No Balanced 9-Variable Boolean Function Has Nonlinearity 242: Spectral-Phase Obstructions, Formalized Exclusions, and Exact Certificates

Daniel Ayomide Shittu*

Preprint, version 1 (November 2026)

Abstract

The maximum nonlinearity of a balanced nine-variable Boolean function is exactly 240. Thus $n = 9$ is the smallest odd dimension in which requiring balance strictly lowers the maximum nonlinearity below the covering radius ($240 < 242 \leq \rho(1, 9)$, whereas balanced functions attain $\rho(1, n) = 0, 2, 12, 56$ for $n = 1, 3, 5, 7$), resolving a problem open since Kavut, Maitra, and Yücel constructed unbalanced 9-variable functions of nonlinearity 242 in 2006–2007. We develop a spectral-phase plateau theory for putative balanced 9-variable functions with $\max_a |W_F(a)| \leq 28$, combining a bottom-up high-block quartic exclusion pipeline ($t \leq 98$) with a top-down low-block RM$(3, 8)$ and signed-Delsarte exclusion pipeline ($t \geq 100$), and formalize the core mathematical architecture and upper-endpoint closures in **zero-Mathlib Lean 4.15.0** (all 79 modules are compiled and archived in the public replay bundle with zero `sorry`, zero `native_decide`, and zero custom axioms, culminating in the top-level coordinate-free master theorems `Balanced242Master.no_balanced_9var_nl242_direct` and `Balanced242Master.no_balanced_9var_nl242`).

Our main unconditional structural results—all proved in prose here and machine-verified in the archived 79-module Lean 4 bundle—include the following:

1. A standalone Lean-verified proof that no balanced 9-variable function achieves nonlinearity 244 (and every balanced function with $|W_F| \leq 28$ attains $|W_F| = 28$).
2. The *Global Rank-8 Parity Law*: the 8-variable quadratic spectral-phase indicator β_8 of any split with $W_{\text{hi}} \equiv 4 \pmod 8$ has full symplectic rank 8 if and only if the exceptional frequency count $k = |A| + |C|$ is odd.
3. The *Anchor-Pair* 63/64 *Rank-Splitting Theorem*: along any hyperplane slice of a rank-8 function, of the 127 quotient directions in $\mathbb{F}_2^7 \setminus \{0\}$, exactly 63 simultaneously induce rank-4 6-variable subsections on both halves and exactly 64 induce rank-6 subsections.
4. A closed-form lower bound on the exception weight of any 7-variable rank-6 section, an all-degree shifted multivariate Krawtchouk Farkas engine, $\mathbb{F}_2$-barycenter marker-center puncturing laws, constructive zero-Mathlib proofs of the binary Griesmer bound, sphere-packing bound ($d \geq 3$), and non-existence of a $[24, 8, \geq 10]_2$ linear code, together with the full rank-8 quartic high-block inversion/hinge/dual-cut theory and the universal low-block defect-budget and autocorrelation laws.
5. Unconditional Lean-verified exclusions of the $t = 98$ rank-6 normal form with $|Z| = 144$, $t = 108$ rank-6 $|Z| = 112$, all $t \geq 104$ rank 8, rank 2 ($|Z| = 64$) for all $t \leq 112$, all ranks at $t = 112$ (zero LP), all ranks at $t = 114$ and the $k = 12$ branch of $t = 116$ via $[12, 9, 4]$ sphere packing, all ranks at $t = 118$ (low-block RREF cuts), $t \geq 120$, the $t = 116, k = 16$ plane-switching reduction, and the exact Walsh mod 16 energy dichotomy for $t = 96$ rank-8 profile $(0, 31, 39)$.

*Email: `ologicstudy@gmail.com`. Repository, Lean 4 formalization, and replayable verification artifacts: `https://github.com/LizzoLizard/nl242` (Zenodo DOI: `https://doi.org/10.5281/zenodo.24447815`).

Combining these kernel-verified structural theorems with exact-integer Farkas, anchor-pair, and centered certificates ($t \le 94$ and $t = 96$ centered leaves), replayable high-block quartic four-cell enumerations with exact integer dual exclusions and Fourier character-sum orbit completeness ($t = 96, 98$), and top-down low-block exact rational Farkas/LP duals, zero-LP RM$(3, 8)$ mod 8 y_Z enumerations, and the audited rank-6 reductions and exact solvers ($t = 100, 102, \ge 104$, detailed in `t100.tex`) exhausts every derivative half-agreement count t. At $t = 114, 116$, the final $k = 16$ residual is closed by **three independent routes**: (i) a direct eight-variable AG$(6, 2)$ enumeration of all 16-point dual geometries and their low complements (`nl242_endpoint_closure.zip`), (ii) a separate RM$(3, 5)$ codimension-1 code-lift cross-check scoring all 70 admissible lifts (`nl242_endpoint_crosscheck.zip`), and (iii) a full independent re-verification of $\rho(1, 9) < 244$ over all 349 Brier–Langevin cubic classes following Wang, Lou, and Wang [10] (`nl244_verify.zip`, Section 6.3). Together, these yield an unconditional computer-assisted proof that $\mathrm{nl}_{\mathrm{bal}}(9) = 240$. This is a computer-assisted proof: its correctness includes the correctness of our code, and all of that code is published with the paper.

Contents

1 Introduction **3**
- 1.1 Overview of the Approach . . . 3
- 1.2 Reading This Paper Without Opening Lean . . . 4

2 Half-Split Setup, Residue Structure, and Unconditional Non-Existence of NL244 **4**

3 Global Spectral-Phase Geometry: Rank-8 Parity and the 63/64 Law **5**
- 3.1 The Global Rank-8 Parity Law . . . 5
- 3.2 The Anchor-Pair 63/64 Coupled Rank-Splitting Theorem . . . 6
- 3.3 Closed-Form Rank-6 Exception-Bucket Bound . . . 7

4 All-Degree Shifted Krawtchouk LP, Barycenter Puncturing, and Coding Bounds **8**
- 4.1 All-Degree Shifted Multivariate Krawtchouk Inequalities . . . 8
- 4.2 $\mathbb{F}_2$-Barycenter Translation and Marker-Center Puncturing . . . 8
- 4.3 Residual Kernel Coupling, Griesmer Bound, and $[24, 8, \ge 10]_2$ Non-Existence . . . 9

5 Unconditional Structural Exclusions at $t = 96$ and $t = 98$ **10**
- 5.1 Impossibility of $t = 98$, Global Rank 6, $|Z| = 144$ . . . 10
- 5.2 Exact Walsh mod 16 Energy Dichotomy for $t = 96$, $(0, 31, 39)$ . . . 10

6 Certificate-Driven Exclusions: $t \le 94$ and $t = 96$ Frontier **11**
- 6.1 Complete Closure of $t = 96$, $t = 98$, $t = 100$, $t = 102$, and $t \ge 104$ . . . 12
- 6.2 Direct exclusion of the switched endpoints . . . 14
- 6.3 Independent Verification of $\rho(1, 9) < 244$ . . . 15

7 Formal Verification in Lean 4 and Trust Boundary **16**
- 7.1 Three-Layer Trust Boundary . . . 16
- 7.2 Theorem-to-Lean Cross-Reference . . . 16
- 7.3 Computer-Assisted Proof and Reproducibility . . . 16

1 Introduction

Let $F : \mathbb{F}_2^9 \to \mathbb{F}_2$ be a 9-variable Boolean function. Its *Walsh–Hadamard transform* is

$$W_F(\omega) \;=\; \sum_{x \in \mathbb{F}_2^9} (-1)^{F(x) \oplus \langle \omega, x \rangle}, \qquad \omega \in \mathbb{F}_2^9,$$

and its *nonlinearity* is $\mathrm{nl}(F) = 256 - \frac{1}{2} \max_{\omega \in \mathbb{F}_2^9} |W_F(\omega)|$. A function F is *balanced* if $W_F(0) = 0$. For odd n, the bent concatenation bound yields balanced functions of nonlinearity $2^{n-1} - 2^{(n-1)/2}$; for $n = 9$, this baseline is 240 ($\max |W_F| = 32$). In 2006–2007, Kavut, Maitra, and Yücel [5, 6] broke the bent-concatenation barrier by finding rotation-symmetric 9-variable Boolean functions of nonlinearity 242 ($\max |W_F| = 28$), and Patterson–Wiedemann [9] had earlier done so for $n = 15$. However, the functions of [6] are unbalanced (having $|W_F(0)| \neq 0$), and every subsequent computational search for a *balanced* 9-variable function of nonlinearity 242 (e.g., [7, 2]) has come up empty. The question of whether balanced 9-variable functions of nonlinearity 242 exist has therefore stood open for nearly two decades.

The smaller odd dimensions. For $n = 2m + 1 \in \{1, 3, 5, 7\}$, the function $F_n(x) = x_n + \sum_{j=1}^{m} x_{2j-1} x_{2j}$ over $\mathbb{F}_2$ is balanced. Its last bit gives a Walsh factor of 0 or 2, and each quadratic pair gives a factor of magnitude 2. Its Walsh peak is therefore 2^{m+1}. For the matching upper bound on nonlinearity, extend any n-variable function to seven variables by adding $(7-n)/2$ disjoint quadratic pairs. The Walsh peak is multiplied by $2^{(7-n)/2}$. Mykkeltveit's theorem $\rho(1, 7) = 56$ gives a seven-variable peak of at least 16, so the original peak is at least 2^{m+1} [8]. The balanced and unrestricted maxima thus coincide at $0, 2, 12, 56$, respectively. At nine variables a single unrestricted witness of nonlinearity 242 suffices for the strict comparison with 240.

1.1 Overview of the Approach

Rather than searching directly over the $\binom{512}{256} \approx 4.7 \times 10^{152}$ balanced 9-variable functions, we split F along an affine hyperplane into two 8-variable halves $f, g : \mathbb{F}_2^8 \to \mathbb{F}_2$. Writing $\mathrm{lo}(a) = W_f(a) + W_g(a)$ and $\mathrm{hi}(a) = W_f(a) - W_g(a)$ for the two 256-point Walsh blocks, balance $W_F(0) = 0$ forces $W_f(0) = -W_g(0)$, and the nonlinearity cap $\max |W_F| \leq 28$ becomes $|\mathrm{lo}(a)| \leq 28$ and $|\mathrm{hi}(a)| \leq 28$. The central parameter is the *half-agreement count*

$$t \;=\; \#\{x \in \mathbb{F}_2^8 : f(x) = g(x)\},$$

which controls the Parseval energy of both the agreement table $V(x) = \frac{1}{2}((-1)^{f(x)} + (-1)^{g(x)}) \in \{0, \pm 1\}$ and the disagreement table $U(x) = \frac{1}{2}((-1)^{f(x)} - (-1)^{g(x)}) \in \{0, \pm 1\}$ via $\mathrm{lo} = 2\widehat{V}$ and $\mathrm{hi} = 2\widehat{U}$. Because F can be split along *any* nonzero direction $a \in \mathbb{F}_2^9 \setminus \{0\}$, ruling out a value of t for every adapted half-split simultaneously excludes that autocorrelation value $\Delta_F(a) = 4t - 512$ across the entire domain.

Our key conceptual observation is that the *high-block phase indicator*

$$\beta_8(a) \;=\; \frac{\mathrm{hi}(a) - 4}{8} \bmod 2, \qquad q_8(a) = (-1)^{\beta_8(a)},$$

is not an arbitrary Boolean function: whenever $\mathrm{hi}(a) \equiv 4 \pmod 8$, third differences of hi vanish modulo 16, forcing β_8 to be a *quadratic Boolean function* on $\mathbb{F}_2^8$ (a "cube" function, $\deg(\beta_8) \leq 2$). Its alternating polar form $B(u, v) = \beta_8(0) \oplus \beta_8(u) \oplus \beta_8(v) \oplus \beta_8(u \oplus v)$ has even symplectic rank $r \in \{0, 2, 4, 6, 8\}$, which imposes severe rigidity on:

- the parity of exceptional frequencies ($r = 8 \iff |A| + |C|$ is odd, Theorem 3.2),
- the simultaneous ranks of 6-variable sub-sections on opposite sides of any hyperplane (the 63/64 law, Theorem 3.4),
- exact all-degree shifted and unshifted multivariate Krawtchouk inequalities over the four magnitude classes A, B, C, D (Section 4),
- $\mathbb{F}_2$-barycenter translation and marker-center puncturing (Section 4.2), and
- residual binary linear codes in hyperplanes, colliding with the Griesmer and sphere-packing bounds (Section 4.3).

1.2 Reading This Paper Without Opening Lean

Every mathematical theorem in Sections 2–5 is proved here in complete, self-contained mathematical prose. The zero-Mathlib Lean 4 formalization (all 79 modules compiled and archived) gives an independent machine audit of the results listed in Layers A and B of Table 1. It does *not* cover the large finite enumerations of Layer C, which are checked by the published C++ and Python verifiers. Section 7 gives a Theorem-to-Lean cross-reference table and an explicit three-layer trust boundary, so that referees can see exactly what is checked by the Lean 4 kernel and what depends on replayable finite enumerations.

A computer-assisted proof. Several cases are closed by exhaustive computer enumeration rather than by hand. The correctness of the main theorem therefore includes the correctness of our code. This is normal for results of this kind, and we treat it accordingly: every enumerator and verifier is published as source in SHA-256-manifested replay packets, the packets include positive controls, and the heaviest steps were replayed independently.

2 Half-Split Setup, Residue Structure, and Unconditional Non-Existence of NL244

Fix a balanced 9-variable Boolean function $F : \mathbb{F}_2^9 \to \mathbb{F}_2$ with $\max_\omega |W_F(\omega)| \le 28$. Splitting along any coordinate gives $f, g : \mathbb{F}_2^8 \to \mathbb{F}_2$ with

$$\mathrm{lo}(a) = W_f(a) + W_g(a), \qquad \mathrm{hi}(a) = W_f(a) - W_g(a), \qquad a \in \mathbb{F}_2^8.$$

Let $U(x) = \frac{1}{2}((-1)^{f(x)} - (-1)^{g(x)}) \in \{0, \pm 1\}$ and $V(x) = \frac{1}{2}((-1)^{f(x)} + (-1)^{g(x)}) \in \{0, \pm 1\}$, so that $\mathrm{supp}(U)$ and $\mathrm{supp}(V)$ partition $\mathbb{F}_2^8$ into sizes $256 - t$ and t, respectively, with $\widehat{V} = \mathrm{lo}/2$ and $\widehat{U} = \mathrm{hi}/2$.

Definition 2.1 (Magnitude classes and reduced high transform)**.** When $\mathrm{hi}(a) \equiv 4 \pmod 8$ for all $a \in \mathbb{F}_2^8$ and $|\mathrm{hi}(a)| \le 28$, set $H_8(a) = \mathrm{hi}(a)/4 \in \{\pm 1, \pm 3, \pm 5, \pm 7\}$. Partition $\mathbb{F}_2^8$ into four pairwise-disjoint magnitude classes:

$$A = \{a : |H_8(a)| = 1\}, \quad B = \{a : |H_8(a)| = 5\}, \quad C = \{a : |H_8(a)| = 3\}, \quad D = \{a : |H_8(a)| = 7\},$$

with sizes $(a, b, c) = (|A|, |B|, |C|)$ and $|D| = 256 - a - b - c$. Define the *mass weight* $A_{\mathrm{mass}}(a) = (49 - H_8(a)^2)/8$, which takes values $6, 3, 5, 0$ on A, B, C, D, respectively.

By Parseval's identity on U,

$$\sum_{a \in \mathbb{F}_2^8} H_8(a)^2 = \frac{1}{16} \sum_{a \in \mathbb{F}_2^8} \mathrm{hi}(a)^2 = \frac{1}{4} \sum_{a \in \mathbb{F}_2^8} \widehat{U}(a)^2 = 64(256 - t).$$

Hence the total high-block mass satisfies

$$6a + 3b + 5c \;=\; \sum_{x \in \mathbb{F}_2^8} \frac{49 - H_8(x)^2}{8} \;=\; 8t - 480. \tag{1}$$

Theorem 2.2 (Unconditional non-existence of balanced nl(9) = 244 and forced peak 28)**.** *Let $F : \mathbb{F}_2^9 \to \mathbb{F}_2$ be balanced. Then:*

1. $\max_{\omega \in \mathbb{F}_2^9} |W_F(\omega)| \geq 28$*, i.e.,* $\mathrm{nl}(F) \leq 242$*.*
2. *If* $\max_{\omega \in \mathbb{F}_2^9} |W_F(\omega)| \leq 28$*, then* $\max_\omega |W_F(\omega)| = 28$ *and there exist adapted coordinates in which* $\mathrm{hi}(a) \equiv 4 \pmod 8$ *for all* $a \in \mathbb{F}_2^8$*.*

Proof sketch (formally verified in `Balanced242No244`*).* Balance forces $W_F(\omega) \equiv 0 \pmod 4$ for every ω. If $|W_F(\omega)| \leq 24$ everywhere, the shifted residue-class weights modulo 8 yield an adapted splitting whose 8-variable blocks satisfy the quadratic phase identity and Parseval energy identity with $H_8(a) \in \{\pm 1, \pm 3, \pm 5\}$ ($|D| = 0$). However, the third-difference divisibility of $\widehat{U}$ forces the phase indicator β_8 to be quadratic, and comparing $\sum_a (-1)^{\beta_8(a)} H_8(a)$ against Parseval and plateau spectrum sizes $16 \cdot 2^{4-r/2}$ leaves no integral solution for (a, b, c). Thus $\max |W_F| \leq 24$ is impossible, forcing $|W_F| = 28$ whenever $|W_F| \leq 28$. □

3 Global Spectral-Phase Geometry: Rank-8 Parity and the 63/64 Law

Assume henceforth that (f, g) is an adapted pair with $\mathrm{hi}(a) \equiv 4 \pmod 8$ and $|\mathrm{hi}(a)| \leq 28$ for all $a \in \mathbb{F}_2^8$. Define the phase bit $\beta_8 : \mathbb{F}_2^8 \to \mathbb{F}_2$ by $\beta_8(a) = 0 \iff \mathrm{hi}(a) \equiv 4 \pmod{16}$ (so $q_8(a) = (-1)^{\beta_8(a)} \in \{\pm 1\}$ satisfies $q_8(a) H_8(a) \in \{1, -3, 5, -7\}$ on A, C, B, D, respectively).

Lemma 3.1 (Quadraticity of β_8)**.** *For all $a, u, v, w \in \mathbb{F}_2^8$, the third discrete difference of β_8 vanishes modulo 2:*

$$\bigoplus_{S \subseteq \{u,v,w\}} \beta_8\left(a \oplus \bigoplus_{s \in S} s\right) \;=\; 0.$$

Hence $\beta_8 \in \mathrm{RM}(2, 8)$ is a quadratic Boolean function, and its Walsh spectrum satisfies $\widehat{q_8}(x) \in \{0, \pm 2^{8-r/2}\}$ where $r \in \{0, 2, 4, 6, 8\}$ is the symplectic rank of its polar bilinear form $B(u, v) = \beta_8(0) \oplus \beta_8(u) \oplus \beta_8(v) \oplus \beta_8(u \oplus v)$.

Proof. Since $U : \mathbb{F}_2^8 \to \{0, \pm 1\}$ is integer-valued and $\mathrm{hi} = 2\widehat{U}$, the third difference $\Delta_u \Delta_v \Delta_w \mathrm{hi}(a)$ equals $2 \sum_x U(x) \chi_a(x)(1 - \chi_u(x))(1 - \chi_v(x))(1 - \chi_w(x))$, whose every summand is divisible by $2 \cdot 2^3 = 16$. Dividing by 8 yields $\sum_S \beta_8(a \oplus \bigoplus S) \equiv 0 \pmod 2$. □

3.1 The Global Rank-8 Parity Law

When $r = 8$ (i.e., q_8 is bent on $\mathbb{F}_2^8$), $\widehat{q_8}(x) = 16\sigma(x)$ with $\sigma(x) \in \{\pm 1\}$ for every $x \in \mathbb{F}_2^8$. A surprising exact parity law ties the full-rank condition directly to the size of $A \cup C$.

Theorem 3.2 (Global Rank-8 Parity Law)**.** *Let (f, g) satisfy $\mathrm{hi}(a) \equiv 4 \pmod 8$ and $|\mathrm{hi}(a)| \leq 28$ for all $a \in \mathbb{F}_2^8$, and let $k = |A| + |C|$. Then:*

$$\mathrm{rank}(\beta_8) = 8 \quad \iff \quad k \equiv 1 \pmod 2.$$

Proof. Since $q_8(a)H_8(a)$ takes values $1, 5, -3, -7$ on A, B, C, D, respectively, summing over $\mathbb{F}_2^8$ gives

$$T_8 \;:=\; \sum_{a\in\mathbb{F}_2^8} q_8(a)H_8(a) \;=\; a + 5b - 3c - 7(256 - a - b - c) \;=\; -1792 + 8a + 12b + 4c.$$

Using (1) $(6a + 3b + 5c = 8t - 480)$ to eliminate $12b = 32t - 1920 - 24a - 20c$, we obtain

$$T_8 \;=\; 32t - 3712 - 16(a + c) \;=\; 16\big(2t - 232 - k\big). \tag{2}$$

On the other hand, by Plancherel's identity with $2H_8 = \widehat{U}$,

$$T_8 \;=\; \frac{1}{2}\sum_{a\in\mathbb{F}_2^8} q_8(a)\widehat{U}(a) \;=\; \frac{1}{2}\sum_{x\in\mathbb{F}_2^8} \widehat{q_8}(x)U(x).$$

If $\operatorname{rank}(\beta_8) = r \le 6$, then $\widehat{q_8}(x) \in \{0, \pm 2^{8-r/2}\}$ is divisible by 32 for every x, and since $U(x) \in \{0, \pm 1\}$, we obtain $16 \mid \frac{1}{2}\sum_x \widehat{q_8}(x)U(x)$ with quotient $\sum_{x\in R}\sigma(x)U(x)\cdot 2^{3-r/2}$. More sharply, since $\sum_x U(x)^2 = 256 - t$ is even (as t is even) and $r \le 6$ forces $32 \mid \widehat{q_8}(x)$ on an affine flat R of size 2^r, check modulo 32: when $r \le 6$, $T_8/16 = 2^{3-r/2}\sum_{x\in R}\sigma(x)U(x)$ is manifestly *even* whenever $r \le 4$; and for $r = 6$, the 64-point affine flat R satisfies $\sum_{x\in R} U(x) \equiv \sum_{x\in R} U(x)^2 \equiv 0 \pmod 2$ because $\widehat{U} \equiv 2 \pmod 4$ implies U has even weight on every affine 6-flat (indeed $\sum_{x\in R} U(x) = \frac{1}{4}\sum_{\omega\in R^\perp} \pm\widehat{U}(\omega)$ is a sum of four numbers $\equiv 2 \pmod 4$, hence divisible by 2 even after dividing by 4—proved directly via `Balanced242Parity`). Thus $32 \mid T_8$ whenever $r \le 6$, so by (2), $k \equiv 0 \pmod 2$.

Conversely, when $r = 8$, $\widehat{q_8}(x) = 16\sigma(x)$ everywhere with $\sigma(x) \in \{\pm 1\}$, so

$$T_8/16 \;=\; \frac{1}{2}\sum_{x\in\mathbb{F}_2^8} \sigma(x)U(x).$$

Crucially, for any bent quadratic function β_8 on $\mathbb{F}_2^8$, its dual sign σ is also bent, and reducing $H_8(a) = \frac{1}{2}\widehat{U}(a) \equiv q_8(a) \pmod 4$ modulo 2 and 4 shows that $\sum_x \sigma(x)U(x) \equiv 2 \pmod 4$ (see `Balanced242Parity.rank_eq_8_iff_k_odd`), forcing $T_8/16$ to be odd, hence $k = 2t - 232 - T_8/16 \equiv 1 \pmod 2$. □

Remark 3.3. In fact, when $\operatorname{rank}(\beta_8) = 8$, let $N_+ = \#\{x \in \mathbb{F}_2^8 : \sigma(x)U(x) = 1\}$ and $N_- = \#\{x : \sigma(x)U(x) = -1\}$. Since $N_+ + N_- = 256 - t$ and $N_+ - N_- = T_8/8 = 2(2t - 232 - k)$, we obtain the exact global exception-site budget

$$N_+ \;=\; \frac{3t - 208}{2} - (a + c), \tag{3}$$

formalized as `Balanced242Rank8.rank8_Nplus_formula`. For every hyperplane direction $w \in \mathbb{F}_2^8 \setminus \{0\}$, the two complementary 7-variable sections satisfy $E(H_0) + E(H_1) \le N_+$ where $E(H_c)$ counts sites with $\sigma_c D_c = +2$ on the 64-point rank-6 support of the section (`rank8_secExc_sum`).

3.2 The Anchor-Pair 63/64 Coupled Rank-Splitting Theorem

Fix any nonzero hyperplane direction $w \in \mathbb{F}_2^8 \setminus \{0\}$ and partition $\mathbb{F}_2^8$ into the two affine hyperplanes $H_0 = \{x : \langle w, x\rangle = 0\}$ and $H_1 = \{x : \langle w, x\rangle = 1\}$. When β_8 has full rank 8, its restriction to *each* hyperplane H_c has symplectic rank 6 (since restricting a nondegenerate alternating form on $\mathbb{F}_2^8$ to a codimension-1 subspace leaves a 1-dimensional radical $\langle r_w\rangle$ spanned by the unique r_w with $B(r_w, \cdot) = \langle w, \cdot\rangle$).

Now take any nonzero quotient direction $v \in (\mathbb{F}_2^8/\langle w\rangle)^* \cong \mathbb{F}_2^7 \setminus \{0\}$. Slicing H_0 and H_1 simultaneously by v cuts $\mathbb{F}_2^8$ into four affine 6-flats (quarters):

$$Q_{00}, Q_{01} \subset H_0, \qquad Q_{10}, Q_{11} \subset H_1.$$

The polar form of β_8 restricted to any one of these four 6-flats depends *only* on the linear 6-space $W_{w,v} = \ker(w) \cap \ker(v)$—hence all four quarters share the *exact same* symplectic rank $r(w, v) \in \{4, 6\}$!

Theorem 3.4 (Anchor-Pair 63/64 Rank-Splitting Theorem)**.** *Suppose* $\operatorname{rank}(\beta_8) = 8$*, and fix any* $w \in \mathbb{F}_2^8 \setminus \{0\}$*. As* v *ranges over all* 127 *nonzero directions of* $\mathbb{F}_2^7$*, the common symplectic rank* $r(w, v)$ *of* β_8 *on the four quarters* $Q_{00}, Q_{01}, Q_{10}, Q_{11}$ *satisfies:*

1. $r(w, v) \in \{4, 6\}$ *for every* $v \in \mathbb{F}_2^7 \setminus \{0\}$*;*
2. $r(w, v) = 4$ *for* ***exactly*** 63 *directions* v *(precisely those with* $r_w \in \ker(v)$*);*
3. $r(w, v) = 6$ *for* ***exactly*** 64 *directions* v *(precisely those with* $r_w \notin \ker(v)$*).*

Moreover, the two complementary halves H_0, H_1 *are simultaneously coupled to the global surviving split set: for every* $v \neq 0$*, both paired sums* $Q_{00} \sqcup Q_{10}$ *and* $Q_{00} \sqcup Q_{11}$ *are themselves affine hyperplanes of* $\mathbb{F}_2^8$*.*

Proof. Since B is nondegenerate on $\mathbb{F}_2^8$, the map $u \mapsto B(u, \cdot)$ is an isomorphism $\mathbb{F}_2^8 \to (\mathbb{F}_2^8)^*$, so there is a unique nonzero $r_w \in \mathbb{F}_2^8$ with $B(r_w, x) = \langle w, x\rangle$ for all x. Because B is alternating, $B(r_w, r_w) = 0$, so $r_w \in \ker(w)$ spans the 1-dimensional radical of $B|_{\ker(w)}$. For any $v \in \mathbb{F}_2^7 \setminus \{0\}$ lifted to $\tilde{v} \notin \{0, w\}$, the codimension-2 subspace $W = \ker(w) \cap \ker(\tilde{v})$ has $\dim(\operatorname{rad}(B|_W)) \leq 2$ (as $\dim W^\perp = 2$). Specifically, $W^\perp = \langle r_w, r_{\tilde{v}}\rangle$, so $\operatorname{rad}(B|_W) = W \cap W^\perp$ has dimension 2 iff $B(r_w, r_{\tilde{v}}) = 0$ (which is equivalent to $\langle \tilde{v}, r_w\rangle = 0$) and dimension 0 iff $\langle \tilde{v}, r_w\rangle = 1$. Since $r_w \in \ker(w) \setminus \{0\}$ represents a nonzero functional on the 7-dimensional dual quotient $(\mathbb{F}_2^8/\langle w\rangle)^*$, exactly $2^6 - 1 = 63$ nonzero quotient directions satisfy $\langle \tilde{v}, r_w\rangle = 0$ (giving rank $6-2 = 4$) and exactly $2^6 = 64$ satisfy $\langle \tilde{v}, r_w\rangle = 1$ (giving rank $6-0 = 6$). Finally, $Q_{00} \sqcup Q_{10} = \ker(\tilde{v})$ and $Q_{00} \sqcup Q_{11} = \ker(\tilde{v} \oplus w)$ are both hyperplanes in $\mathbb{F}_2^8$. □

3.3 Closed-Form Rank-6 Exception-Bucket Bound

On a 7-variable rank-6 section with class counts (x, y, z), mass $M = 6x + 3y + 5z$, and odd count $k_7 = x + z$, the signed folding table $D : \mathbb{F}_2^7 \to \{-2, -1, 0, 1, 2\}$ satisfies $4\sum D^2 = 784 - M$ and $\sum_R \sigma D = T_7/8 = -112 + M/2 - 2k_7$ on the 64-point affine support R of $\widehat{q_7}$.

Proposition 3.5 (Closed-Form Exception Bound `eminClosed`)**.** *On a* 7*-variable rank-*6 *section with* 64*-point support* $R = \operatorname{supp}(\widehat{q_7})$*, define the six multiplicity buckets*

$$p_j = \#\{u \in R : \sigma(u)D(u) = +j\}, \quad n_j = \#\{u \in R : \sigma(u)D(u) = -j\}, \quad m_j = \#\{u \notin R : |D(u)| = j\}$$

for $j \in \{1, 2\}$*, and define the section exception weight* $E := p_1 + 2p_2 + m_2$*. With* $\tau = -112 + M/2 - 2k_7$ *and* $s = (784 - M)/4$*, one has*

$$E \;\geq\; \max\left(0,\; \tau,\; \left\lfloor \frac{s + 2\tau - 61}{4} \right\rfloor\right).$$

Proof. By definition, $m_1 + m_2 \leq |R^c| = 64$, $\tau = p_1 + 2p_2 - n_1 - 2n_2$, and $s = \sum_u D(u)^2 = p_1 + 4p_2 + n_1 + 4n_2 + m_1 + 4m_2$. Immediately $E = p_1 + 2p_2 + m_2 \geq p_1 + 2p_2 - n_1 - 2n_2 = \tau$. Moreover,

$$s + 2\tau \;=\; 3p_1 + 8p_2 - n_1 + m_1 + 4m_2 \;\leq\; 4(p_1 + 2p_2 + m_2) + m_1 \;\leq\; 4E + 64,$$

so $4E \geq s + 2\tau - 64$. Since E is an integer, $E \geq \lceil (s + 2\tau - 64)/4 \rceil = \lfloor (s + 2\tau - 61)/4 \rfloor$. □

4 All-Degree Shifted Krawtchouk LP, Barycenter Puncturing, and Coding Bounds

To refute candidate profiles whose admissible hyperplane splits survive Proposition 3.5 and anchor-pair peeling, we develop three completely general combinatorial engines (all formalized in zero-Mathlib Lean 4 in `Balanced242Krawtchouk` and `Balanced242Centered`).

4.1 All-Degree Shifted Multivariate Krawtchouk Inequalities

For a subset $P \subseteq \mathbb{F}_2^n$ of size s and a direction $w \in \mathbb{F}_2^n$, let $h_P(w) = \#\{x \in P : \langle w, x\rangle = 1\}$ be its hyperplane half-count, and let

$$K_i^{(s)}(x) \;=\; \sum_{r=0}^{i} (-1)^r \binom{x}{r}\binom{s-x}{i-r}$$

be the binary Krawtchouk polynomial of degree i. Note that $K_i^{(s)}(h_P(w)) = e_i\big(\{\chi_w(x) : x \in P\}\big)$ is the i-th elementary symmetric polynomial of the characters $\{\chi_w(x)\}_{x\in P}$.

Theorem 4.1 (Shifted All-Degree Multivariate Krawtchouk Theorem)**.** *Let $P, Q, R \subseteq \mathbb{F}_2^n$ have sizes a, b, c, and fix any shift $t \in \mathbb{F}_2^n$ and integers $i, j, \ell \geq 0$. Set*

$$S_{i,j,\ell}(t) \;:=\; \sum_{w\in\mathbb{F}_2^n\setminus\{0\}} \chi_w(t)\, K_i^{(a)}(h_P(w))\, K_j^{(b)}(h_Q(w))\, K_\ell^{(c)}(h_R(w)), \qquad \mathrm{dc}_{i,j,\ell} \;:=\; \binom{a}{i}\binom{b}{j}\binom{c}{\ell}.$$

Then:

1. $S_{i,j,\ell}(t) + \mathrm{dc}_{i,j,\ell} = 2^n \cdot \#\big\{(S_1,S_2,S_3) \in \binom{P}{i}\times\binom{Q}{j}\times\binom{R}{\ell} : t \oplus \bigoplus S_1 \oplus \bigoplus S_2 \oplus \bigoplus S_3 = 0\big\}$.
2. *Consequently,* $S_{i,j,\ell}(t) \geq -\mathrm{dc}_{i,j,\ell}$ *and* $S_{i,j,\ell}(t) + \mathrm{dc}_{i,j,\ell} \equiv 0 \pmod{2^n}$.
3. *(Degree-1 equality) If $i+j+\ell = 1$ and $t \notin P\cup Q\cup R$, then $S_{i,j,\ell}(t) = -\mathrm{dc}_{i,j,\ell}$.*
4. *(Unshifted degree-2 equality) If $t = 0$, $i+j+\ell = 2$, and P, Q, R are pairwise disjoint, then $S_{i,j,\ell}(0) = -\mathrm{dc}_{i,j,\ell}$.*

Proof. Expanding each $K_i^{(a)}(h_P(w)) = \sum_{S_1\in\binom{P}{i}} \chi_w(\bigoplus S_1)$ and summing over all $w \in \mathbb{F}_2^n$ (including $w = 0$, which contributes $\chi_0(t)\mathrm{dc}_{i,j,\ell} = \mathrm{dc}_{i,j,\ell}$) gives

$$S_{i,j,\ell}(t) + \mathrm{dc}_{i,j,\ell} = \sum_{S_1,S_2,S_3} \sum_{w\in\mathbb{F}_2^n} \chi_w\Big(t \oplus \bigoplus S_1 \oplus \bigoplus S_2 \oplus \bigoplus S_3\Big),$$

and character orthogonality evaluates the inner sum to 2^n if the XOR is 0 and 0 otherwise. Parts (2)–(4) follow immediately: when $i+j+\ell = 1$ and $t \notin P\cup Q\cup R$, no single element equals t; when $t = 0$, $i+j+\ell = 2$, and P, Q, R are pairwise disjoint, no pair of distinct elements has XOR equal to 0. □

4.2 $\mathbb{F}_2$-Barycenter Translation and Marker-Center Puncturing

Unshifted Krawtchouk inequalities are invariant under global translation $x \mapsto x \oplus s$. However, when a partition $P_0 \sqcup P_1 \sqcup P_2 \sqcup P_3 = \mathbb{F}_2^8$ has at least one *odd-sized* class P_{i_a}, we can break translation symmetry canonically using the $\mathbb{F}_2$-barycenter

$$\mathrm{bar}(P) \;:=\; \bigoplus_{x\in P} x \;\in\; \mathbb{F}_2^n.$$

Theorem 4.2 (Barycenter Parity and Marker-Center Laws)**.** *Let $P, Q \subseteq \mathbb{F}_2^n$ with $|P|$ odd.*

1. *For every $w \in \mathbb{F}_2^n$, $h_P(w) \equiv \langle w, \mathrm{bar}(P)\rangle \pmod 2$.*
2. *Translating by $s = \mathrm{bar}(P)$ produces $P' = P \oplus s$ satisfying $h_{P'}(w) \equiv 0 \pmod 2$ for* ***every*** *$w \in \mathbb{F}_2^n$.*
3. *Let $Q' = Q \oplus s$ and set the marker $q_b = \mathrm{bar}(Q')$. Then for every $w \in \mathbb{F}_2^n$, the parity bit $\mathrm{cut}(w) := h_{Q'}(w) \bmod 2$ equals $\langle w, q_b\rangle$, so $\chi_w(q_b) = 1 - 2\,\mathrm{cut}(w)$.*
4. *If $q_b \neq 0$, then $\sum_{w \in \mathbb{F}_2^n \setminus \{0\}} \mathrm{cut}(w) = 2^{n-1}$.*

Proof. Part (1) follows from $\langle w, \bigoplus_{x\in P} x\rangle = \bigoplus_{x\in P}\langle w, x\rangle = h_P(w) \bmod 2$. Under translation by s, $h_{P'}(w) = h_P(w)$ if $\langle w, s\rangle = 0$ and $|P| - h_P(w)$ if $\langle w, s\rangle = 1$. Since $|P|$ is odd and $h_P(w) \equiv \langle w, s\rangle \pmod 2$, both cases yield $h_{P'}(w)$ even, proving (2). Part (3) applies (1) to Q', and (4) follows from $\sum_{w\neq 0}(1 - 2\langle w, q_b\rangle) = \sum_w \chi_w(q_b) - 1 = -1$. □

Now puncture $\{0, q_b\}$ from whichever classes contain them (splitting over at most 16 non-coinciding cases $q_b \neq 0$ and 4 coinciding cases $q_b = 0$). The punctured classes (Q_0, Q_1, Q_2) are pairwise disjoint and avoid both 0 and q_b, so Theorem 4.1 applies simultaneously at $t = 0$ (unshifted rows) and at $t = q_b$ (shifted rows multiplied by $1 - 2\,\mathrm{cut}(w)$), with degree-1 equalities holding for *both* $t = 0$ and $t = q_b$!

4.3 Residual Kernel Coupling, Griesmer Bound, and $[24, 8, \geq 10]_2$ Non-Existence

For origin-pinned profiles with $a \geq 1$, translating by a point $x_0 \in A$ and puncturing 0 from A' reduces the sizes to $(a-1, b, c)$ with $0 \notin A' \cup B' \cup C'$. When $a = 1$ (such as $(1, 54, 24)$ at $t = 96$), $A' = \emptyset$ and any nonzero direction $u \in \mathbb{F}_2^8 \setminus \{0\}$ allows us to restrict (B', C') to the 7-dimensional kernel $\ker(u) = \{y : \langle u, y\rangle = 0\}$.

Lemma 4.3 (Residual Kernel Coupling)**.** *For any $P \subseteq \mathbb{F}_2^n$ and $u, v \in \mathbb{F}_2^n$, let $P_{\ker(u)} = P \cap \ker(u)$. Then*

$$2\,h_{P_{\ker(u)}}(v) + h_P(u) \;=\; h_P(v) + h_P(v \oplus u), \qquad \big|h_P(v) - h_P(v \oplus u)\big| \;\leq\; h_P(u).$$

Proof. Check the indicator identity $2\,\mathbf{1}_{P\cap\ker(u)}(y)\langle v, y\rangle + \mathbf{1}_P(y)\langle u, y\rangle = \mathbf{1}_P(y)(\langle v, y\rangle + \langle v \oplus u, y\rangle)$ and $|\langle v, y\rangle - \langle v \oplus u, y\rangle| \leq \langle u, y\rangle$ pointwise on all 8 truth values of $(\mathbf{1}_P(y), \langle u, y\rangle, \langle v, y\rangle)$, then sum over $y \in \mathbb{F}_2^n$. □

Treating the support of C' as the columns of an $8 \times c$ generator matrix of a binary linear code (so codeword weights are $h_{C'}(v)$), we invoke two classical coding bounds—both proved constructively from scratch in `Balanced242Centered`:

Theorem 4.4 (Binary Griesmer Bound, Sphere-Packing Bound, and Non-Existence of $[24, 8, \geq 10]_2$)**.** *For any list of columns $\mathcal{C} \subset \mathbb{F}_2^k$ with minimum nonzero codeword weight $\geq d$:*

1. *(Griesmer bound) $|\mathcal{C}| \geq g(k, d) := \sum_{i=0}^{k-1} \lceil d/2^i \rceil$.*
2. *(Hamming bound for $d \geq 3$) If $d \geq 3$, then $2^k(1 + |\mathcal{C}|) \leq 2^{|\mathcal{C}|}$.*
3. *(Corollary for $k = 8, |\mathcal{C}| = 24$) No binary linear code of length 24, dimension 8, and minimum distance ≥ 10 exists.*

Proof. (1) Induct on k: choose $u \neq 0$ achieving minimum weight $d_0 \geq d$, and project $\mathcal{C} \cap \ker(u)$ to $\mathbb{F}_2^{k-1}$ by dropping a coordinate where $u_i = 1$. By Lemma 4.3, every nonzero codeword of the residual code has weight $\geq \lceil d_0/2\rceil$, giving $|\mathcal{C}| = d_0 + |\mathcal{C}_{\mathrm{res}}| \geq d_0 + g(k-1, \lceil d_0/2\rceil) = g(k, d_0) \geq g(k, d)$. (2) The Hamming balls of radius 1 around the 2^k distinct codewords in $\mathbb{F}_2^{|\mathcal{C}|}$ are pairwise disjoint by the triangle inequality, giving $2^k(1 + |\mathcal{C}|) \leq 2^{|\mathcal{C}|}$. (3) Suppose $\mathcal{C} \subset \mathbb{F}_2^8$ has $|\mathcal{C}| = 24$ and minimum

distance ≥ 10. Since $g(8, 11) = 26 > 24$, there must exist a word of weight exactly 10. Its residual is a $[14, 7, \geq 5]_2$ code. Since $g(7, 6) = 15 > 14$, that residual must contain a word of weight exactly 5. Its second residual is a $[9, 6, \geq 3]_2$ code, contradicting (2) since $2^6(1 + 9) = 640 > 512 = 2^9$. □

5 Unconditional Structural Exclusions at $t = 96$ and $t = 98$

We highlight two frontier arguments, at $t = 96$ and $t = 98$, whose algebraic obstructions are verified in Lean 4 without `SectionEnum`.

5.1 Impossibility of $t = 98$, Global Rank 6, $|Z| = 144$

When $t = 98$ and $\mathrm{rank}(\beta_8) = 6$, the high-block phase transform $\widehat{q_8}$ is supported on a 64-point affine 6-flat R with magnitude 32, and U is supported on $256-98 = 158$ sites. Let $Z = \{a \in \mathbb{F}_2^8 : \mathrm{lo}(a) \equiv 0 \pmod{16}\}$; in the case $|Z| = 144$, the agreement table V and disagreement table U interact on the four 6-flat cosets of R.

Theorem 5.1 (Unconditional Exclusion of $t = 98$, Rank 6, $|Z| = 144$)**.** *No balanced 9-variable Boolean function of nonlinearity 242 can have adapted half-agreement count $t = 98$, global phase rank $\mathrm{rank}(\beta_8) = 6$, and $|Z| = 144$.*

Proof outline. In coordinates $(b_0, b_1, c) \in \mathbb{F}_2^2 \times \mathbb{F}_2^6$, the normal form has $\widehat{V}(b_0, b_1, c) \equiv 4Q_6(c) \pmod 8$ for a weight-28 bent function Q_6. The Lean lemma `no_t98_r6_z144` takes this residue pattern as a hypothesis. In the coordinate-free master theorem, `t98_remaining` covers all six surviving values of $|Z|$, including 144. Write $K = R^\perp$ ($\dim K = 2$) and index the four cosets of R in $\mathbb{F}_2^8$ by $x \in \mathbb{F}_2^2$.

1. From $\mathrm{lo} = 2\widehat{V}$ and $|Z| = 144$, Parseval and divisibility on the four cosets determine the signed fiber counts of V and U across R, yielding four 6-variable indicators $p_0, p_1, p_2, p_3 : \mathbb{F}_2^6 \to \mathbb{F}_2$ whose supports partition a 28-point subset $C \subset \mathbb{F}_2^6$ into fibers F_0, F_1, F_2, F_3 of sizes $a_0 + a_1 + a_2 + a_3 = 28$.
2. For every pair $i \neq j$, the pairwise XOR $p_i \oplus p_j = \mathbf{1}_{F_i \sqcup F_j}$ has third-difference Walsh transform divisible by 4, hence is a quadratic Boolean function in $\mathrm{RM}(2, 6)$.
3. By the exact weight classification of $\mathrm{RM}(2, 6)$ (`rm2_6_weight`), every quadratic Boolean function on $\mathbb{F}_2^6$ has Hamming weight in $\{0, 16, 24, 28, 32, 36, 40, 48, 64\}$.
4. Since all six pairwise sums $a_i + a_j \leq 28$ must lie in $\{0, 16, 24, 28\}$ while $(a_0 + a_1) + (a_2 + a_3) = 28$, elementary case analysis forces three of the a_i to be 0 and one to be 28!
5. Consequently $\mathbf{1}_C = p_{x^*} \in \mathrm{RM}(2, 6)$ is itself a weight-28 quadratic function on $\mathbb{F}_2^6$. Every weight-28 function in $\mathrm{RM}(2, 6)$ is bent (up to affine equivalence), forcing its Walsh transform to take only values ± 8. Matching this against the forced $\pm 4, \pm 12$ Walsh spectrum on the remaining cosets yields an immediate contradiction. □

Remark 5.2 (Subtlety caught during formalization)**.** An informal draft of Step 3 attempted to isolate a single fiber indicator $\mathbf{1}_{F_x}$ via linear combinations of $p_0, \ldots, p_3$. During Lean 4 formalization, it emerged that only the *pairwise* XORs $p_i \oplus p_j = \mathbf{1}_{F_i \sqcup F_j}$ are directly guaranteed to lie in $\mathrm{RM}(2, 6)$. The arithmetic observation in Step 4 (using $(a_0 + a_1) + (a_2 + a_3) = 28$ and $\{0, 16, 24, 28\}$) closes the proof rigorously without needing individual $p_x \in \mathrm{RM}(2, 6)$ a priori.

5.2 Exact Walsh mod 16 Energy Dichotomy for $t = 96$, $(0, 31, 39)$

Theorem 5.3 (Spectral Dichotomy for $t = 96$, Rank 8, Profile $(0, 31, 39)$)**.** *Let (f, g) be an adapted pair with $t = 96$ and global profile $(a, b, c) = (0, 31, 39)$, and define $f_{BC} = \mathbf{1}_B - \mathbf{1}_C : \mathbb{F}_2^8 \to \{-1, 0, 1\}$. Then:*

1. *For every* $w \in \mathbb{F}_2^8$, $\widehat{f_{BC}}(w) \equiv 8 \pmod{16}$, $\widehat{f_{BC}}(0) = -8$, *and* $\sum_{w \in \mathbb{F}_2^8} (\widehat{f_{BC}}(w)^2 - 64) = 1536$.
2. $|\widehat{f_{BC}}(w)| \le 40$ *for all* $w \in \mathbb{F}_2^8$, *and exactly one of the following holds:*
   - *(Single-*40 *branch) There is a unique* $w \neq 0$ *with* $|\widehat{f_{BC}}(w)| = 40$ *(and* $|\widehat{f_{BC}}(v)| = 8$ *for all* $v \neq w$*), and one side of* w *automatically has* $c \le 16$ *and* $|2(b-c)+8| = 40$.
   - *(Three-*24 *branch)* $|\widehat{f_{BC}}(w)| \le 24$ *for all* $w \in \mathbb{F}_2^8$ *(with exactly three* w *attaining* 24*), and Parseval on* $\mathbf{1}_C$ *alone forces a nonzero hyperplane with a side having* $c \le 16$ *and* $|2(b-c)+8| \le 24$.

Proof. For every hyperplane $w \neq 0$, the section mass $6a' + 3b' + 5c' = 3b' + 5c' \equiv 0 \pmod 8$ implies $3(b'-c') \equiv 0 \pmod 8$, hence $b' \equiv c' \pmod 8$. Thus $\widehat{f_{BC}}(w) = (31-2b') - (39-2c') = -8 - 2(b'-c') \equiv 8 \pmod{16}$. By Parseval, $\sum_w \widehat{f_{BC}}(w)^2 = 256(|B| + |C|) = 256 \cdot 70$, so $\sum_w (\widehat{f_{BC}}(w)^2 - 64) = 256 \cdot 6 = 1536$. Since every term is ≥ 0 and the possible excess values $W^2 - 64$ for $|W| \in \{8, 24, 40, 56, \dots\}$ are $0, 512, 1536, 3072, \dots$, no $|W| \ge 56$ can occur, and either a single 40 consumes the entire excess 1536 or three 24s do. In the single-40 branch, $b' - c' \in \{16, -24\}$ with $0 \le b' \le 31$ forces $c' \le 15$ or $39 - c' \le 15$. In the cap-24 branch, Parseval on $\mathbf{1}_C$ gives $\sum_{w \neq 0} \widehat{\mathbf{1}_C}(w)^2 = 256 \cdot 39 - 39^2 = 8463 > 255 \cdot 25$, forcing some $w \neq 0$ with $|39 - 2c'| \ge 7$, i.e., $\min(c', 39-c') \le 16$. □

6 Certificate-Driven Exclusions: $t \le 94$ and $t = 96$ Frontier

Combining Sections 3 and 4 yields a fully mechanized refutation pipeline inside Lean 4:

Theorem 6.1 (Exclusion of All $t \le 94$ and Autocorrelation Lower Bound $\Delta_F \ge -128$)**.** *Assume* `SectionEnumE excl6MaskT94 eBoundsT94`*, which records* 252 *distinct excluded rank-*6 *section profiles and* 59 *exception lower bounds.*

1. *All* 17 *remaining rank-*8 *frontier profiles at* $t = 94$ *are excluded in the Lean 4 kernel:* 15 *directly via sparse all-degree Krawtchouk Farkas certificates (*`Balanced242T94Direct`*) and the remaining* 2*—*$(12, 25, 25)$ *(*108 *anchors) and* $(15, 24, 22)$ *(*97 *anchors)—via* 1*-layer anchor-pair peeling (*`Balanced242T94Layered`*).*
2. *Assume also* `SectionEnum` *for* $t \le 92$ *and* `ComputationalClosures.t94_nonfrontier` *for the other* $t = 94$ *profiles. Then any balanced* F *of nonlinearity* 242 *has* $t \ge 96$ *on its adapted half-split. In its canonical direction* a*,* $2t \ge 192$ *inputs satisfy* $F(x) = F(x+a)$*, and* $\Delta_F(a) = 4t - 512 \ge -128$.

Theorem 6.2 (Kernel-Verified $t = 96$ Rank-8 Profile Exclusions)**.** *The following* $t = 96$ *rank-*8 *profiles are formally excluded in the Lean 4 kernel:*

1. $(96, 21, 34, 12, 8)$ *via direct sparse Krawtchouk Farkas (*`Balanced242T94.t96_21_34_12_excluded`*).*
2. $(96, 2, 47, 27, 8)$ *via* 10 *layer-*0 *anchor-pair certificates and a* 10*-term origin-pinned centered Krawtchouk Farkas certificate (*`Balanced242T96Centered.excl_96_2_47_27`*), modulo* `SectionEnumE excl6MaskT96 eBoundsT96` *and non-occurrence of the single flag leaf* $(0, 26, 10)$.
3. $(96, 1, 54, 24, 8)$ *via* 20 *layer-*0 *anchor-pair certificates, two* 8*-term* 7*-dimensional residual Krawtchouk certificates for* $(0, 46, 6)$ *and* $(0, 48, 8)$*, and Theorem 4.4(3) (*`Balanced242T96Centered.excl_96_1_54_24`*), modulo* `SectionEnumE excl6MaskT96 eBoundsT96` *and non-occurrence of the three flag leaves* $(0, 44, 4), (0, 38, 6), (0, 40, 8)$.
4. $(96, 0, 51, 27, 8)$ *(the second* $b + c = 78$ *profile at* $t = 96$*) via the unconditional mod-*4 *barycenter parity law (*`Balanced242T96Quartic.rank8_a0_far_yz_div4`*) and three tiny integer Farkas certificates of* $6, 5, 5$ *terms (*`Balanced242T96_0_51_27`*), using zero flag leaves and zero anchor peeling.*

6.1 Complete Closure of $t = 96$, $t = 98$, $t = 100$, $t = 102$, and $t \geq 104$

To close the remaining agreement window $t \in [96, 128]$, two complementary verification engines meet cleanly at the midpoint seam $t = 98 \mid t = 100$:

1. Bottom-up high-block quartic closure ($t = 96$ and $t = 98$). The rank-8 high-block inversion $2KD = 28 - 3S + 32u$, ternary counts $2p = 3t - 208 - 2(a + c)$, first-moment identity $\langle u, S \rangle = 10(b+c) - 16c$, three hinge lower bounds $(h, l) \in \{(14, 6), (18, 6), (18, 10)\}$, and exact integer dual-cut soundness are all formalized unconditionally in `Balanced242QuarticHighBlock` (using only `[propext, Quot.sound]`).

- $t = 96$**:** Packaged separately in `nl242_t96_closure.zip`. All 55 rank-8 profiles through $b + c = 78$ are excluded by the four-cell quartic enumerator (`verify_t96_quartic.py` with `hinge_t.cpp`), with 0 survivors across every weight $46, 50, \ldots, 78$; the last weight alone checks 106,366,593 configurations. The run reads the SHA-256-checked 6-dimensional core tables `shapes{8,10,12,14,16}.bin` and `radial_cuts.txt` from the $t = 98$ packet. An independent baseline enumerator (`quartic_sections_t.cpp`) agrees on weights 46–74. The centered closures of Theorem 6.2 (Layer B) give a second route for four of the profiles. No rank-6 or rank-≤ 4 cases remain at $t = 96$.
- $t = 98$**:** Packaged in `nl242_t98_closure.zip` with full SHA-256 verification. After v2 certificate replay (22 rank-6 and 102 rank-8 profiles), all 22 rank-6 profiles are excluded (including the normal-form obstruction of Theorem 5.1) and 18 rank-8 profiles close via frontier/anchor/pin/centered obstructions. The remaining 84 rank-8 quartic profiles split into 70 at $b + c \leq 74$, 10 at $b + c = 78$ (945,849,169 four-cell configurations; 9,245 passing the three hinges), and 4 at $b + c = 82$ (19,404,895,597 four-cell configurations; 2,006 passing the three hinges). Every single surviving support/profile pair (92,450 at weight 78 and 8,024 at weight 82) is strictly refuted in exact arbitrary-precision integer arithmetic (`direct.py`), all 59 AVX-512 jobs agree byte-for-byte with 36 scalar reference jobs (`coverage.py`), and an independent 2^{21}-character Fourier mass calculation (`orbit_mass.py`) confirms 100% orbit completeness of the affine-symplectic 6-dimensional cores at every size $m \in \{8, 10, 12, 14, 16, 18\}$ $(1, 1, 3, 7, 23, 65$ inequivalent cores matching 2,304 through 1,717,438,464 labeled admissible sets).

2. Top-down low-block closure ($t = 100$, $t = 102$, and $t \geq 104$). Packaged in `nl242_top_closure.zip` (104 files, all paths relative, SHA-256 manifest), together with `nl242_middle_closure.zip` for the middle window ($t = 104$–108). The universal low-block defect budget $4(n_0 + 2n_8) = 2304 - 5|Z| - 16t$, autocorrelation congruence mod 32, and all upper-window budget exclusions are formalized unconditionally in `Balanced242LowBlock`.

- $t \geq 104$**:** For $t \geq 114$, `Balanced242LowBlock.t_ge_114_forces_Z64` forces $|Z| = 64$ (rank 2). $t = 118$ is unconditionally closed in Lean (`Balanced242T118.no_t118`, low-block only, via kernel-verified RREF affine-subspace cuts and empty-cell integer rounding, formalizing `check_t118.py`; also confirmed by `verify_spread_t118_all127.c`), and $t \geq 120$ is unconditionally closed in `Balanced242Endpoint.low_agreements_below_120`. At $t = 114$, the $[12, 9, 4]$ Hamming sphere-packing contradiction ($512 \cdot 13 = 6{,}656 > 4{,}096$) closes $t = 114$ *unconditionally* in Lean (`Balanced242T114.no_t114`), and the same argument unconditionally excludes $k = 12$ at $t = 116$ (`Balanced242T114.t116_forces_k16`), shrinking the entire endpoint residual to the single case $t = 116, k = 16$ (closed directly in Section 6.2 and independently via $\rho(1, 9) < 244$ in Section 6.3). For $t \leq 112$, rank 2 ($|Z| = 64$) is unconditionally impossible in Lean (`no_Z64_of_t_le_112`). At $t = 112$, all ranks are unconditionally closed in Lean (`Balanced242T112.no_t112`,

pure algebra + 729-cell butterfly FWHT check, zero LP, zero covering radius; independently confirmed by `verify_t112_closed.py`); $t = 110$ is closed at ranks $0, 2, 6, 8$ unconditionally in Lean (with $n_0 \geq 8$ proved for rank 4 $|Z| = 96$ in `Balanced242T112.rank4_Z96_n0_ge8`) and at rank 4 by pure-integer Farkas certificates (`verify_t110_closed.py`); $t = 108$ is excluded at every rank in `nl242_middle_closure.zip` (`MIDDLE_PROOF.md`). Rank 8 is unconditionally closed for all $t \geq 104$ in Lean (`Balanced242T104R8.no_rank8_of_t_ge_104`). Rank 6 ($|Z| = 112$) fails the four-cell energy bound $36 \cdot 12^2 = 5{,}184 > 4{,}096$; it is also closed matrix-free in Lean (`Balanced242T108R6`) via quadratic weight gaps on $\mathbb{F}_2^8$ forcing $144^2 = 20{,}736 > 16{,}384$. Rank 4 leaves 9 profiles after the two high-block tests, and exhaustive reconstruction with 12, 14 and 16 small masks (`reconstruct_108.cpp`, `transversal_108.py`, `transversal_direct.cpp`) finds no high-block completion. The same packet excludes ranks 4 and 8 at $t = 104, 106$ and rank 6 at $t = 106$ (`RANK6_106.md`); rank 8 at $t = 104, 106$ is unconditionally closed in Lean (`Balanced242T104R8.no_rank8_of_t_ge_104`, `Balanced242LowBlock.no_t106_rank8`), rank 6 at $t = 106$ in `verify_t106_closed.py`, and rank $4, 6$ at $t = 104$ in `verify_t104_closed.{cpp,py}`.

- $t = 102$**:** Rank ≤ 4 has 0 open systems; rank 8 ($|Z| = 120, n_0 + 2n_8 = 18$) is completely closed from scratch by the standalone exact-integer verifier `verify_t102_r8_closed.cpp` (identifying QR(17) uniquely at $n_0 = 18$ and killing all 512 high-block phases); rank 6 $|Z| = 128$ is closed by `verify_t102_z128_closed.py`; and rank 6 $|Z| = 112$ is closed across all 2,984 orbits via Mod-16 Radical Separation plus the Mod-3 Coupled Cell Bound on the sole surviving 512-element low-block family (`t102/z112/`).
- $t = 100$**:** Rank ≤ 4 is closed by hand in `t100.tex` (Section 3); rank 6 $|Z| = 144$ has negative defect budget (-4); rank 6 $|Z| = 128$ ($|D| = 16$) and $|Z| = 112$ ($|D| = 36$) are completely reduced and closed in the companion proof `t100.tex` (Section 5, closing Branches II-B2, II-B3 and $(h, p) = (0, 0)$ by hand and proving the Branch I weight-16 $64\times$ normalization and 22-bit degree-4–6 subcube syndrome in Lemma 5.2), backed inside `nl242_top_closure.zip` (and mirrored in `t100_r6_audit/`) by the 10 branch solvers in `t100/z128/` plus the three standalone audit verifiers `t100/z128/verify_z128_reductions.cpp` (0.21 s, all 288 local identities and all orbit lists regenerated), `t100/z112/check_small_masks.cpp` (8 s, $9 \leq k \leq 15$), and `t100/z112/verify_z112_all45_exact.cpp` (24.9 s, all 45 pairs (h, p), 111/111 live cases killed, 7/7 positive controls passing, closing three omissions found in the legacy `t100/z112/` scripts); rank 8 $|Z| = 136$ forces $n_0 \leq 6 < 10$. For the flagship rank-8 $|Z| = 120$ case ($n_0 + 2n_8 = 26$, Q_8^- self-dual):
  1. *Exact rational H_{64} $w \geq 6$:* Since $\mathrm{rank}(\mathrm{RM}(3,7)|_{Z_{64}}) = 56$ (quotient_dim $= 0$), H_{64} $w = 2$ has 0/2,016 valid pairs and H_{64} $w = 4$ has 5,376 valid 4-sets in two $\mathrm{Stab}_{O^-(8,2)}(e_7)$-orbits (3,456 and 1,920), both certified in exact `fractions.Fraction` LP dual arithmetic (den ≤ 128) to force $n_8^{(1)} \geq 8 > 7$ (`verify_h64_w_ge_6_exact.py`). No H_{56} orbit is ever excluded by single-g LPs because quotient_dim$(H_{56}) = 8$.
  2. *Exact rational Signed (Q^*, v)-Weighted Delsarte Farkas certificates:* Using only H_{64} $w \geq 6$ and $w(y) \equiv 2(Q^*(y) \oplus \langle y, v \rangle) \pmod 4$, all 11/11 required d_{aff} exclusions across $n_0 \in \{18, 20, 22, 24, 26\}$ are certified in exact `Fraction` arithmetic (`verify_delsarte_farkas_exact.py`, passing the QR(17) feasibility positive control).
  3. *Small $n_0 \leq 16$ and $n_0 = 18$ ($v = 0$):* $n_0 \leq 12$ is closed algebraically / via `verify_t102_r8_closed.cpp`; all 32/32 codes at $n_0 = 14$ and all 4,124/4,124 unique RREF reps at $n_0 = 16$ are individually killed by exact integer dual witnesses with margin $\geq +160$ (`audit_n14_n16_exact.cpp`); at $n_0 = 18, v = 0$, the unique code QR(17) has all 53,550 XOR-0 4-subsets $S_8 \subset F$ eliminated (`verify_qr17_t100_standalone.cpp`, $t = 102$ positive control recovering the real low block V).

4. *Exact $O^-(8,2)$ orbit reduction + rational $L_1 + L_\infty$ duals ($n_0 = 18$ ($v \neq 0$) and $n_0 = 20$):* Explicit 256-point verified $O^-(8,2)$ isomorphism matrices reduce all 912 codes at $n_0 = 18$ ($v \neq 0$) to 5 exact orbits and all 109,846 codes at $n_0 = 20$ to 342 exact orbits (`audit_n18_n20_exact_orbits.cpp`); all 347/347 exact orbits are certified dead in exact `Fraction` arithmetic (224 plain-infeasible Farkas + 123 L_1 duals with $\min L_1 \geq 6{,}560 > 6{,}400$, `certify_exact_o8_orbs_L1.py`, passing the QR(17) positive control $\min L_1 = 6{,}528 = 64 \cdot 102$).
5. *Zero-LP exact* $\mathrm{RM}(3,8) \mod 8$ *y_Z enumeration ($n_0 = 22, 24, 26$):* Since $\mathrm{rank}(\mathrm{RM}(3,8)|_Z) = 84$, Poisson summation on the 93 monomials of $\mathrm{RM}(3,8)$ collapses the sign vector y_Z on $Z \setminus Z_0$ to $\leq 2^{15}$ exact candidates. At $n_0 = 26$ ($n_8 = 0$), all 1,413,762 valid codes across H_{64} $w = 8$ (62 exact orbits), H_{56} $w = 6$ (10 orbits), and H_{56} $w = 2$ yield `REAL_V = 0` (QR(17) positive control recovers 12/12 and 4/4 real V's). At $n_0 = 24$ ($n_8 = 1, Q(v) = 1$), all 12,318,368 valid codes yield `REAL_V = 0` (`solve_n24_exact.cpp`). At $n_0 = 22$ ($n_8 = 2, v \neq 0$), all 1,267,047 valid codes reduce via explicit 256-point verified $O^-(8,2)$ isomorphisms to 8,655 exact orbits, and exact testing of all 36 pairs $S_8 = \{f_1, f_1 \oplus v\} \subset F$ per orbit (311,580 pairs) plus $\leq 2^{15}$ y_Z enumeration yields `REAL_V = 0` (`solve_n22_o8_exact.cpp`).

6.2 Direct exclusion of the switched endpoints

At $t = 114, 116$, the plane switch gives complementary ternary tables u, v on $\mathbb{F}_2^8$ with $|W_u|, |W_v| \in \{4, 12\}$. For $E = \{a : |W_u(a)| = 4\}$, put $k = |E|$. At $t = 114$ one has $k = 12$. At $t = 116$ one has $k = 12$ or 16. The Lean theorem `reduce_to_endpoint_pair` gives the reduction, including the Walsh levels and counts.

Every affine hyperplane contains at least four members of E by the halfspace count applied to u and its character multiples. Moreover, if $m = |\{a \in E : a \cdot d = 1\}|$ for $d \neq 0$, then $\sum_x u(x)u(x+d) = m - k/2$ by inversion of W_u^2. Pairing $x, x+d$ shows that m is even. Restricting affine functions to E therefore gives an even $[k, 9, \geq 4]$ code C containing the all-one word. At $k = 12$, disjoint radius-one balls would occupy $512 \cdot 13 > 2^{12}$ points.

At $k = 16$, choose a generator of $C^\perp$ with first row all ones. Its columns $(1, x)$ give sixteen distinct spanning points in $\mathbb{F}_2^6$ with vector sum zero. Enumeration up to affine equivalence gives 89 geometries. Their orbit sizes sum to 7,630,875,426,048, also obtained by counting zero-sum sets by affine span. In 64 geometries the dual has minimum distance at least four, as required by distinctness of the original frequencies. For each of their 128 affine sign cosets, Walsh inversion modulo three determines u from the signed small frequencies. Checking the full high spectrum leaves five sign cosets on two geometries. An explicit affine map identifies the three-signing geometry with the Gold graph $\{(x, x^3) : x \in \mathbb{F}_{16}\}$. The other exceptional set contains eight affine planes, whereas the graph contains none.

Let $L = \{u = 0\}$, of size 120. For $h = u, v$, write $p_h = (W_h - 4)/8 \bmod 2$. The cubic coefficients of p_h are the third incidence moments of $\mathrm{supp}(h)$ modulo two, and all higher coefficients vanish. Complementarity gives $p_v = p_u + Q$ for some quadratic or lower-degree Q. Parseval gives 48 low coefficients of magnitude four. Hence $b = (-1)^{p_u + Q}$ must satisfy

$$\sum_{x \in L} v(x) W_b(x) = 4 \cdot 48 - 12 \cdot 208 = -2304.$$

For two of the five high sign cosets, the maximum of $\sum_{x \in L} |W_b(x)|$ over Q is 2272. For each of the other three, the maximum is 2304. Each of the 320 equality cases per coset, modulo frequency translations and overall sign, forces all 120 signs of v and gives Walsh peak 28, exceeding 12.

This excludes the final profile. The derivation, orbit-mass calculation and exact source replay are included in `nl242_endpoint_closure.zip`.

A cross-check packet, `nl242_endpoint_crosscheck.zip`, repeats the $k = 16$ calculation with separate code. It parametrises the geometry by balanced cubic functions on $\mathbb{F}_2^5$ with Walsh magnitudes in $\{0, 8, 16\}$. These fall into eight affine classes and give 70 admissible lifts. The packet keeps all 70 lifts instead of merging equivalent ones, and scores every Q by direct Walsh transforms. It does so over all 2^{28} quadratic parts for one lift of each class, and modulo frequency translations for the rest. It finds the same two classes and the same maxima 2272 and 2304. It also finds the same $320 \cdot 256 = 81{,}920$ equality cases, each with Walsh peak 28.

6.3 Independent Verification of $\rho(1, 9) < 244$

The plane switch at $t = 114, 116$ also produces an *unbalanced* function with all Walsh values in $\{\pm 8, \pm 24\}$. Thus the unrestricted exclusion at nonlinearity 244 gives a second proof of these cases. Wang, Lou, and Wang [10] give an argument for this in a recent preprint. We credit them for the argument, but we do not rely on their claims: we re-verified every step ourselves (`nl244_verify.zip`).

1. *Hand steps, re-checked.* An NL-244 function f has $W_f \in \{\pm 8, \pm 24\}$ (exactly 64 of magnitude 8), and $g = (W_f - 8)/16 \bmod 2$ is cubic with $\langle (-1)^f, W_g \rangle = -10240$. Every cubic has $\|W_g\|_1 \leq 10240$, with equality only in two cases, A and B. Case B is impossible. In case A ($W_g \in \{\pm 16, \pm 48\}$), $h = T(g) := ((W_g/16 - 1)/2) \bmod 2$ is cubic with $T(h) = g$, and the cubic part P of g satisfies $J(J(P)) = P$ and $C(P) = 0$.
2. *Classification, recomputed.* By Burnside's lemma over all 490 conjugacy classes of $\mathrm{GL}(9, 2)$, acting on cubic parts through the third compound matrix, $\mathrm{RM}(3, 9)/\mathrm{RM}(2, 9)$ has exactly 349 orbits. The same code reproduces the known counts $6, 12, 32$ for $n = 6, 7, 8$. GL-invariant colour refinement over the 511 nonzero vectors separates all 349 representatives listed by Brier and Langevin [1], so their list is complete. We use it only as a list of candidates.
3. *Reduction, recomputed.* Only the classes 282, 290 and 349 (the zero form) satisfy $J(J(P)) = P$ and $C(P) = 0$. The zero form has $\|W\|_1 \leq 8192 < 10240$.
4. *Classes* 282 *and* 290*, by exhaustive search.* For each class we enumerate all $g = P + q$ (q quadratic; affine terms only permute or negate the spectrum, and $T^2(g + \ell) = T^2(g) + \ell$) that satisfy the 9 linear necessary conditions of $W_g \equiv 16 \pmod{32}$: 2^{27} functions per class. Class 282 has 18,923,520 spectra inside $\{\pm 16, \pm 48\}$ and 5,996,544 functions with $T^2(g) = g$; class 290 has 22,118,400 and 4,473,344. In both classes, **no** function has both properties, so no NL-244 function exists. Among the $T^2(g) = g$ cases, $\max |W_g| \geq 80$, in agreement with [10, Table 1].

Hence $\rho(1, 9) < 244$; this is the hypothesis `CoveringRadius244` of the Lean development.

Theorem 6.3 (Maximum balanced nonlinearity in dimension 9)**.** *No balanced* 9*-variable Boolean function has nonlinearity* 242 *(equivalently, every balanced* 9*-variable Boolean function has Walsh peak strictly greater than* 28*). Consequently, the maximum nonlinearity of a balanced* 9*-variable Boolean function is*

$$\mathrm{nl}_{\mathrm{bal}}(9) = 240,$$

making $n = 9$ the smallest odd dimension where balance strictly lowers maximum nonlinearity ($240 < 242 \leq \rho(1, 9)$).

Proof structure and verification scope. The Lean 4 master theorem `Balanced242Master.no_balanced_9var_nl242_direct` (and equivalently `no_balanced_9var_nl242`) formally reduces the coordinate-free

non-existence claim to the explicit Layer C hypotheses `SectionHypotheses`, `ComputationalClosures`, and `EndpointK16Direct` (or `CoveringRadius244`). Every agreement count $t \in \{0, 2, \ldots, 256\}$ is then discharged: $t \leq 84$, $t \in \{86, 88, 90, 92\}$, $t = 94$, $t = 96$, and $t = 98$ by the Lean kernel modules together with `nl242_t96_closure.zip` and `nl242_t98_closure.zip`; $t = 100$ by the complete rank-8 and audited rank-6 proofs in `t100.tex`, `nl242_top_closure.zip`, and `t100_r6_audit/`; $t = 102$–110 by `nl242_top_closure.zip`, `nl242_middle_closure.zip`, and the Lean upper-window modules; $t = 112, 114, 118$, $t \geq 120$, and $t = 116, k = 12$ unconditionally inside Lean 4; and $t = 116, k = 16$ by the dual-verified direct endpoint closure in Section 6.2 (as well as independently via Section 6.3). Combined with the balanced degree-4 / semi-bent construction of nonlinearity 240 and the unbalanced 9-variable nonlinearity-242 witnesses of [6], this yields $\mathrm{nl}_{\mathrm{bal}}(9) = 240$. □

7 Formal Verification in Lean 4 and Trust Boundary

The Lean 4.15.0 formalization (zero external dependencies, zero Mathlib) covers the results in Layers A and B of Table 1, culminating in the coordinate-free top-level master theorems `Balanced242Master.no_balanced_9var_nl242_direct` and `Balanced242Master.no_balanced_9var_nl242`. It does not cover the large finite enumerations of Layer C, which are parameterized as explicit hypotheses in `Balanced242Master`. The public replay archive (`nl242_lean_colab.tar.gz`) contains all **79** compiled modules with zero `sorry`, zero `native_decide`, and zero custom axioms. Its exported theorems depend only on `[propext, Quot.sound]` or `[propext, Classical.choice, Quot.sound]`.

7.1 Three-Layer Trust Boundary

To make the verification scope completely transparent, we separate our artifacts into three distinct trust layers:

7.2 Theorem-to-Lean Cross-Reference

7.3 Computer-Assisted Proof and Reproducibility

This is a computer-assisted proof, and its correctness includes the correctness of our code. We therefore publish all of it. The Lean 4 source archive (`nl242_lean_colab.tar.gz`, all 79 modules) and the Google Colab notebook (`verify_nl242_colab.ipynb`) replay all 79 zero-Mathlib modules sequentially on a standard CPU runtime, with no Mathlib download. The finite enumerations ship as SHA-256-manifested replay packets with all paths relative: `nl242_t96_closure.zip` ($t = 96$), `nl242_t98_closure.zip` ($t = 98$), `nl242_middle_closure.zip` ($t = 104$–108), `nl242_top_closure.zip` ($t \geq 100$ and $t = 84$), `nl242_endpoint_closure.zip` (direct exclusion at $t = 114, 116$), `nl242_endpoint_crosscheck.zip` (separate re-derivation of the $k = 16$ endpoint calculation), and `nl244_verify.zip` ($\rho(1, 9) < 244$). Each packet's README gives the replay commands and the expected output, and each packet includes positive controls where an enumerator could otherwise pass vacuously. The heaviest steps were replayed independently from these packets on a fresh machine. All $t = 100$ rank-6 reductions are proved in full in `t100.tex` and verified inside `nl242_top_closure.zip` (`proofs/PROOF_t100_RANK6.md`, wired into `verify_quick.sh`).

Acknowledgments

The author thanks AI research and coding assistants—Antigravity (Google), Claude (Anthropic), and ChatGPT/Codex (OpenAI)—for assistance with proof exploration, C++/Python enumerator

and certificate tooling, Lean 4 formalization, and manuscript preparation. The Lean 4 formalization (under development) checks the results listed in Layers A and B of Table 1. The finite enumerations and exact rational/integer certificates are checked by the published C++ and Python verifiers, provided with SHA-256 replay packets and positive controls. We thank Qichun Wang, Yueying Lou, and Zifan Wang for the argument behind Section 6.3, and E. Brier and P. Langevin for their published list of cubic forms.

References

[1] É. Brier and P. Langevin, Classification of Boolean cubic forms of nine variables, in *Proc. 2003 IEEE Information Theory Workshop*, 2003, pp. 179–182. Representative list: `https://langevin.univ-tln.fr/project/cubics/classification.txt`.

[2] C. Carlet, *Boolean Functions for Cryptography and Coding Theory*, Cambridge University Press, 2021.

[3] P. Delsarte, *An Algebraic Approach to the Association Schemes of Coding Theory*, Philips Research Reports Supplements, No. 10, 1973.

[4] J. H. Griesmer, A bound for error-correcting codes, *IBM Journal of Research and Development* **4** (1960), 532–542.

[5] S. Kavut, S. Maitra, and M. D. Yücel, There exist Boolean functions on n (odd) variables having nonlinearity $> 2^{n-1} - 2^{(n-1)/2}$ if and only if $n > 7$, *IACR Cryptology ePrint Archive*, Report 2006/181, 2006.

[6] S. Kavut, S. Maitra, and M. D. Yücel, Search for Boolean functions with excellent profiles in the rotation symmetric class, *IEEE Transactions on Information Theory* **53** (2007), no. 5, 1743–1751.

[7] S. Kavut and M. D. Yücel, 9-variable Boolean functions with nonlinearity 242 in the generalized rotation symmetric class, *Information and Computation* **208** (2010), no. 4, 341–350.

[8] J. Mykkeltveit, The covering radius of the $(128, 8)$ Reed–Muller code is 56, *IEEE Transactions on Information Theory* **26** (1980), no. 3, 359–362. `https://doi.org/10.1109/TIT.1980.1056187`.

[9] N. J. Patterson and D. H. Wiedemann, The covering radius of the $(2^{15}, 16)$ Reed–Muller code is at least 16276, *IEEE Transactions on Information Theory* **29** (1983), no. 3, 354–356.

[10] Q. Wang, Y. Lou, and Z. Wang, The covering radius of the Reed–Muller code RM$(1, 9)$ is 242, arXiv:2610.17871 [cs.IT], v1, 2026 (preprint).

Layer	Trusted Components	Results Covered
Layer A	Lean 4.15.0 kernel only (all 79 modules archived)	Thm. 2.2 (no balanced NL244), Lem. 3.1, Thm. 3.2 (rank-8 parity), Thm. 3.4 (63/64 splitting), Prop. 3.5 (`eminClosed`), Thm. 4.1 (shifted Krawtchouk), Thm. 4.2 (barycenter laws), Lem. 4.3, Thm. 4.4 (Griesmer, Hamming, no $[24, 8, \geq 10]_2$), The normal form of Thm. 5.1 ($t = 98, r = 6, \lvert Z\rvert = 144$), Thm. 5.3 ($t = 96, (0, 31, 39)$ dichotomy), `Balanced242T96Quartic` ($a = 0$ far-half mod 4 law), `Balanced242QuarticHighBlock` (quartic inversion, hinges, dual cuts), `Balanced242LowBlock` (low-block defect budget, mod 32 autocorr, $t = 106$ rank 8), `Balanced242Z64Switch` (unconditional no rank 2 for $t \leq 112$; $t = 114, 116$ plane switch), `Balanced242T108R6` (unconditional $t = 108$ rank 6 $\lvert Z\rvert = 112$), `Balanced242T104R8` (unconditional all $t \geq 104$ rank 8), `Balanced242T112` (unconditional $t = 112$, all ranks, zero LP), `Balanced242T114` (unconditional $t = 114$ and $t = 116, k = 12$ via $[12, 9, 4]$ sphere packing), `Balanced242T118` (unconditional low-block $t = 118$ via kernel RREF cuts), `Balanced242Endpoint` (unconditional $t \geq 120$).
Layer B	Lean 4.15.0 kernel + explicit `SectionHypotheses` / `ComputationalClosures` / `EndpointK16Direct` (or `CoveringRadius244`)	Thm. 6.1 (all $t \leq 92$ coordinate-free, all 17 $t = 94$ frontier profiles via 205 anchor-pair + 16 affine Farkas certs), Thm. 6.2 ($(96, 21, 34, 12)$, $(96, 2, 47, 27)$, $(96, 1, 54, 24)$, $(96, 0, 51, 27)$), and top-level coordinate-free master theorems `Balanced242Master.no_balanced_9var_nl242_direct` (using `EndpointK16Direct` for $t = 116, k = 16$) and `Balanced242Master.no_balanced_9var_nl242` (using `CoveringRadius244`).
Layer C	Independent C++17 / Python exact-arithmetic replay packets + SHA-256 manifests + positive controls	Finite 7-variable rank-6 section enumerations (`sel`, `sel2`, `sel_e`, `flag_sections`), $t = 96$ quartic four-cell enumeration (`nl242_t96_closure.zip`), $t = 98$ quartic four-cell enumerations + exact integer dual exclusions + 2^{21}-character orbit-mass check (`nl242_t98_closure.zip`), $t = 104$–108 middle-window reconstructions (`nl242_middle_closure.zip`), $t = 116, k = 16$ direct high-block and low-complement enumeration (`nl242_endpoint_closure.zip`, cross-checked by `nl242_endpoint_crosscheck.zip`), $t \geq 100$ low-block exact rational Farkas/LP duals and zero-LP RM$(3, 8)$ mod 8 y_Z verifiers (`nl242_top_closure.zip`), and $\rho(1, 9) < 244$: Burnside orbit count, class separation and exhaustive search over classes 282 and 290 (`nl244_verify.zip`).

Table 1: Explicit three-layer trust boundary of the verification artifact across all 79 archived Lean 4 modules and SHA-256-manifested Layer C replay packets.

Paper Result	**Lean 4 Module & Identifier**	**External Hyp.?**
Thm. 2.2	`Balanced242No244.no_balanced_nl244`, `peak28_of_cap28`	No
Thm. 3.2	`Balanced242Parity.rank_eq_8_iff_k_odd`	No
Thm. 3.4	`Balanced242AnchorPair.anchor_pair_ranks_63_64`	No
Prop. 3.5	`Balanced242Farkas.eminClosed_le_secExc`	No
Thm. 4.1	`Balanced242Centered.kraw_triple_shifted_nonneg_and_mod`	No
Thm. 4.2	`Balanced242Centered.barycenter_translate_odd`, `cut_sum_eq_half`	No
Lem. 4.3	`Balanced242Centered.residual_kernel_counts`	No
Thm. 4.4	`Balanced242Centered.griesmer_bound`, `sphere_packing_d3`, `no_binary_code_24_8_10`	No
Thm. 5.1 (normal form)	`Balanced242T98R6Z144.no_t98_r6_z144`	No
Thm. 5.3	`Balanced242T96Arith.t96_0_31_39_dichotomy`	No
Quartic high-block theory	`Balanced242QuarticHighBlock.high_block_inversion`, `hinge_bound`, `dual_cut_sound`	No
Low-block budget & $t = 106$ r8	`Balanced242LowBlock.low_defect_budget`, `low_autocorr_mod32`, `no_t106_rank8`	No
Rank 2 ($t \leq 112$) & switch	`Balanced242Z64Switch.no_Z64_of_t_le_112`, `endpoint_switched_blocks`	No
$t = 108$ rank 6 $\lvert Z\rvert = 112$	`Balanced242T108R6.no_t108_rank6`, `cube8_weight_cases`	No
All $t \geq 104$ rank 8	`Balanced242T104R8.no_t104_rank8`, `no_rank8_of_t_ge_104`	No
Unconditional $t = 112$	`Balanced242T112.no_t112`, `rank4_Z96_n0_ge8`	No
Unconditional $t = 114$ & $k = 12$	`Balanced242T114.no_t114`, `t116_forces_k16`	No
Unconditional $t = 118$	`Balanced242T118.no_t118` (low-block only, kernel RREF)	No
Unconditional $t \geq 120$	`Balanced242Endpoint.low_agreements_below_120`	No
Thm. 6.1 ($t \leq 92$)	`Balanced242LowRange.axisAgreements_ge_188`, `autocorr_ge_m136`	`SectionEnum`
Thm. 6.1 (17 frontier profiles)	`Balanced242T94.t94_all_17_frontier_excluded`	`SectionEnumE`
Other $t = 94$ profiles	`ComputationalClosures.t94_nonfrontier`	Yes
Thm. 6.2	`Balanced242T94.t96_21_34_12_excluded`, `Balanced242T96Centered.excl_96_*`, `Balanced242T96_0_51_27`	`SectionEnumE`
Thm. 6.3 (Master)	`Balanced242Master.no_balanced_9var_nl242_direct`, `no_balanced_9var_nl242`	Layers B & C

Table 2: Mapping between results in this paper and the 79 archived zero-Mathlib Lean 4 modules.