Centered Arithmetic Is Not Leakage-Neutral

Finite-Noise Optimal Attacks and Minimax Laws for Modular Representations

Online Resource 1: Supplementary Proofs

Journal of Cryptographic Engineering

José Luis Delgado
Universitat Oberta de Catalunya, Barcelona, Spain
Corresponding author: jdelgado13@uoc.edu
ORCID: 0009-0009-2771-7507

1 Theorem Proofs

Definitions and Notation

For prime q, $h \in \mathbb{F}_q^2$, and $c_j \in \mathbb{F}_q^2$, set

$$S(h) = \sum_{j=1}^{T} g_j(c_j \cdot h + b_j). \tag{1}$$

For $\mathcal{X}_s = \{x : \sum_i x_i = s\}$, define

$$F_s(z) = \sum_{x\in\mathcal{X}_s} A_x e^{\langle \lambda_x, z\rangle}, \qquad \lambda_{x,i} = w_i(x_i)/\sigma_i^2,$$
$$A_x = \prod_i \exp\left(-\frac{(c_i - w_i(x_i))^2}{2\sigma_i^2}\right),$$

and

$$C_i(s,t,\Lambda) = \sum_{\substack{x\in\mathcal{X}_s, y\in\mathcal{X}_t\\ \lambda_x+\lambda_y=\Lambda}} A_x A_y (\lambda_{x,i} - \lambda_{y,i}). \tag{2}$$

For positive real-analytic location densities, let

$$B_s(\alpha) = \frac{q^{1-d}}{\alpha!} \sum_{x\in\mathcal{X}_s} \prod_i f_i^{(\alpha_i)}(-a_i(x_i)),$$

and let $H_{ij}(s,t;\alpha)$ be the coefficient of z^α in

$$z_i\{(\partial_i p_s)p_t - p_s(\partial_i p_t)\} - z_j\{(\partial_j p_s)p_t - p_s(\partial_j p_t)\}. \tag{3}$$

For the digital-sum quantities $H(n) = \sum_{x<n} \mathrm{HW}(x)$ and $Q(n) = \sum_{x<n} \mathrm{HW}(x)^2$, put

$$S_t = H(q-t) + \beta t - H(t), \tag{4}$$
$$T_t = Q(q-t) + \beta^2 t - 2\beta H(t) + Q(t), \tag{5}$$

so that

$$E_1(t) = \frac{qT_t - S_t^2}{q^2}. \tag{6}$$

Theorem 1.1 (Exact modular likelihood factor)**.** *For arbitrary conditionally independent local channels, put $a_i(x) = \ell_i(y_i \mid x)$. Then all conditional sum likelihoods are*

$$p_s(y_0,\ldots,y_{d-1}) = q^{1-d}(a_0 * \cdots * a_{d-1})(s), \qquad s \in \mathbb{Z}_q.$$

Consequently one observation is scored for every sum hypothesis by

$$q^{1-d}\,\mathrm{IDFT}\left(\prod_i \widehat{a}_i\right)$$

in $O(dq\log q)$ arithmetic operations and $O(q)$ working memory. The same statement holds on every finite abelian group after replacing the DFT by its character transform.

Proof. Because there are exactly q^{d-1} share tuples of prescribed sum, summing the product channel density over that fiber produces the displayed circular convolution; the convolution theorem then yields the algorithm. Iterated character convolution is identical on a finite abelian group. $\square$

Theorem 1.2 (All-pair finite-field ridge scoring). *For every nonzero c_j, write uniquely $c_j = \lambda_j u_j$, where $u_j \in \{(1,s) : s \in \mathbb{F}_q\} \cup \{(0,1)\}$ and $\lambda_j \in \mathbb{F}_q^*$. Define*

$$A_u(z) = \sum_{j:u_j=u} g_j(\lambda_j z + b_j), \qquad C = \sum_{j:c_j=0} g_j(b_j).$$

Then

$$S(h) = C + \sum_{u\in\mathbb{P}^1(\mathbb{F}_q)} A_u(u \cdot h).$$

With unnormalized one- and two-dimensional DFTs,

$$\begin{aligned} \widehat{S}(ku) &= q\widehat{A}_u(k) && (k \neq 0), \\ \widehat{S}(0) &= q\sum_u \widehat{A}_u(0) + q^2 C. \end{aligned}$$

Consequently all q^2 values of S are obtained exactly, up to the chosen floating-point DFT accuracy, in $O(Tq + q^2 \log q)$ arithmetic operations and $O(q^2)$ memory after the tables g_j are available. Adding a new trace block costs $O(\Delta Tq)$; every frozen prefix reuses the previous projective accumulator. For the complete masked likelihood, including construction of all g_j, the total cost is $O(Tdq \log q + q^2 \log q)$, versus $O(Tq^2)$ just to enumerate pair lookups naively.

Proof. The projective decomposition is unique because q is prime: normalize by the first nonzero coordinate. Substitution in Equation (1) produces the spatial identity, including the constant contribution of $c_j = 0$.

With $e_q(x) = \exp(2\pi \mathrm{i} x/q)$, Fourier inversion takes the form

$$A_u(u \cdot h) = q^{-1} \sum_{k\in\mathbb{F}_q} \widehat{A}_u(k) e_q(ku \cdot h).$$

Taking the two-dimensional DFT and using character orthogonality, the term at frequency $\xi \neq 0$ vanishes unless $\xi = ku$. Every nonzero ξ has a unique such projective direction and nonzero k, and the two character sums contribute $q^2/q = q$, proving the first formula. At frequency zero, every direction contributes its $k = 0$ coefficient and the constant contributes $q^2 C$. Thus $q + 1$ length-q DFTs fill a two-dimensional spectrum, which one inverse DFT evaluates at every pair. Constructing the A_u scans each length-q row once. These observations establish the time, memory, and incremental-prefix bounds. Theorem 1.1 constructs one length-q log-likelihood row in $O(dq \log q)$, yielding the final bound. $\square$

Theorem 1.3 (Exact likelihood of the product statistic). *The joint density of* $(\operatorname{sign} T, \log |T|)$*, conditional on $S = s$, is*

$$h_s(\epsilon, u) = q^{1-d} \sum_{\substack{x_0+\cdots+x_{d-1}=s \\ e_0\oplus\cdots\oplus e_{d-1}=\epsilon}} (g_{0,x_0}^{e_0} *_{\mathbb{R}} \cdots *_{\mathbb{R}} g_{d-1,x_{d-1}}^{e_{d-1}})(u).$$

Equivalently, h is convolution on $\mathbb{Z}_q \times \mathbb{Z}_2 \times \mathbb{R}$. Consequently $\sum_j \log h_{s_j(a)}(\epsilon_j, u_j)$ is MAP-optimal among every attack measurable with respect to the product statistics $(T_j)_j$. Equality between the complete and product likelihoods for all priors and decision problems holds exactly on the corresponding sufficiency locus.

Proof. On each orthant, the map $z_i = (-1)^{e_i} e^{u_i}$ has local Jacobian e^{u_i}. Thus $g_{i,x_i}^{e_i}$ is precisely the density of the local signed log-magnitude. Product sign is parity of the local signs and product log-magnitude is the sum of the local log-magnitudes. Conditioning on the modular share sum and summing the resulting real convolutions proves the formula. MAP optimality is the likelihood principle for the pushed-forward experiment. The final assertion is Blackwell data processing, with equality exactly when the product is sufficient. $\square$

Theorem 1.4 (Finite Gaussian product-sufficiency classification)**.** *For $d \geq 2$, the statistic $P(Y-c)$ is sufficient for the complete family $\{p_s : s \in \mathbb{Z}_q\}$ if and only if*

$$C_i(s,t,\Lambda) = 0$$

for every s, t, Λ, and coordinate i. This is a finite necessary-and-sufficient algebraic test in the leakage tables, centers, and noise variances.

Proof. Gaussian expansion factors $p_s(z) = K(z)F_s(z)$, where $K > 0$ is independent of s. Hence every likelihood ratio is F_s/F_t. The product-preserving vector fields $D_{ij} = z_i\partial_i - z_j\partial_j$ show that this ratio is P-measurable exactly when all $D_{ij}(F_s/F_t)$ vanish. Expanding their numerators and grouping equal exponential slopes produces

$$(D_{ij}F_s)F_t - F_s(D_{ij}F_t) = \sum_{\Lambda}(z_iC_i - z_jC_j)e^{\langle\Lambda,z\rangle}.$$

Linear independence of polynomial-exponential functions establishes the coefficient criterion; the almost-everywhere argument and the treatment of disconnected fibers appear in Section 2. ☐

Corollary 1.5 (Global finite-noise strictness)**.** *Suppose every local table is nonnegative, has its unique zero at residue zero, and satisfies $w_i(1) > 0$. For every $q \geq 2$, $d \geq 2$, and finite positive Gaussian variances, the centered product is not sufficient. For the pair $S = 0$ versus $S = 1$, the full observation has strictly larger $D_{\mathrm{KL}}(P_0\|P_1)$ and $D_{\mathrm{KL}}(P_1\|P_0)$, Rényi-$\alpha$ divergence for every $0 < \alpha < 1$, Bhattacharyya distance, and Chernoff information. It therefore has strictly larger Chernoff–Stein type-II exponents in either declared direction and a strictly larger equal-prior Bayesian i.i.d. exponent.*

Proof. Choose the combined slope whose only nonzero coordinate is $\Lambda_i = w_i(1)/\sigma_i^2$. Nonnegativity and the unique zero force the only contributing pair in Equation (2), for $(s,t) = (0,1)$, to be $x = 0$ and $y_i = 1, y_j = 0$ for $j \neq i$. Therefore

$$C_i(0,1,\Lambda) = -A_xA_yw_i(1)/\sigma_i^2 < 0.$$

Theorem 1.4 establishes insufficiency, while strict conditional Jensen establishes the KL and Rényi statements. Minimization over α preserves strictness because a nonidentical transformed pair has an interior affinity minimizer, whereas an identical transformed pair has affinity one. The Chernoff claim and the classical exponent theorems then imply the stated exponents. ☐

Theorem 1.6 (Non-Gaussian product-sufficiency laws)**.** *For positive real-analytic location densities, the centered product is sufficient for $S = s$ versus $S = t$ if and only if $H_{ij}(s,t;\alpha) = 0$ for every $i \neq j$ and every multi-index α. Any one nonzero coefficient is therefore a finite strictness certificate. In particular, when $q = 2$, every local leakage table is nonconstant, and $d \geq 2$, the centered product is never sufficient at any finite noise scale.*

For independent Laplace location noise with arbitrary positive share scales, let

$$N_s(v) = \#\{x \in \mathcal{X}_s : (w_0(x_0), \ldots, w_{d-1}(x_{d-1})) = v\}.$$

Product sufficiency for s versus t holds if and only if $N_s(v) = N_t(v)$ for every leakage-value vector v, equivalently $p_s = p_t$. Hence a nonnegative local table with residue zero as its unique zero implies strict insufficiency for $S = 0$ versus $S = 1$, for every $q \geq 2$ and $d \geq 2$. In each strict case the full observation has strictly larger KL, Rényi-α, Bhattacharyya, and Chernoff information than the centered product.

Proof. Equation (3) is the numerator of $D_{ij}(p_s/p_t)$, where $D_{ij} = z_i\partial_i - z_j\partial_j$. An analytic function is constant on the fibers of $P(z) = \prod_i z_i$ exactly when every D_{ij} annihilates it; Taylor expansion at the origin also joins the disconnected fibers. Coefficient extraction proves the first equivalence. For $q = 2$, parity factorization reduces the likelihood ratio to an invertible function of $\prod_i u_i(z_i)$, where each u_i is bounded and analytic. Product measurability forces every u_i to be constant, which would make a density proportional to a nonzero translate of itself, impossible for an integrable density.

For Laplace noise, distinct translates are linearly independent because the Laplace Fourier transform has no real zero. On every cell cut out by the translate locations, product measurability and the fields D_{ij} make a ratio of finite product mixtures constant along product fibers. In a tail cell all mixture terms have the same slopes, so the ratio is constant; analytic continuation across adjacent cells and continuity across coordinate hyperplanes make that constant global. Normalization forces $p_s = p_t$, and translate independence establishes the histogram criterion. The all-zero leakage vector has counts one and zero under $S = 0$ and $S = 1$, respectively, under the unique-zero hypothesis. Strict conditional Jensen and interior Chernoff minimization establish the information statements, with the rigidity and cell-continuation arguments developed in Section 3. ☐

Theorem 1.7 (High-noise Chernoff design law)**.** *Let $Y_i = w_i(X_i) + \tau N_i$, where the N_i are independent standard Gaussians, the local tables are bounded, and $d \geq 1$. Let $C_\tau(s,t)$ be the Chernoff information between the complete conditional laws given $S = s$ and $S = t$, and put*

$$K(s) = \mathbb{E}\left[\prod_i \big(w_i(X_i) - \mathbb{E}w_i(X_i)\big) \,\middle|\, S = s \right].$$

Uniformly over the finite set of pairs s, t,

$$C_\tau(s,t) = \frac{(K(s) - K(t))^2}{8\tau^{2d}} + O(\tau^{-2d-1}).$$

Consequently the mean pairwise Chernoff information satisfies

$$\frac{1}{q^2} \sum_{s,t} C_\tau(s,t) = \frac{1}{4\tau^{2d}} \sum_{k \neq 0} \left| q^{-1} \widehat{w}(k) \right|^{2d} + O(\tau^{-2d-1}) \tag{7}$$

for identical local tables $w_i = w$. Thus minimizing E_d is exactly the leading-order minimax problem for the average pairwise Chernoff information of the complete finite-noise modular experiment. For any finite family of representations, every strict E_d minimizer is the unique Chernoff minimizer for all sufficiently large finite τ.

The same statement holds for randomized encodings and independent local encoding randomness after replacing $w_i(x)$ by the conditional mean $g_i(x) = \mathbb{E}[W_i \mid x]$. If every g_i is constant, the τ^{-2d} coefficient vanishes and finite-noise information, if any, starts at a strictly higher order.

Proof. After scaling the observations by τ, put $\varepsilon = \tau^{-1}$. Relative to the standard Gaussian product measure, the conditional density ratio has the Hermite expansion

$$r_s(z;\varepsilon) = \mathbb{E}\left[\prod_i \sum_{n \geq 0} \frac{\varepsilon^n w_i(X_i)^n H_n(z_i)}{n!} \,\middle|\, S = s \right].$$

Every moment involving fewer than all d coordinates is independent of s, because every proper subset of arithmetic shares is uniform. Hence

$$r_s - r_t = \varepsilon^d (K(s) - K(t)) \prod_i H_1(z_i) + O_{L^2}(\varepsilon^{d+1});$$

centering changes only proper-subset terms. Because the finite leakage alphabets make the remainder uniform, the Rényi affinity has, uniformly for $\alpha \in [0,1]$, the quadratic-mean expansion

$$-\log \int p_s^\alpha p_t^{1-\alpha} = \frac{\alpha(1-\alpha)}{2} \varepsilon^{2d} (K(s) - K(t))^2 + O(\varepsilon^{2d+1}),$$

since the Hermite product has unit L^2 norm. Maximizing $\alpha(1-\alpha)$ establishes the first display. Averaging squared differences produces $2q^{-1} \sum_s K(s)^2$, while Parseval and Theorem 1.8 imply $q^{-1} \sum_s K(s)^2 = E_d$. Finiteness converts a strict leading-coefficient gap into an eventual strict ordering. Conditioning also averages over independent encoding randomness; its first Hermite coefficient is $g_i(x)$, proving the randomized extension. $\square$

Theorem 1.8 (Exact d-share spectrum)**.** *For every $d \geq 2$, leakage function $w : G \to \mathbb{R}$, and $k \neq 0$,*

$$K_d^w = q^{1-d} \underbrace{\bar{w} * \cdots * \bar{w}}_{d \text{ times}}, \qquad \widehat{K}_d^w(k) = q^{1-d} \widehat{w}(k)^d.$$

Proof. There are q^{d-1} share tuples with a prescribed sum. Therefore

$$K_d^w(s) = q^{1-d} \sum_{x_0 + \cdots + x_{d-1} = s} \prod_i \bar{w}(x_i),$$

which is the stated circular convolution. The unnormalized Fourier transform maps convolution to multiplication. Mean subtraction affects only mode zero, so $\widehat{\bar{w}}(k) = \widehat{w}(k)$ for $k \neq 0$. $\square$

Corollary 1.9 (Heterogeneous share spectrum)**.** *Let* $w_0, \dots, w_{d-1} : G \to \mathbb{R}$*, put* $\bar{w}_i = w_i - \mathbb{E}w_i$*, and define*

$$K^{w_0,\dots,w_{d-1}}(s) = \mathbb{E}\left[\prod_i \bar{w}_i(X_i) \,\middle|\, \sum_i X_i = s\right].$$

Then, for every $k \neq 0$*,*

$$K^{w_0,\dots,w_{d-1}} = q^{1-d}\bar{w}_0 * \cdots * \bar{w}_{d-1}, \qquad \widehat{K}^{w_0,\dots,w_{d-1}}(k) = q^{1-d}\prod_i \widehat{w}_i(k).$$

Proof. The tuple count and convolution argument in Theorem 1.8 does not require the factors to be equal. Fourier transformation converts the heterogeneous convolution into the stated product. □

Lemma 1.10 (Two's-complement complement identity)**.** *For* $0 \leq n < 2^\beta$*,*

$$\mathrm{HW}(\mathrm{tc}_\beta(-(n+1))) = \beta - \mathrm{HW}(n).$$

Proof. The word of $-(n+1)$ is $2^\beta - (n+1) = (2^\beta - 1) - n$, the bitwise complement of the β-bit word of n. □

Theorem 1.11 (Centered-minus-unsigned law)**.** *For* $0 \leq n < h$*, the residue* $x = q - 1 - n$ *lies in the negative centered half and*

$$\Delta(q-1-n) = \mathsf{s} + \mathrm{HW}(r+n) - \mathrm{HW}(n).$$

On the nonnegative half, $\Delta = 0$*.*

Proof. The centered value of $q - 1 - n$ is $-(n+1)$, so Lemma 1.10 implies $w_c = \beta - \mathrm{HW}(n)$. Moreover

$$q - 1 - n = 2^m - (r + n + 1) = (2^m - 1) - (r + n),$$

and hence $w_u = m - \mathrm{HW}(r + n)$. Subtraction proves the identity. □

Lemma 1.12 (Section-complement symmetry)**.** *For* $1 \leq t < q$ *and* $0 \leq x < q$*,*

$$w_{q-t}(q - 1 - x) = \beta - w_t(x).$$

Consequently $E_d(t) = E_d(q - t)$ *for every* $d \geq 1$*.*

Proof. The word representing $-1 - x$ is the bitwise complement of the word representing x. Reversal, cyclic shift, sign, and an additive constant preserve all non-DC Fourier magnitudes. □

Theorem 1.13 (Global order-one representation minimax)**.** *For every* $q \geq 2$*,* $\beta \geq \lceil \log_2 q \rceil$*, and contiguous section* $0 \leq t < q$*,*

$$E_1(t) \geq E_1(0).$$

If $q < 2^\beta$*, equality holds only at* $t = 0$*, so unsigned storage is the unique minimizer. If* $q = 2^\beta$*, every section induces the same word table and all sections tie. These are the only degeneracies.*

Proof. At exact width $M = 2^m$, put $r = M - q < M/2$. If $r = 0$, every section is a cyclic permutation of the complete m-bit word table and all energies tie. Assume henceforth $r > 0$. The section words are the complement of the interval $I_t = [q - t, M - t)$ of length r. For $y(x) = \mathrm{HW}(x) - m/2$, write $a_I = \sum_I y$ and $b_I = \sum_I y^2$. Its variance numerator is

$$V(I) = q(Mm/4 - b_I) - a_I^2.$$

The binary interval bounds in Section 4 imply, for every interval of length r,

$$|a_I| \leq |a_{[0,r)}|, \qquad b_I \leq b_{[0,r)}.$$

The terminal interval removed by unsigned storage is the bit complement of the prefix, so the two inequalities imply $V(I_t) \geq V(I_0)$. To obtain strictness, write an interior interval as $I = [a, a + r)$, with $a > 0$ and $a + r < M$. Graham's digital-sum inequality states that

$$\sum_{x \in I} \mathrm{HW}(x) - H(r) = H(a + r) - H(a) - H(r) \geq \min\{a, r\} > 0.$$

Bit complementation establishes the symmetric strict inequality relative to the terminal interval. Hence $|a_I| < |a_{[0,r)}|$ away from the two endpoints; together with $b_I \le b_{[0,r)}$, this makes $V(I) > V(I_0)$. The initial interval corresponds to the excluded duplicate section $t = q$, leaving $t = 0$ as the unique minimizer.

For additional word bits, direct differentiation of the integer numerator in Equation (6) results in

$$\frac{\partial V_\beta(q,t)}{\partial \beta} = 2t(q-t)\left(\beta - \frac{H(t)}{t} - \frac{H(q-t)}{q-t}\right) > 0$$

for $0 < t < q$, using $H(n)/n \le \log_2(n)/2$. The unsigned variance is unchanged by leading zero bits. This proves the theorem and its equality classification. □

Theorem 1.14 (Exact all-order representation phase change)**.** *At $(q, \beta) = (14, 4)$, unsigned is the unique minimizer of E_d for $d = 1, 2, 3$, whereas the centered section $t = 7$ is the unique minimizer for every $d \ge 4$.*

Proof. For integer leakage tables, Parseval and cyclic convolution establish the exact integer identity

$$q^{2d} E_d(t) = q \sum_s \left(w_t^{*d}(s)\right)^2 - \left(\sum_x w_t(x)\right)^{2d}.$$

Through $d = 8$, exact integer convolution identifies the unique minimizers just stated. In particular, the unsigned/centered numerators at $d = 3, 4$ are respectively

$$(131{,}473, 134{,}848), \qquad (5{,}950{,}145, 4{,}275{,}348).$$

For all larger orders, the centered spectral radius completes the argument. Let Y be the unnormalized squared Fourier magnitude for centered storage, with $Y = 36$ at the Nyquist mode. At the other even modes and at the primitive odd modes, respectively, exact reduction in $\mathbb{Z}[z]/(\Phi_7)$ and $\mathbb{Z}[z]/(\Phi_{14})$ produces

$$Y^3 - 35Y^2 + 378Y - 1183 = 0,$$
$$Y^3 - 45Y^2 + 402Y - 169 = 0.$$

Both cubics are positive at 36 and strictly increasing on $[36, \infty)$, so every such mode has $Y < 36$. Hence the centered spectral radius is exactly $9/49$. Direct evaluation of the Nyquist coefficient for every other section has squared magnitude at least $1/4$. Therefore, for every $t \ne 7$,

$$E_d(7) \le 13(9/49)^d, \qquad E_d(t) \ge (1/4)^d.$$

The inequality $13(36/49)^9 = 1320279436689408/1628413597910449 < 1$, together with the decrease of the ratio in d, proves unique centered optimality for all $d \ge 9$. The integer cases $4 \le d \le 8$ complete the proof. □

Theorem 1.15 (Lexicographic high-order minimax law)**.** *Let $\mathcal{R}$ be any finite family of real leakage tables on $\mathbb{Z}_q$. For $w \in \mathcal{R}$, put*

$$E_d(w) = \sum_{k=1}^{q-1} |q^{-1}\widehat{w}(k)|^{2d},$$

and sort its squared non-DC spectrum as $\Lambda_w = (\lambda_{w,1} \ge \cdots \ge \lambda_{w,q-1})$. The minimizers of $E_d(w)$ over $\mathcal{R}$ for all sufficiently large d are exactly the tables with lexicographically smallest Λ_w. If a lexicographic minimizer and a competitor first differ at index j, with $a = \lambda_{\min,j} < b = \lambda_{w,j}$, then the former defeats the competitor for every

$$d > \frac{\log(q-j)}{\log(b/a)}.$$

If $a = 0 < b$, it defeats it at every order. Tables with identical sorted spectra tie permanently. Moreover, any spectral-radius minimizer satisfies

$$E_d(w_\infty) \le (q-1) \min_{w \in \mathcal{R}} E_d(w)$$

for every $d \ge 1$.

Proof. The first $j-1$ powers cancel in a pairwise comparison, while sortedness implies

$$\sum_{r=j}^{q-1} \lambda_{\min,r}^d \le (q-j)a^d < b^d \le \sum_{r=j}^{q-1} \lambda_{w,r}^d.$$

Taking the maximum of the finitely many pairwise thresholds proves the first claims. Put

$$M = \min_{w\in\mathcal{R}} \max_k |q^{-1}\widehat{w}(k)|^2.$$

A radius minimizer has energy at most $(q-1)M^d$, while every table has energy at least M^d, which proves the approximation bound. $\square$

Theorem 1.16 (Constant-weight finite-noise minimax)**.** *Let $q \ge 2$, $d \ge 2$, and $\beta \ge 1$. Put*

$$M_\beta = \binom{\beta}{\lfloor \beta/2 \rfloor}.$$

The four equivalent conditions are:

(i) there is $c \in \mathcal{C}_\beta(q)$ for which the complete conditional laws P_s^c are identical for all s;

(ii) there is $c \in \mathcal{C}_\beta(q)$ with $\mathfrak{C}_\tau(c) = 0$ for one, equivalently every, $\tau > 0$;

(iii) there is an injective constant-Hamming-weight codebook of size q;

(iv) $q \le M_\beta$.

Whenever these conditions hold, the global minimum of $\mathfrak{C}_\tau$ over $\mathcal{C}_\beta(q)$ is zero, and its minimizers are exactly the injective constant-weight codebooks. The exact minimum width for zero complete finite-noise leakage is therefore

$$\beta_{\mathrm{cw}}(q) = \min\left\{\beta : \binom{\beta}{\lfloor \beta/2 \rfloor} \ge q\right\}. \tag{8}$$

Moreover, with

$$F(q) = \log_2 q + \frac{1}{2}\log_2\log_2 q + \frac{1}{2}\log_2(\pi/2),$$

$$\beta_{\mathrm{cw}}(q) = F(q) + O(1), \qquad -\varepsilon(q) \le \beta_{\mathrm{cw}}(q) - F(q) \le 1 + \varepsilon(q), \quad \varepsilon(q) \longrightarrow 0. \tag{9}$$

For the two public moduli in the manuscript,

$$\beta_{\mathrm{cw}}(3329) = 14, \qquad \beta_{\mathrm{cw}}(8380417) = 26.$$

If $q > M_\beta$, zero leakage is impossible and every codebook obeys the sharp spectral gap

$$E_d(c) := \sum_{k\ne 0} \left|q^{-1}\widehat{w}_c(k)\right|^{2d} \ge \frac{q-1}{q^{2d}}. \tag{10}$$

Equality holds if and only if $q = M_\beta + 1$ and, up to adding a constant and changing sign, w_c is the indicator of one residue. Equivalently, the codebook uses $q-1 = M_\beta$ words from one largest Hamming layer and one word from an adjacent layer.

Proof. If all codewords have Hamming weight t, then every local noiseless leakage is the constant t. The complete noisy observation law is therefore the same product Gaussian translate for every share tuple and every modular sum. Thus (iii) implies (i), and (i) implies (ii). Since Chernoff information is nonnegative and vanishes exactly for identical probability laws, (ii) implies (i).

Conversely, suppose the complete conditional laws are identical, set $\mu = q^{-1}\sum_x w_c(x)$, center each local observation, and take the conditional expectation of the product over all d shares. Independence and zero-mean noise remove every term containing a noise factor, so this conditional moment is $K_d^{w_c}(s)$. It is constant in s; its average over uniform s is zero, hence $K_d^{w_c} \equiv 0$. By Theorem 1.8,

$$0 = \widehat{K}_d^{w_c}(k) = q^{1-d}\widehat{w}_c(k)^d \qquad (k \ne 0).$$

All non-DC Fourier coefficients of w_c vanish, so w_c is constant. Thus every minimizer is a constant-weight codebook, proving the equivalence of (i)–(iii).

Since a Hamming layer of weight t contains $\binom{\beta}{t}$ words, the ratio $\binom{\beta}{t+1}/\binom{\beta}{t} = (\beta - t)/(t+1)$ shows that the largest layer is central, proving (iii) equivalent to (iv) and Equation (8).

We use Robbins' exact factorial inequalities: for every integer $n \geq 1$,

$$\sqrt{2\pi}\, n^{n+1/2} e^{-n+1/(12n+1)} < n! < \sqrt{2\pi}\, n^{n+1/2} e^{-n+1/(12n)}.$$

Applied to the even and odd central layers, they imply uniformly

$$\binom{\beta}{\lfloor \beta/2 \rfloor} = 2^{\beta} \sqrt{\frac{2}{\pi\beta}} \left(1 + O(\beta^{-1})\right).$$

If $\beta = \beta_{\mathrm{cw}}(q)$, then $M_{\beta-1} < q \leq M_{\beta}$. Taking base-two logarithms in these two inequalities and substituting $\log_2 \beta = \log_2 \log_2 q + o(1)$ shows that $F(q) \leq \beta + o(1)$ and $F(q) > \beta - 1 + o(1)$, which is exactly Equation (9). The exact comparisons

$$\binom{13}{6} = 1716 < 3329 \leq 3432 = \binom{14}{7},$$

$$\binom{25}{12} = 5200300 < 8380417 \leq 10400600 = \binom{26}{13},$$

establish the two stated widths.

For the spectral gap, let $y_x = w_c(x)$. The pairwise variance identity

$$q^2 \operatorname{Var}(y) = \sum_{x<t} (y_x - y_t)^2$$

shows that every nonconstant integer table satisfies $\operatorname{Var}(y) \geq (q-1)/q^2$. Equality holds exactly when $q-1$ entries are one integer and the remaining entry is an adjacent integer. Parseval reads

$$\sum_{k \neq 0} |\widehat{w}_c(k)|^2 = q^2 \operatorname{Var}(w_c) \geq q - 1.$$

Applying the power-mean inequality to the $q-1$ nonnegative numbers $|\widehat{w}_c(k)|^2$ yields Equation (10). Equality in both steps forces the one-exception table; its nonzero Fourier coefficients all have modulus one, so it also suffices. Such a codebook needs $q-1$ words in one layer, which under $q > M_\beta$ is possible exactly when $q = M_\beta + 1$; an adjacent layer supplies the exceptional word. ☐

Proposition 1.17 (Zero deterministic weighted-HW robustness)**.** *Let $q \geq 2$ and let $c : \mathbb{Z}_q \to \{0,1\}^\beta$ be injective and constant-weight. For weighted leakage $w_a(x) = \sum_j a_j c_j(x)$, the set of exactly neutral weights is the linear subspace*

$$\mathcal{A}_c = \left(\operatorname{span}\{c(x) - c(0) : x \in \mathbb{Z}_q\}\right)^{\perp}.$$

It contains the ordinary-HW vector $\mathbf{1}$, but its largest ℓ_∞ ball about $\mathbf{1}$ has radius zero. Explicitly, for every $\rho > 0$ there is a positive weight vector a with $\|a - \mathbf{1}\|_\infty < \rho$ for which w_a is nonconstant.

Proof. The displayed subspace is exactly the system of linear equations $a \cdot (c(x) - c(0)) = 0$. Equal ordinary Hamming weights put $\mathbf{1}$ in this subspace. Choose distinct codewords and a coordinate j on which they differ. For $0 < \varepsilon < \min(\rho, 1)$, the positive vector $a = \mathbf{1} + \varepsilon e_j$ separates their weighted leakages. Hence no nontrivial open norm ball is neutral. ☐

Theorem 1.18 (Universal complement-tag neutralization)**.** *If $\Pr[B = 1] = p$, the conditional weighted-HW mean is*

$$\mathbb{E}[W \mid x] = (1 - 2p) w_a(x) + pA.$$

Consequently an unbiased tag makes this mean exactly $A/2$ for every residue, every base representation, and every fixed vector of bit sensitivities. If w_a is nonconstant, $p = 1/2$ is the unique neutral choice in this two-branch family. With independent tags and centered, independent noise at the shares, every conditional centered-product mean is zero at every order. This tagged realization uses one stored tag bit and one unbiased random bit per word; Theorem 1.20 determines when the tag can be absorbed into existing word slack.

Proof. The two noiseless branch leakages are $w_a(x)$ and $A - w_a(x)$, which proves the displayed affine identity. Its coefficient of the nonconstant function w_a vanishes exactly when $1 - 2p = 0$. Conditional independence of tags and noise factors each centered-product mean into a product of zero conditional means. □

Corollary 1.19 (Global optimum for conditional-mean product kernels)**.** *Let $g(x) = \mathbb{E}[W \mid x]$ be the effective local first-moment table of a randomized encoding. The unbiased complement-tag encoder satisfies*

$$\sum_{k \neq 0} |q^{-1}\widehat{g}(k)|^{2d} = 0 \qquad \text{for every } d \geq 1.$$

It is therefore a simultaneous global minimizer, with the smallest possible value, over every randomized encoding under this conditional-mean product-kernel objective. This objective ignores conditional variance and mixture shape. Within the two-branch complement family and for nonconstant w_a, the minimizer is unique in the selector probability.

Proof. Theorem 1.18 makes g constant, so all of its non-DC Fourier coefficients vanish. The objective is a sum of nonnegative terms; uniqueness of the selector follows from the same theorem. □

Theorem 1.20 (Sharp randomness and redundancy lower bounds)**.** *Every such code satisfies*

$$n \geq \lceil \log_2 q \rceil + 1, \qquad \max_x H_\infty(P_x) \geq 1, \qquad \max_x H(P_x) \geq 1.$$

If the encoder is a deterministic function of the message and an independent seed R, then also $H_\infty(R), H(R) \geq 1$. Simultaneous equality in the entropy bounds holds exactly when all P_x are uniform on distinct two-word supports with a common coordinatewise midpoint: on some $r \geq 1$ varying coordinates they are antipodal pairs, and all remaining coordinates are fixed. Consequently $q \leq 2^{r-1}$. Conversely every such family attains equality. The total-width and seed bounds are tight for every q.

Proof. Fix a codeword u of probability a under message x, and another message y. Equality of bit marginals makes the expected Hamming distance from u identical under P_x, P_y. Disjoint decoding makes the latter at least one, whereas the former is at most $n(1-a)$. Thus every atom is at most $1 - 1/n$, every support has at least two words, and $2q \leq 2^n$.

At most one message can have an atom heavier than $1/2$: thresholding the common mean vector would identify that atom uniquely, contradicting disjoint supports for a second message. Some message therefore has both min-entropy and Shannon entropy at least one. If the maxima equal one, a message attaining the bound is uniform on two words. Every differing coordinate has common mean $1/2$, which bounds every atom of every other message by $1/2$; equality then forces all of them to be uniform pairs. Two binary pairs have the same coordinatewise midpoint exactly when they are antipodal on the same varying coordinates and agree elsewhere, giving 2^{r-1} available pairs and the equality classification. A deterministic map cannot increase Shannon or min-entropy, proving the seed bound; equality forces an unbiased binary seed. Choosing $r = \lceil \log_2 q \rceil + 1$ and any q antipodal pairs proves tightness. □

Proposition 1.21 (Exact finite-noise boundary)**.** *Under an unbiased tag, residues x, y have identical noiseless leakage laws if and only if*

$$|w_a(x) - A/2| = |w_a(y) - A/2|\,.$$

The same equivalence holds after adding Gaussian noise of any finite positive variance. Hence this encoder is locally leakage-independent precisely when the displayed absolute deviation is constant over all residues.

Proof. The conditional law is $\frac{1}{2}(\delta_{w_a(x)} + \delta_{A - w_a(x)})$, so equality is exactly equality of the two unordered, midpoint-$A/2$ pairs. Gaussian convolution is injective on finite signed measures because its Fourier transform has no zeros. □

Lemma 1.22 (Representation criterion)**.** *If $D(q,\beta) > 2U(q)$, then $|\widehat{w}_c(1)| > |\widehat{w}_u(1)|$. If $U(q) > 0$, then for every $d \geq 2$,*

$$\frac{|\widehat{K}_d^{w_c}(1)|}{|\widehat{K}_d^{w_u}(1)|} \geq \left(\frac{D(q,\beta) - U(q)}{U(q)}\right)^d > 1.$$

If $U(q) = 0$, the unsigned kernel coefficient is zero and the centered coefficient is nonzero.

Proof. The reverse triangle inequality implies $|\widehat{w}_c(1)| \geq q(D-U) > qU = |\widehat{w}_u(1)|$. For $U > 0$, divide and apply Theorem 1.8. For $U = 0$, the same strict inequality establishes the stated zero-versus-nonzero case. □

Theorem 1.23 (Continuous phase-aware margin)**.**

$$\inf_{\lambda \in [1/2,1]} G_3(\lambda) \geq 0.24003482888485672142.$$

Proof. Section 5 establishes the uniform tail bounds

$$\sum_{a>A} |U_a| \leq (\pi + \tfrac{1}{2})2^{1-A}, \qquad \sum_{a>A} |C_a| \leq (1+\pi)2^{1-A},$$

together with Lipschitz constants for the truncations. With $A_U = A_C = 10$, a 20,000-interval midpoint cover produces the outward-rounded mesh lower bound 0.31835504282967676110. The Lipschitz radius is at most 0.05600619449019234493, while the unsigned and carry tail penalties are at most 0.01422497130308512984 and 0.00808904815154256492. Subtracting these interval upper bounds proves the stated value. All midpoint coordinates are exact rationals and every complex exponential and modulus is evaluated in Sage's 160-bit complex interval field. □

Theorem 1.24 (Sharp dominance threshold)**.** *For every fixed integer $s \geq 0$, let $\beta_s(q) = \lceil \log_2 q \rceil + s$. Then*

$$D(q, \beta_s(q)) > 2U(q) \quad \textit{for every } q \geq 3 \quad \iff \quad s \geq 3.$$

Proof. For $\mathsf{s} = 3$, the finite transfer in Section 6 applies Theorem 1.23 to every $m \geq 17$. At the threshold band, the finite margin is at least 0.23365624115347768. An independent 160-bit interval sweep evaluates the exact geometric formula for all $3 \leq q \leq 2^{16}$; its minimum is greater than 0.318309886305. The same transfer covers every larger width: the dyadic factor and endpoint penalties decrease with m, while the fixed Abel tails are already included in the threshold-band bound.

For $\mathsf{s} \geq 4$, the phase-blind lower bound

$$D \geq \mathsf{s}|S_q|/q - |C_{q,r}|/q$$

is already positive above $2U$ for every q, and is monotone in s. This establishes sufficiency for all $\mathsf{s} \geq 3$.

For necessity, take $q = 2^m$. Then $r = 0$, hence $C_{q,0} = 0$. Every unsigned bit below the most significant bit completes an integer number of periods and has zero first Fourier coefficient. The most significant bit is the half-interval indicator, so $U(2^m) = |S_q|/q$. Therefore

$$D_s(2^m) = sU(2^m),$$

and the strict inequality fails for every $s \leq 2$. □

Theorem 1.25 (Exact weighted law)**.** *On the negative half,*

$$\Delta_a(q-1-n) = \sum_{j=m}^{\beta-1} a_j + \sum_{j=0}^{m-1} a_j(\mathrm{bit}_j(r+n) - \mathrm{bit}_j(n)).$$

Proof. For $j < \beta$, the centered negative word has bit $1 - \mathrm{bit}_j(n)$. For $j < m$, the unsigned word has bit $1 - \mathrm{bit}_j(r+n)$; for $j \geq m$, it has bit zero. Subtracting the weighted bit sums proves the formula. □

Theorem 1.26 (Exact phase-aware weighted radius)**.** *Assume $D_{\mathbf{1}} > 2U_{\mathbf{1}}$, and define*

$$\rho_\infty(q,\beta) = \inf\{\|a - \mathbf{1}\|_\infty : D_a \leq 2U_a\}.$$

For $\phi \in [0, 2\pi)$, let

$$v_j(\phi) = \delta_j - 2e^{\mathrm{i}\phi}u_j, \qquad V(\phi) = \sum_j v_j(\phi).$$

Then the exact radius is

$$\rho_\infty(q,\beta) = \min_\phi \sup_{y:\, \sum_j |\Re(\bar{y}v_j(\phi))|>0} \frac{|\Re(\bar{y}V(\phi))|}{\sum_j |\Re(\bar{y}v_j(\phi))|}. \tag{11}$$

If the $v_j(\phi)$ span $\mathbb{R}^2$, the inner maximum is

$$\max_{j:v_j\neq 0} \frac{|\Im(\overline{v_j}V)|}{\sum_\ell |\Im(\overline{v_j}v_\ell)|}. \tag{12}$$

If $v_j = c_j\xi$ with $c_j \in \mathbb{R}$, it is

$$\frac{|\sum_j c_j|}{\sum_j |c_j|}.$$

Every open box of radius below ρ_∞ satisfies $D_a > 2U_a$; the closed box at ρ_∞ contains equality, and every larger box contains a violating vector.

Proposition 1.27 (Finite algebraic phase elimination)**.** *Write $[z,w] = \Im(\bar{z}w)$. Every determinant in Equation* (12) *has the form*

$$[v_j(\phi), v_\ell(\phi)] = A_{j\ell} + B_{j\ell}\cos\phi + C_{j\ell}\sin\phi.$$

Under $t = \tan(\phi/2)$, the common denominator cancels and each active branch is a ratio of absolute quadratic polynomials. The projective phase line partitions at their real roots. On each cell, a minimizer of the upper envelope occurs only at a branch stationary point, an intersection of two active branches, a rank drop, or a cell endpoint. The stationary equations have degree at most two and branch intersections degree at most four. Consequently $\rho_\infty(q,\beta)$ is determined by a finite algebraic candidate set together with the rank-one formula.

Corollary 1.28 (Exact dyadic weighted radius)**.** *For $q = 2^m$ and $\beta = m + s$,*

$$\rho_\infty(2^m, m+s) = \begin{cases} 0, & s \le 2, \\ (s-2)/(s+2), & s \ge 3. \end{cases}$$

For $s \ge 3$, equality is attained by $a_{m-1} = 2s/(s+2)$ and $a_m = \cdots = a_{m+s-1} = 4/(s+2)$. In particular, no uniform excess-width-three box around HW can have radius greater than $1/5$.

Theorem 1.29 (Robust weighted-HW separation)**.** *For every $q \ge 3$, $\beta = m+3$, $\alpha \neq 0$, and bit sensitivities satisfying $|a_j/\alpha - 1| \le 0.077$,*

$$D_a(q,\beta) > 2U_a(q).$$

For $(q,\beta) = (3329, 16)$ and $(8380417, 32)$, the corresponding rigorous perturbation radii are at least 0.3163292868 *and* 0.6346298702.

Proof. The exact finite interval sweep establishes a radius of at least 0.1523421023 through 2^{16}. For every larger width, Lemma 6.1 replaces a generic bounded-variation transfer by an exact geometric-to-integral factor on every high bit-plane. The independent 160-bit interval verifier identifies the smallest analytic radius at $m = 17$:

$$\frac{0.23365624115347768}{3.0215785299117311} = 0.07732919692155194.$$

The explicit band formulas increase thereafter, so 0.077 is uniform with a strict reserve exceeding $3.29 \cdot 10^{-4}$. Single-parameter interval evaluations establish the two stated NIST radii. □

Corollary 1.30 (Certified exact NIST radii)**.** *The exact minimax radii of Theorem 1.26 satisfy*

$$0.47444937499013233220 < \rho_\infty(3329, 16) < 0.47444937499013233221,$$

$$0.63672613882631207518 < \rho_\infty(8380417, 32) < 0.63672613882631207519.$$

Proof. The 192-bit outward-rounded verifier recomputes every bit-plane coefficient and covers the complete phase circle by 32,768 closed cells. Each cell has a uniform rank-two determinant witness. Fixed-sign branches are ratios of trigonometric linear forms; the verifier certifies their derivative signs and isolates every competitive decreasing/increasing crossing by interval bisection. The minimizing dual-bit crossings are 10:0 for ML-KEM and 14:0 for ML-DSA. Point intervals at the isolated crossings bound every branch from above. The complete certificate widths are below $9.49 \cdot 10^{-44}$ and $4.41 \cdot 10^{-44}$, respectively. For each upper bound, the verifier also serializes every coordinate of the supporting-face primal perturbation, encloses its box norm strictly inside the displayed decimal endpoint, and encloses the complex boundary residual in a ball of radius below $1.4 \cdot 10^{-44}$. □

Theorem 1.31 (Transition-leakage law)**.** *For each share i, let (P_i, N_i) be a random β-bit previous-word/noise pair. Assume that these pairs are mutually independent, independent of all shares, and satisfy $\mathbb{E}N_i = 0$. Put*

$$p_{i,j} = \Pr[\mathrm{bit}_j(P_i) = 1], \qquad a_{i,j} = 1 - 2p_{i,j},$$

and, for $\rho \in \{u, c\}$, let

$$L_{\rho,i} = \mathrm{HD}(Y_\rho(X_i), P_i) + N_i.$$

Define its mean-free conditional product kernel by

$$K_{\mathrm{HD},\rho}(s) = \mathbb{E}\left[\prod_i (L_{\rho,i} - \mathbb{E}L_{\rho,i}) \,\middle|\, \sum_i X_i = s\right].$$

The mean-free order-d transition numerator has the exact spectrum

$$\widehat{K}_{\mathrm{HD},\rho}(k) = q^{1-d} \prod_{i=0}^{d-1} \widehat{w}_{\rho,a_i}(k), \qquad k \neq 0.$$

If $D_{a_i}(q,\beta) > 2U_{a_i}(q,\beta)$ for every share, then the centered first-mode factor is strictly larger at every share. Consequently, whenever the unsigned product is nonzero,

$$\frac{|\widehat{K}_{\mathrm{HD},c}(1)|}{|\widehat{K}_{\mathrm{HD},u}(1)|} = \prod_i \frac{|\widehat{w}_{c,a_i}(1)|}{|\widehat{w}_{u,a_i}(1)|} > 1;$$

if an unsigned factor vanishes, its centered counterpart is nonzero.

In particular, the condition holds for every $q \geq 3$, $\beta = m + 3$, if for every share there is an $\alpha_i \neq 0$ satisfying

$$\left| \frac{1 - 2p_{i,j}}{\alpha_i} - 1 \right| \leq 0.077 \quad \text{for every bit } j,$$

and it also holds at $(q, \beta) = (3329, 16)$ and $(8380417, 32)$ throughout the corresponding 31.6% and 63.4% boxes. At the opposite boundary, if every $p_{i,j} = 1/2$, both conditional-mean numerators vanish exactly.

Proof. For a fixed current word y, linearity of expectation implies

$$\mathbb{E}[\mathrm{HD}(y, P_i)] = \sum_j \bigl(p_{i,j} + (1 - 2p_{i,j}) \,\mathrm{bit}_j(y)\bigr).$$

Thus, up to a constant removed by mean centering, the conditional leakage function is exactly w_{ρ,a_i}. Conditional on the shares, independence across i factors the product of centered leakages. Corollary 1.9 establishes the spectrum. The reverse-triangle argument of Lemma 1.22, applied to each weighted leakage, establishes the strict comparison under $D_{a_i} > 2U_{a_i}$. The stated uniform and NIST boxes are Theorem 1.29, applied independently at every share. When all previous bits are unbiased, every $a_{i,j} = 0$, so each conditional mean is constant and every nonzero Fourier numerator is zero. ☐

Proposition 1.32 (Abel carry bound)**.** *For every mode k,*

$$\frac{|C_{q,r}(k)|}{q} \leq \frac{m\left(2\min(r,h) + h|1 - e^{2\pi i k r/q}|\right)}{q}.$$

Proof. Let $a_n = \mathrm{HW}(n)$. Shift the first sum in $C = \sum_{n<h}(a_{n+r} - a_n)z^n$, then add and subtract the unshifted length-h interval. The phase mismatch contributes $mh|1 - z^r|$. The two intervals have symmetric difference at most $2\min(r,h)$, and $0 \leq a_n \leq m$. ☐

Theorem 1.33 (Exact near-dyadic carry law)**.** *Let*

$$N = 2^m, \qquad q = N - r, \qquad 1 \leq r < N/2, \qquad L = m - 1, \qquad M = 2^L,$$

and put

$$c = \lceil r/2 \rceil, \qquad f = \lfloor r/2 \rfloor, \qquad h = M - c.$$

For a mode $k \in \mathbb{Z}_q$, put $z_k = e^{2\pi \mathrm{i} k/q}$. Write $s_2(n) = \mathrm{HW}(n)$, $H(t) = \sum_{0 \le n < t} s_2(n)$, and define

$$F_L(z) = \sum_{t=0}^{M-1} s_2(t) z^t,$$

$$A_c(z) = \sum_{u=0}^{c-1} z^{M-c+u}, \qquad E_c(z) = \sum_{u=0}^{c-1} s_2(c-1-u) z^{M-c+u},$$

$$D_f(z) = \sum_{v=0}^{f-1} (1+s_2(v)) z^v, \qquad B_r(z) = \sum_{v=0}^{r-1} s_2(v) z^v.$$

Then every carry mode has the exact finite identity

$$C_{q,r}(k) = (z_k^{-r} - 1) F_L(z_k) + L A_c(z_k) - E_c(z_k) + z_k^{-r} \big(z_k^M D_f(z_k) - B_r(z_k) \big). \tag{13}$$

At mode one, write $z = z_1$. Then

$$F_L(z) = \frac{1 - z^M}{1 - z} \left(\frac{L}{2} + \frac{\mathrm{i}}{2} T_{m,r} \right), \qquad T_{m,r} = \sum_{j=0}^{L-1} \tan\Big(\frac{\pi 2^j}{q} \Big). \tag{14}$$

Define the universal constant

$$\kappa = \sum_{a=1}^{\infty} \tan\Big(\frac{\pi}{2^{a+1}} \Big) = 1.8098372016777826781584884053 \ldots. \tag{15}$$

If $N \ge 4r$, let

$$\ell_r = \lceil \log_2(r+1) \rceil, \qquad \mathcal{E}_{m,r} = \frac{2\pi r^2}{q} \left(\frac{7}{4} L + 6 + \frac{7}{4} \ell_r \right).$$

Then

$$|C_{q,r} - (fL - 2f - 2H(f) + \mathrm{i} r\kappa)| \le \mathcal{E}_{m,r}. \tag{16}$$

Consequently the analytic carry bound

$$\frac{|C_{q,r}|}{q} \le \frac{\sqrt{\big(fL - 2f - 2H(f)\big)^2 + (r\kappa)^2} + \mathcal{E}_{m,r}}{q} \tag{17}$$

holds, and the matching lower bound is obtained by replacing the plus sign before $\mathcal{E}_{m,r}$ by a minus sign and truncating at zero.

Proof. Fix k and write $z = z_k$. Shift the first digit-sum interval in the definition of the carry:

$$\begin{aligned} C_{q,r}(k) &= z^{-r} \sum_{t=r}^{r+h-1} s_2(t) z^t - \sum_{t=0}^{h-1} s_2(t) z^t \\ &= (z^{-r} - 1) \sum_{t=0}^{h-1} s_2(t) z^t + z^{-r} \left(\sum_{t=h}^{h+r-1} s_2(t) z^t - \sum_{t=0}^{r-1} s_2(t) z^t \right). \end{aligned}$$

Since $h = M - c$ and $h + r = M + f$, binary complementation below M and the identity $s_2(M+v) = 1 + s_2(v)$ imply

$$\begin{aligned} \sum_{t=0}^{h-1} s_2(t) z^t &= F_L(z) - L A_c(z) + E_c(z), \\ \sum_{t=h}^{h+r-1} s_2(t) z^t &= L A_c(z) - E_c(z) + z^M D_f(z). \end{aligned}$$

Substitution cancels the phase-mismatched copies of $LA_c - E_c$ and proves Equation (13) for every mode. For the remainder of the proof set $k = 1$.

For Equation (14), mark selected binary digits in

$$P_L(u, z) = \prod_{j=0}^{L-1} (1 + uz^{2^j}).$$

Then $F_L(z) = \partial_u P_L(1, z)$, while

$$\prod_{j=0}^{L-1} (1 + z^{2^j}) = \frac{1 - z^M}{1 - z}$$

and, for $x \not\equiv \pi \pmod{2\pi}$,

$$\frac{e^{\mathrm{i}x}}{1 + e^{\mathrm{i}x}} = \frac{1}{2} + \frac{\mathrm{i}}{2} \tan(x/2).$$

Here no denominator vanishes because $M < q < 2M$. Applying the last identity with $x = 2\pi 2^j/q$ proves Equation (14).

Put

$$A = z^{-r} \left(\sum_{t=0}^{r-1} z^t \right) (1 - z^M), \qquad Q = A/2 + A_c(z),$$

and

$$R = -E_c(z) + z^{-r} \big(z^M D_f(z) - B_r(z) \big).$$

Equations (13) and (14) become

$$C_{q,r} = LQ + \frac{\mathrm{i}}{2} A T_{m,r} + R. \tag{18}$$

Let $\theta = 2\pi/q$. Since $z^M = -e^{\mathrm{i}\theta r/2}$, the elementary inequality $|e^{\mathrm{i}x} - 1| \le |x|$ implies

$$|A - 2r| \le \frac{5}{2} \theta r^2. \tag{19}$$

Indeed, with $X = \sum_{j=1}^{r} z^{-j}$ and $Y = 1 - z^M$,

$$|X - r| \le \theta \sum_{j=1}^{r} j \le \theta r^2, \quad |Y| \le 2, \quad |Y - 2| \le \theta r/2.$$

Also

$$A_c(z) = -\sum_{u=0}^{c-1} e^{\mathrm{i}\theta(r/2 - c + u)},$$

so

$$|A_c(z) + c| \le \theta \sum_{u=0}^{c-1} |r/2 - c + u| \le \frac{1}{2} \theta r^2. \tag{20}$$

Because $f = r - c$, Equations (19) and (20) imply

$$|Q - f| \le \frac{7}{4} \theta r^2. \tag{21}$$

Reindex the tangent sum by $j = L - a$:

$$T_{m,r} = \sum_{a=1}^{L} \tan x_{a,m}, \qquad x_{a,m} = \frac{\pi}{2^{a+1}(1 - r/N)}.$$

When $N \ge 4r$, $x_{1,m} \le \pi/3$. The function $\tan x/x$ is increasing on $(0, \pi/2)$, and $2\sqrt{3}/\pi < 6/5$, so $\tan x \le (6/5)x$ on $[0, \pi/6]$. Hence

$$T_{m,r} \le \sqrt{3} + \frac{6}{5} \sum_{a=2}^{\infty} x_{a,m} \le \sqrt{3} + \frac{2\pi}{5} < 3,$$

where the last strict inequality follows, for example, from $\sqrt{3} < 26/15$ and $\pi < 22/7$. Thus

$$0 < T_{m,r} < 3. \tag{22}$$

Let $x_a = \pi/2^{a+1}$. The mean-value theorem, with $\sec^2 x \le 4$ on $[0, \pi/3]$, and the geometric tail imply

$$|T_{m,r} - \kappa| \le \frac{8\pi r}{3N} + \frac{6\pi}{5N}. \tag{23}$$

Indeed, the finite-part error is at most $4(r/N)(1-r/N)^{-1}\sum_a x_a \le 8\pi r/(3N)$, while the omitted tail is at most $6\pi/(5N)$. In particular,

$$r|T_{m,r} - \kappa| \le \frac{\pi r^2}{N}\left(\frac{8}{3} + \frac{6}{5r}\right) \le \frac{58\pi r^2}{15N} < \frac{4\pi r^2}{N} \le 2\theta r^2,$$

so

$$r|T_{m,r} - \kappa| \le 2\theta r^2. \tag{24}$$

The residual term R is controlled by noting that every digit sum in E_c, D_f, B_r is at most ℓ_r. Comparing the phases in the three finite boundary sums with their limits $-1, -1, +1$, respectively, produces

$$\begin{aligned}
|E_c(z) + H(c)| &\le \frac{1}{2}\theta \ell_r r^2,\\
\left|z^{-r} z^M D_f(z) + f + H(f)\right| &\le \frac{1}{4}\theta(1+\ell_r) r^2,\\
|z^{-r} B_r(z) - H(r)| &\le \theta \ell_r r^2.
\end{aligned}$$

Therefore

$$\left|R - \big(H(c) - f - H(f) - H(r)\big)\right| \le \theta r^2\left(\frac{1}{4} + \frac{7}{4}\ell_r\right). \tag{25}$$

The binary prefix recurrences

$$H(2t) = 2H(t) + t, \qquad H(2t+1) = 2H(t) + t + s_2(t)$$

show in both parity cases that

$$H(c) - f - H(f) - H(r) = -2f - 2H(f). \tag{26}$$

Subtracting $fL - 2f - 2H(f) + \mathrm{i}r\kappa$ from Equation (18) and applying Equations (21), (19), (22), (24), and (25) produces

$$\begin{aligned}
|\text{difference}| &\le L\frac{7}{4}\theta r^2 + \frac{1}{2}\frac{5}{2}\theta r^2 \cdot 3 + 2\theta r^2 + \theta r^2\left(\frac{1}{4} + \frac{7}{4}\ell_r\right)\\
&= \theta r^2\left(\frac{7}{4}L + 6 + \frac{7}{4}\ell_r\right),
\end{aligned}$$

which is Equation (16). The two modulus bounds follow from $\big||u| - |v|\big| \le |u - v|$. ☐

Corollary 1.34 (Sharp near-Mersenne obstruction)**.** *Let $q = 2^m - 1$ and $\beta = m$. Put $L = m - 1$, $\phi = \pi/q$, and*

$$T_{m,1} = \sum_{j=0}^{L-1} \tan\left(\frac{\pi 2^j}{q}\right).$$

Then

$$C_{q,1} = \frac{L}{2}\left(e^{-2\mathrm{i}\phi} - e^{-\mathrm{i}\phi}\right) + \mathrm{i}e^{-3\mathrm{i}\phi/2}\cos(\phi/2) T_{m,1}. \tag{27}$$

In particular,

$$|qD(q,m) - \kappa| \le \frac{\pi(7m+24)}{2q}, \qquad D(q,m) = \frac{\kappa}{q} + O\left(\frac{m}{q^2}\right). \tag{28}$$

Thus the previous Abel estimate $D(q,m) = O(m/q)$ is loose by a factor $\Theta(m)$ on the extremal near-Mersenne sequence.

Proof. For $r = 1$, one has $c = 1$, $f = 0$, $E_c = D_f = B_r = 0$, and Equation (18) reduces to Equation (27). In Equation (16), the comparison center is $\mathrm{i}\kappa$, $\ell_1 = 1$, and

$$\mathcal{E}_{m,1} = \frac{2\pi}{q}\left(\frac{7}{4}(m-1) + \frac{31}{4}\right) = \frac{\pi(7m+24)}{2q}.$$

Since $qD(q,m) = |C_{q,1}|$, the result follows. □

Corollary 1.35 (Fixed-defect and thin-defect classification)**.** *Let $q_m = 2^m - r_m$, $r_m \geq 0$, and let the excess width $s = \beta - m$ be fixed.*

If $r_m = 0$, then $C_{q_m,0} = 0$. At exact width $s = 0$, $\Delta \equiv 0$; for $s \geq 1$, $\Delta = s1_{\text{negative}}$ and

$$\widehat{w}_c(1) = (s+1)\widehat{w}_u(1). \tag{29}$$

For every finite pair with $N \geq 4r$, if

$$G_{m,r} = \sqrt{\left(fL - 2f - 2H(f)\right)^2 + (r\kappa)^2} + \mathcal{E}_{m,r},$$

then

$$\left|D(q, m+s) - \frac{s}{\pi}\right| \leq \frac{s(\pi+1)}{N} + \frac{G_{m,r}}{q}. \tag{30}$$

Suppose now that $mr_m/2^m \to 0$. Then

$$D(q_m, m) \longrightarrow 0, \qquad D(q_m, m+s) \longrightarrow \frac{s}{\pi}, \qquad U(q_m) \longrightarrow \frac{1}{\pi}, \tag{31}$$

and, including the phase,

$$\frac{\widehat{w}_u(1)}{q_m} \longrightarrow \frac{\mathrm{i}}{\pi}, \qquad \frac{\widehat{\Delta}(1)}{q_m} \longrightarrow \frac{\mathrm{i}s}{\pi}, \qquad \frac{|\widehat{w}_c(1)|}{|\widehat{w}_u(1)|} \longrightarrow s+1. \tag{32}$$

Consequently, for every fixed masking order $d \geq 2$,

$$\frac{|\widehat{K}_d^{w_c}(1)|}{|\widehat{K}_d^{w_u}(1)|} \longrightarrow (s+1)^d. \tag{33}$$

For fixed $r \geq 2$, the sharper exact-width expansion is

$$qD(q,m) = \sqrt{\left(f(m-3) - 2H(f)\right)^2 + (r\kappa)^2} + O_r\left(\frac{m}{q}\right), \qquad f = \lfloor r/2 \rfloor, \tag{34}$$

so $qD(q,m)/(m-1) \to \lfloor r/2 \rfloor$. The defects $r = 2t$ and $r = 2t+1$ have the same real leading term; parity changes the limiting imaginary term by exactly κ.

Proof. The dyadic statement follows because at mode one every unsigned bit below the most significant bit completes a whole number of periods, while the most significant bit is the negative-half indicator.

Assume $r_m > 0$. Theorem 1.33, the trivial bounds $H(f) \leq fL$ and $f \leq r/2$, show $|C_{q_m,r_m}|/q_m = O(mr_m/2^m) + O(mr_m^2/2^{2m}) = o(1)$. The half-interval sum satisfies

$$\left|\frac{S_q}{q} - \frac{\mathrm{i}}{\pi}\right| \leq \frac{\pi+1}{2^m}. \tag{35}$$

Indeed, compare its left Riemann sum with $\int_0^{1/2} e^{2\pi \mathrm{i}t}\,dt = \mathrm{i}/\pi$: the Lipschitz error is at most $\pi/(2q)$, and the unmatched endpoint interval costs at most $1/(2q)$. The reverse triangle inequality in modulus, together with $|C_{q,r}| \leq G_{m,r}$, establishes Equation (30) and the two limits for D.

For the unsigned phase, let $\omega_q = e^{-2\pi\mathrm{i}/q}$, $\omega_N = e^{-2\pi\mathrm{i}/N}$, and $a_x = s_2(x)$. On the complete dyadic word table only the most significant bit survives at mode one, hence

$$\frac{1}{N}\sum_{x=0}^{N-1} a_x \omega_N^x = \frac{\mathrm{i}e^{\mathrm{i}\pi/N}}{N\sin(\pi/N)}. \tag{36}$$

Put $x = \pi/N \le \pi/4$. The distance from the right-hand side to i/π is at most

$$\frac{1}{N\cos(x/2)} + \frac{1}{\pi}\Big(\frac{x}{\sin x} - 1\Big) \le \frac{2}{N}.$$

Here $1/\cos(\pi/8) < 11/10$, while $x/\sin x - 1 \le x^2$, the latter following from $1 - \cos x \le x^2/2$ and $\cos x \ge 1/\sqrt{2}$. For the non-dyadic word table, add and subtract the first q terms at the dyadic frequency. Since $0 \le a_x \le m$,

$$\begin{aligned}\left|\frac{\widehat{w}_u(1)}{q} - \frac{1}{N}\sum_{x=0}^{N-1} a_x\omega_N^x\right| &\le \frac{mr}{N} + \frac{m}{N}\sum_{x=0}^{q-1} 2\pi x\left|\frac{1}{q} - \frac{1}{N}\right| + \frac{mr}{N} \\ &\le \frac{(\pi+2)mr}{N}.\end{aligned}$$

Thus

$$\left|\frac{\widehat{w}_u(1)}{q} - \frac{\mathrm{i}}{\pi}\right| \le \frac{2 + (\pi+2)mr}{N}. \tag{37}$$

The actual delta coefficient is $\widehat{\Delta}(1) = z(sS_q + C_{q,r})$, so its phase limit follows from Equation (35) and $|C|/q = o(1)$. Addition proves the coefficient ratio, and Theorem 1.8 raises it to the d-th power. Equation (34) is the modulus form of Equation (16) with fixed r. ☐

Theorem 1.36 (Centered-versus-unsigned regime classification)**.** *(i) $\Delta \equiv 0$ iff $q = 2^m$ and $\beta = m$.*

(ii) If $q = 2^m$ and $\beta > m$, then $\Delta = \mathsf{s}\,1_{\text{negative}}$ exactly.

(iii) If $q < 2^m$ and $\beta = m$, the difference is carry-only.

(iv) At exact width, $q = 2^m - 1$ satisfies $qD(q, m) \to \kappa$. More generally, fixed defect $q = 2^m - r$ has $qD(q, m)/(m-1) \to \lfloor r/2 \rfloor$, with the exceptional $r = 1$ limit $qD(q, m) \to \kappa$.

(v) In the thin-defect regime $mr/2^m \to 0$, exact width is asymptotically degenerate, every fixed positive slack restores the macroscopic limit $D \to \mathsf{s}/\pi$, and the centered/unsigned coefficient ratio tends to $\mathsf{s} + 1$.

(vi) For the fixed-excess-width family $\beta_s(q) = \lceil \log_2 q \rceil + s$, the uniform condition $D > 2U$ holds iff $s \ge 3$.

(vii) At $s = 3$, the weighted condition $D_a > 2U_a$ holds throughout the box $|a_j/\alpha - 1| \le 0.077$.

(viii) Under Hamming-distance leakage from independent previous words, the effective bit weights are exactly $a_{i,j} = 1 - 2p_{i,j}$. The preceding weighted separation applies share by share, while $p_{i,j} = 1/2$ for every bit is an exact zero-numerator boundary.

Proof. Items (i)–(iii) follow directly from Theorem 1.11: at $n = 0$, $\Delta(q-1) = \mathsf{s} + \mathrm{HW}(r)$, which vanishes exactly when $\mathsf{s} = r = 0$. Items (iv)–(v) are Corollary 1.34 and Corollary 1.35; item (vi) is Theorem 1.24; item (vii) is Theorem 1.29; and item (viii) is Theorem 1.31. ☐

Theorem 1.37 (Centered-Reduction Leakage Law)**.** *Fix $q \ge 3$, $\beta \ge m = \lceil \log_2 q \rceil$, and $d \ge 2$. For centered HW leakage, there exists $k_c \in \{1, \dots, q-1\}$ such that*

$$|\widehat{w}_c(k_c)| \ge 1, \qquad |\widehat{K}_d^{w_c}(k_c)| \ge q^{1-d},$$

and its phase-aligned predictor has mean magnitude at least q^{-d}. Moreover, exactly one of the following holds.

(a) If $q = 2^m$ and $\beta = m$, then $w_c = w_u$ and $\Delta \equiv 0$.

(b) Otherwise, for the isolated additive representation component Δ, there exists $k \in \{1, \dots, q-1\}$ such that

$$|\widehat{\Delta}(k)| \ge 1, \qquad |\widehat{K}_d^{\Delta}(k)| \ge q^{1-d},$$

and the corresponding phase-aligned predictor has mean magnitude at least q^{-d}.

Thus the centered and unsigned leakage tables are identical if and only if the modulus is an exact power of two stored at exact width.

Proof. The centered leakage is integer-valued and nonconstant because $w_c(0) = 0$ and $w_c(1) = 1$. Theorem 1.36(i) proves case (a). Otherwise Δ is likewise integer-valued and nonconstant: it is zero on the nonnegative half and nonzero somewhere on the negative half. For either nonconstant integer-valued function f on q points, the mean-free energy is at least $(q-1)/q$, because

$$\sum_x |f(x) - \mathbb{E}f|^2 = \frac{1}{q}\sum_{x<y} |f(x) - f(y)|^2 \geq \frac{q-1}{q}.$$

The last inequality follows since a nontrivial integer partition has at least $q-1$ cross-pairs, each differing by at least one. Parseval therefore implies

$$\sum_{k=1}^{q-1} |\widehat{f}(k)|^2 = q\sum_x |f(x) - \mathbb{E}f|^2 \geq q-1.$$

At least one of the $q-1$ nonzero modes has magnitude at least one. Applying this argument first to w_c and, in case (b), to Δ, then using Theorem 1.8, establishes both kernel bounds. Dividing a kernel Fourier coefficient by q converts it into the expectation of the unit-magnitude phase-aligned predictor. $\square$

Lemma 1.38 (Noise does not cancel the numerator)**.**

$$\mathbb{E}\left[\prod_i (L_i - \mathbb{E}L_i) \,\middle|\, \sum_i X_i = s\right] = \left(\prod_i \alpha_i\right) K_d^w(s).$$

Proof. Expand the product. Every term containing a noise factor has conditional expectation zero by independence and zero mean. The noise-free term is the stated gain product times K_d^w. $\square$

Theorem 1.39 (Exact all-mode optimal predictor)**.** *For every $k \neq 0$,*

$$\widehat{\mu}(k) = q^{1-d}\prod_i \alpha_i \widehat{w}_i(k), \qquad \widehat{\nu}(k) = q^{1-d}\prod_i \widehat{v}_i(k).$$

If $\mu \neq 0$, the unique correlation-optimal mean-zero unit predictor, up to sign, is $g_{\mathrm{corr}} = \pm\mu/\|\mu\|_2$, and

$$\rho^2_{\mathrm{corr},*} = \frac{q^{-2d}\sum_{k\neq 0}\prod_i \alpha_i^2 |\widehat{w}_i(k)|^2}{\prod_i (\alpha_i^2 \operatorname{Var}(w_i) + \sigma_i^2)}.$$

If $\mu = 0$, every predictor has zero correlation. When $\mu \neq 0$, for $1 \leq k \leq \lfloor q/2 \rfloor$, let $c_k = 1$ when q is even and $k = q/2$, and $c_k = 2$ otherwise. The exact fraction of optimal-predictor energy in the real k-th harmonic is

$$\eta_k = \frac{c_k \prod_i \alpha_i^2 |\widehat{w}_i(k)|^2}{\sum_{\ell \neq 0}\prod_i \alpha_i^2 |\widehat{w}_i(\ell)|^2}.$$

Thus mode one is the global optimum exactly when all product coefficients outside $\{1, q-1\}$ vanish.

Theorem 1.40 (Power-optimal product score)**.** *If $\mu \neq 0$, then among scores $Y_g = g(S)T$,*

$$g_{\mathrm{pow}}(s) \propto \frac{\mu(s)}{\nu(s)}, \qquad \Lambda_* = \max_g \frac{\mathbb{E}[Y_g]^2}{\mathbb{E}[Y_g^2]} = \frac{1}{q}\sum_s \frac{\mu(s)^2}{\nu(s)},$$

where a zero-over-zero summand is zero. If $\mu = 0$, then $\Lambda_* = 0$ and every score with nonzero second moment ties. If $\nu > 0$ and $\mathbb{E}g(S) = 0$ is imposed, the optimum is

$$g_{\mathrm{pow},0}(s) \propto \frac{\mu(s) - c}{\nu(s)}, \quad c = \frac{\mathbb{E}[\mu/\nu]}{\mathbb{E}[1/\nu]},$$

with

$$\Lambda_{0,*} = \mathbb{E}[\mu^2/\nu] - \frac{\mathbb{E}[\mu/\nu]^2}{\mathbb{E}[1/\nu]}.$$

When $\mu \neq 0$, the correlation- and power-optimal directions coincide exactly when ν is constant on the support of μ. For n independent traces and $\Lambda_* > 0$, the oriented optimal score obeys

$$\Pr[\bar{Y}_n \le 0] \le \frac{1-\Lambda_*}{1+(n-1)\Lambda_*}.$$

This bound is at most $\eta \in (0,1)$ whenever

$$n \ge \max\left\{1, \left\lceil \frac{(1-\eta)(1-\Lambda_*)}{\eta\Lambda_*} \right\rceil\right\}.$$

For a matched-second-moment null satisfying $\mathbb{E}_0[T \mid S] = 0$ and $\mathbb{E}_0[T^2 \mid S] = \nu(S)$, a threshold with one-sided Cantelli bounds at most η_0 and η_1 exists whenever

$$n \ge \left\lceil \frac{1}{\Lambda_*}\left(\sqrt{\frac{1-\eta_0}{\eta_0}} + \sqrt{(1-\Lambda_*)\frac{1-\eta_1}{\eta_1}}\right)^2 \right\rceil. \tag{38}$$

Within the score class $g(S)T$, g_{pow} minimizes this requirement. If every $\sigma_i > 0$ and the noises are scaled by τ, then

$$\Lambda_*(\tau) = \frac{q^{-2d}\sum_{k\neq 0}\prod_i \alpha_i^2|\widehat{w}_i(k)|^2}{\tau^{2d}\prod_i \sigma_i^2}(1+O(\tau^{-2})).$$

Define the continuous Cantelli sample factor $F_r^{\mathrm{opt}}(\tau) = (1-\Lambda_{*,r}(\tau))/\Lambda_{*,r}(\tau)$. Consequently, for equal shares and common gain/noise, provided both nonzero spectral energy sums exist,

$$\lim_{\tau\to\infty} \frac{F_u^{\mathrm{opt}}(\tau)}{F_c^{\mathrm{opt}}(\tau)} = \frac{\sum_{k\neq 0}|\widehat{w}_c(k)|^{2d}}{\sum_{k\neq 0}|\widehat{w}_u(k)|^{2d}}. \tag{39}$$

Theorem 1.41 (Exact finite-noise variance)**.**

$$|\mathbb{E}Z_w| = |\alpha|^d \left(\frac{|\widehat{w}(k)|}{q}\right)^d,$$

$$\mathrm{Var}(Z_w) = (\alpha^2\,\mathrm{Var}(w) + \sigma^2)^d - \alpha^{2d}\left(\frac{|\widehat{w}(k)|}{q}\right)^{2d}.$$

Whenever both representation means are nonzero, define the exact deflection sample factor $F_w = \mathrm{Var}(Z_w)/|\mathbb{E}Z_w|^2$. Then

$$\frac{F_u}{F_c} = \frac{\mathrm{Var}(Z_u)/|\mathbb{E}Z_u|^2}{\mathrm{Var}(Z_c)/|\mathbb{E}Z_c|^2}.$$

Proof. The numerator is Theorem 1.8 projected onto Φ_k. Since $|\Phi_k| = 1$, the shares are independent before conditioning, and the noise is independent,

$$\mathbb{E}|Z_w|^2 = \prod_i \mathbb{E}(\alpha\bar{w}(X_i) + N_i)^2 = (\alpha^2\,\mathrm{Var}(w) + \sigma^2)^d.$$

Subtracting $|\mathbb{E}Z_w|^2$ proves the variance formula and the sample-factor identity follows by definition. □

Corollary 1.42 (Exact phase-aligned real variance)**.** *If $\mathbb{E}Z_w \neq 0$, let $\theta = \arg\mathbb{E}Z_w$ and $Y_w = \Re(e^{-\mathrm{i}\theta}Z_w)$. Then*

$$\mathrm{Var}(Y_w) = \frac{\mathbb{E}|Z_w|^2 + \Re(e^{-2\mathrm{i}\theta}\mathbb{E}Z_w^2)}{2} - |\mathbb{E}Z_w|^2,$$

where

$$\mathbb{E}Z_w^2 = \left[\frac{1}{q}\sum_{x\in\mathbb{Z}_q} e^{4\pi\mathrm{i}kx/q}\left(\alpha^2\bar{w}(x)^2 + \sigma^2\right)\right]^d.$$

Thus the exact aligned variance is identical for every independent zero-mean noise law with the same variance.

Proof. Use $(\Re V)^2 = (|V|^2 + \Re V^2)/2$. Independence across shares factors $\mathbb{E}Z_w^2$; the mixed leakage–noise term vanishes by zero mean. $\square$

Theorem 1.43 (Exponential-confidence harmonic test)**.** *Fix $q \geq 3$, $d \geq 2$, and a mode $k \neq 0$ with $2k \not\equiv 0 \pmod q$. For each trace sample S uniformly in $\mathbb{Z}_q$, sample uniform arithmetic shares conditioned on $\sum_i X_i = S$, and observe*

$$L_i = \alpha \bar{w}(X_i) + \tau N_i,$$

where $\alpha \neq 0$, the noises are independent across shares and traces, independent of the shares, mean zero, and have common finite variance $\sigma^2 > 0$. Let

$$Z_w = e^{2\pi \mathrm{i} kS/q} \prod_{i=0}^{d-1} L_i, \qquad \theta_w = \arg \mathbb{E}Z_w, \qquad Y_w = \Re(e^{-\mathrm{i}\theta_w} Z_w).$$

Then

$$\mu_w := \mathbb{E}Y_w = |\alpha|^d \left(\frac{|\widehat{w}(k)|}{q} \right)^d. \tag{40}$$

Let H_0 be the matched independent-label null obtained by replacing S in the phase factor by an independent uniform S'. Write $V_{1,w} = \mathrm{Var}_{H_1}(Y_w)$, $V_{0,w} = \mathrm{Var}_{H_0}(Y_w)$, and $V_w^\star = \max(V_{0,w}, V_{1,w})$. For $\mu_w > 0$, define

$$b_w = \max\left\{1, \left\lceil \frac{32 V_w^\star}{\mu_w^2} \right\rceil \right\}, \qquad J_\delta = 2\left\lceil \frac{16}{9} \log \frac{2}{\delta} \right\rceil + 1.$$

With $n_w = b_w J_\delta$ independent traces, split the scores into J_δ blocks of length b_w, take the median of the block means, and accept H_1 iff this median exceeds $\mu_w/2$. Both type-I and type-II errors are at most $\delta/2$.

For fixed δ, let $n_w(\tau, \delta)$ denote this integer sufficient budget. If the centered and unsigned coefficients at mode k are nonzero, then

$$\lim_{\tau \to \infty} \frac{n_u(\tau, \delta)}{n_c(\tau, \delta)} = \left(\frac{|\widehat{w}_c(k)|}{|\widehat{w}_u(k)|} \right)^{2d}. \tag{41}$$

At mode one, if $D(q, \beta) > 2U(q) > 0$, this limit is at least

$$\left(\frac{D(q,\beta) - U(q)}{U(q)} \right)^{2d}. \tag{42}$$

Along every thin-defect family of Corollary 1.35, it converges to $(s+1)^{2d}$.

Proof. Equation (40) is the phase-aligned form of Theorem 1.8; every noise term in the product expansion has zero expectation. Under the independent-label null, the phase has zero mean, so the aligned score has mean zero.

For one block average $\overline{Y}$, Chebyshev's inequality implies, under H_1,

$$\Pr(\overline{Y} \leq \mu_w/2) \leq \frac{4V_{1,w}}{b_w \mu_w^2} \leq \frac{1}{8},$$

and under H_0,

$$\Pr(\overline{Y} > \mu_w/2) \leq \frac{4V_{0,w}}{b_w \mu_w^2} \leq \frac{1}{8}.$$

Because the block failures are independent, Hoeffding's inequality for their Bernoulli indicators bounds the probability that at least half the blocks fail by

$$\exp\bigl[-2J_\delta(1/2 - 1/8)^2\bigr] = \exp(-9J_\delta/32) \leq \delta/2.$$

Thus the finite sample statement requires only finite variance, with Gaussian and sub-Gaussian noise as special cases.

Because $2k \not\equiv 0 \pmod q$, the constant noise term vanishes from $\mathbb{E}Z_w^2$. Corollary 1.42 therefore implies, uniformly for every fixed bounded leakage table,

$$V_{0,w}(\tau) = \frac{1}{2}\tau^{2d}\sigma^{2d}(1 + o(1)), \qquad V_{1,w}(\tau) = \frac{1}{2}\tau^{2d}\sigma^{2d}(1 + o(1)).$$

The common factor J_δ cancels, the block lengths diverge, and their ceilings are asymptotically negligible. Substituting Equation (40) proves Equation (41). The reverse triangle inequality establishes Equation (42), and Corollary 1.35 yields the final limit. $\square$

2 Likelihood Sufficiency Details

After $z_i = y_i - c_i$, expansion of the Gaussian density produces

$$p_s(z) = K(z)F_s(z), \qquad K(z) = q^{1-d} \prod_i \frac{e^{-z_i^2/(2\sigma_i^2) - c_i z_i/\sigma_i^2}}{\sqrt{2\pi}\sigma_i}.$$

For necessity in Theorem 1.4, let $\phi_{ij,u}$ multiply z_i by e^u and z_j by e^{-u}. If a likelihood ratio R is P-measurable, then $R \circ \phi_{ij,u} = R$ almost everywhere. The flow is a diffeomorphism, and positivity of the Gaussian mixtures identifies statistical and Lebesgue null sets. Continuity upgrades the equality to every point; differentiation at $u = 0$ yields $D_{ij}R = 0$.

Conversely, expand an analytic solution at the origin as $R(z) = \sum_\alpha r_\alpha z^\alpha$. Since $D_{ij}z^\alpha = (\alpha_i - \alpha_j)z^\alpha$, all surviving multi-indices are $(n, \dots, n)$, and locally $R = h_0(P)$. If $P(z) = P(z')$, the real-analytic functions $u \mapsto R(uz)$ and $u \mapsto R(uz')$ agree near zero because both products equal $u^d P(z)$. The one-variable identity theorem extends equality to $u = 1$. This includes disconnected nonzero fibers and the zero-product fiber.

Finite sums of distinct real exponentials with polynomial coefficients are linearly independent: restrict to a generic line that makes the scalar exponents distinct and successively annihilate all but one term by constant-coefficient differential operators. Thus each polynomial $z_iC_i - z_jC_j$ in the proof vanishes identically, which for $i \neq j$ is equivalent to $C_i = C_j = 0$.

For the information comparison, under a binary pair P, Q, let $R = dP/dQ$ and $\bar{R} = \mathbb{E}_Q[R \mid P(Y - c)]$. Non-sufficiency means that strict conditional Jensen applies on a set of positive measure. It is strict for $x \log x$ and for $-x^\alpha$, $0 < \alpha < 1$. If the transformed pair differs, its Chernoff affinity has an interior minimizer α_*, and

$$\min_\alpha A_T(\alpha) = A_T(\alpha_*) > A(\alpha_*) \geq \min_\alpha A(\alpha).$$

If it coincides, the left side is one while the distinct full pair has minimum affinity below one, establishing all strictness statements in Corollary 1.5.

3 Non-Gaussian Rigidity Details

The two rigidity arguments underlying Theorem 1.6 begin with a positive function R that is real analytic on $\mathbb{R}^d$. If $R = h(P)$ almost everywhere, invariance under the flow

$$(z_i, z_j) \longmapsto (e^u z_i, e^{-u} z_j)$$

holds everywhere by continuity, hence $D_{ij}R = 0$. Conversely write $R(z) = \sum_\alpha r_\alpha z^\alpha$ at the origin. Since $D_{ij}z^\alpha = (\alpha_i - \alpha_j)z^\alpha$, simultaneous annihilation leaves only $\alpha = (n, \dots, n)$, so locally $R = h_0(P)$. For any $P(z) = P(z')$, the analytic functions $u \mapsto R(uz)$ and $u \mapsto R(uz')$ agree near zero and therefore at $u = 1$. Thus R is globally constant on every product fiber. Expanding $D_{ij}(p_s/p_t)$ now proves the coefficient criterion in the theorem.

For the binary claim, write the two shifted local densities as g_{i0}, g_{i1}, and set

$$u_i(z) = \frac{g_{i0}(z) - g_{i1}(z)}{g_{i0}(z) + g_{i1}(z)}.$$

Direct parity expansion produces

$$p_s(z) = 2^{-d} \prod_i (g_{i0} + g_{i1}) \left(1 + (-1)^s \prod_i u_i(z_i) \right).$$

If $\prod_i u_i(z_i)$ were product-measurable, then on open intervals avoiding the isolated zeros one could fix all but two coordinates and obtain $u(x)v(y) = h(xy) \neq 0$. Logarithmic differentiation implies $xu'(x)/u(x) = yv'(y)/v(y) = \kappa$. Analyticity at zero forces κ to be a nonnegative integer, while boundedness on the real line forces $\kappa = 0$. Hence every u_i is constant. This would make two distinct translates of the same density proportional. Such proportionality is impossible: masses on consecutive intervals of the translation length would form a two-sided geometric sequence whose total diverges.

For Laplace noise, distinct translates are linearly independent because Fourier transformation turns a linear relation into a finite relation among distinct exponentials, because $(1 + b^2\xi^2)^{-1} \neq 0$ on $\mathbb{R}$. Tensoring proves independence of the multivariate product translates.

Product sufficiency forces proportional densities through a cellwise continuation argument. Partition each axis at zero and at all translate locations. On every open Cartesian cell, both conditional densities are finite exponential sums with fixed signs in the absolute values and their ratio is analytic. The fields D_{ij} make the ratio a function of P on that cell. Cells within a fixed sign orthant form a connected adjacency graph. Across a common nonzero face, varying a second coordinate makes P cover an open interval; continuity and the one-variable identity theorem therefore identify the two cell functions. In the outer tail cell all mixture components have the same coordinate slopes, so the ratio is constant. Propagation makes it constant on the orthant, and continuity at a coordinate hyperplane identifies adjacent orthant constants. Thus $p_s = Cp_t$ globally. Both integrate to one, so $C = 1$, and tensor independence makes this equality equivalent to equality of every finite histogram coefficient $N_s(v) = N_t(v)$.

4 Binary Interval Minimax Proof

Let $s(x) = \mathrm{HW}(x)$, $M = 2^m$, and

$$A_m(n) = \sum_{x<n}(s(x) - m/2), \qquad B_m(n) = \sum_{x<n}(s(x) - m/2)^2.$$

The proof uses the two interval bounds

$$A_m(r) \leq \sum_{x\in I}(s(x) - m/2) \leq -A_m(r), \tag{43}$$

$$\sum_{x\in I}(s(x) - m/2)^2 \leq B_m(r), \tag{44}$$

for every interval $I \subset [0, M)$ of length r.

With $S(n) = \sum_{x<n} s(x)$ and $D(a, b) = S(a + b) - S(a) - S(b)$, Graham's inequality states that $D(a, b) \geq \min(a, b)$ for positive arguments, proving the left side of Equation (43); bit complementation proves the right side. The accompanying elementary upper bound is

$$D(a, b) \leq a + b - 1. \tag{45}$$

Indeed, let $F_j(n)$ count ones in bit j among $0, \ldots, n - 1$, and choose h with $2^h < a + b \leq 2^{h+1}$. For $j < h$, the superadditivity defect of two prefixes of the periodic word $0^{2^j}1^{2^j}$ is at most 2^j. At bit h it is at most $a + b - 2^h$, and all higher terms vanish. Summation establishes $D(a, b) \leq \sum_{j<h} 2^j + a + b - 2^h = a + b - 1$.

Equation (44) is proved simultaneously with

$$B_m(u + v) - A_m(u + v) \geq B_m(u) + A_m(u) + B_m(v) + A_m(v) \tag{46}$$

for $u+v \leq M$. The base case is immediate. Split the next cube into two halves. On the lower half the new centered weight is $y_m - 1/2$, and on the upper half it is $y_m + 1/2$. An interval contained in one half is bounded by the induction hypothesis and Equation (43). If it crosses the midpoint with fragment lengths u, v, bit complementation makes its square sum

$$B_m(u) + A_m(u) + u/4 + B_m(v) + A_m(v) + v/4,$$

which Equation (46) bounds by the next-dimensional prefix.

Propagation of the glue inequality uses its gap $Q_m(u, v)$; assume $u \leq v$. The exact prefix recurrences produce three cases. If $n = u + v \leq M$,

$$Q_{m+1}(u, v) = Q_m(u, v) + n - D(u, v) \geq 0.$$

If $n = M + r$ and $v \leq M$, then $u \geq r$, and square complement symmetry implies

$$Q_{m+1}(u, v) = B_m(r) - B_m(u) + B_m(u - r) + M \geq M,$$

by the interval-square hypothesis on $[u-r,u)$. In the third case, if $v = M+a$, put $r = u+a$ and $b = M-a-u$. Prefix complement symmetry $A_m(M-x) = A_m(x)$ implies

$$\begin{aligned} Q_{m+1}(u,v) &= Q_m(u,a) + M - a + A_m(r) + A_m(u) - A_m(a) \\ &= Q_m(u,a) + u + b - D(u,b) \geq 0 \end{aligned}$$

by Equation (45). This closes the simultaneous induction for intervals of length at most half the cube. For a longer interval, its complement is a prefix and suffix of total length below half; applying the short-interval result once more to the gap between them proves the same bound. Hence Equation (44) holds at every length.

5 Continuous Tail and Mesh Bounds

5.1 Unsigned tail

Write $b_a = 1/2 + s_a$, where s_a is zero mean, period $p_a = 2^{1-a}$, and has a periodic primitive S_a satisfying $\|S_a\|_\infty \leq p_a/4$. The constant half has zero first coefficient. Integration by parts on $[0,\lambda]$ establishes

$$|U_a(\lambda)| \leq \lambda^{-1}\left(2\|S_a\|_\infty + \|S_a\|_\infty \frac{2\pi}{\lambda}\lambda\right) \leq (\pi + \tfrac{1}{2})p_a.$$

Summing p_a over $a > A$ proves

$$\sum_{a>A} |U_a| \leq (\pi + \tfrac{1}{2})2^{1-A}.$$

5.2 Carry tail

The difference $d_{a,\lambda}(t) = b_a(t+1-\lambda) - b_a(t)$ has zero mean and a periodic primitive bounded by $p_a/2$. Integration by parts on $[0,\lambda/2]$, using $\lambda \geq 1/2$, establishes

$$|C_a(\lambda)| \leq (1+\pi)p_a, \qquad \sum_{a>A} |C_a| \leq (1+\pi)2^{1-A}.$$

5.3 Lipschitz bounds

For

$$f_a(\lambda) = \lambda^{-1}\int_0^\lambda b_a(t)e^{-2\pi i t/\lambda}\,dt,$$

differentiation away from square-wave boundaries implies $|f_a'| \leq (2+\pi)/\lambda$. Continuity at boundaries extends the one-sided bound globally. Hence

$$\mathrm{Lip}(U_A; [\lambda_0, 1]) \leq A(2+\pi)/\lambda_0.$$

Separating normalization, endpoint, phase, and shift motion for the carry term produces

$$\mathrm{Lip}(C_A) \leq 2^{A+2} - 4 + A(12 + 2\pi).$$

The complex modulus is 1-Lipschitz, so their sum, with a factor two on the unsigned term, bounds $G_{3,A}$.

6 Finite-to-Continuous Transfer

6.1 Exact dyadic transfer

Lemma 6.1 (Aligned exponential quadrature)**.** *Let $N \geq 1$, $q \geq 3$, $\lambda = q/N$, $\delta_\sigma = \sigma 2\pi i/q$ for $\sigma \in \{-1,1\}$, and*

$$\eta_\sigma(q) = \frac{\delta_\sigma}{e^{\delta_\sigma} - 1}.$$

If f is constant on every cell $[n/N,(n+1)/N)$, then for every integer $J \geq 0$,

$$\frac{1}{q}\sum_{n=0}^{J-1} f(n/N)e^{\sigma 2\pi \mathrm{i} n/q} = \frac{\eta_\sigma(q)}{\lambda}\int_0^{J/N} f(t)e^{\sigma 2\pi \mathrm{i} t/\lambda}\,dt.$$

Moreover, writing $x = \pi/q$,

$$|\eta_\sigma(q)| \leq \frac{1}{1-x^2/6}, \qquad |\eta_\sigma(q)-1| \leq x + \frac{x^2}{6-x^2}.$$

Proof. On cell n, the integral evaluates to

$$\int_{n/N}^{(n+1)/N} e^{\sigma 2\pi \mathrm{i} t/\lambda}\,dt = e^{\delta_\sigma n}\frac{e^{\delta_\sigma}-1}{\delta_\sigma N}.$$

Multiplying by the cell value of f, summing, and using $N/q = 1/\lambda$ proves the identity. Also

$$\eta_\sigma(q) = e^{-\sigma \mathrm{i} x}\frac{x}{\sin x}.$$

The inequalities follow from $\sin x \geq x - x^3/6$, $|e^{-\sigma \mathrm{i} x}-1| \leq x$, and the triangle inequality. □

This lemma applies simultaneously to every retained bit-plane: their dyadic transitions, as well as the carry shift r/N, lie exactly on the N^{-1} grid. Only the endpoint h/N differs from $\lambda/2$, and only for odd q, by $1/(2N)$.

For the ten-term interval certificate, let M_A be the truncated phase-aware margin and let $\mathcal{U}_A, \mathcal{C}_A$ bound the sums of the unsigned and carry bit-plane magnitudes. The 160-bit interval computation establishes

$$\begin{aligned} M_A &\geq 0.26234884833948441617, \\ \mathcal{U}_A &\leq 0.65715489659645633631, \\ \mathcal{C}_A &\leq 0.72378995228602846574. \end{aligned}$$

For a bit band $m \geq 17$, put $N = 2^m$,

$$x_m = \frac{\pi}{2^{m-1}}, \quad a_m = \frac{1}{1-x_m^2/6}, \quad e_m = x_m + \frac{x_m^2}{6-x_m^2},$$

and set

$$T_U = (\pi + \tfrac{1}{2})2^{-9}, \qquad T_C = (2\pi+1)2^{-9}.$$

Lemma 6.1, the odd endpoint bound, and the finite Abel tails below imply, uniformly over the entire band,

$$\begin{aligned} D_3(q) - 2U(q) \geq M_A &- 2(e_m\mathcal{U}_A + T_U) \\ &- \big[e_m(\mathcal{C}_A + 10/N) + 10/N + T_C\big] \\ &- 3\big[e_m(1/\pi + 1/N) + 1/N\big]. \end{aligned}$$

For weighted HW, the corresponding perturbation denominator is at most

$$B_m = 3a_m(1/\pi + 1/N) + a_m(\mathcal{C}_A + 10/N) + T_C + 2(a_m\mathcal{U}_A + T_U).$$

At $m = 17$, these expressions evaluate respectively to

$$\begin{aligned} D_3(q) - 2U(q) &\geq 0.23365624115347768, \\ B_{17} &\leq 3.0215785299117311, \\ \frac{0.23365624115347768}{B_{17}} &\geq 0.07732919692155194. \end{aligned}$$

Every variable error term and B_m is nonincreasing with m, so this first analytic band is the global analytic minimum. A separate direct oracle checks the exact quadrature identity for every bit-plane of every $3 \leq q \leq 257$, with maximum complex error below $6.83 \cdot 10^{-15}$.

6.2 Bounded-variation fallback

Lemma 6.2 (Bounded-variation grid error)**.** *If $F : [0, L] \to \mathbb{C}$ has bounded variation and $J = \lfloor NL \rfloor$, then*

$$\left| \frac{1}{N} \sum_{n=0}^{J-1} F(n/N) - \int_0^{J/N} F(t)\, dt \right| \le \frac{\operatorname{Var}_{[0,J/N]}(F)}{N}.$$

Moreover, replacing J/N by L costs at most $\|F\|_\infty/N$.

Proof. On each grid cell, the difference between the left endpoint value and the cell average is at most the oscillation of F on that cell. Multiplying by the cell width and summing bounds the error by $N^{-1} \operatorname{Var}(F)$. The final interval has length less than N^{-1}. □

High-bit Riemann errors. Let $N = 2^m$, $\lambda = q/N$, $p_a = 2^{1-a}$, and $\phi_\lambda(t) = e^{-2\pi i t/\lambda}$. For the high bit $j = m - a$, the normalized finite unsigned sum is the left Riemann sum for

$$U_a(\lambda) = \lambda^{-1} \int_0^\lambda b_a(t)\phi_\lambda(t)\, dt.$$

On $[0, \lambda]$, b_a has at most $2\lambda/p_a + 2$ unit jumps and $\operatorname{Var}(\phi_\lambda) = 2\pi$. The product-variation inequality and $\lambda \ge 1/2$ therefore imply

$$\begin{aligned} |U_a^{(m)} - U_a| &\le \frac{2\lambda/p_a + 2 + 2\pi}{\lambda N} \\ &\le \frac{2^a + 4 + 4\pi}{2^m}. \end{aligned}$$

Summing over $a = 1, \dots, A$ produces

$$R_U(m, A) = \frac{2^{A+1} - 2 + A(4 + 4\pi)}{2^m}.$$

For the carry term, put

$$d_{a,\lambda}(t) = b_a(t + 1 - \lambda) - b_a(t).$$

The finite sum ends at h/N, where $h = \lfloor q/2 \rfloor$, while the integral ends at $\lambda/2$; their distance is at most $1/(2N)$. On this half interval, each square wave has at most $\lambda/p_a + 2$ jumps. Hence

$$\operatorname{Var}(d_{a,\lambda}) \le 2\lambda/p_a + 4, \quad \|d_{a,\lambda}\|_\infty \le 1, \quad \operatorname{Var}(e^{2\pi i t/\lambda}) = \pi.$$

Lemma 6.2, the product-variation inequality, and endpoint completion now imply the explicit per-bit estimate

$$\begin{aligned} |C_a^{(m)} - C_a| &\le \frac{2\lambda/p_a + 4 + \pi + 1/2}{\lambda N} \\ &\le \frac{2^a + 9 + 2\pi}{2^m}. \end{aligned}$$

Consequently,

$$R_C(m, A) = \frac{2^{A+1} - 2 + A(9 + 2\pi)}{2^m}.$$

Finite low-bit tails. For a zero-mean P-periodic sequence g_n, let $G_t = \sum_{n=0}^t g_n$ and $z = e^{-2\pi i/q}$. Abel summation establishes, for every $J \le q$,

$$\left| \sum_{n=0}^{J-1} g_n z^n \right| \le \|G\|_\infty \big(1 + J|1 - z|\big).$$

For a centered bit square wave, $\|G\|_\infty \le P/4$; for the difference of a shifted and an unshifted bit, $\|G\|_\infty \le P/2$. Since $|1 - z| \le 2\pi/q$, $P = N2^{1-a}$, and $N/q \le 2$, normalization and summation over $a > A$ yield

$$T_U(A) = (\pi + \tfrac{1}{2})2^{1-A}, \qquad T_C(A) = (2\pi + 1)2^{1-A}.$$

Step and total error. Comparing the finite half sum with $\int_0^{1/2} e^{2\pi i t}\,dt$, the left-Riemann error is at most $\pi/(2q)$, and the unmatched endpoint interval has length at most $1/(2q)$. Since $q > 2^{m-1}$,

$$|S_q/q - \mathrm{i}/\pi| \le (\pi+1)2^{-m}.$$

As an independent, coarser transfer that does not use grid alignment, combining a common continuous truncation with its rigorous mesh lower bound produces

$$D_3(q) - 2U(q) \ge \inf G_3 - [R_C + T_C] - 2[R_U + T_U] - 3(\pi+1)2^{-m}.$$

At $m = 17, A_U = 9, A_C = 10$, the right-hand side is 0.16261983851311484, already positive; the exact dyadic transfer above raises the certified margin to 0.23365624115347768.

7 Alternating Carry Obstruction

For odd m, set

$$q_m = (2^{m+1} - 1)/3, \quad r_m = (2^m + 1)/3, \quad h_m = r_m - 1.$$

Under $n = 2^m t$, $q_m/2^m \to 2/3$, $r_m/2^m \to 1/3$, and $h_m/2^m \to 1/3$. For each fixed high-bit offset a, Riemann convergence away from finitely many dyadic discontinuities implies

$$\Gamma_a = \frac{3}{2}\int_0^{1/3} (b_a(t+1/3) - b_a(t))e^{3\pi \mathrm{i} t}\,dt.$$

The zero-mean primitive argument establishes

$$|\Gamma_a| \le \tfrac{3}{4}(1+\pi)2^{1-a}, \qquad \sum_{a>A} |\Gamma_a| \le \tfrac{3}{4}(1+\pi)2^{1-A}.$$

The bound is summable, so finite-prefix convergence followed by dominated tail control proves

$$\lim_{\substack{m\to\infty \\ m \text{ odd}}} \frac{|C_{q_m,r_m}(1)|}{q_m} = \left|\sum_{a\ge 1} \Gamma_a\right| = 0.35933157036826014648\ldots.$$

8 Exact Weighted-Radius Proof

Proof of Theorem 1.26. Let

$$\mathcal{B} = \{a \in \mathbb{R}^\beta : D_a \le 2U_a\}.$$

This set is closed and nonempty because it contains zero. Its distance from $\mathbf{1}$ is attained: any minimizing sequence lies in a common compact ℓ_∞ box, and a convergent subsequence has its limit in $\mathcal{B}$. A nearest point cannot satisfy a strict inequality. Indeed, on the segment from $\mathbf{1}$, where $D_a - 2U_a > 0$, to such a point, continuity would produce an equality point at a strictly smaller distance.

Put $A = \sum_j a_j\delta_j$ and $B = \sum_j a_j u_j$. The equality $|A| = 2|B|$ holds exactly when

$$A = 2e^{\mathrm{i}\phi}B$$

for some ϕ. If $B = 0$, equality also forces $A = 0$, and every phase works, so no simultaneous-zero case is lost. The equality boundary is therefore the union over $\phi \in [0, 2\pi)$ of

$$H_\phi = \left\{a \in \mathbb{R}^\beta : \sum_j v_j(\phi)a_j = 0\right\}.$$

Writing $a = \mathbf{1} + e$, its fixed-phase distance is

$$\rho_\phi = \min\left\{\|e\|_\infty : \sum_j v_j(\phi)e_j = -V(\phi)\right\}. \tag{47}$$

Consider the centrally symmetric zonotope

$$Z_\phi = \sum_j [-v_j(\phi), v_j(\phi)] \subset \mathbb{R}^2.$$

Equation (47) is the gauge of $-V$ in Z_ϕ. If Z is compact, convex, and centrally symmetric, and $b \in \operatorname{span} Z$, then

$$\inf\{t \geq 0 : b \in tZ\} = \sup_{y: h_Z(y)>0} \frac{|\langle y, b\rangle|}{h_Z(y)}. \tag{48}$$

If $b \in tZ$, the support inequality implies $|\langle y, b\rangle| \leq t h_Z(y)$, proving one direction. Conversely, if t is below the gauge, $b \notin tZ$. Strict finite-dimensional separation supplies $y \neq 0$ with $\langle y, b\rangle > t h_Z(y)$, after reversing y if necessary. Taking the supremum and then all t below the gauge proves Equation (48).

Here

$$h_{Z_\phi}(y) = \sum_j |\Re(\bar{y} v_j(\phi))|.$$

Applying Equation (48) and then minimizing over the phase proves Equation (11).

For the finite formula, suppose the generators span $\mathbb{R}^2$. Quotient the nonzero dual directions by real scaling and partition the resulting projective circle at directions perpendicular to a nonzero v_j. On each open cell, every sign in the denominator is fixed, so the denominator is one real linear form. Subdivide once at a possible zero of the numerator. On the resulting cell the quotient has the form

$$R(\theta) = \frac{\langle y(\theta), p\rangle}{\langle y(\theta), w\rangle}, \qquad y(\theta) = (\cos\theta, \sin\theta), \quad \langle y(\theta), w\rangle > 0.$$

Direct differentiation shows

$$R'(\theta) = \frac{\det(p, w)}{\langle y(\theta), w\rangle^2}.$$

Thus R is monotone or constant and its maximum occurs at a cell endpoint. At such an endpoint $y = \mathrm{i}v_j$, up to a nonzero real scalar. Substitution produces exactly Equation (12).

An explicit primal witness follows by setting $n_j = [v_j, V]$, $h_j = \sum_\ell |[v_j, v_\ell]|$, and $r_j = |n_j|/h_j$. For every $\ell \neq j$ with $[v_j, v_\ell] \neq 0$, set

$$e_\ell = -r_j \operatorname{sgn}(n_j) \operatorname{sgn}([v_j, v_\ell]),$$

and set zero-projection coordinates to zero. Then $[v_j, -V - \sum_{\ell \neq j} e_\ell v_\ell] = 0$, so the residual is a real multiple $e_j v_j$. Choosing this e_j solves $\sum_\ell e_\ell v_\ell = -V$. When j is the sole zero-projection coordinate and its branch is maximal, gauge optimality implies $|e_j| \leq r_j$. With additional parallel coordinates, distribute the residual among them using any attained gauge minimizer. The former case holds at both certified NIST phases and is the supporting-face vector enclosed by the interval certificate.

If the generators have real rank one, write $v_j = c_j \xi$. Every feasible perturbation satisfies

$$\left|\sum_j c_j\right| = \left|\sum_j c_j e_j\right| \leq \|e\|_\infty \sum_j |c_j|.$$

Equality is attained by assigning each nonzero coordinate the corresponding signed face value. Rank zero would imply $\sum_j \delta_j = 2e^{\mathrm{i}\phi} \sum_j u_j$, contradicting strict nominal separation. These arguments also show attainment of the primal and dual values. The open-box, equality, and larger-box statements now follow from the definition of the nearest bad point.

For a calibrated center c and anisotropic radii r_j, write $a_j = c_j + r_j e_j$. Repeating the zonotope proof with generators $r_j v_j$ establishes the stated formula. For an ℓ_p ball, the support function of its linear image is the ℓ_{p^*} norm of the projected generators by Hölder duality. $\square$

Proof of Proposition 1.27. Expanding $v_j = \delta_j - 2e^{\mathrm{i}\phi} u_j$ produces

$$\begin{aligned}
A_{j\ell} &= [\delta_j, \delta_\ell] + 4[u_j, u_\ell], \\
B_{j\ell} &= -2\big([\delta_j, u_\ell] + [u_j, \delta_\ell]\big), \\
C_{j\ell} &= 2\big(\Re(\bar{u}_j \delta_\ell) - \Re(\bar{\delta}_j u_\ell)\big).
\end{aligned}$$

With $t = \tan(\phi/2)$, define

$$P_{j\ell}(t) = (A_{j\ell} + B_{j\ell}) + 2C_{j\ell}t + (A_{j\ell} - B_{j\ell})t^2.$$

The determinant is $P_{j\ell}(t)/(1+t^2)$; the common positive factor cancels from every branch ratio. Also $P_{jV} = \sum_\ell P_{j\ell}$. Partition the real projective line at all real roots of the nonzero polynomials $P_{j\ell}$, P_{jV}, and the rank-drop minors. On an open cell every absolute-value sign is fixed, so an active branch is a ratio of two quadratic polynomials.

An interior minimum of their pointwise maximum is either a stationary point of one active branch or a crossing of two active branches; all other minima occur at the partition boundary or a rank change. For a ratio of quadratics, the cubic derivative terms cancel, leaving degree at most two. Equality of two ratios produces degree at most four after cross multiplication. The point $t = \infty$, corresponding to $\phi = \pi$, is included projectively. Identically zero minors are handled by the rank-one stratum already proved. This yields the stated finite candidate set. $\square$

Proof of Corollary 1.28. For $q = 2^m$, every unsigned bit below $m - 1$ completes an integer number of periods and has zero first coefficient. Let $H \neq 0$ be the coefficient of bit $m - 1$, the half-interval indicator. The carry vanishes, and each of the s sign-extension planes has coefficient H. Hence

$$U_a = |Ha_{m-1}|, \qquad D_a = \left| H \sum_{j=m}^{m+s-1} a_j \right|.$$

For $s \leq 2$, the nominal vector is already bad. For $s \geq 3$ and $\|a - \mathbf{1}\|_\infty \leq \varepsilon < 1$,

$$D_a \geq |H|s(1 - \varepsilon), \qquad 2U_a \leq 2|H|(1 + \varepsilon).$$

Separation is therefore strict for $\varepsilon < (s - 2)/(s + 2)$. At equality, choose

$$a_{m-1} = \frac{2s}{s + 2}, \qquad a_m = \cdots = a_{m+s-1} = \frac{4}{s + 2}.$$

Then $D_a = 2U_a$, proving both optimality and the equality classification. $\square$

9 All-Mode Optimality Proof

Proof of Theorem 1.39. Conditioned on $S = s$, exactly q^{d-1} share tuples are compatible. Expanding T, every term containing a noise factor has zero conditional expectation. Therefore

$$\mu = q^{1-d}(\alpha_0 \bar{w}_0) * \cdots * (\alpha_{d-1} \bar{w}_{d-1}),$$

and Fourier transformation establishes the first claimed spectrum. Conditioned on a fixed tuple, independence across shares implies

$$\mathbb{E}[T^2 \mid X_0 = x_0, \ldots, X_{d-1} = x_{d-1}] = \prod_i v_i(x_i).$$

Averaging the compatible tuples produces $\nu = q^{1-d} v_0 * \cdots * v_{d-1}$, proving the second spectrum. Averaging without conditioning also produces

$$\mathbb{E}T^2 = \prod_i (\alpha_i^2 \operatorname{Var}(w_i) + \sigma_i^2).$$

For every predictor g, $\mathbb{E}[g(S)T] = \langle g, \mu \rangle_q$, and $\mathbb{E}\mu = 0$. Cauchy–Schwarz shows that a mean-zero unit predictor has absolute numerator at most $\|\mu\|_2$, with equality exactly for $g = \pm\mu/\|\mu\|_2$ when $\mu \neq 0$. When $\mu = 0$, every numerator is zero. Parseval for the unnormalized DFT implies

$$\|\mu\|_2^2 = q^{-2d} \sum_{k \neq 0} \prod_i \alpha_i^2 |\widehat{w}_i(k)|^2,$$

and division by $\mathbb{E}T^2$ proves the correlation formula.

Since μ is real, frequencies k and $q - k$ form one real harmonic orbit. Parseval assigns twice the single-frequency energy to this orbit, except at the self-conjugate even Nyquist mode. Dividing the orbit energy by the total nonzero-mode energy proves the formula for η_k. The mode-one projection is the entire predictor exactly when all product coefficients outside its conjugate orbit vanish. $\square$

Proof of Theorem 1.40. Conditional Jensen implies $\mu(s)^2 \le \nu(s)$, so $\nu(s) = 0$ implies $\mu(s) = 0$. Weighted Cauchy–Schwarz establishes

$$\langle g, \mu\rangle_q^2 \le \langle g^2, \nu\rangle_q \left\langle \frac{\mu^2}{\nu}, 1 \right\rangle_q ,$$

with equality for $g \propto \mu/\nu$ on the support of ν. This proves the first optimality formula and also $0 \le \Lambda_* \le 1$. When $\mu = 0$, both the optimum and every score numerator are zero, giving the stated tie classification.

The correlation direction is μ, whereas the unconstrained power direction is μ/ν. Conditional Jensen makes $\nu > 0$ wherever $\mu \neq 0$; hence these directions are proportional if and only if ν is constant on the support of μ.

For the mean-zero constraint, use the weighted inner product $(f, g)_\nu = \langle fg, \nu\rangle_q$. The objective vector is μ/ν, while $\langle g, 1\rangle_q = 0$ is orthogonality to $1/\nu$. Projection produces $g \propto (\mu - c)/\nu$ with the stated c, and the squared norm of that projection is

$$\mathbb{E}[\mu^2/\nu] - \frac{\mathbb{E}[\mu/\nu]^2}{\mathbb{E}[1/\nu]}.$$

Orient the unconstrained optimum so $M = \mathbb{E}Y_g > 0$, and put $B = \mathbb{E}Y_g^2$. The mean of n independent scores has variance $(B - M^2)/n$. Cantelli's inequality therefore implies

$$\Pr[\bar{Y}_n \le 0] \le \frac{(B - M^2)/n}{(B - M^2)/n + M^2} = \frac{1 - \Lambda_*}{1 + (n-1)\Lambda_*}.$$

Solving this inequality for a target η determines the stated integer sample bound.

Under the matched null, choose $t \in (0, M)$; the score has mean zero and second moment B, so Cantelli bounds type I by $B/(B + nt^2)$. Under the alternative, type II is bounded by $(B - M^2)/(B - M^2 + n(M - t)^2)$. Thus a common threshold exists if

$$M\sqrt{n} \ge \sqrt{B\frac{1-\eta_0}{\eta_0}} + \sqrt{(B - M^2)\frac{1-\eta_1}{\eta_1}}.$$

Substitution of $M^2/B = \Lambda_*$ establishes Equation (38). Its right-hand side is strictly decreasing in Λ_*, so the power-optimal score also minimizes this moment-only testing requirement. Cantelli is sharp over distributions specified only through these first two moments.

After replacing N_i by τN_i,

$$v_{i,\tau}(x) = \tau^2\sigma_i^2 \left(1 + \frac{\alpha_i^2 \bar{w}_i(x)^2}{\tau^2\sigma_i^2}\right).$$

Because $\mathbb{Z}_q$ is finite, multiplication and conditional averaging produce, uniformly in s, $\nu_\tau(s) = \tau^{2d}\prod_i \sigma_i^2(1 + O(\tau^{-2}))$. The conditional mean is independent of τ. Substituting the Parseval formula for $\|\mu\|_2^2$ proves the expansion for $\Lambda_*(\tau)$. Every fixed-deflection or fixed-Cantelli target is asymptotic to a common constant times $1/\Lambda_*$; taking the centered/unsigned ratio proves Equation (39). $\square$