 - \rho$, and $\Delta(f, \mathrm{RS}_k) > \delta$:*

$$|\{\alpha \in F : \Delta(f_{\mathrm{even}} + \alpha f_{\mathrm{odd}}, \mathrm{RS}_{k/2}) \le \delta/2\}| \le 1.$$

Proof. By Lemma 13: $\Delta((f_{\mathrm{even}}, f_{\mathrm{odd}}), \mathrm{RS}^{=2}_{k/2}) \ge \Delta(f, \mathrm{RS}_k) > \delta$. The isomorphism $\mathrm{RS}_k \xrightarrow{\sim} \mathrm{RS}^{=2}_{k/2}$ (Remark 4) gives $\mathrm{RS}^{=2}_{k/2} = \mathrm{RS}_{k/2} \times \mathrm{RS}_{k/2}$, so $\Delta_{\mathrm{joint}}((f_{\mathrm{even}}, f_{\mathrm{odd}}), \mathrm{RS}_{k/2} \times \mathrm{RS}_{k/2}) > \delta$. Apply Theorem 5 to $(f_1, f_2) = (f_{\mathrm{even}}, f_{\mathrm{odd}})$ on $C = \mathrm{RS}_{k/2}$ over L': at most 1 value of α has $\Delta(f_{\mathrm{even}} + \alpha f_{\mathrm{odd}}, \mathrm{RS}_{k/2}) \le \delta/2$. □

6 FRI Soundness

Definition 15 (Standard FRI query model)**.** *The verifier samples q base-level points $z_1, \ldots, z_q \in L$ **i.i.d. uniform with replacement**; for each z_j it queries the folding-tree path $z_j \to \pi_1(z_j) \to \pi_2(z_j) \to \cdots \to \pi_R(z_j)$, where $\pi_i : L \to L^{(i)}$ is the iterated squaring map, and at every round $i = 1, \ldots, R$ checks the consistency relation $f^{(i)}(\pi_i(z_j)) = f^{(i-1)}_{\mathrm{even}}(\pi_i(z_j)) + \alpha_i \cdot f^{(i-1)}_{\mathrm{odd}}(\pi_i(z_j))$. (Equivalent statements for additive folding (App. B) and circle folding (App. C) replace π_i with the corresponding round-i projection.)*

Remark 16 (Implementation models: with-replacement vs. without-replacement)**.** *This paper's bounds are stated for the i.i.d. with-replacement model (Definition 15), which is standard in the FRI/STIR/WHIR analyses [18, 5, 8, 9]. Deployed implementations typically draw q distinct base-level indices via Fiat–Shamir hashing (without replacement, or with domain separation that makes collisions cryptographically negligible). The without-replacement miss probability $\binom{|L|-|\pi_i^{-1}(D_i)|}{q} / \binom{|L|}{q}$ is bounded above by $(1-\tau)^q$ for the same $\tau = |D_i|/|L^{(i)}|$ (negative association of without-replacement indicators), so the bounds in Theorems 18, 46, 52 apply unchanged to deployed protocols.*

Lemma 17 (Catch probability under i.i.d. base sampling)**.** *Under Definition 15, fix any round i and any deviation set $D_i \subseteq L^{(i)}$ with $|D_i|/|L^{(i)}| \geq \tau$. Then*

$$\Pr_{z_1,\ldots,z_q \overset{\text{i.i.d.}}{\sim} L}\left[\, \pi_i(z_j) \notin D_i \text{ for all } j = 1,\ldots,q \,\right] \;\leq\; (1-\tau)^q.$$

Proof. Let $X_j := \mathbf{1}\{\pi_i(z_j) \in D_i\}$. The map $\pi_i : L \to L^{(i)}$ is 2^i-to-1 (a composition of i squarings, each 2-to-1 on the relevant subgroup), so $|\pi_i^{-1}(D_i)| = 2^i \cdot |D_i|$ and

$$\Pr_{z_j \sim L}[X_j = 1] \;=\; \frac{2^i \cdot |D_i|}{2^i \cdot |L^{(i)}|} \;=\; \frac{|D_i|}{|L^{(i)}|} \;\geq\; \tau.$$

The variables $X_1, \ldots, X_q$ are i.i.d. Bernoulli with parameter $\geq \tau$ because $z_1, \ldots, z_q$ are i.i.d. and each $X_j = h(z_j)$ for the deterministic function $h(z) = \mathbf{1}\{\pi_i(z) \in D_i\}$. Independence under deterministic maps does *not* require the realized values $\pi_i(z_j)$ to be distinct: even when $|L^{(i)}| \leq q$ (late rounds, where collisions among induced positions are guaranteed by pigeonhole), the joint distribution of $(z_1, \ldots, z_q)$ remains the product measure on L^q, hence $(X_1, \ldots, X_q)$ remains a product of Bernoullis. Therefore

$$\Pr[X_1 = \cdots = X_q = 0] \;=\; \prod_{j=1}^{q} \Pr[X_j = 0] \;\leq\; (1-\tau)^q. \quad \square$$

Theorem 18 (FRI Soundness Above Johnson)**.** *For* $\mathrm{RS}[F, L, k]$ *with* $\mathrm{char}(F) \neq 2$*, $k = 2^m$, L closed under negation, $\delta_J < \delta < 1 - \rho$, and $\Delta(f^{(0)}, \mathrm{RS}_k) > \delta$: FRI with $R = m$ rounds and q queries per round satisfies*

$$\Pr[\text{FRI accepts}] \leq \frac{nR}{|F|} + (1 - \delta/2)^q.$$

Proof. We partition all adaptive prover strategies into two cases according to whether the committed oracles match the honest fold. The prover sees α_i before committing $f^{(i)}$; the analysis below covers every such adaptive strategy.

Strategy A: the prover commits the honest fold at every round. In this case, $f^{(i)}$ is the function $f_{\text{even}}^{(i-1)} + \alpha_i \cdot f_{\text{odd}}^{(i-1)}$ (determined by $f^{(0)}$ and $\alpha_1, \ldots, \alpha_i$). No adaptive freedom remains: the prover's only choice was to fold honestly, and the resulting $f^{(R)}$ is a *fixed* function of $f^{(0)}$ and $(\alpha_1, \ldots, \alpha_R)$.

For any fixed evaluation point $y \in L^{(R)}$, $f^{(R)}(y)$ is multilinear in $(\alpha_1, \ldots, \alpha_R)$: each round introduces α_i with degree exactly 1 (via the even/odd decomposition). Consequently, every syndrome $s_j(f^{(R)}) = \sum_y c_y \cdot f^{(R)}(y)$ is a multilinear polynomial of total degree $\leq R$ in the α_i.

Some syndrome is nonzero as a polynomial. Suppose all syndromes vanished identically (as polynomials in $\alpha_1, \ldots, \alpha_R$). Then $f^{(R)} \in \mathrm{RS}_{k/2^R}$ for *every* choice of $(\alpha_1, \ldots, \alpha_R) \in F^R$. We extract $f^{(0)} \in \mathrm{RS}_k$ from this universal hypothesis as a polynomial-identity argument (*not* a protocol-level reconstruction; FRI does not expose intermediate $f_{\text{even}}^{(i-1)}, f_{\text{odd}}^{(i-1)}$, and an attacker cannot reconstruct $f^{(0)}$ from $f^{(R)}$ at runtime). For each round $i = R, R-1, \ldots, 1$: by the multilinearity in α_i, evaluating at $\alpha_i = 0$ yields $f_{\text{even}}^{(i-1)} \in \mathrm{RS}_{k/2^i}$, and the difference at $\alpha_i = 1$ minus $\alpha_i = 0$ yields $f_{\text{odd}}^{(i-1)} \in \mathrm{RS}_{k/2^i}$ (each as a function of $\alpha_1, \ldots, \alpha_{i-1}$, with $|F| \geq 2$ ensuring two distinct evaluation points). The isomorphism of Remark 4 — which requires $k/2^{i-1}$ even, guaranteed by $k = 2^m$ — then gives $f^{(i-1)} \in \mathrm{RS}_{k/2^{i-1}}$. Iterating to $i = 1$ yields $f^{(0)} \in \mathrm{RS}_k$, contradicting the hypothesis. Hence there exists a syndrome s_j that is a nonzero multilinear polynomial of degree $\leq R$.

By Schwartz–Zippel applied to this single nonzero polynomial:

$$\Pr_{\alpha_1,\ldots,\alpha_R}[s_j(f^{(R)}) = 0] \leq R/|F|.$$

A fortiori, $\Pr[f^{(R)} \in \mathrm{RS}_{k/2^R}] \le R/|F|$.

Strategy B: the prover deviates from the honest fold at some round. We first analyze the *single-deviation* strategy (optimal, as argued below) and then show that multi-deviation does not improve the bound. Let i be the unique round where the committed $f^{(i)}$ differs from $f^{(i-1)}_{\mathrm{even}} + \alpha_i \cdot f^{(i-1)}_{\mathrm{odd}}$, and suppose the prover commits a codeword $g^{(i)} \in \mathrm{RS}_{k/2^i}$: from round $i{+}1$ onward, honest folds of a codeword produce codewords, so $f^{(R)} \in \mathrm{RS}_{k/2^R}$ and the final check passes.

The risk: the consistency check at round i in the standard query model (Definition 15). The deviation set $D_i = \{y \in L^{(i)} : g^{(i)}(y) \neq f^{(i-1)}_{\mathrm{even}}(y) + \alpha_i \cdot f^{(i-1)}_{\mathrm{odd}}(y)\}$ satisfies $|D_i|/|L^{(i)}| \ge \Delta(g^{(i)}, \mathrm{fold}(f^{(i-1)}, \alpha_i)) \ge \delta/2$ for good α_i (justified next). By Lemma 17 with $\tau = \delta/2$: the probability that none of the q i.i.d. base queries catches the round-i deviation is at most $(1-\delta/2)^q$.

We show that $\Delta(f^{(i-1)}, \mathrm{RS}_{k/2^{i-1}}) > \delta/2$ at every round i (for good $\alpha_1, \ldots, \alpha_{i-1}$), so the deviation at round i is always $\ge \delta/2$:

- *Round* 1: $\Delta(f^{(0)}, \mathrm{RS}_k) > \delta$. By Theorem 14: for all but ≤ 1 value of α_1, $\Delta(\mathrm{fold}(f^{(0)}, \alpha_1), \mathrm{RS}_{k/2}) > \delta/2$.
- *Rounds* $2, \ldots, R$ *(distance preserved via BCIKS proximity gap)*: Since $\delta < 1-\rho$, we have $\delta/2 < (1-\rho)/2$, so $\delta/2$ is strictly below the unique decoding radius $(1-\rho)/2$ for every $\mathrm{RS}_{k/2^j}$ (rate ρ is invariant under folding). At any round $j \ge 2$ with $\Delta(f^{(j-1)}, \mathrm{RS}_{k/2^{j-1}}) > \delta/2$: by Lemma 13, $\Delta_{\mathrm{joint}} > \delta/2$. The BCIKS proximity gap [5] (Theorem 1.2, Eq. (1.1), applied at threshold $\gamma = \delta/2 < (1-\rho)/2$) gives: for all but $\le n$ values of α_j, $\Delta(\mathrm{fold}(f^{(j-1)}, \alpha_j), \mathrm{RS}_{k/2^j}) > \delta/2$. Distance is *preserved*, not further reduced.

For a single deviation at round i: the honest fold is $> (\delta/2)$-far from $\mathrm{RS}_{k/2^i}$ for good α_i, and the consistency check catches with probability $\ge 1-(1-\delta/2)^q$. The number of bad α's is ≤ 1 at round 1 (Theorem 14) and $\le n$ at rounds $2, \ldots, R$ (BCIKS above), totaling $\le 1 + (R{-}1)n \le nR$.

Over all R rounds (the prover picks the best round):

$$\Pr[\text{single-deviation Strategy B succeeds}] \le \frac{nR}{|F|} + (1-\delta/2)^q.$$

Multi-deviation does not improve the bound. Suppose the prover deviates at a set of rounds $I = \{i_1, \ldots, i_t\}$, $t \ge 2$, committing arbitrary oracles $g^{(i_j)}$ (not necessarily codewords). Two cases:

(a) *All α_{i_j} are unlucky* (fold is $> \delta/2$-far at every deviation round). Each deviation creates a $(\delta/2)$-far discrepancy on $L^{(i_j)}$. The preimage $\pi_{i_j}^{-1}(D_{i_j}) \subseteq L$ covers $\ge (\delta/2)|L|$ base-level positions, and the union $\bigcup_j \pi_{i_j}^{-1}(D_{i_j})$ is at least as large. The q base queries must miss this entire union, so $\Pr[\text{all checks pass}] \le (1-\delta/2)^q$, the same as single deviation.

(b) *Some α_{i_j} is lucky* (fold is $\le \delta/2$-close). This requires α_{i_j} to be one of the $\le n$ bad values at round i_j. If the prover deviates to a non-codeword $g^{(i_j)} = c + e$ with $c \in \mathrm{RS}_{k/2^{i_j}}$ and $e \neq 0$, subsequent honest folds carry the error: $f^{(R)} = c^{(R)} + e^{(R)}$, where $c^{(R)} \in \mathrm{RS}_{k/2^R}$. The error $e^{(R)}$ has at least one nonzero syndrome (by the Strategy A unfolding argument applied to e), so $\Pr[f^{(R)} \in \mathrm{RS}_{k/2^R}] \le R/|F|$ by Schwartz–Zippel.

By union over the $\le nR$ lucky events and the consistency-check evasion:

$$\Pr[\text{multi-deviation succeeds}] \le \frac{nR}{|F|} + (1-\delta/2)^q.$$

Combining. Every adaptive prover either folds honestly (Strategy A) or deviates at some round (Strategy B, single or multi). The acceptance probability is:

$$\Pr[\text{accept}] \le \max(R/|F|,\ nR/|F| + (1-\delta/2)^q) \le nR/|F| + (1-\delta/2)^q.$$

□

Remark 19 (Why $\delta/2$ does not compound across rounds)**.** *A common concern: the reduced threshold $\delta/2$ could degrade further at each round (to $\delta/4$, $\delta/8$, …), eventually making the consistency check trivial. This does not happen. After round 1 reduces the effective distance from δ to $\delta/2$, this threshold sits strictly below the unique decoding radius $(1-\rho)/2$ (since $\delta/2 < (1-\rho)/2$ for any $\delta < 1-\rho$). The BCIKS proximity gap [5] (Theorem 1.2, Eq. (1.1)) at threshold $\gamma = \delta/2 < (1-\rho)/2$ bounds the number of bad challenges to $\leq n$ per round. The distance remains $\delta/2$ (not $\delta/4$) because we invoke the equal-threshold proximity gap rather than re-applying the two-threshold Theorem 5 iteratively. In short: Theorem 5 enters the unique decoding regime at round 1; the BCIKS proximity gap [5] then locks the distance at $\delta/2$ for all subsequent rounds.*

7 List-Size Landscape and OP2

The results in this section are independent of the FRI soundness theorem (Theorem 18), which uses only the half-threshold CA (Theorem 5) and classical list decoding. They address Open Problem 2 (§8): can one show $M = 0$ on all RS-compatible flats? We give the first p-dependent moment bound on the list size (Theorem 23 for $c = 1$; Corollary 28 for general c), prove the $M = 0$ conjecture is false at deployment code lengths (Proposition 33), and reduce the remaining gap to a worst-case-from-average problem.

Remark 20 (p-independent list-size bounds)**.** *Classical bounds give $M \leq n^2/(t(t-1))$ for $k = 2$ (Fisher double-counting; see Sudan [13], Guruswami–Sudan [14] for general list decoding) and $M \leq \sum_{i=0}^{k-2c} \binom{n}{i}$ for general k (Frankl–Wilson [12] linear algebra method), where $c = n-k-w$ is the codimension excess. Both are independent of p and cannot capture the $p^{-\alpha}$ decay observed computationally [36]: for* RS$[16, 8]$ *at $c = 1$, $M_{\max}(p) \approx 4200 \cdot p^{-3/4}$ across 19 primes from $p = 17$ to 131.*

7.1 Syndrome hyperplane structure and second moment

Classical p-independent bounds (Remark 20) cannot capture the $p^{-\alpha}$ decay observed computationally. The following results explain this decay via the hyperplane geometry of syndrome space and provide the first p-dependent moment bounds on the list size. The results hold for *all* k.

Fix RS$[n, k]$ over $\mathbb{F}_p$ with $p > n$, evaluation domain $L = \{\alpha_1, \ldots, \alpha_n\}$, and codimension excess $c = 1$ (error weight $w = n-k-1$, syndrome dimension $D = n-k$). For each w-subset $E \subset [n]$, the syndrome equations $\sum_{i \in E} e_i \alpha_i^j = s_j$ $(j = k, \ldots, n-1)$ form a $D \times w$ Vandermonde system. With $D = w + 1$, the compatibility condition is a single hyperplane $H_E = \{s \in \mathbb{F}_p^D : \langle n_E, s \rangle = 0\}$, and $M(s) = |\{E \in \binom{[n]}{w} : s \in H_E\}|$.

Lemma 21 (Error-locator normal)**.** *The normal vector of H_E is the coefficient vector of the error-locator polynomial:*

$$n_E = (\Lambda_{E,0}, \ldots, \Lambda_{E,D-1}), \qquad \Lambda_E(x) = \prod_{e \in E} (x - \alpha_e).$$

Proof. Since α_e is a root of Λ_E: $\sum_{\ell=0}^{w} \Lambda_{E,\ell}\, \alpha_e^{k+\ell} = \alpha_e^k \Lambda_E(\alpha_e) = 0$ for each $e \in E$. So the vector $(\Lambda_{E,0}, \ldots, \Lambda_{E,w})$ annihilates every column of the $D \times w$ Vandermonde submatrix V_E. Since V_E has rank w (distinct evaluation points, $p > n$), its left kernel is 1-dimensional and spanned by n_E. □

Lemma 22 (Pairwise independence)**.** *For distinct w-subsets $E_1 \neq E_2$, the normals n_{E_1} and n_{E_2} are linearly independent over $\mathbb{F}_p$.*

Proof. Both Λ_{E_1} and Λ_{E_2} are monic of degree w. Monic polynomials of the same degree are proportional only if equal, requiring identical root sets: $E_1 = E_2$. □

Theorem 23 (Exact second moment of the list size)**.** *For* $\mathrm{RS}[n,k]$ *over* $\mathbb{F}_p$ *with* $p > n$ *and* $c = 1$*, the list size* M *over uniform* $s \in \mathbb{F}_p^D$ *satisfies:*

$$\mathbb{E}[M] = \frac{\binom{n}{w}}{p}, \qquad \mathrm{Var}[M] = \mathbb{E}[M]\left(1 - \frac{1}{p}\right).$$

Proof. Write $M = \sum_E \mathbf{1}_{s \in H_E}$. Each indicator has $\Pr[s \in H_E] = 1/p$. For $E_1 \neq E_2$: by Lemma 22, $H_{E_1} \cap H_{E_2}$ has codimension 2 in $\mathbb{F}_p^D$, so $\Pr[s \in H_{E_1} \cap H_{E_2}] = 1/p^2 = \Pr[s \in H_{E_1}] \cdot \Pr[s \in H_{E_2}]$. The indicators are pairwise uncorrelated, giving $\mathrm{Var}[M] = \binom{n}{w} \cdot \frac{1}{p}(1 - \frac{1}{p})$. □

Remark 24 (Implications for the list-size landscape)**.**
(i) Near-Poisson concentration. For large p*:* $\mathrm{Var}[M] \approx \mathbb{E}[M]$*, matching the Poisson dispersion relation. This justifies the heuristic that* M *concentrates around* $\binom{n}{w}/p$ *and is typically* 0 *once* $p > \binom{n}{w}$*. For* $\mathrm{RS}[16,8]$ *at* $c = 1$*:* $\binom{16}{7} = 11440$*, so* $\mathbb{E}[M] < 1$ *for* $p \geq 11441$*. The power-law decay of* $M_{\max}$ *(see Remark 20) is substantially sharper than the first-moment prediction, suggesting algebraic structure beyond pairwise independence governs the worst case.*
(ii) The $c = 1$ *threshold and* $c \geq 2$ *generalization. At* $c = 1$ *(near capacity,* $w = n-k-1$*) with rate* $\rho = 1/2$*:* $\binom{n}{w} = \binom{n}{n/2-1} \sim 2^n/\sqrt{n}$*, so* $\mathbb{E}[M] < 1$ *requires* $p > 2^n/\sqrt{n}$ *— far beyond any practical field once* n *exceeds a few dozen. At the FRI-relevant distance* $\delta > \delta_J$ *with codimension* $c = n/2 - \lfloor \delta n \rfloor \gg 1$ *(for* $\rho = 1/2$*, just above Johnson gives* $c \approx 0.2n$*): Theorem 26 and Corollary 28 establish* $\mathbb{E}[M] = \binom{n}{w}/p^c$ *with pairwise independence for all pairs with* $|E_1 \cap E_2| \leq w - c$*, giving* $\mathrm{Var}[M] \approx \mathbb{E}[M]$ *(Poisson dispersion). At deployment parameters,* $\mathbb{E}[M] = \binom{n}{w}/p^c \leq 2^{-5.3n} \approx 0$.
(iii) Average vs. worst case. By Markov: $\Pr_s[M \geq 1] \leq \binom{n}{w}/p$*. However, §8, item 2, asks for* $M_{\max} = 0$ *over all syndromes (or all RS-compatible flats). The second-moment bound* $\mathrm{Var}[M] \leq \mathbb{E}[M]$ *does not improve on Markov for this worst-case question. Bridging the average-to-worst-case gap requires exploiting the algebraic structure of the error-locator hyperplane arrangement beyond pairwise independence — for example, showing that the* $\binom{n}{w}$ *hyperplanes, despite their number, avoid all RS-compatible flats when* p *is sufficiently large relative to* n*.*

7.2 Generalization to codimension excess $c \geq 2$

Theorem 23 treats codimension excess $c = 1$, where each w-subset E defines a single hyperplane H_E in syndrome space. For general $c \geq 2$ (error weight $w = n - k - c$, syndrome dimension $D = n - k$), the compatibility condition is $s \in W_E := \mathrm{Im}(V_E)$, the column span of the $D \times w$ Vandermonde submatrix V_E. Since V_E has rank w (MDS, $p > n$), W_E is a w-dimensional subspace of $\mathbb{F}_p^D$, equivalently a codimension-c subspace. Its orthogonal complement $W_E^\perp$ has dimension c and is the left null space of V_E. The following results extend the error-locator normal structure and pairwise independence to this regime. The case $c \geq 2$ is the *FRI-relevant* regime: at rate $\rho = 1/2$ and distance $\delta \approx 0.3$, the codimension excess is $c = n/2 - \lfloor 0.3n \rfloor \approx 0.2n$.

Lemma 25 (Generalized normal space)**.** *For codimension excess* $c \geq 1$*, the left null space of the* $D \times w$ *Vandermonde submatrix* V_E *has dimension* c *and is spanned by the coefficient vectors of*

$$\Lambda_E(x) \cdot x^r, \qquad r = 0, 1, \ldots, c-1,$$

where $\Lambda_E(x) = \prod_{e \in E}(x - \alpha_e)$ *is the error-locator polynomial.*

Proof. Since α_e is a root of Λ_E, the polynomial $\Lambda_E(x) \cdot x^r$ vanishes at every α_e for $e \in E$. Its coefficient vector (as a polynomial of degree $w + r < D = w + c$, embedded in $\mathbb{F}_p^D$) therefore annihilates every column of V_E. The c polynomials $\{\Lambda_E \cdot x^r\}_{r=0}^{c-1}$ have degrees $w, w+1, \ldots, w+c-1$; since Λ_E is monic of degree w, these have leading terms $x^w, x^{w+1}, \ldots, x^{w+c-1}$, so they are linearly independent. The left null space of V_E has dimension $D-w = c$ (distinct evaluation points, $p > n$), and c linearly independent vectors span it. □

Theorem 26 (Pairwise independence at $c \geq 2$)**.** *For distinct w-subsets $E_1 \neq E_2$ with $j = |E_1 \cap E_2| \leq w - c$: the $2c$ normal vectors from Lemma 25 (i.e., the coefficient vectors of $\{\Lambda_{E_1} x^r, \Lambda_{E_2} x^r : 0 \leq r < c\}$) are linearly independent over $\mathbb{F}_p$. Consequently:*

$$W_{E_1} \cap W_{E_2} \text{ has codimension } 2c, \qquad \Pr_s[s \in W_{E_1} \cap W_{E_2}] = \frac{1}{p^{2c}} = \Pr_s[s \in W_{E_1}] \cdot \Pr_s[s \in W_{E_2}].$$

The membership indicators $\mathbf{1}_{s \in W_{E_1}}$ and $\mathbf{1}_{s \in W_{E_2}}$ are pairwise uncorrelated for uniform s.

Proof. Suppose $\sum_{r=0}^{c-1} \big(a_r \Lambda_{E_1}(x) x^r + b_r \Lambda_{E_2}(x) x^r\big) = 0$ as polynomials of degree $< D = w + c$. Set $P(x) = \sum a_r x^r$ and $Q(x) = \sum b_r x^r$ (both of degree $< c$). Then $\Lambda_{E_1} P + \Lambda_{E_2} Q = 0$.

Let $G = \gcd(\Lambda_{E_1}, \Lambda_{E_2}) = \Lambda_{E_1 \cap E_2}$ of degree j, and write $\Lambda_{E_i} = G \cdot A_i$ where $\gcd(A_1, A_2) = 1$ and $\deg A_i = w - j$. Dividing by G: $A_1 P = -A_2 Q$. Since $\gcd(A_1, A_2) = 1$, we need $A_1 \mid Q$. But $\deg A_1 = w - j \geq w - (w - c) = c > \deg Q$, so $Q = 0$, hence $P = 0$.

The $2c$ normals are linearly independent, so $W_{E_1}^{\perp} + W_{E_2}^{\perp}$ has dimension $2c$. Since $\dim(W_{E_1}^{\perp}) = \dim(W_{E_2}^{\perp}) = c$, the intersection $W_{E_1} \cap W_{E_2}$ has codimension $2c$, giving exact independence. □

Remark 27 (Sharpness of the threshold $j \leq w - c$)**.** *The threshold is sharp. For $j > w - c$, set $\ell = j - (w - c) \geq 1$. In the notation of the proof above, $\deg A_i = w - j = c - \ell$. For any polynomial R of degree $< \ell$, setting $Q = A_1 R$ (degree $< c$) and $P = -A_2 R$ (degree $< c$) gives $\Lambda_{E_1} P + \Lambda_{E_2} Q = G(-A_1 A_2 R + A_2 A_1 R) = 0$: a nontrivial relation among the $2c$ normals. The space of such relations has dimension $\ell = j - (w - c)$, so the rank drops from $2c$ to $2c - \ell$.*

In particular, at $j = w - c + 1$ ($\ell = 1$): rank $= 2c - 1$, verified computationally for all tested pairs [36].

Corollary 28 (Exact moments of the list size at $c \geq 2$)**.** *For $\mathrm{RS}[n, k]$ over $\mathbb{F}_p$ with $p > n$, codimension excess $c \geq 1$, and $w = n - k - c$, define $M = |\{E \in \binom{[n]}{w} : s \in W_E\}|$ for uniform $s \in \mathbb{F}_p^D$. Then:*

$$\mathbb{E}[M] = \frac{\binom{n}{w}}{p^c}.$$

For the variance, decomposing by intersection size $j = |E_1 \cap E_2|$:

$$\mathrm{Var}[M] = \underbrace{\sum_{j=0}^{w-c} \binom{n}{w}\binom{w}{j}\binom{n-w}{w-j} \cdot 0}_{\text{pairwise independent: } j \leq w-c} + \underbrace{\sum_{j=w-c+1}^{w} \binom{n}{w}\binom{w}{j}\binom{n-w}{w-j}\Big(\frac{1}{p^{2c-\ell(j)}} - \frac{1}{p^{2c}}\Big)}_{\text{correlated: } j > w-c}, \qquad (2)$$

where $\ell(j) = j - (w - c) \in \{1, \ldots, c\}$ is the rank deficiency from Remark 27.

Proof. Write $M = \sum_E \mathbf{1}_{s \in W_E}$. Each indicator has $\Pr[s \in W_E] = 1/p^c$ since W_E has codimension c. For $\mathbb{E}[M]$: linearity of expectation gives $\binom{n}{w}/p^c$.

For the second moment: $\mathbb{E}[M^2] = \sum_{E_1, E_2} \Pr[s \in W_{E_1} \cap W_{E_2}]$. By Theorem 26, pairs with $j \leq w - c$ contribute $1/p^{2c}$ each — exactly the product of marginals, so their contribution to $\mathrm{Var}[M]$ is zero. Pairs with $j > w - c$ have $\ell(j) = j - (w - c)$ dependent normal vectors (Remark 27), so $W_{E_1} \cap W_{E_2}$ has codimension $2c - \ell(j)$, giving the stated excess probability $1/p^{2c-\ell(j)} - 1/p^{2c}$. □

Corollary 29 (Expected true list size)**.** *Under the hypotheses of Corollary 28, the expected true list size (counting only supports whose Vandermonde error values are all nonzero) satisfies*

$$\mathbb{E}[M_{\text{true}}] = \frac{\binom{n}{w}(p-1)^w}{p^{w+c}} = \frac{\binom{n}{w}}{p^c}\Big(1-\frac{1}{p}\Big)^w.$$

Proof. Write $M_{\text{true}} = \sum_E \mathbf{1}_{s\in W_E,\,\text{all } v_i\neq 0}$. For fixed E, the map $v \mapsto V_E v$ is a bijection $\mathbb{F}_p^w \to W_E$ (the $D \times w$ Vandermonde matrix V_E has rank w since $p > n$, so it has a w-dimensional image). Under uniform $s \in \mathbb{F}_p^D$: $\Pr[s \in W_E] = p^w/p^D = 1/p^c$. Conditional on $s \in W_E$, the unique preimage v with $V_E v = s$ is uniform in $\mathbb{F}_p^w$, so $\Pr[\text{all } v_i \neq 0 \mid s \in W_E] = ((p-1)/p)^w$. By linearity:

$$\mathbb{E}[M_{\text{true}}] = \binom{n}{w}\cdot\frac{1}{p^c}\cdot\Big(\frac{p-1}{p}\Big)^w = \frac{\binom{n}{w}(p-1)^w}{p^{w+c}}. \quad \square$$

Remark 30 (Poisson dispersion at general c)**.** *The variance decomposition* (2) *has a clean structure: (i) Diagonal dominance. The $j = w$ term (same subset, $\ell = c$) contributes $\binom{n}{w}(1/p^c - 1/p^{2c}) \approx \mathbb{E}[M]$ to $\mathrm{Var}[M]$. All other correlated terms ($w - c < j < w$) involve $\binom{w}{j}\binom{n-w}{w-j}$ which is polynomially bounded in n, while the excess probability $1/p^{2c-\ell} \leq 1/p^{c+1}$ decays exponentially in c. For $c \geq 2$ and $p \gg n$: $\mathrm{Var}[M] \approx \mathbb{E}[M]$, matching the Poisson dispersion relation exactly as in the $c = 1$ case. (ii) Average vs. worst case. The Poisson dispersion means that for random syndromes, M concentrates sharply around $\binom{n}{w}/p^c$. At deployment ($c \approx 0.2n$, $p = 2^{31}$): $\mathbb{E}[M] \leq 2^{-5.3n} \approx 0$, so $M = 0$ for overwhelmingly most syndromes. The worst case is governed by evaluation syndromes and the incidence bound; see Remark 31 for the precise structure.*

Remark 31 (Syndrome-compatible supports vs. true list size)**.** *The error-locator normal conditions (the key equations of Berlekamp) count supports E where $\langle \Lambda_E \cdot x^r, s\rangle = 0$ for $0 \leq r < c$, but do not verify that the error values (the Vandermonde solution $V_E^{-1}s$) are all non-zero. Define:*

$$M_{\text{compat}}(s) = |\{E \in \tbinom{L}{w} : \langle \Lambda_E \cdot x^r, s\rangle = 0 \;\forall\, r < c\}|,$$
$$M_{\text{true}}(s) = |\{E \in \tbinom{L}{w} : V_E^{-1}s_{0:w} \text{ has all entries } \neq 0 \text{ and extension conditions hold}\}|.$$

The list-decoding list size is M_{true}; the count $M_{\text{compat}} \geq M_{\text{true}}$ overcounts by including supports where some error values vanish (actual error weight $< w$).

Evaluation syndromes: $M_{\text{compat}} = \binom{n-1}{w-1}$ but $M_{\text{true}} = 0$. Fix $\alpha \in L$ and set $s_\alpha := (1, \alpha, \alpha^2, \ldots, \alpha^{D-1})$. For every w-subset E and $0 \leq r < c$: $\langle \Lambda_E \cdot x^r, s_\alpha\rangle = \alpha^r \Lambda_E(\alpha)$, which vanishes iff $\alpha \in E$. Hence $M_{\text{compat}}(s_\alpha) = \binom{n-1}{w-1}$. However, for any $E \ni \alpha$: the Vandermonde system $V_E \cdot v = (1, \alpha, \ldots, \alpha^{w-1})^T$ has the unique solution $v_\alpha = 1$, $v_\beta = 0$ for $\beta \in E \setminus \{\alpha\}$ — a weight-1 error, not weight-w. Therefore $M_{\text{true}}(s_\alpha) = 0$ for $w \geq 2$. The $\binom{n-1}{w-1}$ "compatible" supports are all false positives: the error-locator conditions are satisfied, but the actual error has weight 1.

True list size: bounded by incidence, p-dependent for large n. Exhaustive computation over all $s \neq 0$ yields $\max_{s\neq 0} M_{\text{true}}(s)$:

n	w	c	d	$\max M_{\text{compat}}$	$\max M_{\text{true}}$	p tested
5	2	1	1	4	2	5, 7, 11, 13, 31
6	3	1	2	10	4	7, 11, 13
7	3	1	2	15	5	7, 11, 13
8	3	1	2	21	6	11, 13
6	3	2	1	10	2	7, 11
7	3	2	1	15	2	7, 11
8	3	2	1	21	2	11, 13
9	3	2	1	$\binom{8}{2}$	3 ($p\leq$23); 2 ($p\geq$29)	11–97
12	3	2	1	$\binom{11}{2}$	4 ($p\leq$17); 3 ($p\geq$19)	13–31
15	3	2	1	$\binom{14}{2}$	5 (p=17); M_∞=3	17; $\text{rank}_{\mathbb{Q}}$
7	3	3	0	15	1	7, 11

Here $d := w - c$ and $\max M_{\text{true}} \leq \binom{n}{d}/\binom{w}{d}$ (the incidence bound [22]) in every case. The table is generated by exhaustive enumeration over $\mathbb{F}_p^D$; scripts and outputs are in [36].

p-dependence. For $n \leq 8$, $\max M_{\text{true}}$ is identical across all tested $p > n$. At $n \geq 9$, M_{true} becomes p-dependent: the $\mathbb{Q}$-rank of the combined normal matrix (Proposition 34) determines whether a configuration persists for all large p, or only for finitely many primes. The asymptotic value M_∞ (Definition 35) is computed in Remark 36.

Provable bounds.

(a) Unique decoding ($c \geq w$): $M_{\text{true}} \leq 1$, since $\min \Delta(\text{RS}[n, n-D]) = D+1 = w+c+1 > 2w$.

(b) Incidence bound ($c < w$): $M_{\text{true}} \leq \binom{n}{d}/\binom{w}{d}$ with $d = w - c$ ([22]), tight at $d = 1$ for small p.

(c) Average: $\mathbb{E}[M_{\text{true}}] = \binom{n}{w}(1 - 1/p)^w/p^c$ (Corollary 29), exponentially small at deployment.

Consequence for the error-locator normal approach. The error-locator normal machinery (Theorems 23 and 26) gives tight average bounds ($\mathbb{E}[M] = \binom{n}{w}/p^c$) but operates on M_{compat}, which overcounts M_{true} by including false-positive supports. The worst-case M_{true} is bounded by the incidence bound (polynomial in n) and can depend on p, but this does not affect the FRI soundness theorem (Theorem 18), which uses only the half-threshold CA (Theorem 5).

Remark 32 (Coset extraction on multiplicative subgroups)**.** *When $L = \langle\omega\rangle$ and $k = 2$, the Vieta conditions on an agreement set S force the characteristic polynomial $\prod(x - \alpha_i)$ to be lacunary (only leading, sub-leading, and constant terms nonzero). For $n = 2^K$ and even t: $\gcd(t-1, n) = 1$, so the lacunary polynomial's root set on L is a single coset of the subgroup of order $\gcd(t-1, n)$ (Lidl–Niederreiter [11]; Golomb–Gong [16] Ch. 3–4), giving $M = O(1)$. The power-of-2 domain structure that FRI uses eliminates the coset-union complication entirely.*

7.3 OP2 obstruction and the average-to-worst-case gap

The strongest form of OP2 asks: does $M_{\max} = 0$ for all centers u when p is sufficiently large? The following elementary construction shows this is *false* at FRI-relevant code lengths.

Proposition 33 (List-size obstruction)**.** *For $\text{RS}[n, k]$ over $\mathbb{F}_p$ with $p \geq 3$, Johnson radius w_J, and codimension excess $c_J = n - k - w_J$: if $c_J \leq \lfloor(n-k-1)/2\rfloor$, then $M_{\max} \geq 2$ for all p.*

Proof. Let $c_1, c_2 \in \text{RS}_k$ with $d(c_1, c_2) = d_{\min} = n-k+1$, and let $S = \{i : c_1(i) \neq c_2(i)\}$ with $|S| = d_{\min}$. The hypothesis gives $2w_J \geq d_{\min}$, so we can choose $B_1, B_2 \subset S$ with $|B_1| = |B_2| = w_J$ and $B_1 \cup B_2 = S$. Define u by: $u = c_1 = c_2$ on $L \setminus S$; $u = c_2$ on $B_1 \setminus B_2$; $u = c_1$ on $B_2 \setminus B_1$; $u \notin \{c_1, c_2\}$ on $B_1 \cap B_2$ (feasible since $p \geq 3$). Then $d(u, c_1) = |B_1| = w_J$ and $d(u, c_2) = |B_2| = w_J$, so $M(u) \geq 2$. □

For rate $\rho = 1/2$: $c_J \approx 0.207n$ and $\lfloor (n-k-1)/2 \rfloor \approx n/4$. At $\rho = 1/2$ with n a power of 2 (the FRI setting), the obstruction applies for $n \geq 2^5 = 32$ (verified computationally up to $n = 64$ [36]; fails at $n = 16$ due to ceiling effects).

What this means. The strongest form of OP2 ($M_{\max} = 0$ for all p) is false at deployment-scale FRI. The average-case bounds ($\mathbb{E}[M] = \binom{n}{w}/p^c$, Corollary 28) and practical $M = 0$ at deployment parameters are unaffected. See §8, item 2, for the precise remaining questions.

7.4 Asymptotic stabilization of M_{true}

The following result explains the p-dependent behavior observed in Remark 31: for large p, the worst-case list size stabilizes at a value determined by the rational rank of a certain normal matrix.

Proposition 34 ($\mathbb{Q}$-rank characterization)**.** *Fix n, w, c with $1 \leq c < w \leq n$ and $D = w + c$. Let $E_1, \ldots, E_k$ be pairwise-compatible w-subsets of $L = \{0, 1, \ldots, n-1\}$ (i.e., $|E_i \cap E_j| \leq w - c$ for $i \neq j$). Form the combined normal matrix*

$$\mathcal{N} = \begin{pmatrix} \Lambda_{E_1} \cdot x^0 \\ \vdots \\ \Lambda_{E_1} \cdot x^{c-1} \\ \hline \Lambda_{E_2} \cdot x^0 \\ \vdots \\ \Lambda_{E_k} \cdot x^{c-1} \end{pmatrix} \in \mathbb{Z}^{kc \times D},$$

where row $\Lambda_{E_i} \cdot x^r$ is the coefficient vector of $\Lambda_{E_i}(x) \cdot x^r$ in the basis $\{1, x, \ldots, x^{D-1}\}$.

(a) Upper bound. If $\operatorname{rank}_{\mathbb{Q}}(\mathcal{N}) = D$, then for all primes p not dividing any $D \times D$ minor of $\mathcal{N}$, the intersection $\bigcap_{i=1}^{k} W_{E_i} = \{0\}$ over $\mathbb{F}_p$. In particular, no nonzero syndrome can simultaneously have all k supports in its list, so $M_{\text{true}}(s) < k$ for every $s \neq 0$.

(b) Lower bound. If $\operatorname{rank}_{\mathbb{Q}}(\mathcal{N}) < D$ and the configuration is nondegenerate (see proof), then for all primes $p > \max(n, kw)$, there exists $s \in \mathbb{F}_p^D \setminus \{0\}$ such that $s \in W_{E_i}$ and the Vandermonde solution $V_{E_i}^{-1} s$ has all entries nonzero, for every $i = 1, \ldots, k$. Hence $M_{\text{true}}(s) \geq k$.

Proof. (a) If $\operatorname{rank}_{\mathbb{Q}}(\mathcal{N}) = D$, then $\mathcal{N}$ has a $D \times D$ submatrix with nonzero determinant $\Delta \in \mathbb{Z}$. For any prime $p \nmid \Delta$: $\operatorname{rank}_{\mathbb{F}_p}(\mathcal{N}) = D$, so the only $s \in \mathbb{F}_p^D$ satisfying $\mathcal{N}s = 0$ is $s = 0$. Since $s \in W_{E_i}$ iff the c normals of E_i annihilate s, the intersection $\bigcap W_{E_i} = \ker(\mathcal{N}) = \{0\}$.

(b) If $\operatorname{rank}_{\mathbb{Q}}(\mathcal{N}) = D - \delta$ with $\delta \geq 1$, the rational kernel $K = \ker_{\mathbb{Q}}(\mathcal{N})$ has dimension δ. Clearing denominators gives δ linearly independent integer vectors in $\ker(\mathcal{N})$. For any prime p larger than the largest prime factor of the denominators (which depends only on n), these reduce to δ independent vectors in $\ker_{\mathbb{F}_p}(\mathcal{N})$, so $\dim_{\mathbb{F}_p}(\bigcap W_{E_i}) \geq \delta$.

It remains to show that some s in the kernel gives all-nonzero Vandermonde solutions. For each support E_i, the map $s \mapsto V_{E_i}^{-1} s$ restricted to $\bigcap W_{E_i}$ is a linear map to $\mathbb{F}_p^w$. Call the configuration *nondegenerate* if for every (i, j), the linear functional $s \mapsto (V_{E_i}^{-1} s)_j$ does not vanish identically on $\bigcap W_{E_i}$. When nondegenerate: each of the kw "bad" sets $\{s : (V_{E_i}^{-1} s)_j = 0\}$ is a proper hyperplane in the p^δ-element kernel, and the union of kw proper hyperplanes omits at least $p^\delta - kw \cdot p^{\delta-1} = p^{\delta-1}(p - kw) > 0$ points when $p > kw$. Nondegeneracy is verified computationally for all configurations in Remark 36; see [36]. □

Definition 35 (Asymptotic worst-case list size)**.**

$$M_\infty(n, w, c) := \max\{k : \exists\ k \text{ pairwise-compatible } w\text{-subsets of } [n] \text{ with } \operatorname{rank}_{\mathbb{Q}}(\mathcal{N}) < D\}.$$

By Proposition 34, $M_{\text{true}}(n, w, c, p) = M_\infty(n, w, c)$ for all sufficiently large primes p, and $M_{\text{true}} \le M_\infty + 1$ for all $p > n$.

Remark 36 (Computed values of M_∞)**.** *For $w = 3$, $c = 2$ ($d = 1$):*

n	M_∞	$\lfloor n/3 \rfloor$	*rank-deficient triples*	*verification*
≤ 8	*2*	≤ 2	*0*	*exhaustive, nondegenerate*
9–10	*2*	*3*	*0*	*exhaustive; 280 (n=9), 2800 (n=10)*
11	*3*	*3*	*1 of 15400*	*nondeg.; no rank-def. 4-tuple*
12	*3*	*4*	*2 of 61600*	*nondeg.; no rank-def. 4-tuple*
13	*3*	*4*	*4 of 200200*	*no rank-def. 4-tuple*
14	*3*	*4*	*6 of 560560*	*no rank-def. 4-tuple*
15	*3*	*5*	*18 of 1401400*	*no rank-def. 4-tuple*

Rank-deficient triples include both translates of

$$\mathcal{C}_a = \big\{(a,\ a{+}2,\ a{+}7),\ \ (a{+}3,\ a{+}8,\ a{+}10),\ \ (a{+}4,\ a{+}5,\ a{+}6)\big\}$$

and non-translate configurations (e.g., $\{(0, 7, 8), (2, 6, 10), (4, 5, 12)\}$ at $n = 13$; 13 of 18 configurations at $n = 15$ are non-translates).

Algebraic obstruction to $M_\infty \ge 4$. For any rank-deficient triple, the kernel of the 6×5 normal matrix is one-dimensional. A fourth disjoint triple $\{d, e, f\}$ preserves rank 4 iff all three roots satisfy a compatibility cubic. Polynomial division reduces integer compatibility to a divisibility condition $(Q(u)) \mid (R(u))$ after centering $u = d - \bar{d}$. For the translate family: $(u^2 - 5) \mid 4u$ (independent of a), yielding exactly the 9 used points as the only integer solutions. Verified computationally [36]: for every rank-deficient triple at $n \le 15$ (including all non-translate configurations), the analogous divisibility condition has solutions only among the 9 used points.

No rank-deficient 4-tuple was found in exhaustive search up to $n = 15$ or in 10,000 random trials at $n \in \{18, 21, 30\}$ [36]. Conjecture: $M_\infty(n, 3, 2) = 3$ for all $n \ge 11$ and $M_\infty = 2$ for $6 \le n \le 10$.

Additional structural results — fiber bounds, pinned-flat uniqueness, Bézout bounds, codimension bounds ($\dim V(r_0{-}1, r_1, r_2) \le w{-}3$), incidence bounds ($M \le \binom{n}{d}/\binom{w}{d}$), and irreducibility of the remainder hypersurface — are developed in a companion note [22] and provide the correct order-of-magnitude picture of M as a function of field size and code parameters.

7.5 Möbius structure and worst-case bound at $c = 2$

The FRI-relevant case is codimension excess $c = 2$ at rate $\rho = 1/2$ (so $w = n/2 - 2$). Here a clean algebraic structure controls how M_{true} depends on the syndrome line. Fix $s_1, s_2 \in \mathbb{F}_p^D$ defining a line $s(\gamma) = s_1 + \gamma s_2$, and let

$$M_{\text{true}}(s_1, s_2) := \#\big\{\gamma \in \mathbb{F}_p : \exists\, E \in \tbinom{L}{w} \text{ s.t. } V_E^{-1}\big(s(\gamma)\big) \text{ has all entries } \neq 0\big\}.$$

Lemma 37 (Möbius structure on cores)**.** *Let $C \subset L$ with $|C| = w - 1$ (a core) and $E = C \cup \{x\}$ for $x \in L \setminus C$. As x varies, the unique γ_E for which the first compatibility equation $\langle \Lambda_E,\ s(\gamma)\rangle = 0$ holds is a degree-1 rational function (Möbius transform) of the evaluation point α_x:*

$$\gamma_E \ = \ \frac{A_C \cdot \alpha_x + B_C}{C_C \cdot \alpha_x + D_C}$$

for constants $A_C, B_C, C_C, D_C \in \mathbb{F}_p$ depending on (s_1, s_2) and C but not on x.

Proof. Factor $\Lambda_E(z) = \Lambda_C(z) \cdot (z - \alpha_x)$. The coefficient of z^j in Λ_E is $\varepsilon_{j-1}(C) - \alpha_x \cdot \varepsilon_j(C)$, where $\varepsilon_j(C)$ are the coefficients of Λ_C (with $\varepsilon_{-1}(C) := 0$). Hence the coefficient vector of Λ_E is linear in α_x. For any $s \in \mathbb{F}_p^D$, $\langle \Lambda_E, s \rangle = \sum_j \varepsilon_{j-1}(C) s_j - \alpha_x \sum_j \varepsilon_j(C) s_j$ is linear in α_x. Setting $\langle \Lambda_E, s_1 + \gamma s_2 \rangle = 0$ and solving for γ gives a ratio of two linear functions of α_x. □

Theorem 38 (Worst-case bound at $c = 2$)**.** *At codimension excess $c = 2$ over any $\mathbb{F}_p$ with $p > n$,*

$$M_{\text{true}}(s_1, s_2) \;\le\; M_{\text{compat}}(s_1, s_2) \;\le\; \min\Big(p,\; 2\binom{n}{w-1}\Big)$$

for every line (s_1, s_2).

Proof. A pair (E, γ) contributes to M_{compat} iff $\langle \Lambda_E, s(\gamma) \rangle = 0$ *and* $\langle \Lambda_E \cdot x, s(\gamma) \rangle = 0$ (the two normals at $c = 2$). Decompose $E = C \cup \{x\}$ with $|C| = w - 1$ and $x \in L \setminus C$. Both inner products are bilinear in (α_x, γ): linear in α_x by the shift identity used in Lemma 37, linear in γ by the line parameterization. Eliminating γ between the two bilinear equations yields a polynomial in α_x of degree ≤ 2.

Two cases. If this polynomial is not identically zero, it has at most two roots, so at most two values of α_x (and hence at most two pairs (α_x, γ_E)) arise from this core. Summing over $\binom{n}{w-1}$ cores: $M_{\text{compat}} \le 2\binom{n}{w-1}$. If the polynomial vanishes identically, the line $\ell(\gamma) = s_1 + \gamma s_2$ lies inside the union $\bigcup_{x \in L \setminus C} W_{C \cup \{x\}}$, but $M_{\text{compat}} \le p = |\mathbb{F}_p|$ holds trivially. Either way the stated bound holds. Since $M_{\text{true}} \le M_{\text{compat}}$, the bound passes to M_{true}. □

The bound complements the incidence bound: $M_{\text{true}} \le \min(p,\; 2\binom{n}{w-1},\; \binom{n}{d}/\binom{w}{d})$ with $d = w - c$. The $2\binom{n}{w-1}$ factor is loose at small n (e.g. $n = 15$, $w = 6$, $c = 2$: bound 210 vs. empirical $\max M_{\text{true}} = 5$ [36]) but it is the only p-independent upper bound currently known at the FRI-relevant codimension $c = 2$ that is universal in (s_1, s_2).

Sub-Poisson concentration on bin sizes

Theorem 38 bounds the Möbius-image bin size $M_\gamma(s_1, s_2) := |\{E \in \binom{L}{w} : \gamma_E(s_1, s_2) = \gamma\}|$ by $2\binom{n}{w-1}/w$ universally, but this ignores the negative correlation induced by Lemma 37. A sharper *in-expectation* bound recovers the birthday rate.

Theorem 39 (Sub-Poisson tail on M_γ at $c = 2$)**.** *Fix $\rho = k/n$, $w = n - k - 2$, $D = n - k$. For any prime $p > \max(n^2, 2(n - w + 1))$, with probability $\ge 1 - 1/p$ over uniform $(s_1, s_2) \in \mathbb{F}_p^D \times \mathbb{F}_p^D$:*

$$\max_{\gamma \in \mathbb{F}_p} M_\gamma(s_1, s_2) \;\le\; \frac{\binom{n}{w}}{p} \;+\; 4\sqrt{2 \log(p) \cdot \frac{\binom{n}{w}}{p}}.$$

Combined with $M_\gamma \le p$ trivially, in the same probability event:

$$\max_{\gamma \in \mathbb{F}_p} M_\gamma(s_1, s_2) \;\le\; C_0 \cdot \sqrt{\binom{n}{w} \cdot \log p}, \qquad C_0 \le 16.$$

Proof sketch. Set $X_E := \mathbf{1}[\gamma_E(s_1, s_2) = \gamma]$, with marginal $\Pr_s[X_E = 1] = 1/p$ (Lemma 37 gives a Möbius bijection on the leaf α_x, so for fixed core C and uniform (s_1, s_2), the value γ_E is uniform on $\mathbb{F}_p$). Pairs (E_1, E_2) classify by overlap r: at $r \le w - 2$, the four normals are linearly independent (Theorem 26), so $\text{Cov}_s(X_{E_1}, X_{E_2}) = 0$; at $r = w - 1$, the Möbius bijectivity excludes equal images from distinct leaves except on a degenerate-Möbius locus of measure $\le 1/p$,

giving $\mathrm{Cov}_s(X_{E_1}, X_{E_2}) \le 0$. Hence $\{X_E\}$ are pairwise non-positively correlated indicators with $\mathrm{Var}[M_\gamma] \le \mathbb{E}[M_\gamma] = \binom{n}{w}/p =: \mu$.

The standard sub-Poisson tail bound for indicators with non-positive pairwise covariance (Janson [26], Dubhashi–Ranjan [27]; the Bernstein argument requires only the non-positive correlation hypothesis) gives

$$\Pr_s\big[M_\gamma \ge \mu + t\big] \;\le\; \exp\!\left(-\frac{t^2}{2(\mu + t/3)}\right).$$

Setting $t = 4\sqrt{2\log(p)\cdot\mu}$ and applying a union bound over $\gamma \in \mathbb{F}_p$:

$$\Pr_s\big[\max_\gamma M_\gamma \ge \mu + t\big] \;\le\; p\cdot\exp\!\left(-\frac{16\cdot 2\log p\cdot\mu}{2(\mu + (4/3)\sqrt{2\log p\cdot\mu})}\right) \;\le\; \frac{1}{p}$$

for $\mu \ge \log p$, which holds for $p > p_{\min}(n)$ as stated. The combined bound with $M_\gamma \le p$ follows from the AM–GM step $\min(\mu, \sqrt{\binom{n}{w}\log p}) \le \sqrt{\binom{n}{w}\log p}$. Full proof in companion note [24]. □

The high-probability bound matches the empirical $0.66\cdot 1.36^n$ envelope (Remark 40) up to $\sqrt{\log p}$ factors. A *worst-case* extension over all (s_1, s_2) is open (§8, item 2); the natural higher-moment Markov approach is structurally blocked by configurations where three supports share a common $(w-c)$-element core (companion note [24], script in [36]).

Remark 40 (Empirical exponential growth at the worst case)**.** *At $c = 2$ and $w = n/2-2$ (the FRI-relevant regime), exhaustive search over (s_1, s_2) pairs, prime fields, and supports up to $n = 24$ [36] gives*

$$\max_{s_1,s_2,p} M_{\mathrm{true}}(s_1, s_2; p) \;\approx\; 0.66\cdot 1.36^n \qquad (R^2 > 0.99,\; n \in \{8, 10, \ldots, 24\}).$$

The peak prime tracks $p_{\mathrm{peak}} \approx \sqrt{\binom{n}{w}}$, the regime where bin sizes (the number of supports projecting to each $\mathbb{F}_p$ value of the linear part of γ) become $\Theta(\sqrt{\binom{n}{w}})$ on average and concentrate within a constant factor of this average. This rules out any polynomial-in-n worst-case bound on M_{true} at $c = 2$, settling the asymptotic order. The exponential growth is, however, slower than $\sqrt{\binom{n}{w}} \sim 1.41^n$, consistent with a constant-factor concentration of the bin distribution. The corresponding ratio M_{true}/p, the FRI-relevant quantity, decays exponentially: $0.676^n \to 0$.

7.6 Phase diagram across codimension excess c

Theorem 38 and Remark 40 characterize the worst-case list size M_{true} at $c = 2$. The picture across all c is a clean three-phase diagram:

- $c = 1$ *(covering radius)*: $M_{\mathrm{true}} = \min(p, \binom{n}{w})$ asymptotically (the saturating regime; trivially $M_{\mathrm{true}} \le \binom{n}{w}$ from one constraint).
- $c = 2$: exponential, $M_{\mathrm{true}} \sim 0.66\cdot 1.36^n$ (Remark 40); governed by the Möbius coupling on $(w-1)$-cores (Lemma 37).
- $c \ge 3$: empirically *linear* in n at sufficiently large p, $M_{\mathrm{true}} \le \lfloor (2D-1)/c \rfloor$ where $D = n - k$.

The deployment-relevant regime for FRI is $c = \Theta(n)$ (specifically $c \approx 0.2n$ above the Johnson radius at rate 1/2), so the third phase is the one that controls FRI soundness analysis.

Conjecture 41 (Open-Set Rank Lemma at $c \ge 3$)**.** *For supports $E_1, \ldots, E_m \subset L$ each of size $w = D - c$ with $c \ge 3$, and distinct $\gamma_1, \ldots, \gamma_m \in \mathbb{F}_p^*$, let $A \in \mathbb{F}_p^{mc\times 2D}$ be the constraint matrix with*

row blocks $[N_{E_i} \mid \gamma_i N_{E_i}]$, where N_{E_i} collects the c error-locator normals from E_i (Lemma 25). For all p above an effective threshold $p_0(n, k, c)$ (polynomial in n):

$$\textit{either} \quad \operatorname{rank}(A) = \min(mc, 2D), \quad \textit{or} \quad \exists\, i \textit{ with } \langle n_0(E_i), s_2 \rangle = 0 \textit{ for all } (s_1, s_2) \in \ker A.$$

Equivalently, every (s_1, s_2) with $M_{\text{true}}(s_1, s_2) = m$ satisfies $m \le \lfloor (2D-1)/c \rfloor$.

Remark 42 (Empirical evidence and small-p counterexamples)**.** *Conjecture 41 has been verified at large p for (n, k) up to $n = 40$ with c up to 9 via random-syndrome search [25, 36]; selected data points (rate $1/2$, ≥ 14000 random configs each):*

n	k	c	bound $\lfloor (2D-1)/c \rfloor$	observed $\max M_{\text{true}}$	p tested
20	10	5	3	≤ 2	large
24	12	5	4	≤ 2	large
28	14	6	4	≤ 4	large
32	16	7	4	≤ 4	large
36	18	8	4	≤ 4	large
40	20	9	4	≤ 4	large

The observed M_{true} is constant (≤ 4) at the Johnson-radius codimension for rate $1/2$, well within the $O(n/c)$ envelope.

At small primes, Conjecture 41 can fail due to mod-p algebraic coincidences. Two structural counterexamples are known:

- *Triangle obstruction at $c = 2$. For $n = 8$, $w = 2$, $p = 113$: the four supports $\{(3,6), (5,6), (3,5), (0,1)\}$ realize $M_{\text{true}} = 4$, exceeding the would-be bound $\lfloor (2D-1)/c \rfloor = 3$. The first three supports form a K_3 on vertices $\{3, 5, 6\}$, providing a row dependency that the fourth disjoint support extends. Consistent with the $c = 2$ phase being exponential, not linear.*
- *Tetrahedron obstruction at $c = 3$, small p. For $n = 12$, $w = 3$, $p = 61$: the four supports $\{(1,4,5), (1,4,8), (4,5,8), (1,5,8)\}$ are exactly the size-3 subsets of $\{1, 4, 5, 8\}$, forming a K_4 tetrahedron, and realize $M_{\text{true}} = 4 > \lfloor (2D-1)/c \rfloor = 3$. The lemma fails at this small p but holds for $p > p_0(n, k, c)$.*

The general obstruction at general c is the $(w+1)$-clique — all size-w subsets of a $(w+1)$-vertex set. These configurations are realizable only at primes below an explicit threshold; the conjecture asserts an effective Schwartz–Zippel-style bound on this threshold.

The connection to FRI deployment: the FRI verifier accepts a low-distance codeword on the evaluation domain only when the syndrome line hits one of M_{true} many error patterns. At deployment scale, $c \approx 0.2n$ exceeds 3, putting us in the third phase of the diagram. At fixed rate, $\lfloor (2D-1)/c \rfloor$ is a *constant* in n (equal to 4 at rate $1/2$ near the Johnson radius); Conjecture 41 therefore predicts $M_{\text{true}} = O(1)$ uniformly in code length at deployment parameters, consistent with the empirical row of the table.

8 Open Problems

We list open problems in order of deployment relevance.

1. **The half-threshold ratio is optimal; equal threshold has exponential constant.** Theorem 5 proves $\varepsilon_{\text{ca}}(C, \delta/2, \delta) \le 1/|F|$ for any linear code — the 1:2 ratio is optimal for the "two bad γ" argument type (Remark 12).

At *equal* threshold ($\delta_{\text{hd}} = \delta$), the erasure-pattern analysis of Remark 12 gives $\varepsilon_{\text{ca}}(C, \delta, \delta) = \binom{n}{w}/|F|$: each of the $\binom{n}{w}$ erasure patterns contributes exactly one decodable γ via a linear equation. This is $O(1)/|F|$ for fixed n, but $\binom{n}{w} \sim 2^{nH(\delta)}$ grows exponentially in the code length. At FRI scale ($n \gg \log|F|$): $\binom{n}{w} \gg |F|$, so the equal-threshold bound is vacuous and the $2\times$ query overhead is intrinsic. Reducing it below $2\times$ requires a non-CA proof technique.

The remaining open question is whether the 1:2 ratio can be improved by arguments that use *more than two* bad γ's or bypass the CA framework entirely. Three bad γ's give an overdetermined system (three equations in two unknowns (f_1, f_2)), yielding polynomial constraints on h_3 given h_1, h_2; but the additional constraint does not tighten the inclusion-exclusion bound beyond the two-γ result. Bypassing CA is pursued in companion work via an elementary action–orbit theorem on cyclic FRI domains; rigorous toy bounds and a refined "one-orbit" conjecture are established there, with universal extension still open.

2. **Worst-case M_{true}: structure and bounds.** Proposition 33 shows $M_{\max} \geq 2$ at FRI parameters ($n \geq 32$, rate 1/2) for all p; the strongest form of the $M = 0$ conjecture is false. The average-case bounds $\mathbb{E}[M] = \binom{n}{w}/p$ (Theorem 23) for $c = 1$ and $\mathbb{E}[M] = \binom{n}{w}/p^c$ (Corollary 28) for general c, both with Poisson dispersion $\text{Var}[M] \approx \mathbb{E}[M]$ (via pairwise independence of error-locator normals, Theorem 26), show $M = 0$ for a random syndrome once $p^c > \binom{n}{w}$. At deployment ($c \approx 0.2n$, $p = 2^{31}$): $\mathbb{E}[M] \leq 2^{-5.3n} \approx 0$.

   Evaluation syndromes are false positives. Remark 31 shows that the error-locator normal *compatibility* count M_{compat} reaches $\binom{n-1}{w-1}$ at evaluation syndromes $s_\alpha = (1, \alpha, \ldots, \alpha^{D-1})$. But these $\binom{n-1}{w-1}$ "compatible" supports all produce weight-1 errors ($v_\alpha = 1$, $v_\beta = 0$ for $\beta \neq \alpha$), so $M_{\text{true}}(s_\alpha) = 0$ for $w \geq 2$. The $\binom{n-1}{w-1}$ is a failure of the error-locator normal machinery to enforce non-zero error values, not a genuine list-size obstruction.

   Provable bounds on M_{true}. At $c \geq w$ (within the unique decoding radius): $M_{\text{true}} \leq 1$. At $c < w$: $M_{\text{true}} \leq \binom{n}{d}/\binom{w}{d}$ with $d = w - c$ (the incidence bound [22]).

   p-dependence and asymptotic stabilization. Proposition 34 proves that $\max M_{\text{true}}$ stabilizes at the value M_∞ (Definition 35) for all sufficiently large p, and Remark 36 tabulates computed values. Corollary 29 gives the exact formula $\mathbb{E}[M_{\text{true}}] = \binom{n}{w}(1 - 1/p)^w/p^c$.

   What remains open: (i) a closed-form for $M_\infty(n, w, c)$ (the data at $w = 3$, $c = 2$ suggest $M_\infty = 3$ for all $n \geq 11$; see Remark 36); (ii) whether any worst-case list-size bound can improve the proximity gap beyond the half-threshold $\delta/2$; (iii) the *birthday-bound conjecture* : extending Theorem 39 from in-expectation to a uniform worst-case bound. The in-expectation form

$$\max_{\gamma} M_\gamma(s_1, s_2) \;\leq\; \frac{\binom{n}{w}}{p} + 4\sqrt{2\log(p) \cdot \frac{\binom{n}{w}}{p}} \qquad \text{w.p. } \geq 1 - 1/p$$

   is proved in Theorem 39 via a Bernstein-style sub-Poisson tail on the variance bound from Lemma 37. The conjectural worst-case version is

$$\max_{s_1, s_2} \max_{\gamma \in \mathbb{F}_p} M_\gamma(s_1, s_2) \;\leq\; C_1 \cdot \frac{\binom{n}{w}}{p} \qquad \text{(uniform in } (s_1, s_2)\text{)}.$$

   If true, this combines with $M_\gamma \leq p$ to yield $\max M_{\text{true}} = O(\sqrt{\binom{n}{w}}) = O(2^{n/2})$, matching the 1.41^n envelope of Remark 40 and bridging the gap between the universal bound $2\binom{n}{w-1}$ (Theorem 38) and the empirical $0.66 \cdot 1.36^n$.

Status of the worst-case extension. The natural higher-moment + Markov + union-bound approach for promoting Theorem 39 to all (s_1, s_2) is structurally blocked. The k-wise extension of the pairwise-independence theorem (Theorem 26) is empirically false at $c = 2$ for $k \geq 3$: three supports E_1, E_2, E_3 sharing a common $(w-c)$-element core C^* have all pairwise overlaps $|E_i \cap E_j| = w - c$ (saturating the pairwise-independence hypothesis) but their six normals all lie in the principal ideal $(\Lambda_{C^*}) \subset \mathbb{F}_p[x]/(x^D)$ of codimension $w - c$, hence dimension $2c$, so the rank is at most $2c < 3c$. Joint compatibility probability is then $1/p^{2c}$ rather than the independent prediction $1/p^{3c}$, making the third moment super-Gaussian and ruling out sub-Gaussian Markov tails. Empirical verification: companion script [36] finds thousands of rank-deficient triples for $n \in \{12, 14, 16\}$. This is a higher-order analog of Lemma 37 (the two-leaf case at $c = 2$): the multi-leaf coupling at $c \geq 2$ requires sequence-school cross-correlation tools beyond what the error-locator normal machinery alone provides.

The $c \geq 3$ regime is qualitatively different (§7.6). The same rank-deficiency that creates exponential growth at $c = 2$ disappears at $c \geq 3$, and Conjecture 41 (Open-Set Rank Lemma) predicts $M_{\text{true}} \leq \lfloor (2D-1)/c \rfloor$ at sufficiently large p, linear in n. The deployment regime $c = \Theta(n)$ falls in this third phase, where the conjecture predicts $M_{\text{true}} = O(1)$ at the Johnson radius and the empirical data confirms $M_{\text{true}} \leq 4$ for rate $1/2$ across $n = 20$ to 40. Resolving Conjecture 41 for $c \geq 3$ would close the worst-case worry at deployment scale even without progress on the $c = 2$ birthday-bound conjecture above.

3. **Odd k.** The $k = 2^m$ condition excludes odd k (Remark 4). Standard FRI uses $k = 2^m$, so this is not a deployment blocker.

Resolved: characteristic 2 was originally listed as an open problem. Appendix B resolves it completely: replacing the multiplicative $\pm\sqrt{y}$ decomposition with the additive subspace-polynomial fold $s(x) = x^2 + \beta x$ yields FRI soundness (Theorem 61) with the same $nR/|F| + (1-\delta/2)^q$ bound over binary tower fields. The key observation is that the half-threshold CA (Theorem 5) is characteristic-independent — it uses only linearity and Hamming distances. Only the coupling lemma requires the fold replacement, and the additive fold preserves the isomorphism $\text{RS}_k \xrightarrow{\sim} \text{RS}^{=2}_{k/2}$ for even k. This means the safety net now covers Binius-style deployments on $\mathbb{F}_{2^{128}}$ tower fields with no additional assumptions.

9 Deployment Roadmap

Everything in this paper is deployable today. The FRI, STIR, and WHIR protocol structure — folding maps, Merkle commitments, prover messages, and verifier checks — remains unchanged. The deployment parameter that changes is the query count q, and increasing q increases proof size, verifier hashing, calldata, and wrapper-circuit cost linearly in the query-heavy portion of the verifier. The theorem bounds IOPP soundness of the interactive proximity test; deployment-grade non-interactive knowledge-soundness (via Fiat–Shamir, DEEP-quotienting, polynomial commitment wrapping) requires standard but separate compilation arguments. We summarize what our results provide in concrete parameters.

	Conjectured (zero-loss CA)	**Proved (this paper)**
Status	Unproven	Theorem
Security budget ($q = 128$, $R = 20$, $\delta = 0.4$, rate $\rho = 1/2$):		
Commit-phase (BabyBear base, $\lvert F\rvert = 2^{31}$)	$\sim$7 bits	$\sim$7 bits
Commit-phase (BabyBear$_4$, $\lvert F\rvert = 2^{124}$)	$\sim$100 bits	$\sim$100 bits
Query-phase ($q = 128$, δ vs $\delta/2$)	$(1-\delta)^{128} \approx 2^{-94}$	$(1-\delta/2)^{128} \approx 2^{-41}$
Total (BabyBear base)	$\min(7, 94) = 7$	$\min(7, 41) = 7$
Total (BabyBear$_4$, extension)	$\min(100, 94) = 94$	$\min(100, 41) = 41$
Query overhead	baseline	2–2.2$\times$ baseline
Protocol-structure changes	none	**none** (q recalibrated)
Protocols covered	FRI, STIR, WHIR	**FRI, STIR, WHIR** (proved; see matrix)†
Characteristic	odd only	**any (incl. char 2)**†
Survives (δ, δ)-CA exponential constant (§8)	no	**yes**

† FRI soundness is proved as a formal theorem in both characteristics (Theorem 18 for char $\neq 2$; Theorem 61 for char $= 2$). STIR's commit term is $(k+2)R/|F|$ (Theorem 46), not $nR/|F|$; the security rows above are for FRI. The STIR/WHIR extension to characteristic 2 follows from the same argument (the batch-CA theorem is characteristic-independent and the additive coupling lemma replaces the multiplicative one), as stated in Remark 63; we do not formalize it as a separate theorem.

Proof-coverage matrix. The precise scope of what is formally proved vs. deferred:

Protocol / composition step	Status	Reference
FRI above Johnson, char$(F) \neq 2$	**fully proved**	Thm 18
FRI above Johnson, char$(F) = 2$	**fully proved**	Thm 61
Half-threshold CA, any linear code (RVW13 [28])	classical	Thm 5
Batch CA, any field	**fully proved**	Thm 44
STIR soundness, char$(F) \neq 2$	**fully proved**	Thm 46
STIR soundness, char$(F) = 2$	**sketched**	Rem. 63
WHIR, linear-subcode CA core	**fully proved**	Cor 48
WHIR, affine-subcode composition	**fully proved**	Thm 52
Circle-FRI variant (Stwo)	**fully proved**	Thm 69
Fiat–Shamir / non-interactive soundness	**fully proved**	Thm 73; Appendix D
DEEP-FRI knowledge soundness	**fully proved**	Thm 75; Cor 76
Composition with Plonkish lookups (logUp)	**not addressed**	future work

The "safety net" covers the full proximity-testing stack: information-theoretic IOPP soundness (Theorems 18, 61, 69, 46, 52), Fiat–Shamir compilation to non-interactive soundness (Theorem 73), and DEEP-FRI knowledge soundness for the polynomial commitment scheme (Theorem 75, Corollary 76). The remaining "not addressed" item (Plonkish/logUp composition) is the arithmetization layer, which is protocol-specific and orthogonal to proximity testing.

Adoption requires no changes to protocol structure: the FRI/STIR/WHIR prover algorithm, round structure, folding equations, and proof format are identical for all multiplicative-FRI and additive-FRI deployments. For Binius-style deployments over binary tower fields ($\mathbb{F}_{2^{128}}$), the additive FRI folding (Appendix B) achieves the same soundness bound with the same proof structure; no protocol modification beyond using subspace-polynomial folding $s(x) = x^2 + \beta x$ in place of $y = x^2$ is required — this is the canonical characteristic-2 analogue introduced by Ben-Sasson–Bentov–Horesh–Riabzev [18] and further developed in additive FRI and Binius [19]; our contribution is a *soundness theorem* for this fold above the Johnson bound, not the fold itself. The verifier implementation must still be parameterized by the chosen query count q, so changing q changes the number of Merkle openings and consistency checks. Remaining engineering tasks:

1. **Interactive security parameter tuning**: the commit-phase term $nR/|F|$ matches BCHKS up to constants; on BabyBear_4 ($|F| = p^4 \approx 2^{124}$), this gives ∼100 bits. The binding constraint is the query-phase term $(1-\delta/2)^q$: at $q = 128$ this gives ∼41 bits for $\delta = 0.4$. Reaching 128-bit interactive query-phase security at $\delta = 0.4$ requires $q \approx 398$ (a ∼3.1× increase from typical $q = 128$), but BabyBear_4 remains capped at ∼100 total interactive bits by the commit term; the non-interactive setting is stricter because of the Fiat–Shamir Q-loss in the next item. Multi-rate reference for deployers:

ρ	$\delta_J = 1-\sqrt{\rho}$	δ	q (conj. zero-loss)	q (ours)
1/2	0.293	0.30	~ 195	~ 427
1/2	0.293	0.40	~ 136	~ 311
1/4	0.500	0.55	~ 87	~ 216
1/8	0.646	0.70	~ 58	~ 161

   All rows target $\lambda = 100$ bits of query-phase security, with δ chosen strictly above the rate-dependent Johnson radius $\delta_J = 1-\sqrt{\rho}$ so that our theorem applies (the open zone is $\delta_J < \delta < 1-\rho$). Below δ_J, BCHKS's zero-loss bound is the relevant baseline and our reduced-threshold result offers no improvement; the table therefore omits below-Johnson operating points.
2. **Non-interactive parameter tuning**: Fiat–Shamir compilation multiplies the interactive RBR error by the adversary's random-oracle query budget Q and adds the Merkle/random-oracle binding term $3(Q^2+1)/2^\kappa$ (Theorem 73). Table 3 gives the resulting deployment budget for $Q = 2^{64}$ and $\delta = 0.4$. The entries are generated by `notes/scripts/deployment_params.py`; "commit" means no finite q can reach the target because the $Q \cdot nR/|F|$ term already exceeds the target.

Config	$\lvert F\rvert$ bits	q now	Interactive	NI now	q_{64}	q_{100}	q_{128}
BabyBear base	31	128	6.7	−57.3	commit	commit	commit
BabyBear_4	124	128	41.2	−22.8	commit	commit	commit
BabyBear_5	155	128	41.2	−22.8	398	commit	commit
BabyBear_7	217	128	41.2	−22.8	398	510	597
Mersenne31_4	124	128	41.2	−22.8	commit	commit	commit
Goldilocks_2	128	128	41.2	−22.8	commit	commit	commit
Binary tower 2^{128}	128	128	41.2	−22.8	commit	commit	commit

   Table 3. Non-interactive FRI security under the loose Theorem 73 accounting, with $n = 2^{20}$, $R = 20$, $\delta = 0.4$, $Q = 2^{64}$, $\kappa = 256$ except BabyBear_7 where $\kappa = 384$ so the hash-binding term is not the bottleneck. "Interactive" and "NI now" are proved bits at $q = 128$; negative NI bits mean the theorem gives no non-trivial non-interactive bound at that query count after the Q-loss. The q_λ columns give the required query count for λ-bit non-interactive security when the commit and hash terms permit it.
3. **Verifier cost update**: increasing q is a parameter change with real cost. It leaves the IOP structure intact but increases Merkle path verifications, FRI consistency checks, proof bytes, calldata, and any recursive/wrapper verifier circuit work roughly linearly in q.
4. **Reproducibility**: the deployment table above is generated by `notes/scripts/deployment_params.py`; the paper artifact, Lean status, proximity-gap figure, and saved computational outputs are documented in `REPRODUCING.md`. Companion notes still marked "in preparation" are not required for Table 3; central claims that depend on them should either be promoted into the paper or linked to fixed public artifacts before prize submission.
5. **Formal verification**: the half-threshold CA bound (Theorem 5), the FRI coupling chain, and the batch-CA bound (Theorem 44) have been formally verified in Lean 4 with Mathlib [36]. Remaining: fill in the RS isomorphism axioms (standard polynomial degree bounds) and extend to the full probability-level soundness statement.

10 Conclusion and Outlook

We have proved an unconditional FRI soundness theorem above the Johnson bound for plain Reed–Solomon codes on the evaluation domains used in deployment — in every characteristic, under the structural hypotheses ($k = 2^m$, L admitting a fixed-point-free involution) that are standard for deployed FRI. The bound $\varepsilon_{\mathrm{FRI}} \le nR/|F| + (1-\delta/2)^q$ is a theorem, not a conjecture, covering multiplicative FRI, circle FRI, additive FRI, STIR, and WHIR (Section 9). Deployment requires changing only the query count q at the protocol-parameter level; the prover/verifier equations remain unchanged, while verifier cost scales with the chosen q.

The CA framework is settled. The half-threshold CA bound is the classical 1:2 proximity-gap inequality of RVW13 [28]. Our contribution is applying it *above the Johnson bound*: the observation $\delta/2 < (1-\rho)/2$ places the protocol into the unique-decoding regime after round 1, where BCIKS [5] locks the distance. The 1:2 ratio is optimal at deployment scale (Remark 12), and the $M = 0$ approach is obstructed at FRI-relevant code lengths (Proposition 33). The 2× query overhead is intrinsic to any CA-based proof at current parameters.

Beyond the CA framework. Every published FRI soundness proof factors through a CA lemma (§1.7). Reducing the 2× overhead requires proving FRI soundness *without* this intermediate step — for example, via a direct algebraic analysis of the fold map on multiplicative subgroups, or via character-sum techniques exploiting that $L = \langle\omega\rangle$ is a cyclic group. Cross-correlation distributions from the Helleseth–Golomb–Gong tradition [15, 16] and Weil-type bounds on incomplete character sums [17] are natural candidates. Companion work in preparation pursues this direction: an elementary action–orbit theorem on cyclic FRI domains gives a tight rigorous bound for two-round FRI at toy parameters $(n_0, k_0) = (32, 8)$, with a refined "one-orbit" conjecture targeting universal extension. Whether FRI admits a non-CA soundness proof with overhead $< 2\times$ at deployment scale is a key remaining question for FRI soundness theory.

Practical bottom line. Prior to this work, above the Johnson bound: zero bits of proven FRI soundness. After: $\varepsilon_{\mathrm{FRI}} \le nR/|F| + (1-\delta/2)^q$, giving ∼41 interactive query bits at $q = 128$ or ∼128 interactive query bits at $q \approx 398$ for $\delta = 0.4$. Total security is still capped by the commit term (for example ∼100 interactive bits on BabyBear$_4$ at $n = 2^{20}, R = 20$). The 2× query overhead vs. the conjectured zero-loss baseline is optimal within the CA framework; against what is actually proved, our bound is the tightest available.

Roadmap. Three follow-ups: (a) complete the Lean 4 formalization [36] (the core CA bound and coupling chain are proved; the RS isomorphism axioms and probability-level soundness statement remain); (b) tight bounds on worst-case M_{true} and whether they can improve the proximity gap beyond $\delta/2$ (§8, item 2); (c) the non-CA route: companion work in preparation establishes an elementary action–orbit theorem on cyclic FRI domains and proves a tight rigorous toy bound, with a one-orbit conjecture targeting universal extension — preliminary evidence that the $< 2\times$ regime is reachable outside the CA framework, though deployment-scale rigor remains open.

A Extension to STIR and WHIR

We show that the half-threshold CA bound (Theorem 5) extends to the batch setting required by STIR [8] and to the constrained RS subcodes required by WHIR [9]. This makes the safety net

directly applicable to deployed STIR/WHIR-based systems.

A.1 Batch Correlated Agreement

STIR and its variants test proximity of multiple functions simultaneously by taking a random linear combination and checking the result. This requires a *batch* CA bound.

Definition 43 (Batch Correlated Agreement)**.** *For code C over domain L, B functions $f_1, \ldots, f_B : L \to F$, and independent uniform $\gamma_1, \ldots, \gamma_B \in F$:*

$$\varepsilon_{\mathrm{ca}}^{(B)}(C, \delta_{\mathrm{hd}}, \delta_{\mathrm{nt}}) = \max_{f_1,\ldots,f_B} \Pr_{\gamma_1,\ldots,\gamma_B}\left[\Delta\left(\sum_{i=1}^{B} \gamma_i f_i,\, C\right) \le \delta_{\mathrm{hd}} \text{ and } \exists\, i :\ \Delta(f_i, C) > \delta_{\mathrm{nt}}\right].$$

Theorem 44 (Half-Threshold Batch CA)**.** *For any linear code C over F:*

$$\varepsilon_{\mathrm{ca}}^{(B)}(C, \delta/2, \delta) \le \frac{B}{|F|}.$$

Proof. Fix $f_1, \ldots, f_B$ with $\Delta(f_j, C) > \delta$ for some index j. We bound the probability over $(\gamma_1, \ldots, \gamma_B)$ that $g = \sum_i \gamma_i f_i$ is $(\delta/2)$-close to C.

Step 1: Conditioning. Fix all γ_i for $i \neq j$, so that $h = \sum_{i \neq j} \gamma_i f_i$ is a fixed function. Then $g = \gamma_j f_j + h$, which is linear in the single random variable γ_j.

Step 2: Pairwise argument. Call γ_j "bad" if $\exists\, c_\gamma \in C$ with $|S_\gamma| \ge (1 - \delta/2)n$ where $S_\gamma = \{x \in L : \gamma_j f_j(x) + h(x) = c_\gamma(x)\}$.

For two distinct bad values γ_j, γ_j': on $S_\gamma \cap S_{\gamma'}$ (size $\ge (1-\delta)n$), we have $(\gamma_j - \gamma_j')f_j = c_\gamma - c_{\gamma'}$, so f_j agrees with the codeword $(c_\gamma - c_{\gamma'})/(\gamma_j - \gamma_j') \in C$ on $\ge (1-\delta)n$ points, giving $\Delta(f_j, C) \le \delta$. This contradicts $\Delta(f_j, C) > \delta$. Hence at most 1 bad γ_j.

Step 3: Union bound. The bound in Step 2 holds for each fixed h (i.e., each fixed choice of γ_i, $i \neq j$). Averaging over the other γ_i:

$$\Pr_{\gamma_1,\ldots,\gamma_B}[\Delta(g, C) \le \delta/2 \mid f_j \text{ is } \delta\text{-far}] \le \frac{1}{|F|}.$$

By the union bound over all indices j with $\Delta(f_j, C) > \delta$ (at most B of them):

$$\varepsilon_{\mathrm{ca}}^{(B)}(C, \delta/2, \delta) \le \frac{B}{|F|}. \quad \square$$

Remark 45 (Individual vs. joint distance premise)**.** *Definition 43 uses the individual distance condition $\exists\, i :\ \Delta(f_i, C) > \delta$ rather than the joint distance $\Delta_{\mathrm{joint}}((f_1, \ldots, f_B), C^{=B}) > \delta$. The individual condition is a weaker premise (easier to satisfy), making Theorem 44 a stronger statement.*

In the STIR proof (Theorem 46, Strategy B), the relevant condition is $\Delta(f^{(i-1)}|_{L^{(i)}}, \mathrm{RS}_{k_{i+1}}) > \delta$: this is the individual-distance condition applied to one specific function, which is what Theorem 44 requires. The pairwise argument in the proof of Theorem 5 is invoked directly.

A.2 STIR Soundness Above Johnson

We show that our half-threshold CA is a drop-in replacement for the conjectured zero-loss CA in the STIR soundness proof [8].

STIR protocol recap. STIR [8] is an IOPP for Reed–Solomon proximity testing. Fix a degree schedule $k_0 = k, k_1, \ldots, k_R$ with $k_i \mid k_{i-1}$ and domains $L = L^{(0)} \supset L^{(1)} \supset \cdots \supset L^{(R)}$ with

$|L^{(i)}| = n/2^i$ (the standard radix-2 schedule sets $k_i = k_{i-1}/2$, matching FRI). At each round i, the verifier sends a random evaluation point $z_i \in F \setminus L$ and a random batching coefficient γ_i. The prover responds with $v_i = f^{(i-1)}(z_i)$ and commits the quotient oracle $q_i(x) = (f^{(i-1)}(x) - v_i)/(x - z_i)$, then sets $f^{(i)} = q_i + \gamma_i \cdot f^{(i-1)}\restriction_{L^{(i)}}$ (the batched next-round function). The key soundness step is: if $f^{(i-1)} \notin \mathrm{RS}_{k_i}$, then with high probability over γ_i, the batched $f^{(i)}$ is far from $\mathrm{RS}_{k_{i+1}}$ — this is precisely the CA bound applied to the pair $(q_i,\, f^{(i-1)}\restriction_{L^{(i)}})$ with randomness γ_i.

Theorem 46 (STIR Soundness Above Johnson)**.** *For* $\mathrm{RS}[F, L, k]$ *with* $\mathrm{char}(F) \neq 2$, $k = 2^m$, L *closed under negation,* $\delta_J < \delta < 1 - \rho$*: STIR with* R *rounds and* q *queries per round satisfies*

$$\Pr[\text{STIR accepts}] \leq \frac{(k+2)R}{|F|} + \left(1 - \frac{\delta}{2}\right)^q.$$

Proof. The proof follows the same Strategy A / Strategy B dichotomy as Theorem 18, adapted to STIR's round structure.

Strategy A: the prover commits honest quotients and batching at every round. The prover computes honest quotients $q_i(x) = (f^{(i-1)}(x) - f^{(i-1)}(z_i))/(x - z_i)$ and honest batching $f^{(i)} = q_i + \gamma_i \cdot f^{(i-1)}\restriction_{L^{(i)}}$. No adaptive freedom remains beyond the randomness of (z_i, γ_i). We bound the probability that $f^{(R)} \in \mathrm{RS}_{k_R}$ by a per-round union bound.

At round i, suppose $f^{(i-1)} \notin \mathrm{RS}_{k_{i-1}}$. Then:

- Over $z_i \in F \setminus L$: by the STIR out-of-domain soundness lemma [8] (Theorem 4.5 / Lemma 6.2), if $f^{(i-1)} \notin \mathrm{RS}_{k_{i-1}}$ then the quotient $q_i(x) = (f^{(i-1)}(x) - f^{(i-1)}(z_i))/(x - z_i)$ inherits non-codeword status except on a set of $\leq k_{i-1}$ degenerate values of z_i (roots of a nonzero degree-$(\leq k_{i-1})$ polynomial in z_i arising from the interpolation identity).
- Over γ_i: for a non-degenerate z_i, the batched syndrome of $f^{(i)} = q_i + \gamma_i \cdot f^{(i-1)}\restriction_{L^{(i)}}$ is degree-1 in γ_i (nonzero), so it vanishes for at most 1 value of γ_i.

Per round: $\Pr[f^{(i)} \in \mathrm{RS}_{k_i} \mid f^{(i-1)} \notin \mathrm{RS}_{k_{i-1}}] \leq (k_{i-1} + 1)/|F| \leq (k+2)/|F|$.

By the union bound over R rounds: $\Pr[f^{(R)} \in \mathrm{RS}_{k_R}] \leq (k+2)R/|F|$.

Note on the k-factor. The $k_{i-1}/|F|$ term is intrinsic to STIR's out-of-domain evaluation: the quotient involves evaluating a degree-$(<k_{i-1})$ polynomial at the random point z_i, producing $\leq k_{i-1}$ degenerate values. On extension fields (e.g., $\mathrm{BabyBear}_4$ with $|F| = p^4 \approx 2^{124}$), this term is negligible ($k/|F| \approx 2^{-105}$). On the base field ($|F| = p \approx 2^{31}$), it yields ~ 7 bits — comparable to the BCHKS commit-phase bound. The improvement of this paper over BCHKS for STIR is therefore in the *query phase*: $(1 - \delta/2)^q$ replaces the conjectured $(1 - \delta)^q$.

Strategy B: the prover deviates from honest computation at some round i. The optimal adversarial strategy: send a codeword $q_i \in \mathrm{RS}_{k_{i+1}}$ as the quotient, then compute honestly from round $i{+}1$ onward (all subsequent functions land in the code). Two checks threaten at round i:

1. *Out-of-domain (OOD) check*: the verifier sends $z_i \in F \setminus L^{(i)}$ and receives the claimed value $v_i = f^{(i-1)}(z_i)$, then checks $f^{(i-1)}(y) - v_i = (y - z_i)q_i(y)$ at q random points $y \in L^{(i)}$. In Strategy B, $f^{(i-1)}$ is an arbitrary prover-committed oracle (not necessarily a polynomial). The soundness of the OOD step follows from the STIR analysis [8] (Theorem 4.5): for all but $\leq k_{i-1}$ "lucky" values of z_i, any fixed adversarial pair (v_i, q_i) with $q_i \in \mathrm{RS}_{k_{i+1}}$ satisfies the OOD identity on at most $k_{i-1}{-}1$ points of $L^{(i)}$, giving a catch probability $\geq 1{-}(k_{i-1}/|L^{(i)}|)^q \geq 1 - \rho_i^q$. Since $\rho_i \leq 1/2 < 1 - \delta/2$: the OOD failure ρ_i^q is dominated by $(1{-}\delta/2)^q$.
2. *Batch-CA at the deviation round*: $f^{(i)} = q_i + \gamma_i \cdot f^{(i-1)}\restriction_{L^{(i)}}$ must be $(\delta/2)$-close to $\mathrm{RS}_{k_{i+1}}$. Since $q_i \in \mathrm{RS}_{k_{i+1}}$, this requires $f^{(i-1)}\restriction_{L^{(i)}}$ to be $(\delta/2)$-close to $\mathrm{RS}_{k_{i+1}}$ (up to the γ_i linear combination). By the pairwise argument of Theorem 5: if $\Delta(f^{(i-1)}\restriction_{L^{(i)}}, \mathrm{RS}_{k_{i+1}}) > \delta$, at most 1 value of γ_i succeeds. Failure: $\leq 1/|F|$.

As in Theorem 18 (Remark 19), the effective distance $\delta/2$ is preserved at rounds $2, \ldots, R$: since $\delta/2 < (1-\rho)/2$ (unique decoding radius), the BCIKS proximity gap [5] (Theorem 1.2, Eq. (1.1)) at threshold $\gamma = \delta/2$ bounds the bad challenges to $\leq n$ per round. The OOD shift-query bound $(k_{i-1}/|L^{(i)}|)^q$ in item (1) and the consistency-check bound $(1-\delta/2)^q$ both invoke Lemma 17: STIR samples q i.i.d. shift queries from $L^{(i)}$ at each round, so the per-round catch indicators are independent.

Combining items (1) and (2) per round, and taking the union over R rounds:

$$\Pr[\text{Strategy B}] \leq \frac{(k+2)R}{|F|} + (1-\delta/2)^q.$$

Combining. Every adaptive prover either commits honest quotients (Strategy A) or deviates (Strategy B):

$$\Pr[\text{accept}] \leq \max\big((k+2)R/|F|,\ (k+2)R/|F| + (1-\delta/2)^q\big) = (k+2)R/|F| + (1-\delta/2)^q.$$

□

Remark 47 (Comparison with conjectured STIR bound)**.** *The original STIR analysis [8] assumes zero-loss CA ($\varepsilon_{\text{ca}}(C, \delta, \delta) = O(1)/|F|$, conjectured) and achieves $\varepsilon = O(R)/|F| + (1-\delta)^q$. Our proved STIR bound has two differences:*

- *Commit phase: $(k+2)R/|F|$ instead of $O(R)/|F|$. The k-factor comes from the OOD polynomial evaluation at the random point z_i (degree $\leq k_{i-1}$, giving k_{i-1} degenerate values). This is intrinsic to STIR's quotient structure and absent in plain FRI. On extension fields ($|F| = p^4$), this term is negligible; on the base field ($|F| = p$), it dominates.*
- *Query phase: $(1-\delta/2)^q$ instead of $(1-\delta)^q$, the same $\sim 2\times$ overhead as for FRI.*

For BabyBear$_4$ ($|F| \approx 2^{124}$, $k = 2^{19}$, $R = 20$): the commit term gives $\approx 2^{-100}$ bits (negligible), and the query term is the bottleneck as for FRI. On the base field BabyBear ($|F| \approx 2^{31}$): the commit term gives only ~ 7 bits, comparable to BCHKS. The improvement over BCHKS for base-field STIR is entirely in the query phase.

A.3 WHIR Extension

WHIR [9] achieves super-fast verification by working with *constrained* Reed–Solomon codes: subcodes $C' \subseteq \text{RS}_k$ cut out by additional linear or affine constraints (multilinear-evaluation constraints, out-of-domain (OOD) evaluation constraints). The CA bound (Theorem 5) uses only one property of the code:

Linearity (closure under F-linear combinations).

This is preserved by every *linear* subcode of RS_k over the same domain L. We make the linear-subcode reduction explicit (Corollary 48), then build up to the full WHIR soundness theorem (Theorem 52) via a per-iteration analysis that handles affine constraints through the sumcheck weight polynomial rather than through CA on affine codes.

Corollary 48 (Linear-subcode CA)**.** *For any **linear** subcode $C' \subseteq \text{RS}_k$ over domain L and the same δ regime as Theorem 5:*

$$\varepsilon_{\text{ca}}(C', \delta/2, \delta) \leq 1/|F|, \qquad \varepsilon^{(B)}_{\text{ca}}(C', \delta/2, \delta) \leq B/|F|.$$

Proof. The proofs of Theorems 5 and 44 invoke linearity in exactly one place: constructing $g_1, g_2 \in C$ from $h_1, h_2 \in C$ via F-linear combinations. Linear subcodes are closed under such combinations, so the proofs go through verbatim. □

Remark 49 (Why this does not extend to affine subcodes)**.** *For an affine subcode $C' = g_0 + C_0$ with $g_0 \notin C_0$: the construction $g_2 = (c_\gamma - c_{\gamma'})/(\gamma - \gamma')$ in the proof of Theorem 5 produces an element of C_0 (a difference, not a sum), not of C'. Hence we cannot conclude $g_2 \in C'$. WHIR addresses this by enforcing affine constraints at a different layer (the sumcheck weight polynomial), not by extending CA to affine codes. The CA bound is only invoked on the ambient $\mathrm{RS}_{k/2^i}$ at each folding step.*

WHIR protocol structure: what we replace

We adopt the notation of [9] (Construction 5.1). One WHIR iteration consists of:

(i) **OOD sample.** Verifier sends $z \in F \setminus L$; prover responds with a claimed evaluation $y = f(z)$ and continues with the augmented constraint set $\{(z_j, y_j)\}_j \cup \{(z, y)\}$.

(ii) **Sumcheck folding.** Verifier sends folding randomness $\alpha \in F$; the constrained RS code $C^{(i)}$ folds via the sumcheck weight-polynomial machinery (Definition 4.5 of [9]) to a constrained RS code $C^{(i+1)}$ of half the dimension on the half-domain $L^{(i+1)}$.

(iii) **Shift queries.** Verifier samples t i.i.d. points from $L^{(i+1)}$ and checks consistency of the folded function with the prover's commitment.

After M iterations the protocol terminates with q i.i.d. final queries on the smallest folded domain.

Where MCA enters in [9]: only at step (ii), the sumcheck folding. WHIR's analysis (Theorem 4.20 of [9]) bounds the probability that a δ-far function f folds to a δ-close folded function via the MCA error of the constrained RS code at threshold δ. Above the Johnson bound, this requires Conjecture 4.12 of [9] (zero-loss MCA in the open zone). *This is the conjecture we eliminate.* The OOD sample (i) and shift queries (iii) do not invoke MCA; their soundness is bounded directly by polynomial-identity testing and Lemma 17 respectively.

Substitution: half-threshold CA + proved MCA in unique-decoding

We replace WHIR's MCA at threshold δ with a two-stage argument:

- *Iteration 1, first folding round only:* apply our half-threshold CA (Theorem 5) on the linear ambient RS_k via Corollary 48. This drops the effective distance from δ (premise) to $\delta/2$ (conclusion), with $\leq 1/|F|$ probability of escape. *This is the only invocation of half-threshold CA in the entire WHIR proof.*
- *All subsequent folding rounds (iteration 1 rounds $2, \ldots, r_1$ and all folds of iterations $2, \ldots, M$):* the effective distance $\delta/2$ satisfies $\delta/2 < (1-\rho)/2 \leq \delta_J$, so the protocol enters the unique-decoding regime. WHIR's MCA Corollary 4.11 (proved unconditionally in [9] for $\delta < \delta_J$) applies, giving $\leq n/|F|$ bad challenges per round. We deliberately do *not* re-apply half-threshold CA at later folds: doing so would compound the halving to $\delta/4, \delta/8, \ldots$.

The key structural point: *WHIR's Conjecture 4.12 is invoked only above Johnson; below Johnson, MCA is proved in [9].* Our one-time half-threshold step at iteration 1's first fold brings the effective distance below $(1-\rho)/2 \leq \delta_J$, so all subsequent folds (in iteration 1 and all later iterations) use the proved MCA, never the conjecture.

Affine constraints: handled by sumcheck, not by CA

WHIR's OOD constraints define an affine subcode $C' = g_0 + C_0$ where $C_0 \subseteq \mathrm{RS}_k$ is linear and g_0 encodes the constraint values. The CA bound (Corollary 48) applies only to the linear part, but in WHIR the constraints are not enforced via CA: they are absorbed into the sumcheck weight polynomial $\mathrm{wt}(X)$ (Definition 4.5 of [9]). The folding step (ii) then asks the prover to produce a folded function f' such that the sumcheck identity

$$\sum_{x \in L^{(i)}} \mathrm{wt}(x) \cdot f^{(i)}(x) \;=\; (\text{wt-folded sumcheck claim})$$

holds round by round. This identity is degree-≤ 2 in the sumcheck round variable, with soundness $\leq 2/|F|$ per round (Schwartz–Zippel). The CA bound enters only at the proximity-preserving step that maps $f^{(i)} \to f^{(i+1)}$, where the underlying code is the linear $\mathrm{RS}_{k/2^i}$. Affine vs. linear is therefore handled by separation of concerns: **sumcheck verifies constraint satisfaction, CA verifies proximity**.

Per-iteration error budget

Lemma 50 (Per-iteration commit-phase error)**.** *Under the above setup, one WHIR iteration with r folding rounds has commit-phase soundness error (failures from random alpha/sumcheck challenges, not from query sampling)*

$$\varepsilon^{\mathrm{iter}}_{\mathrm{commit}} \;\leq\; \frac{k}{|F|} \;+\; \frac{r\,n}{|F|} \;+\; \frac{2r}{|F|},$$

with terms: OOD ($k/|F|$), folding rounds at $\leq n$ bad challenges each ($rn/|F|$), and sumcheck rounds ($2r/|F|$). Shift queries are handled separately (Remark 51).

Proof. OOD ($k/|F|$). The verifier's $z \in F \setminus L$ is uniform; an adversarial prover with non-codeword f can fool the OOD check only if z is a root of a fixed degree-$<k$ polynomial (the difference between f and the prover's claimed evaluation extended as a low-degree interpolant). At most k such roots, giving $\leq k/|F|$.

Folding rounds ($rn/|F|$). Iteration 1's first fold uses half-threshold CA: the input is δ-far, Corollary 48 gives at most 1 bad α for $(\delta/2)$-closeness ($\leq 1/|F|$, tighter than $n/|F|$). *All other folds* (iteration 1's rounds $2, \ldots, r$ and *all* folds of iterations ≥ 2) use the BCIKS proximity gap [5] at threshold $\delta/2 < (1-\rho)/2 \leq \delta_J$ (the input is $(\delta/2)$-far by distance preservation across iterations, Remark 19; we deliberately do *not* re-apply the half-threshold CA at iterations ≥ 2, which would compound the distance halving). BCIKS gives $\leq n$ bad challenges per round, $\leq rn/|F|$ per iteration. We state $rn/|F|$ as a uniform upper bound covering all r folds (loose by $\frac{n-1}{|F|}$ for iteration 1's first fold).

Sumcheck ($2r/|F|$). Per-round Schwartz–Zippel on the degree-≤ 2 sumcheck identity gives $2/|F|$; total $2r/|F|$.

Union over the three commit-phase failure modes gives the stated bound. □

Remark 51 (Shift queries and query-phase max-over-strategies)**.** *The prover commits to a strategy before seeing verifier randomness. Adversarial strategies (mutually exclusive choices) and their query-phase failure probabilities:*

- *Strategy A: honest fold throughout. The cumulative oracle $f^{(M)}$ is $(\delta/2)$-far by distance preservation. Final q queries miss with probability $\leq (1-\delta/2)^q$ (Lemma 17).*

- *Strategy B: commit a codeword $f^{(i+1)} \in C^{(i+1)}$ in iteration i instead of the honest fold; then continue honestly. The discrepancy between the committed codeword and the honest fold is $(\delta/2)$-far on $L^{(i+1)}$. Iteration-i shift queries miss this with probability $\leq (1-\delta/2)^t$. (Subsequent iterations' shift queries pass automatically: codewords fold to codewords, so consistency is preserved from iteration i+1 onward; final queries also pass since $f^{(M)}$ is a codeword.)*
- *Strategy B-multi: deviate at multiple iterations $\{i_1, \ldots, i_t\}$. The events "shift queries miss in iteration i_j" are independent across iterations (different randomness), so joint miss $\leq (1-\delta/2)^{t\cdot|\{i_j\}|} \leq (1-\delta/2)^t$.*

The prover picks ONE strategy; the verifier's failure probability is the maximum over strategies, not the sum:

$$\Pr[\textit{query-phase miss}] \;\leq\; \max\big((1-\delta/2)^q,\, (1-\delta/2)^t\big) \;=\; (1-\delta/2)^q,$$

using $t \geq q$ to make the shift-query bound tighter. No factor of M arises: a single chosen strategy can only escape its single failure mode.

Theorem 52 (WHIR Soundness Above Johnson)**.** *For $\mathrm{RS}[F, L, k]$ with $k = 2^m$, L admitting a fixed-point-free involution, $\delta_J < \delta < 1 - \rho$, and an initial constrained-RS code $C^{(0)} \subseteq \mathrm{RS}_k$ defined by linear and affine constraints (multilinear-evaluation and OOD): **provided the prover-committed initial oracle** $f^{(0)} : L \to F$ **satisfies** $\Delta(f^{(0)}, C^{(0)}) > \delta$, WHIR [9] (Construction 5.1) with M iterations, r_i folding rounds in iteration i, $t \geq q$ shift queries per iteration, and q final queries satisfies*

$$\Pr[\text{WHIR accepts}] \;\leq\; \frac{Mk + (n{+}2)R}{|F|} \;+\; \left(1 - \frac{\delta}{2}\right)^q,$$

where $R = \sum_{i=1}^{M} r_i$ is the total number of FRI-folding rounds across all iterations. The query term is $(1-\delta/2)^q$ (no M factor: max over prover strategies, not sum; see Remark 51). This is a theorem; it does not require WHIR Conjecture 4.12.

Proof. We separate three failure modes: *constraint-violation failures* (committed $f^{(0)}$ is close to RS_k but violates the OOD/multilinear constraints, so $\Delta(f^{(0)}, C^{(0)}) > \delta$ even though $\Delta(f^{(0)}, \mathrm{RS}_k) \leq \delta$); *commit-phase failures* (verifier's random α/sumcheck challenges land in bad sets); and *query-phase failures* (verifier's queries miss a deviation). The hypothesis $\Delta(f^{(0)}, C^{(0)}) > \delta$ is used to drive a case split on whether the ambient distance $\Delta(f^{(0)}, \mathrm{RS}_k)$ is also $> \delta$.

Case split on ambient vs. constrained distance. Two cases:

Case (a): $\Delta(f^{(0)}, \mathrm{RS}_k) > \delta$. The ambient RS distance is $> \delta$, so Corollary 48 applies directly: half-threshold CA on the linear ambient RS_k at iteration 1's first fold drops the effective distance to $\delta/2$ with $\leq 1/|F|$ escape (subsumed by the $rn/|F|$ uniform bound of Lemma 50).

Case (b): $\Delta(f^{(0)}, \mathrm{RS}_k) \leq \delta$ *but* $\Delta(f^{(0)}, C^{(0)}) > \delta$. Some witness $g \in \mathrm{RS}_k$ with $d(f^{(0)}, g) \leq \delta n$ exists, but every such witness violates at least one constraint defining $C^{(0)} \subsetneq \mathrm{RS}_k$. We must show that, even though the proximity layer cannot reject (the ambient distance is small), the sumcheck/OOD layer rejects.

Pre-query case-determination. The case split is made on the *committed* oracle $f^{(0)} : L \to F$ before any verifier randomness. The Merkle commitment binds $f^{(0)}$ as a function on all of L at commit time (see [18, 9] for the standard binding analysis); the ambient distance $\Delta(f^{(0)}, \mathrm{RS}_k) = \min_{g \in \mathrm{RS}_k} d(f^{(0)}, g)/n$ is then a property of the committed function, independent of which positions are subsequently queried. The prover cannot retroactively shift $f^{(0)}$ between Case (a) and Case (b) based on query outcomes (modulo the standard collision term $3(Q^2{+}1)/2^\kappa$ for κ-bit Merkle binding under Q random-oracle queries, absorbed into the Fiat–Shamir budget).

Black-box reduction to WHIR's sumcheck/OOD soundness. Conditional on Case (b), we invoke *as black boxes* the WHIR sumcheck and OOD soundness statements [9] (Theorems 5.4–5.5), which establish the per-iteration affine/constrained-code consistency bounds $k/|F|$ (OOD) and $2r/|F|$ (sumcheck). These statements are part of WHIR's published soundness analysis and are *independent of MCA / Conjecture 4.12*: the sumcheck soundness uses only Schwartz–Zippel on the prover-supplied round polynomials over independently sampled verifier randomness, and the OOD soundness uses standard polynomial-identity / interpolation arguments under the augmented constraint set. We do not re-derive these here. Our contribution replaces the MCA/proximity-preservation step (which [9] bounds via Conjecture 4.12 above Johnson) with the proven half-threshold CA + BCIKS unique-decoding chain of §5; the affine-constraint and OOD layers are imported unchanged.

Reduction to Case (a). If the prover survives iteration 1's OOD/sumcheck checks in Case (b), then by WHIR Theorems 5.4–5.5 the augmented constraint set $\{(z_j, y_j)\}_j \cup \{(z, y)\}$ is consistent with $f^{(0)}$ on a $(1-\delta)$-fraction of L. Iteration 2 then operates on a folded oracle $f^{(1)}$ that is either $(\delta/2)$-far from the augmented code $C^{(1)}$ (Case (a) at the next level, handled by the proximity layer) or again close-to-ambient-but-constraint-violating (Case (b) at the next level, recursively absorbed into iteration 2's OOD/sumcheck budget by the same black-box invocation). After M iterations the cumulative augmented constraint set forces $f^{(M)}$ to be a codeword in the final folded code or to be caught by the final q queries; either way, the per-iteration budget $k/|F| + r_i n/|F| + 2r_i/|F|$ of Lemma 50 applies at every iteration.

Query/commit interaction. The Merkle binding (above) ensures the case split survives query interaction: queries reveal positions of the pre-committed $f^{(i)}$, but cannot retroactively change which case applies. The query-phase miss probability $(1-\delta/2)^q$ from Remark 51 is therefore additive with the commit-phase OOD/sumcheck budget; no joint conditioning is needed.

Commit phase: induction over iterations. We show by induction on $i = 0, 1, \ldots, M-1$ that, conditional on the i preceding iterations having all succeeded, the commit-phase escape probability of iteration $i{+}1$ is bounded by $\varepsilon^{\text{iter}}_{\text{commit}}$ (Lemma 50).

Base case ($i = 0$). By the case split above (Case (a) for proximity; Case (b) absorbed into sumcheck/OOD terms), iteration 1 applies Lemma 50: the OOD term covers Case (b)'s constraint violations, the folding term applies half-threshold CA at the first fold (using Case (a)'s ambient hypothesis $\Delta(f^{(0)}, \mathrm{RS}_k) > \delta$), and the sumcheck term covers the multilinear-evaluation constraint identity. Total: $k/|F| + r_1 n/|F| + 2r_1/|F|$.

Inductive step. Suppose iterations $1, \ldots, i$ have succeeded. Then the surviving oracle $f^{(i)}$ on the folded domain $L^{(i)}$ satisfies $\Delta(f^{(i)}, \mathrm{RS}_{k/2^{R_i}}) > \delta/2$, where $R_i = \sum_{j \le i} r_j$ and the effective distance is preserved at $\delta/2$ by the BCIKS distance-locking argument (Remark 19; valid because $\delta/2 < (1-\rho)/2$). Iteration $i{+}1$ then applies Lemma 50 starting from a $(\delta/2)$-far oracle. The half-threshold CA does *not* re-apply here (which would compound the halving to $\delta/4$); instead, every fold in iteration $i{+}1$ uses the BCIKS proximity gap at threshold $\delta/2 < \delta_J$, giving $\le n$ bad α per fold and contributing $r_{i+1} \cdot n/|F|$ to the commit budget. The lemma's uniform bound $rn/|F|$ already accommodates this case.

Query phase. By Remark 51: max over adversarial strategies (honest fold caught by final queries OR codeword commit in some iteration caught by iteration-local shift queries) gives $\le (1-\delta/2)^q$ (using $t \ge q$, which makes shift-query catch tighter than final-query catch).

Total error. Summing per-iteration commit-phase budgets (Lemma 50) and the query-phase term:

$$\Pr[\text{accept}] \;\le\; \sum_{i=1}^{M} \varepsilon^{\text{iter}}_{\text{commit},i} \;+\; (1-\delta/2)^q \;\le\; \frac{Mk + (n+2)R}{|F|} \;+\; (1-\delta/2)^q,$$

where $\sum_i \big(k + r_i n + 2r_i\big) = Mk + (n{+}2)R$ for the per-round terms. □

Remark 53 (Result type)**.** *Theorem 52 is a* ***WHIR soundness theorem that bypasses MCA above Johnson****: it does not prove a new MCA statement, nor does it require WHIR's Conjecture 4.12. Above-Johnson MCA appears nowhere in the proof. The half-threshold CA acts as a one-time "crossing" from above-Johnson to below-Johnson at iteration* 1*'s first fold; from that point onward (all subsequent folds in iteration 1 and all folds in iterations* ≥ 2*), the protocol stays in the unique-decoding regime, where MCA is proved (WHIR Corollary 4.11). The* $2\times$ *query overhead is the price of this single crossing.*

Remark 54 (Deployment status)**.** *The reduced-threshold safety net covers all proximity testing protocols deployed or proposed for Ethereum: FRI (Theorem 18, Theorem 61), STIR (Theorem 46; char* 2 *via Remark 63), and WHIR (Theorem 52; char* 2 *via Remark 63). This covers the odd-characteristic multiplicative-FRI deployments (BabyBear, Goldilocks), circle-FRI deployments (Stwo on Mersenne31, Theorem 69), and the characteristic-*2 *proposals (Binius on binary tower fields). No protocol changes are required — only the soundness analysis is updated.*

B Extension to Characteristic 2

We show that the reduced-threshold CA bound (Theorem 5) — which is characteristic-independent — extends to the FRI folding structure over fields of characteristic 2, yielding FRI soundness above the Johnson bound for Binius-style deployments on binary tower fields. The key replacement is the *additive* (subspace-polynomial) fold in place of the multiplicative $\pm\sqrt{y}$ fold.

B.1 Additive FRI Folding

Definition 55 (Additive FRI Folding)**.** *Let* $F = \mathbb{F}_{2^s}$ *and* $L \subset F$ *an* $\mathbb{F}_2$*-linear subspace of dimension* d*, so* $|L| = n = 2^d$*. Fix* $\beta \in L$*,* $\beta \neq 0$*. The* ***subspace polynomial*** *of* $\langle\beta\rangle$ *is*

$$s(x) = x^2 + \beta x.$$

Since s *is* $\mathbb{F}_2$*-linear (*$s(x{+}y) = s(x) + s(y)$ *in characteristic* 2*, using* $(x{+}y)^2 = x^2 + y^2$*) and* $\ker(s|_L) = \{0, \beta\}$*:*

$$L' := s(L) \subset F, \qquad |L'| = n/2,$$

and L' *is again an* $\mathbb{F}_2$*-linear subspace (of dimension* $d{-}1$*). Each* $y \in L'$ *has exactly two preimages* $\{x, x{+}\beta\}$ *in* L*.*

For $f : L \to F$*, the* ***additive even/odd decomposition*** *is:*

$$f_1(y) = \frac{f(x) + f(x{+}\beta)}{\beta},$$
$$f_0(y) = f(x) + x \cdot f_1(y),$$

where x *is either preimage of* y *(the expressions are invariant under* $x \mapsto x{+}\beta$*: swapping adds* $\beta \cdot f_1(y)$ *to the* f_0 *formula's second term and adds* $\beta \cdot f_1(y)$ *to* $f(x)$*, which cancel). The* ***additive FRI fold*** *is:*

$$f'(y) = f_0(y) + \alpha \cdot f_1(y).$$

Remark 56 (Reconstruction)**.** $f(x) = f_0(s(x)) + x \cdot f_1(s(x))$ *for all* $x \in L$*: at* x*, the right side is* $f_0(y) + x \cdot f_1(y) = (f(x) + x \cdot f_1(y)) + x \cdot f_1(y) = f(x)$*. This is the characteristic-*2 *analogue of the odd-characteristic reconstruction* $f(x) = f_{\text{even}}(x^2) + x \cdot f_{\text{odd}}(x^2)$*.*

B.2 Even/Odd Isomorphism in Characteristic 2

Lemma 57 (Additive Even/Odd Isomorphism)**.** *For k even: $g \mapsto (g_0, g_1)$ is a linear isomorphism* $\mathrm{RS}[F, L, k] \xrightarrow{\sim} \mathrm{RS}[F, L', k/2]^{=2}$.

Proof. **Forward direction** ($g \in \mathrm{RS}_k \Rightarrow g_0, g_1 \in \mathrm{RS}_{k/2}$). Let $g(x) = \sum_{i=0}^{k-1} a_i x^i$. In the polynomial ring $F[x]$, the relation $x^2 = s + \beta x$ (from $s(x) = x^2 + \beta x$) allows recursive reduction: every monomial x^i reduces to $p_i(s) + x \cdot q_i(s)$ where $\deg p_i, \deg q_i \le \lfloor i/2 \rfloor$. Collecting terms: $g(x) = g_0(s) + x \cdot g_1(s)$ with

$$\deg g_0 \le \lfloor (k-1)/2 \rfloor \le k/2 - 1, \qquad \deg g_1 \le \lfloor (k-2)/2 \rfloor \le k/2 - 1.$$

Hence $g_0, g_1 \in \mathrm{RS}_{k/2}(L')$.

Dimension count. The map $\varphi : \mathrm{RS}_k(L) \to \mathrm{RS}_{k/2}(L')^{=2}$ is F-linear and maps a k-dimensional space into a $(k/2 + k/2) = k$-dimensional space.

Injectivity. If $(g_0, g_1) = (0, 0)$, then $g(x) = g_0(s(x)) + x \cdot g_1(s(x)) = 0$ for all $x \in L$. Since g has degree $< k \le |L|$, this implies $g = 0$. So $\ker(\varphi) = \{0\}$.

A linear injection between spaces of equal dimension is an isomorphism. □

Remark 58 (Why k even is needed)**.** *For k odd in characteristic 2: the monomial $f(x) = x^k$ satisfies $\Delta(x^k, \mathrm{RS}_k) = 1 - \rho$ (since $x^k - g$ has $\le k$ roots for any g of degree $< k$). Its additive decomposition gives f_0, f_1 of degree $\le \lceil k/2 \rceil - 1$, so both belong to $\mathrm{RS}_{\lceil k/2 \rceil}$ and the fold $f' = f_0 + \alpha \cdot f_1 \in \mathrm{RS}_{\lceil k/2 \rceil}$ for every α. The isomorphism $\mathrm{RS}_k \to \mathrm{RS}^{=2}_{k/2}$ fails because $\mathrm{RS}^{=2}_{\lceil k/2 \rceil}$ has dimension $2\lceil k/2 \rceil = k+1 > k = \dim \mathrm{RS}_k$: the map is injective but not surjective, so the coupling lemma breaks (not every pair $(g_0, g_1) \in \mathrm{RS}^{=2}_{\lceil k/2 \rceil}$ corresponds to a codeword in RS_k). The fold of a far-from-code function lands in the larger code $\mathrm{RS}_{\lceil k/2 \rceil}$ for all α, and FRI is unsound. This is the same failure mode as in odd characteristic (Remark 4), confirming that $k = 2^m$ is a genuine structural requirement.*

B.3 Coupling Lemma and Proximity Gap

Lemma 59 (Additive FRI Coupling)**.** *For even k and any $f : L \to F$ (with $\mathrm{char}(F) = 2$, L an $\mathbb{F}_2$-linear subspace, $\beta \in L \setminus \{0\}$):*

$$\Delta(f, \mathrm{RS}_k(L)) \le \Delta\big((f_0, f_1), \mathrm{RS}_{k/2}(L')^{=2}\big).$$

Proof. By Lemma 57: $g \mapsto (g_0, g_1)$ is an isomorphism $\mathrm{RS}_k \to \mathrm{RS}^{=2}_{k/2}$. For any $g \in \mathrm{RS}_k$ and $y \in L'$: if $f_0(y) = g_0(y)$ and $f_1(y) = g_1(y)$, then by reconstruction, $f(x) = g(x)$ for both x in the fiber $\{x, x+\beta\}$ of y. So each joint-error position y (where $f_0(y) \ne g_0(y)$ or $f_1(y) \ne g_1(y)$) contributes ≥ 1 error to $f - g$ on L, and each y contributes at most 2. Let $E = |\{y \in L' : (f_0(y), f_1(y)) \ne (g_0(y), g_1(y))\}|$. Then $\Delta(f, g) \le 2E/n = E/|L'| = \Delta_{\mathrm{joint}}$. Taking $\min_g$:

$$\Delta(f, \mathrm{RS}_k) \le \Delta_{\mathrm{joint}}. \quad \square$$

Theorem 60 (Half-Threshold Proximity Gap, Characteristic 2)**.** *For $\mathrm{char}(F) = 2$, even k, $\delta_J < \delta < 1 - \rho$, and $\Delta(f, \mathrm{RS}_k(L)) > \delta$:*

$$|\{\alpha \in F : \Delta(f_0 + \alpha f_1, \mathrm{RS}_{k/2}(L')) \le \delta/2\}| \le 1.$$

Proof. By Lemma 59: $\Delta((f_0, f_1), \mathrm{RS}^{=2}_{k/2}) \ge \Delta(f, \mathrm{RS}_k) > \delta$. Instantiate Theorem 5 with the pair (f_0, f_1) on $\mathrm{RS}_{k/2}$ over L' at threshold δ. Theorem 5 uses only linearity and Hamming distances, with no characteristic hypothesis. The proof is identical to Theorem 14. □

B.4 FRI Soundness in Characteristic 2

Theorem 61 (FRI Soundness Above Johnson, Characteristic 2)**.** *For* $\mathrm{RS}[F, L, k]$ *with* $\mathrm{char}(F) = 2$*,* $k = 2^m$*,* $L \subset F$ *an* $\mathbb{F}_2$*-linear subspace of dimension* $d \geq m$*, and* $\delta_J < \delta < 1 - \rho$*,* $\Delta(f^{(0)}, \mathrm{RS}_k) > \delta$*: FRI with additive folding (Definition 55),* $R = m$ *rounds, and* q *queries per round satisfies*

$$\Pr[\text{FRI accepts}] \leq \frac{nR}{|F|} + (1 - \delta/2)^q.$$

Proof. The proof follows the same Strategy A / Strategy B structure as Theorem 18, with the multiplicative fold replaced by the additive fold. We verify the three components and note where the argument differs from odd characteristic:

Iterability. At each round i, the domain $L^{(i)} = s_i(L^{(i-1)})$ is an $\mathbb{F}_2$-linear subspace of dimension $d-i$, using $\beta_i \in L^{(i-1)} \setminus \{0\}$ as the folding direction. The code dimension halves: $k_i = k/2^i$. Since $k = 2^m$, we have k_i even for all $i < m$ (guaranteeing the isomorphism), and $k_m = 1$ (so RS_1 is the space of constants, and the final check verifies $f^{(R)}$ is constant).

Strategy A. Each round's fold $f^{(i)}(y) = f_0^{(i-1)}(y) + \alpha_i \cdot f_1^{(i-1)}(y)$ is linear in α_i, where $f_0^{(i-1)}, f_1^{(i-1)}$ are determined by $f^{(0)}$ and $(\alpha_1, \ldots, \alpha_{i-1})$ via the additive decomposition on domains $L \supset s_1(L) \supset s_2(s_1(L)) \supset \cdots$ (each $L^{(i)} = s_i(L^{(i-1)})$ is an $\mathbb{F}_2$-subspace of dimension $d-i$, so the chain is well-defined for $R \leq d$). It follows that $f^{(R)}(y)$ is multilinear in $(\alpha_1, \ldots, \alpha_R)$ with total degree $\leq R$. The *nonvanishing*: if all syndromes vanished identically, then $f^{(R)} \in \mathrm{RS}_1$ for every $(\alpha_1, \ldots, \alpha_R)$. As in the proof of Theorem 18, this is a polynomial-identity argument (not a runtime reconstruction); evaluating at $\alpha_i = 0$ and $\alpha_i = 1$ recovers $f_0^{(i-1)}, f_1^{(i-1)} \in \mathrm{RS}_{k/2^i}$, and Lemma 57 (requiring $k/2^{i-1}$ even, guaranteed by $k = 2^m$) gives $f^{(i-1)} \in \mathrm{RS}_{k/2^{i-1}}$; iterating yields $f^{(0)} \in \mathrm{RS}_k$, a contradiction. Hence some syndrome is a nonzero multilinear polynomial, and $\Pr[f^{(R)} \in \mathrm{RS}_1] \leq R/|F|$ by Schwartz–Zippel.

Strategy B. The distance-preservation argument uses Theorem 60: at round 1, all but ≤ 1 value of α_1 gives $\Delta(f^{(1)}, \mathrm{RS}_{k/2}) > \delta/2$. At subsequent rounds, $\delta/2 < (1-\rho)/2$ (unique decoding radius), and the BCIKS proximity gap [5] (Theorem 1.2, Eq. (1.1)) at threshold $\gamma = \delta/2$ gives $\leq n$ bad α per round; distance is preserved at $\delta/2$. The consistency check catches a $(\delta/2)$-far deviation with probability $\geq 1-(1-\delta/2)^q$ by Lemma 17 (the additive query model is the analogue of Definition 15 with π_i replaced by the iterated subspace polynomial $s_i \circ \cdots \circ s_1$, which is 2^i-to-1 on its image; the i.i.d. argument transfers verbatim). The multi-deviation argument proceeds identically to Theorem 18.

Combining: $\Pr[\text{accept}] \leq nR/|F| + (1 - \delta/2)^q$. □

Remark 62 (Domain structure)**.** *In the multiplicative (odd-characteristic) setting, FRI requires* L *closed under negation. In the additive (characteristic-2) setting, the analogous requirement is that* L *is an* $\mathbb{F}_2$*-linear subspace, which is the standard domain choice in additive-FRI systems. The folding direction* β_i *at round* i *must satisfy* $\beta_i \in L^{(i-1)} \setminus \{0\}$*; it is chosen dynamically from the current domain* $L^{(i-1)} = s_{i-1}(\cdots(s_1(L)))$*, which is an* $\mathbb{F}_2$*-subspace of dimension* $d-i+1$*. A concrete prescription: fix an* $\mathbb{F}_2$*-basis* $\{e_1, \ldots, e_d\}$ *of* L*; at round* i*, use* $\beta_i =$ *any nonzero element of* $L^{(i-1)}$ *(e.g., the image* $s_{i-1}(\cdots(s_1(e_i)))$ *if nonzero, or any other convenient choice).*

Remark 63 (STIR/WHIR in characteristic 2)**.** *The batch-CA theorem (Theorem 44) is characteristic-independent (it uses only the pairwise argument from Theorem 5). Combining with Lemma 59 and Theorem 60, Theorem 46 and Corollary 48 extend to characteristic* 2 *with the same bounds. The safety net now covers FRI, STIR, and WHIR over fields of any characteristic.*

C Extension to Circle FRI (Stwo)

We extend the soundness theorem to circle FRI [20], the variant used by Stwo [21] on Mersenne31 fields ($p = 2^{31} - 1$). The half-threshold CA (Theorem 5) is code-agnostic; only the coupling lemma requires adaptation to the circle evaluation domain. The first FRI round (bivariate $\to$ univariate) uses a circle-specific coupling lemma proved below; all subsequent rounds reduce to standard univariate FRI.

C.1 Circle Codes and Folding

Definition 64 (Circle Group and Circle RS Code)**.** *Let p be an odd prime. The **circle group** is $\mathcal{C}(\mathbb{F}_p) = \{(x, y) \in \mathbb{F}_p^2 : x^2 + y^2 = 1\}$ with group law*

$$(x_0, y_0) \oplus (x_1, y_1) = (x_0x_1 - y_0y_1,\ x_0y_1 + y_0x_1),$$

identity $(1, 0)$, and inversion $(x, y)^{-1} = (x, -y)$. The order is $|\mathcal{C}(\mathbb{F}_p)| = p+1$ when -1 is a quadratic non-residue ($p \equiv 3 \pmod 4$) and $p - 1$ otherwise; for Mersenne31 ($p = 2^{31} - 1 \equiv 3 \pmod 4$), the order is 2^{31}.

*Let $D \subset \mathcal{C}(\mathbb{F}_p)$ be an evaluation domain of even size N on which the **involution** $\iota(x, y) = (x, -y)$ is fixed-point-free (equivalently: $y \neq 0$ for all $(x, y) \in D$, i.e., $(\pm 1, 0) \notin D$). Let $D_x = \{x : (x, y) \in D\}$ denote the x-projection, with $|D_x| = N/2$. The involution partitions D into $N/2$ pairs $\{(x, y), (x, -y)\}$, one for each $x \in D_x$.*

*The **circle RS code**[5] is*

$$\mathrm{CRS}_N(D) = \big\{(x, y) \mapsto a(x) + y \cdot b(x) : \deg a < N/2,\ \deg b < N/2\big\}\restriction_D,$$

a linear code of dimension $\dim(\mathrm{CRS}_N) = N$ and rate $\rho = N/|D| = 1$ on D. In the FRI setting, the code is evaluated on D and the rate is $\rho = k/N$ where k is the polynomial degree bound (cf. Stwo's "blowup factor"); by abuse of notation, we write $\mathrm{CRS}_k(D)$ for the subcode with $\deg a < k/2$ and $\deg b < k/2$, having dimension k and rate $\rho = k/N$.

Definition 65 (Circle FRI Folding)**.** *For $f : D \to F$, the **circle even/odd decomposition** is:*

$$f_0(x) = \frac{f(x, y) + f(x, -y)}{2}, \qquad f_1(x) = \frac{f(x, y) - f(x, -y)}{2y},$$

*where $(x, y) \in D$ is either representative of the fiber over x (both expressions are invariant under $y \mapsto -y$). The **circle FRI fold** is $f'(x) = f_0(x) + \alpha \cdot f_1(x)$, a function on D_x.*

C.2 Coupling Lemma and Soundness

Lemma 66 (Circle Even/Odd Isomorphism)**.** *For even k: the map $g \mapsto (g_0, g_1)$ is a linear isomorphism $\mathrm{CRS}_k(D) \xrightarrow{\sim} \mathrm{RS}_{k/2}(D_x)^{=2}$.*

Proof. For $g(x, y) = a(x) + y \cdot b(x) \in \mathrm{CRS}_k$: $g_0 = a$ and $g_1 = b$, both of degree $< k/2$, hence both in $\mathrm{RS}_{k/2}(D_x)$. The map is F-linear, and the source and target both have dimension k. *Injectivity*: if $(a, b) = (0, 0)$ on all of D_x, then $\deg a, \deg b < k/2 \leq |D_x|$ forces $a = b = 0$ as polynomials, so $g = 0$ on D. □

[5] In the notation of [20], this is the FFT subspace $L'_N(\mathbb{F}_p)$. The full space L_N has dimension $N+1$; the extra monomial $x^{N/2}$ is handled by the prover sending its coefficient before the first fold (Remark 70).

Lemma 67 (Circle FRI Coupling)**.** *For even k and any $f : D \to F$:*

$$\Delta(f,\, \mathrm{CRS}_k(D)) \;\le\; \Delta\big((f_0, f_1),\, \mathrm{RS}_{k/2}(D_x)^{=2}\big).$$

Proof. By Lemma 66, each $g \in \mathrm{CRS}_k$ corresponds to $(g_0, g_1) \in \mathrm{RS}^{=2}_{k/2}$. If $f_0(x) = g_0(x)$ and $f_1(x) = g_1(x)$ at some $x \in D_x$, then $f(x, \pm y) = a(x) \pm y \cdot b(x) = g(x, \pm y)$ (by reconstruction: $f(x, y) = f_0(x) + y \cdot f_1(x)$), so *both* points in the fiber over x are error-free. Equivalently, each joint-error position x (where $f_0(x) \neq g_0(x)$ or $f_1(x) \neq g_1(x)$) contributes ≥ 1 base-level error on D and at most 2. With $E = |\{x \in D_x : (f_0(x), f_1(x)) \neq (g_0(x), g_1(x))\}|$:

$$\Delta(f, g) = \frac{|\{P \in D : f(P) \neq g(P)\}|}{N} \leq \frac{2E}{N} = \frac{E}{|D_x|} = \Delta_{\text{joint}}.$$

Taking $\min_g$ (using the isomorphism, the minimum over $\mathrm{RS}^{=2}_{k/2}$ equals the minimum over CRS_k): $\Delta(f, \mathrm{CRS}_k) \leq \Delta_{\text{joint}}$. □

Theorem 68 (Circle FRI Proximity Gap)**.** *For even k, $\delta_J < \delta < 1-\rho$, and $\Delta(f, \mathrm{CRS}_k(D)) > \delta$:*

$$|\{\alpha \in F : \Delta(f_0 + \alpha f_1,\, \mathrm{RS}_{k/2}(D_x)) \leq \delta/2\}| \leq 1.$$

Proof. By Lemma 67: $\Delta((f_0, f_1), \mathrm{RS}^{=2}_{k/2}) \geq \Delta(f, \mathrm{CRS}_k) > \delta$. Apply Theorem 5 to (f_0, f_1) on $\mathrm{RS}_{k/2}$ over D_x: at most 1 value of α has $\Delta(f_0 + \alpha f_1, \mathrm{RS}_{k/2}) \leq \delta/2$. □

Theorem 69 (Circle FRI Soundness Above Johnson)**.** *For $\mathrm{CRS}_k(D)$ over F with even $k = 2^m$, $D \subset \mathcal{C}(\mathbb{F}_p)$ of size N with ι fixed-point-free, $\delta_J < \delta < 1 - \rho$, and $\Delta(f^{(0)}, \mathrm{CRS}_k) > \delta$: circle FRI with $R = m$ rounds and q queries per round satisfies*

$$\Pr[\text{circle FRI accepts}] \;\le\; \frac{NR}{|F|} + \left(1 - \frac{\delta}{2}\right)^q.$$

Proof. The proof follows the Strategy A / Strategy B structure of Theorem 18.

Round 1 (circle → univariate). By Theorem 68: for all but ≤ 1 value of α_1, $\Delta(f^{(1)}, \mathrm{RS}_{k/2}(D_x)) > \delta/2$.

Rounds $2, \ldots, R$ (univariate FRI on D_x). After round 1, the function $f^{(1)} : D_x \to F$ lives on a univariate evaluation domain. Subsequent circle-FRI rounds use the **doubling map** $\pi(x) = 2x^2 - 1$ (the x-coordinate of the point-doubling endomorphism on $\mathcal{C}$) with involution $x \mapsto -x$. This involution is fixed-point-free on D_x because $0 \notin D_x$ (since $(0, \pm 1) \in \mathcal{C}$ has order 4 and $|D| \geq 8$ in deployment). The even/odd decomposition

$$f_{\text{even}}(\pi(x)) = \frac{f(x) + f(-x)}{2}, \qquad f_{\text{odd}}(\pi(x)) = \frac{f(x) - f(-x)}{2x}$$

uses the relation $x^2 = (\pi(x) + 1)/2$ to reduce: for $g \in \mathrm{RS}_{k/2}(D_x)$, the monomials $x^{2j} = ((\pi + 1)/2)^j$ and $x^{2j+1} = x \cdot ((\pi + 1)/2)^j$ give $g_{\text{even}}, g_{\text{odd}} \in \mathrm{RS}_{k/4}(\pi(D_x))$, establishing the isomorphism $\mathrm{RS}_{k/2}(D_x) \xrightarrow{\sim} \mathrm{RS}_{k/4}(\pi(D_x))^{=2}$. The coupling lemma follows identically to Lemma 13.

Since $\delta/2 < (1-\rho)/2 \leq \delta_J$ (the inequality $(1-\rho)/2 \leq 1 - \sqrt{\rho}$ holds for all $\rho \in (0, 1)$), the effective distance $\delta/2$ is strictly below the Johnson bound. At each round $j \geq 2$: the coupling lemma gives $\Delta_{\text{joint}} > \delta/2$, and the BCIKS proximity gap [5] at threshold $\delta/2 \leq \delta_J$ bounds the number of bad challenges to $\leq N/2^{j-1}$ (per round). Distance is *preserved* at $\delta/2$ (not further halved), as in Remark 19. The BCIKS result applies because it requires only the RS code structure and a valid coupling lemma, both of which hold for the π-based folding.

Strategy A. Each round introduces α_i with degree 1, so $f^{(R)}(y)$ is multilinear in $(\alpha_1, \ldots, \alpha_R)$. If all syndromes vanished identically, then $f^{(R)} \in \mathrm{RS}_1$ for every $(\alpha_1, \ldots, \alpha_R) \in F^R$. As in the proof of Theorem 18, we extract $f^{(0)} \in \mathrm{CRS}_k$ via a polynomial-identity argument (not a runtime reconstruction): evaluating at $\alpha_i = 0$ and $\alpha_i = 1$ at each round recovers the round-$(i{-}1)$ even/odd projections in the appropriate codes, and Lemma 66 (round 1) and the round-≥ 2 isomorphisms iterate this back to $f^{(0)} \in \mathrm{CRS}_k$, a contradiction. By Schwartz–Zippel: $\Pr[f^{(R)} \in \mathrm{RS}_1] \leq R/|F|$.

Strategy B. Combining the round-1 proximity gap (≤ 1 bad α), the round-≥ 2 distance preservation ($\leq N$ bad α per round), and the consistency check ($\geq 1{-}(1{-}\delta/2)^q$ catch probability via Lemma 17 applied with the round-i projection $\pi_i =$ (involution at round 1) $\circ$ (doubling map iterated at rounds ≥ 2), each 2^i-to-1 from D to its round-i image): $\Pr[\text{accept}] \leq NR/|F| + (1 - \delta/2)^q$. $\square$

Remark 70 (The dimension gap)**.** *The full circle polynomial space L_N of [20] has dimension $N{+}1$; the FFT subspace CRS_N used above has dimension N. The extra direction is the degree-$(N/2)$ monomial $x^{N/2}$, corresponding to the vanishing polynomial of the half-domain. In Stwo's protocol [21], the prover sends this coefficient before the first fold, reducing the proximity test to CRS_N. This adds at most $1/|F|$ to the commit-phase error (absorbed into the $NR/|F|$ term).*

Remark 71 (STIR and WHIR on circle domains)**.** *The batch-CA theorem (Theorem 44) is code-agnostic. Combined with Lemma 67, Theorem 46 extends to circle-FRI-based STIR with the same bound $(k{+}2)R/|F| + (1{-}\delta/2)^q$, and Corollary 48 extends to circle-code subcodes. The safety net now covers all deployed proximity-testing protocols: multiplicative FRI, additive FRI (char 2), and circle FRI.*

D Fiat–Shamir Compilation and Knowledge Soundness

The theorems in this paper (Theorems 18, 61, 69, 46, 52) establish *information-theoretic* soundness of the interactive proximity-testing protocols. Deployed STARK systems use the non-interactive versions obtained via the Fiat–Shamir transform (replacing verifier randomness with hash evaluations). This appendix states the non-interactive soundness and knowledge-soundness theorems that follow from our bounds via standard compilation.

D.1 Round-by-Round Soundness of FRI

The Fiat–Shamir security of FRI requires not just end-to-end soundness but *round-by-round (RBR) soundness* [33, 34]: the property that at each protocol round, the prover's probability of escaping a "doomed" state is bounded, regardless of the transcript so far.

Theorem 72 (RBR Soundness of Half-Threshold FRI)**.** *For $\mathrm{RS}[F, L, k]$ with $k = 2^m$, L admitting a fixed-point-free involution, and $\delta_J < \delta < 1 - \rho$: the FRI protocol with R rounds has round-by-round soundness with per-round error*

$$\varepsilon_{\mathrm{rbr}}^{(i)} \leq \begin{cases} n/|F| & i = 1 \text{ (commit round: } \leq 1 \text{ bad } \alpha), \\ n/|F| & 2 \leq i \leq R \text{ (commit rounds: BCIKS at } \delta/2 \leq \delta_J), \\ (1 - \delta/2)^q & i = R + 1 \text{ (query round).} \end{cases}$$

Proof. Define the "doomed" predicate at round i: the partial transcript is doomed if $\Delta(f^{(i)}, \mathrm{RS}_{k/2^i}) > \delta/2$ (the current oracle is far from the reduced code at the halved threshold). Initially ($i = 0$): $\Delta(f, \mathrm{RS}_k) > \delta > \delta/2$, so every false statement starts doomed.

Round 1 *(commit).* By the proximity gap (Theorem 14): at most 1 value of α_1 makes the fold $(\delta/2)$-close to $\mathrm{RS}_{k/2}$. Escaping the doomed state requires hitting this value: probability $\leq 1/|F| \leq n/|F|$.

Rounds $2, \ldots, R$ *(commit).* The effective distance is $\delta/2 < (1-\rho)/2 \leq \delta_J$. By the BCIKS proximity gap [5] at threshold $\delta/2$: at most n values of α_i reduce the distance below $\delta/2$. Escape probability: $\leq n/|F|$.

Round $R{+}1$ *(query).* The consistency check samples q i.i.d. base-level positions (Definition 15); by Lemma 17, the probability that all q miss the $(\delta/2)$-far deviation set is $\leq (1-\delta/2)^q$.

These per-round bounds are independent of the transcript prefix, so the protocol is RBR-sound with the stated errors. □

D.2 Non-Interactive FRI via Fiat–Shamir

The BCS transformation [32] compiles a public-coin IOP with RBR soundness into a non-interactive argument in the random oracle model. The security loss is *linear* in the number of random-oracle queries (not exponential in rounds), precisely because RBR soundness enables straight-line simulation.

Theorem 73 (Non-Interactive FRI Soundness)**.** *Let $H : \{0,1\}^* \to \{0,1\}^\kappa$ be a random oracle. For any adversary making at most Q queries to H, the Fiat–Shamir-compiled FRI protocol satisfies:*

$$\varepsilon^{\mathrm{NI}}_{\mathrm{FRI}} \;\leq\; Q \cdot \left(\frac{nR}{|F|} + \left(1 - \frac{\delta}{2}\right)^q \right) \;+\; \frac{3(Q^2+1)}{2^\kappa}.$$

Proof. By the BCS compilation theorem [32, 34]: for a public-coin IOP with RBR soundness error $\varepsilon_{\mathrm{rbr}}$, the Fiat–Shamir-compiled protocol has soundness error at most $Q \cdot \varepsilon_{\mathrm{rbr}} + 3(Q^2{+}1)/2^\kappa$, where the second term is the Merkle-tree binding error. Substituting $\varepsilon_{\mathrm{rbr}} \leq nR/|F| + (1{-}\delta/2)^q$ from Theorem 72 (summing the per-round errors: $R \cdot n/|F| + (1{-}\delta/2)^q$) gives the stated bound. □

Remark 74 (Practical parameters)**.** *For BabyBear$_4$ ($|F| = p^4 \approx 2^{124}$), $n = 2^{20}$, $R = 20$, $\kappa = 256$ (SHA-256), and $Q = 2^{64}$ (a generous query bound):*

- *Interactive FRI error: $nR/|F| + (1{-}\delta/2)^q \approx 2^{-100} + (1{-}\delta/2)^q$.*
- *Fiat–Shamir overhead: $Q \cdot \varepsilon \approx 2^{64} \cdot \varepsilon$, reducing security by $\sim$64 bits.*
- *Merkle binding: $3Q^2/2^\kappa \approx 2^{129}/2^{256} \approx 2^{-127}$ (negligible).*
- *Net: the interactive bound loses $\sim$64 bits in the non-interactive setting. At $q \approx 398$ (128-bit interactive query-phase security for $\delta = 0.4$), the non-interactive query contribution is only $\sim$64 bits against $Q = 2^{64}$ adversaries, and the total is further capped by the commit term unless the challenge field is enlarged.*

This matches the standard parameter regime: deployed systems set $\kappa \geq 2\lambda$ and target λ-bit security against $Q \leq 2^\lambda$ adversaries.

Tighter bound. The BCS compilation [34] applies Q-loss to the per-round max of the RBR error, not the total. Since each commit round has escape probability $\leq n/|F|$ (Theorem 72), a tighter bound is $\varepsilon^{\mathrm{NI}}_{\mathrm{FRI}} \leq Q \cdot \max(n/|F|,\ (1{-}\delta/2)^q) + 3(Q^2{+}1)/2^\kappa$, saving a factor of R in the commit term. We state the looser bound above for consistency with the interactive Theorem 18.

D.3 Knowledge Soundness and DEEP Polynomial Commitments

FRI is deployed not as a standalone proximity test but as the core of a *polynomial commitment scheme (PCS)* via the DEEP technique [35]: to verify a claimed evaluation $\hat{f}(\zeta) = v$ at a random point $\zeta \in F$, the verifier checks that the quotient $q(x) = (f(x) - v)/(x - \zeta)$ is close to RS_{k-1}.

Theorem 75 (DEEP-FRI Knowledge Soundness)**.** *For* $\mathrm{RS}[F, L, k]$ *with* $k = 2^m$*,* L *admitting a fixed-point-free involution, and* $\delta_J < \delta < 1-\rho$*: the DEEP-FRI polynomial commitment scheme [35] has knowledge soundness*

$$\varepsilon_{\mathrm{PCS}} \;\le\; \frac{k}{|F|} \;+\; \frac{nR}{|F|} \;+\; \left(1 - \frac{\delta}{2}\right)^q,$$

where the $k/|F|$ *term is the DEEP quotient error (probability that the random evaluation point* ζ *is degenerate).*

Proof. The DEEP technique [35] reduces polynomial commitment to proximity testing: given oracle access to f and a claimed evaluation $f(\zeta) = v$ at a random $\zeta \in F \setminus L$:

1. If $f \in \mathrm{RS}_k$ and $f(\zeta) = v$: the quotient $q(x) = (f(x) - v)/(x - \zeta)$ satisfies $q \in \mathrm{RS}_{k-1}$, and FRI accepts.
2. If $f \in \mathrm{RS}_k$ but $f(\zeta) \neq v$: the polynomial $f(x) - v$ is not divisible by $(x - \zeta)$ (since $f(\zeta) - v \neq 0$), so the function $q(x) = (f(x) - v)/(x - \zeta)$ on L (well-defined since $\zeta \notin L$) is not the evaluation of any degree-$<(k-1)$ polynomial. Hence $\Delta(q, \mathrm{RS}_{k-1}) \geq 1 - (k-1)/n = 1 - \rho + 1/n > \delta$, and FRI rejects with high probability.
3. If $f \notin \mathrm{RS}_k$: the verifier sends random $\zeta \in F \setminus L$ and the prover provides a claimed value v. By the DEEP worst-case-to-average-case reduction [35]: for all but $\leq k$ degenerate values of ζ, the quotient function $q(x) = (f(x) - v)/(x - \zeta)$ on L is far from RS_{k-1}, regardless of the prover's choice of v. Probability of a degenerate ζ: $k/|F|$. Conditional on non-degenerate ζ: FRI catches the far quotient via Theorem 18.

Taking the union bound: $\varepsilon_{\mathrm{PCS}} \leq k/|F| + nR/|F| + (1-\delta/2)^q$.

Knowledge extraction. By the round-by-round framework [34]: the extractor runs the interactive protocol with the adversary, catching the adversary's committed polynomial via the DEEP evaluation at the random point ζ. The RBR knowledge error is $\varepsilon_{\mathrm{rbr}} + k/|F|$, and the BCS compilation preserves knowledge soundness with the same Q-linear loss as for plain soundness. □

Corollary 76 (Non-Interactive PCS Knowledge Soundness)**.** *The Fiat–Shamir-compiled DEEP-FRI PCS has knowledge soundness against* Q*-query adversaries:*

$$\varepsilon^{\mathrm{NI}}_{\mathrm{PCS}} \;\le\; Q \cdot \left(\frac{k + nR}{|F|} + \left(1 - \frac{\delta}{2}\right)^q\right) + \frac{3(Q^2 + 1)}{2^\kappa}.$$

Remark 77 (Scope of the knowledge-soundness claim)**.** *Theorem 75 and Corollary 76 give knowledge soundness for the DEEP-FRI polynomial commitment scheme: the extractor recovers a degree-*$<k$ *polynomial consistent with the committed oracle. The full SNARK (e.g., Plonk with FRI-PCS) additionally requires showing that the extracted polynomial satisfies the circuit constraints; this is the arithmetization step, which is protocol-specific and orthogonal to the proximity-testing soundness. Our contribution is the proximity-testing layer; the arithmetization-to-PCS composition follows standard arguments from [34].*

References

[1] Ethereum Foundation, *Proximity Prize*, `https://proximityprize.org`, 2026.

[2] G. Arnon, D. Boneh, G. Fenzi, *Open problems in list decoding and correlated agreement*, ePrint 2026/680.

[3] E. Crites, A. Stewart, *On Reed–Solomon proximity gaps conjectures*, ePrint 2025/2046.

[4] E. Ben-Sasson, D. Carmon, U. Haböck, S. Kopparty, S. Saraf, *On proximity gaps for Reed–Solomon codes*, ePrint 2025/2055, 2025.

[5] E. Ben-Sasson, D. Carmon, Y. Ishai, S. Kopparty, S. Saraf, *Proximity gaps for Reed–Solomon codes*, FOCS 2020.

[6] R. Goyal, V. Guruswami, *Optimal proximity gaps for subspace-design codes and (random) Reed–Solomon codes*, ePrint 2025/2054.

[7] F. G. Jeronimo, L. Liu, P. Rajpal, *Optimal proximity gaps for folded Reed–Solomon codes via subspace designs*, arXiv:2601.10047.

[8] G. Arnon, A. Chiesa, G. Fenzi, E. Yogev, *STIR: Reed–Solomon proximity testing with fewer queries*, CRYPTO 2024; ePrint 2024/390.

[9] G. Arnon, A. Chiesa, G. Fenzi, E. Yogev, *WHIR: Reed–Solomon proximity testing with super-fast verification*, ePrint 2024/1586.

[10] A. Kambiré, *Proximity gaps conjecture fails near capacity over prime fields*, arXiv:2604.09724, 2026.

[11] R. Lidl, H. Niederreiter, *Finite Fields*, Cambridge, 2nd ed., 1997.

[12] P. Frankl, R. M. Wilson, *Intersection theorems with geometric consequences*, Combinatorica 1(4):357–368, 1981.

[13] M. Sudan, *Decoding of Reed–Solomon codes beyond the error-correction bound*, J. Complexity 13(1):180–193, 1997.

[14] V. Guruswami, M. Sudan, *Improved decoding of Reed–Solomon and algebraic-geometric codes*, IEEE Trans. Inform. Theory 45(6):1757–1767, 1999.

[15] T. Helleseth, *Some results about the cross-correlation function between two maximal linear sequences*, Discrete Math. 16(3):209–232, 1976.

[16] S. W. Golomb, G. Gong, *Signal Design for Good Correlation: For Wireless Communication, Cryptography, and Radar*, Cambridge University Press, 2005.

[17] O. Moreno, C. J. Moreno, *Exponential sums and Goppa codes: I*, Proc. Amer. Math. Soc. 111(2):523–531, 1991.

[18] E. Ben-Sasson, I. Bentov, Y. Horesh, M. Riabzev, *Fast Reed–Solomon interactive oracle proofs of proximity*, in ICALP 2018, LIPIcs vol. 107, 2018. Full version: ePrint 2017/134.

[19] B. E. Diamond, J. Posen, *Proximity testing with logarithmic randomness*, IACR ePrint 2023/1784, 2023.

[20] U. Haböck, D. Levit, S. Papini, *Circle STARKs*, ePrint 2024/278, 2024.

[21] D. Carmon, S. Goldberg, U. Haböck, V. Lerer, S. Lesokhin, *Stwo whitepaper*, ePrint 2026/532, 2026.

[22] R. Chai, X. Fan, *Structural results on per-word list size for Reed–Solomon codes*, companion note, in preparation, 2026.

[23] R. Chai, X. Fan, *Equal-threshold CA bound: $\varepsilon_{\mathrm{ca}} = \binom{n}{w}/|F|$ via erasure-pattern counting*, companion note, in preparation, 2026.

[24] R. Chai, X. Fan, *Birthday-bound at codimension excess $c = 2$: Bernstein concentration via Möbius exclusion, and the k-wise common-core obstruction*, companion note, in preparation, 2026.

[25] R. Chai, X. Fan, *Open-Set Rank Lemma at codimension excess $c \geq 3$: phase diagram, triangle and tetrahedron obstructions at small primes*, companion note, in preparation, 2026.

[26] S. Janson, *On concentration of probability*, in Contemporary Combinatorics, Bolyai Soc. Math. Stud., vol. 10, Springer, 2002, pp. 289–301.

[27] D. Dubhashi, D. Ranjan, *Balls and bins: A study in negative dependence*, Random Structures & Algorithms, vol. 13, no. 2, 1998, pp. 99–124.

[28] G. N. Rothblum, S. Vadhan, A. Wigderson, *Interactive proofs of proximity: delegating computation in sublinear time*, in STOC 2013. Full version: ECCC TR12-182.

[29] S. Ames, C. Hazay, Y. Ishai, M. Venkitasubramaniam, *Ligero: Lightweight sublinear arguments without a trusted setup*, in Proceedings of the 24th ACM Conference on Computer and Communications Security (CCS), 2017.

[30] R. M. Roth and G. Zémor, personal communication, 2018. Cited in [5].

[31] E. Ben-Sasson, S. Kopparty, S. Saraf, *Worst-case to average-case reductions for the distance to a code*, in 33rd Computational Complexity Conference (CCC), pages 24:1–24:23, 2018.

[32] E. Ben-Sasson, A. Chiesa, N. Spooner, *Interactive oracle proofs*, in TCC 2016-B, LNCS 9986, pages 31–60, 2016.

[33] R. Canetti, Y. Chen, J. Holmgren, A. Lombardi, G. N. Rothblum, R. D. Rothblum, D. Wichs, *Fiat–Shamir: from practice to theory*, in STOC 2019, pages 1082–1090, 2019.

[34] A. Block, A. Garreta, J. Katz, J. Thaler, P. N. Tiwari, M. Zajac, *Fiat–Shamir security of FRI and related SNARKs*, in Asiacrypt 2023, LNCS 14444, pages 3–40, 2023.

[35] E. Ben-Sasson, L. Goldberg, S. Kopparty, S. Saraf, *DEEP-FRI: sampling outside the box improves soundness*, ePrint 2019/336, 2019.

[36] R. Chai, X. Fan, *Computational companion: Lean 4 formalization and verification scripts for FRI soundness above the Johnson bound*, `https://github.com/iotexproject/rs-proximity-gaps`, 2026.

Acknowledgments

We thank Professor Guang Gong (University of Waterloo) for reviewing an earlier draft and for the sequence-school perspective, and the Ethereum Foundation for posing the Proximity Prize that motivated this work.