BEYOND ALL-OR-NOTHING CORRUPTION: ADDITIVE SECRET SHARING UNDER HAMMING-WEIGHT LEAKAGE

ANIRUDDHA BISWAS, JIHUN HWANG, HEMANTA K. MAJI, ILYA D. SHKREDOV, AND XIUYU YE

Abstract. Side-channel attacks threaten secret sharing, as well as the threshold protocols and masking defenses built on it. Motivated by power analysis, Faust et al. (EUROCRYPT 2024) initiated the theoretical study of additive secret sharing over the cyclic group $\mathbb{Z}/N\mathbb{Z}$ when each share leaks its Hamming weight. They asked whether these weights, taken together, reveal vanishingly little about the secret as the modulus grows.

For Mersenne $N = 2^n - 1$ with prime exponent n, they proved insecurity $\ll_k n^{-k/4+1}$, establishing security for $k \geqslant 5$ shares. Two shares were known to be insecure there, but the practically important three- and four-share cases remained open. So did non-Mersenne moduli, including the dyadic modulus $N = 2^n$, and the question of how quickly insecurity could decay.

We prove security for every near-dyadic modulus $|N - 2^n| \leqslant 2^n/n$ and every fixed $k \geqslant 3$, with insecurity $\ll_k n^{-k/2+5/4}$. This gives the first security guarantees for three and four shares and for non-Mersenne moduli, including the dyadic moduli. The leading share-dependent term improves to $n^{-k/2}$, and our lower bounds show that this dependence on k is optimal. Roughly speaking, doubling the share length buys $k/2$ additional bits of statistical security, twice the $k/4$ guaranteed before. At this scale, a fixed target insecurity requires only the square root of the share length prescribed by the earlier analysis.

Our proof uses Fourier energy to bound insecurity, contracting it by $n^{-1/2}$ per share. We identify Bernoulli dynamics as the broader mechanism underlying Faust et al.'s exact binary rotation phenomenon and extend the analysis to near-dyadic moduli. For the lower bounds, we identify Fourier modes that retain information about the secret. The estimates rely on spectral estimates obtained from generating functions and Cauchy's coefficient formula.

Contents

1. Introduction

Secret sharing hides a secret by splitting it into shares. Its classical privacy model is all-or-nothing: too few shares reveal nothing, provided the rest remain hidden [Sha79, Bla79]. This principle underlies threshold cryptography and many masking countermeasures [ISW03, NRR06, NRS11]. Physical leakage need not follow these rules. Even coarse *side-channel* measurements of individual shares can, when taken together, reveal information about the secret [Koc96, KJJ99, Mes00, BCO04]. When an attacker learns a little about every share, rather than everything about a few, *how much secrecy survives?*

Motivated by timing and power-analysis attacks, Faust et al. [FMM+24] studied this question for *additive sharing* under *Hamming-weight* leakage. Here, the shares sum to the secret modulo N, and the attacker learns only the number of ones in each share's binary representation. They showed that, for moduli $N = 2^n - 1$ with prime exponent n, five or more shares suffice: the leaked weights together reveal vanishingly little about the secret as n grows. Yet basic questions remained unresolved: do three or four shares suffice, does security extend beyond prime-exponent Mersenne moduli, and how quickly does increasing share size hide the secret?

Neither five shares nor Mersenne arithmetic is necessary. We prove, for moduli near powers of two, three shares suffice; two do not. For n-bit shares, every additional share actually adds roughly $\frac{1}{2}\log_2 n$ bits of security against Hamming-weight leakage, twice what was previously known.

Leakage experiment. We work with the numbers $0, 1, \ldots, N-1$, adding modulo N; this is the cyclic group $G = \mathbb{Z}/N\mathbb{Z}$. Write $\mathrm{wt}(x)$ for the number of ones in the binary representation of x. We use n for the approximate bit length of each share.[1] To share a secret $s \in G$ into $k \geqslant 2$ shares, sample $(s_1, s_2, \ldots, s_k) \in G^k$ uniformly subject to

$$s_1 + s_2 + \cdots + s_k = s.$$

The adversary receives only the leakage vector

$$\Big(\mathrm{wt}(s_1), \mathrm{wt}(s_2), \ldots, \mathrm{wt}(s_k) \Big).$$

Let $D(s)$ denote this leakage distribution, and $D(U)$ its counterpart when the secret s is uniform over G. Following Faust et al. [FMM+24] and the *local leakage-resilience* framework [BDIR18, BDIR21], we measure their difference by total variation distance:

$$E_k(s) \;\coloneqq\; \Big\| D(s) - D(U) \Big\|_{\mathrm{TV}} \;\coloneqq\; \frac{1}{2} \sum_{\mathbf{w} \in W^k} \Big| D(s)(\mathbf{w}) - D(U)(\mathbf{w}) \Big|, \text{ and} \tag{1}$$

$$E_k \;\coloneqq\; \max_{s \in G} E_k(s), \tag{2}$$

where $W \coloneqq \mathrm{wt}(G)$ is the set of possible Hamming weights.

The *insecurity* E_k bounds the distinguishing advantage of every statistical test. For fixed k, *asymptotic security* means $E_k \to 0$ as the modulus grows. Conversely, an infinite family where E_k stays bounded away from zero is *insecure*.

Why coarse leakage need not be harmless. Revealing just the least significant bit of each share can leave a non-vanishing distinguishing advantage even as the modulus grows [AMN+21, MNP+22, FMM+24]. Hamming weight is not harmless by default either. In an $\mathbb{F}_2$ vector space, where shares are added bit by bit without carries:

$$\mathrm{wt}(s) \equiv \sum_{j=1}^{k} \mathrm{wt}(s_i) \pmod 2.$$

[1] Precisely, 2^n denotes a power of two nearest to N. The fixed-width binary representation uses $\lceil \log_2 N \rceil$ bits: n when $N \leqslant 2^n$, and $n+1$ when $N > 2^n$.

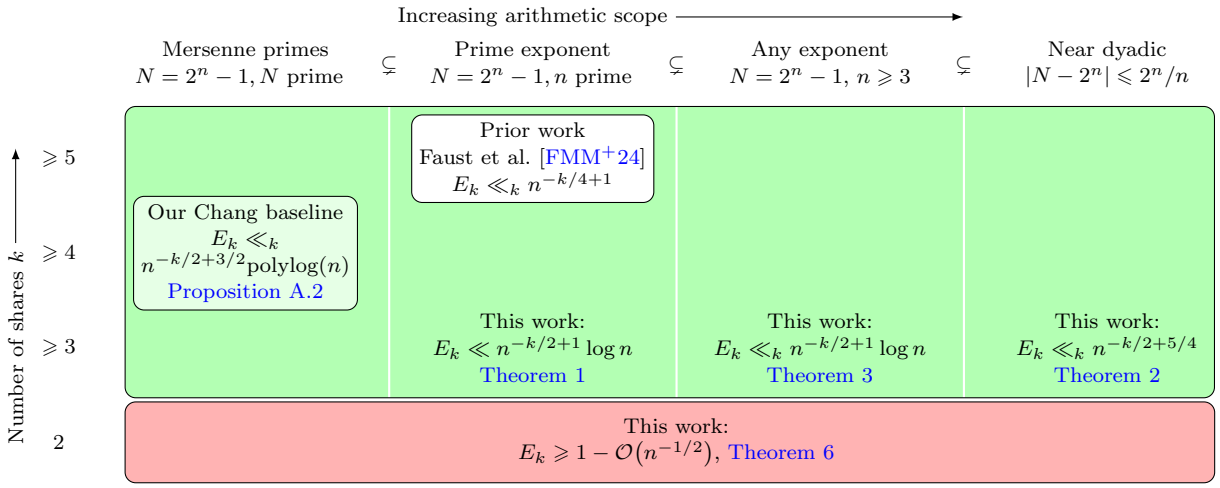

FIGURE 1. **Our contribution: the threshold & tight bounds for leakage resilience.** Columns list modulus families in increasing generality; rows give the number of shares k. Each displayed upper bound also holds in every cell above and to its left. The two-share lower bound holds across all columns. Green indicates asymptotic security; red indicates insecurity. Here $\ll$ suppresses constant factors, and $\ll_k$ denotes a constant exponential in k; see Section 2.

The leaked weights therefore reveal whether the secret has an even or an odd number of ones, regardless of the number or length of the shares.

Carries, while adding in $\mathbb{Z}/N\mathbb{Z}$, break this identity in general, but the least-significant-bit example shows that carries alone do not ensure security. Unlike leakage at a fixed bit position, Hamming weight conceals the locations of the ones, and hence where carries start and how far they propagate. Turning this positional uncertainty into a quantitative security guarantee is the central challenge.

Prior work and the gap. Faust et al. [FMM+24] proved:

$$E_k \ll_k n^{-k/4+1}, \qquad N = 2^n - 1,\ n \text{ prime}.$$

Here $\ll$ and $\gg$ suppress constant factors, and an additional subscript k denotes a constant (at most exponential) depending on k (see Section 2).

Faust et al.'s bound vanishes only for $k \geqslant 5$ shares. Their proof relied on a special arithmetic symmetry: multiplication by two modulo $2^n - 1$ cyclically rotates the binary digits and preserves Hamming weight. Three questions about security remained open.

(1) *How many shares suffice?* Two shares are insecure: for secret 0 over a Mersenne modulus, their weights almost always sum to n, a pattern rarely seen for independent shares. Do three or four shares suffice for asymptotic security?
(2) *Is Mersenne arithmetic essential?* Does security extend to other moduli, especially the dyadic group $\mathbb{Z}/2^n\mathbb{Z}$? In the dyadic group, addition is ordinary n-bit addition with overflow discarded, a natural fit for digital hardware.
(3) *How fast does insecurity decay?* After leakage, each share retains about $(n - \log_2 n)$ bits of entropy, in the sense of [DRS04, DORS08]). Can additive sharing achieve exponential decay in n, as recently proved for *inner-product* sharing [FMM+25], or is polynomial decay unavoidable? In the latter case, what is the optimal dependence of the exponent on k?

1.1. **Our Contributions.** For fixed k and $n \to \infty$, we establish the sharp near-dyadic security threshold and prove upper and lower bounds that match in their leading k-dependent exponent, all without primality assumptions on either N or n; see Figure 1.

Security threshold (Theorem 2 and Theorem 6)

For every modulus satisfying $|N - 2^n| \leqslant 2^n/n$ and every fixed integer $k \geqslant 3$,

$$E_k \ll_k n^{-k/2+5/4}, \qquad \text{whereas} \qquad E_2 = 1 - \mathcal{O}\Big(n^{-1/2}\Big).$$

In particular,

$$E_3 \ll n^{-1/4} \qquad \text{and} \qquad E_4 \ll n^{-3/4},$$

including the dyadic group $\mathbb{Z}/2^n\mathbb{Z}$. The failure of two shares is more general: even outside the near-dyadic window, Theorem 6 gives

$$E_2 \geqslant \frac{2}{3} - \mathcal{O}\Big((\log N)^{-1/2}\Big).$$

Thus, increasing the modulus cannot make two shares asymptotically secure. Table 1 and Table 2 collect our upper and lower bounds, respectively.

Longer shares help, but only polynomially. For fixed $k \geqslant 3$, Theorem 5 complements the upper bound above with

$$E_k \gg_k n^{-k/2}\sqrt{\log n} \qquad (\text{if } |N - 2^n| \leqslant 2^n/n)$$
$$E_k \gg_k \sqrt{\varepsilon} \cdot n^{-k/2+1/2} \qquad (\text{if } |N - 2^n| \leqslant 2^{(1-\varepsilon)n}, \quad 0 < \varepsilon < 1 \text{ fixed.})$$

The first estimate follows by taking $m = n - \log_2 n$ in the finite-parameter statement of Theorem 5, which gives $L = (\frac{1}{2}\log_2 n - C)_+$. Thus the estimate holds for sufficiently large n throughout the near-dyadic window.

Mersenne arithmetic sharpens the bounds. For $N = 2^n - 1$, Theorem 3 gives

$$E_k \ll_k n^{-k/2+1}\log n.$$

Neither n nor N need be prime. For even $k \geqslant 4$, Theorem 4 gives

$$E_k(0) \gg_k n^{-k/2+1},$$

matching the upper bound up to $\log n$.[2] These estimates determine the optimal coefficient $1/2$ for k in the decay exponent, but not the entire exponent for every modulus: at $N = 2^n$, the remaining gap is $n^{3/4}$.

TABLE 1. Our upper bounds on the insecurity E_k, for $k \geqslant 3$.

Parameter	$\boldsymbol{E_k}$ **upper bound**	**Theorem**
$\lvert N - 2^n\rvert \leqslant 2^n/n$	$(cn)^{-k/2+5/4}$	Theorem 2
$N = 2^n - 1$, $n \geqslant 3$	$(cn)^{-k/2+1}\log n$	Theorem 3
$N = 2^n - 1$, $n \geqslant 3$ prime	$c\,n^{-k/2+1}\log n$	Theorem 1

[2]When $n \geqslant 3$ is prime, Theorem 1 gives the same upper bound with an absolute implied constant. This strengthens the dependence on k, not the power of n; the modulus N may still be composite. Both Mersenne refinements use the fourth-moment estimate recorded in Imported Lemma 3.5.

TABLE 2. Our lower bounds on the insecurity E_k.

Parameter	E_k lower bound	Theorem
$\lvert N-2^n\rvert \leqslant 2^{(1-\varepsilon)n}$, $0<\varepsilon<1$	$\varepsilon^{1/2}(dn)^{-k/2+1/2}$	Theorem 5
$N=2^n-1$, $n\geqslant 3$, and even $k\geqslant 4$	$(dn)^{-k/2+1}$	Theorem 4
$\lvert N-2^n\rvert=\varepsilon\cdot N$ and $k=2$	$1-\varepsilon-d\cdot n^{-1/2}\gtrsim 2/3$	Theorem 6

Remark 1.1 (Leaky preprocessed one-time pads)**.** Our analysis also applies to an *online-offline random self-reduction pattern* recurring in cryptography. For example, establish independent uniform one-time-pad keys $s_1,\dots,s_{k-1}$ offline and encrypt s online by publishing the ciphertext $c=s-(s_1+\cdots+s_{k-1})$. Even with Hamming-weight leakage from the keys and full disclosure of c, we show that the insecurity remains $n^{-k/2+5/4}$ when $|N-2^n|\leqslant 2^n/n$. Roughly put, handing the ciphertext c to the adversary costs nothing beyond what Hamming weights of keys already reveal. The quantity Δ_{k-1} of Lemma 3.1 measures exactly this stronger adversary, and it is controlled by the same quantity that drives our upper bounds, which is bounded for all near-dyadics in Corollary 1. Remark 3.2 gives the information-theoretic explanation. △

Remark 1.2 (Asymptotic scope & concrete parameters)**.** Our techniques (for example, in Theorem 2) are tailored to a modular asymptotic analysis; the resulting constants are not optimized and may be large. Together, our upper and lower bounds establish the sharp security threshold and optimal leading dependence on the number of shares. These results provide benchmarks for future work on sharper, group-specific estimates; concrete security guarantees at practical parameters require an explicit finite-parameter analysis, which we do not provide here. △

1.2. Comparison with Prior Work and Broader Context.

The strongest prior baseline. Faust et al. [FMM+24] consider Mersenne moduli $N=2^n-1$, with prime exponent n. They used the binary-rotation property and the primality of the exponent n to prove asymptotic security for $k\geqslant 5$; see Proposition A.1. Against this baseline, our results settle the three- and four-share cases, extend security to the near-dyadic window, and improve the leading k-dependent exponent from $-k/4$ to $-k/2$.

By further restricting to Mersenne primes and explicitly leveraging the primality of N, we prove: four shares suffice, with the optimal leading exponent. This benchmark uses Chang's large spectrum bound [Cha02, Lemma 3.1, p. 17] (specifically, the version stated in [Shk11, Theorem 1.2] which we state here as Imported Lemma A.3). But this benchmark comes at a steep price: this modulus family may not even be infinite. Can we match it for general moduli? For all moduli near powers of two, we prove that we can. Our energy argument retains the leading exponent while proving security with just three shares. Appendix A.2 gives the precise comparison.

Inherent resilience of a standard scheme. Leakage-resilient cryptography addresses probing, noisy, bounded-output, and related leakage [ISW03, DDF14, NRS11, KR19]. One approach designs encodings or sharing schemes for this purpose [BPRW16, GK18a, KMS19, BS19, BIS19, SV19, GK18b, ADN+19, FY19, MSV20, CGG+20, HVW20, CKOS21, CKOS22, CG25, IS24]. We instead quantify the Hamming-weight resilience of unchanged standard additive sharing.

Benhamouda et al. [BDIR18, BDIR21] initiated the study of standard linear sharing under bounded-output local leakage; later work refined attacks and security bounds for additive sharing and analyzed physical-bit leakage in Shamir sharing [AMN+21, MNP+21, MNP+22, MNPCW22, MNPY24, FMM+24, HMNY25, Ngu25]. Hamming weight differs both in its growing output range and in its arithmetic interaction from leakage at fixed bit positions. Our guarantees concern this specific leakage function, not arbitrary local leakage or every physical side channel.

A learning-problem perspective. *Learning with Physical Rounding* (LWPR) brings the learning-problem viewpoint of LWE and LWR into the side-channel setting: it asks whether a secret key can be recovered from coarse physical measurements of linear computations involving it [DMMS21, HMM+23, HMRS26]. Our single-sum problem isolates a basic aspect of this broader challenge: what do coarse measurements reveal about a hidden linear relation? Quantifying leakage in this simpler setting provides an information-theoretic benchmark and a step toward understanding the arithmetic that drives LWPR.

Digital structure and carries. The problem connects to carries and binary digit sums in additive combinatorics [DF09, MR10, SW23]. Binary notation makes Hamming weight simple but addition nonlocal through carries; cyclic Fourier analysis simplifies addition but transfers the digital structure to doubling orbits. We use the latter viewpoint for uniform spectral estimates near powers of two; moduli far from powers of two remain outside our guarantees.

1.3. **High-level Technical Overview.** Upper bounds control the total leakage; lower bounds isolate Fourier modes that retain information about the secret.

1.3.1. *Upper bound strategy: a stronger attacker, a simpler question.* We begin by giving the attacker more information. Reveal the final share $s - Y$ fully, where $Y = s_1 + \cdots + s_{k-1}$, and reveal only the Hamming weights of the independent uniform first $(k-1)$ shares. It suffices that Y remain nearly uniform conditional on those weights, averaged over the leakage outcomes, not for every outcome separately. Data processing inequality then bounds the original leakage experiment (see Lemma 3.1 and Remark 3.2).

What each additional share contributes. Adding independent variables multiplies their Fourier coefficients. The *energy* F_m aggregates squared Fourier magnitudes over the leaked weights of m independent uniform shares and controls the average conditional distance of their sum from uniform. The reduction gives

$$2 \cdot E_k \leqslant F_{k-1}.$$

Crucially, this energy does not depend on the secret: one estimate controls all secrets.[3]

Estimating F_m directly is difficult, so we simplify the problem one share at a time without sacrificing the leading exponent in the upper bound. The key is to control the combined energy across leakage outcomes, rather than bound each Fourier coefficient separately.

For $N = 2^n - 1$ with prime exponent n, we prove

$$F_m \leqslant n^{-1/2} \cdot F_{m-1}, \qquad m \geqslant 2.$$

Each additional share thus reduces this upper bound by a factor of $\sqrt{n}$, explaining the coefficient $1/2$ of k in our decay exponent. For moduli near powers of two, Section 4 proves the analogous energy contraction with $\leqslant$ replaced by $\ll$; only an absolute constant is lost per share.

From exact rotations to nearby moduli. For Mersenne moduli, doubling a number cyclically rotates its binary digits without changing their Hamming weight. Faust et al. [FMM+24] used this symmetry. When n is prime, each nonzero frequency belongs to an n-element orbit of frequencies with equal Fourier coefficients. We use this repetition at the energy level to prove the contraction above.

For other moduli, this exact rotation disappears. We first work with the full n-bit range $\{0, 1, \ldots, 2^n - 1\}$ (the full-range model), keeping addition modulo N. A generating function organizes the calculation bit by bit. Each factor records the choice between a zero and a one; the

[3]The proxy F_m is also the V_2-correlation measure of Mojahedian et al. [MBG+19, Equation (9)] between the leakage vector and the aggregate sum; see Remark 3.2. In particular, this interpretation yields the leaky preprocessed one-time-pad application of Remark 1.1.

exponent of an auxiliary variable z counts the ones. The unnormalized Fourier sum over all weight-w n-bit strings at frequency t is

$$[z^w]\prod_{j=0}^{n-1}\left(1+z\cdot\exp\left(-2\pi\imath\cdot\frac{2^j t}{N}\right)\right), \tag{3}$$

where $[z^w]$ means "take the coefficient of z^w."

The contribution from the j-th bit depends only on $\{2^j t/N\}$, where $\{x\}$ denotes the fractional part of x. We track these fractions using the *Bernoulli doubling map* $T(x)=\{2x\}$, which shifts the binary expansion left by one place; see, for example, [EW11, Chapter 2]. Its successive iterates are

$$\{t/N\}\to\{2t/N\}\to\{2^2t/N\}\to\cdots.$$

For $N=2^n-1$, this sequence repeats after n steps, recovering the cyclic rotation above. For other moduli, this n-step repetition need not hold, but the estimates in Section 4.4 still give the same order of improvement per share.

The technical challenge: bounding the coefficients. The challenge is to turn Cauchy's coefficient-extraction formula for Equation 3 into estimates strong enough to yield this improvement. This is a Krawtchouk-type problem, but with a crucial difference. When binary strings are added without carries (by bit-wise XOR on $\{0,1\}^n$) the Fourier coefficients of a fixed-weight slice are given by Krawtchouk polynomials [Pol19, Section 2.2, Equations (2.7) and (2.17)]. Refer to Figure 2 for the discussion below. The coefficients fall into clear-cut levels: frequencies t with the same number of ones $\mathrm{wt}(t)$ have exactly the same coefficient, regardless of where those ones appear. In their generating function, each bit contributes either $(1+z)$ or $(1-z)$ factor to the product, unlike our case.

In our cyclic setting, those two signs are replaced by complex phases determined by t/N. The coefficient now depends on the full pattern of successive doublings $\{2^j t/N\}$, not just the number of ones in the frequency. Our estimates leave open whether a different, equally simple hierarchy of magnitudes exists.

For security, we only need a uniform separation from the zero-frequency coefficient $\binom{n}{w}$. Lemma 4.10 gives an $\mathcal{O}\big(n^{-1/2}\big)$ relative bound at every $t\neq 0$ when $|w-n/2|=\mathcal{O}(\sqrt{n})$. The exact nonzero maximum and its maximizers remain undetermined; in particular, we do not prove that $t=1$ attains it. Moreover, $t=1$ is not isolated in our experiments: $t=3$ carries comparable aggregate Fourier mass across the weights for Mersenne moduli (see Figure 6 for a visual illustration). Aggregating the weight-dependent coefficient bounds and controlling the binomial tails yields the energy contraction.

Finally, our transference lemma (Lemma 4.8) carries the bound from the full n-bit range back to the actual range $\{0,\dots,N-1\}$. The discrepancy $|N-2^n|$ controls the transference error.

When the iteration stops. The one-share estimate, $F_1\ll n^{1/4}$, and its full-range counterpart complete the basic argument. Starting with $k-1$ shares requires $k-2$ contractions. For prime-exponent Mersenne moduli, the calculation is:

$$\begin{aligned}
2\cdot E_k &\leqslant F_{k-1} && \text{(reduce to energy)},\\
&\leqslant n^{-1/2}\cdot F_{k-2}\leqslant\cdots\leqslant n^{-(k-2)/2}\cdot F_1 && \text{(iterate the contraction)},\\
&\ll n^{-(k-2)/2+1/4}=n^{-k/2+5/4} && \text{(use the one-share estimate)}.
\end{aligned}$$

The full-range argument and transfer step give the same exponent for near-dyadic moduli. For Mersenne moduli with exponent n, a separate fourth-moment estimate lets us stop instead at the stronger two-share energy estimate

$$F_2\ll n^{-1/2}\log n.$$

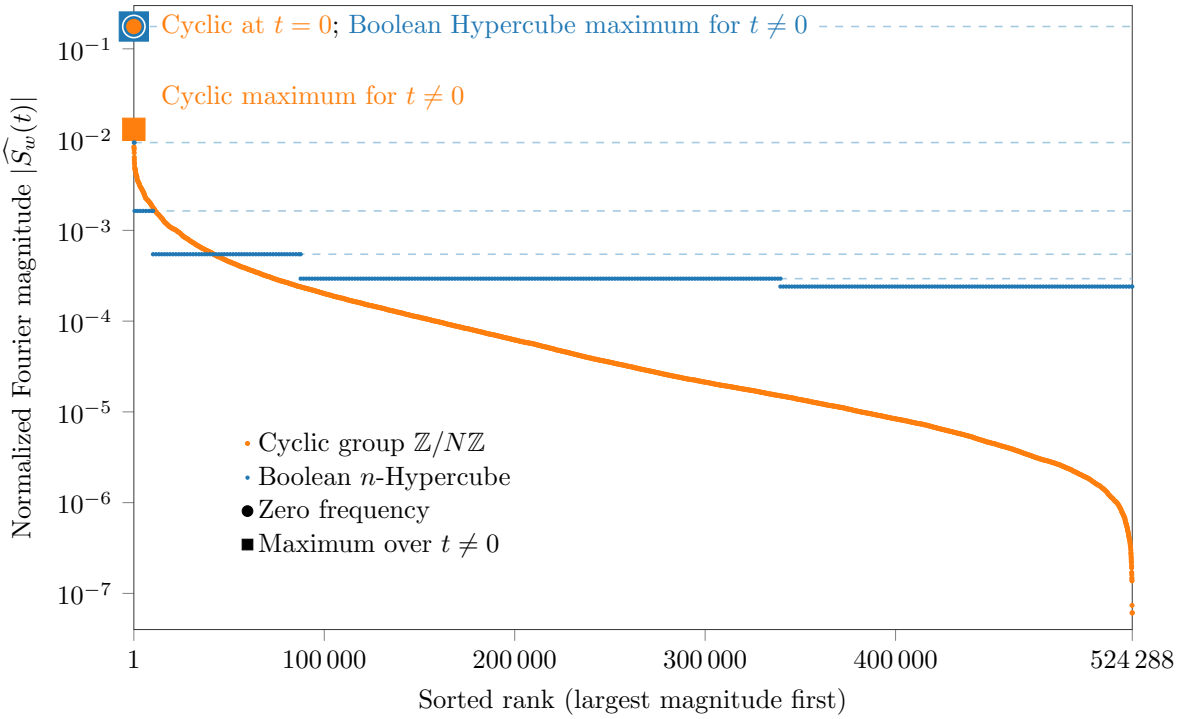

FIGURE 2. **One Hamming slice, two Fourier spectra.** Fourier magnitudes $\left|\widehat{S_w}(t)\right|$ sorted in descending order and plotted on a logarithmic vertical scale. Blue points represent the Boolean hypercube $(\{0,1\}^n, \oplus)$; orange points represent the cyclic group $\mathbb{Z}/N\mathbb{Z}$, where $N = 2^n - 1$. Here $w = \lfloor n/2 \rfloor$ and $n = 19$.

This gives the sharper Mersenne estimates, including the extension to arbitrary n in Section 4.5. The external fourth-moment input, Imported Lemma 3.5, is needed only for the sharper Mersenne bounds (Theorem 1 and Theorem 3), not for the general near-dyadic theorem (Theorem 2).

1.3.2. *Lower bound strategy: finding the information that survives.* Our lower bounds look for traces of the secret that remain in the leaked weights. The starting point is the *dual formulation of total variation distance*. Instead of adding absolute probability differences over leakage outcomes, we first multiply each difference $D(s)(\mathbf{w}) - D(U)(\mathbf{w})$ by a complex number $\phi_{\mathbf{w}}$ of magnitude one – a *phase* – and then add. For a fixed secret s, the largest possible magnitude of this sum is exactly $2E_k(s)$. Finding one sufficiently large weighted sum is therefore enough. We use this freedom to bring out Fourier frequencies that carry enough magnitude. The challenge is to ensure that their contributions do not disappear through cancellation.

The key frequencies are powers of two, up to a sign. For $N = 2^n - 1$, with an even number of shares $k \geqslant 4$, binary-rotation symmetry makes their contributions repeat. We choose these phases so that, at the secret $s = 0$, every frequency contributes a non-negative quantity. We can then keep only the doubling frequencies: the remaining terms cannot cancel them. Their combined contribution proves Theorem 4.

Near-dyadic moduli. For other moduli in our window around 2^n, this exact structure is no longer available. We instead fix the phases, square the magnitude of the resulting expression, and average over all secrets. Fourier orthogonality removes the cross-terms, leaving a sum of squared magnitudes. Different frequencies can no longer cancel one another in this average. Again, we identify a subset of the frequencies $\{1, 2, \ldots, 2^{n-1}\}$ that retain enough magnitude. We estimate their contribution using the bit-by-bit generating function introduced above. This proves Theorem 5: at

least one secret leaves a detectable trace. For a fixed number of shares, this rules out exponentially small worst-case insecurity as shares grow longer.

The two-share obstruction. is relatively straightforward. For any modulus N, let 2^n be a nearest power of two. For a suitable secret, the two shares are complementary binary strings with probability at least $2/3$; their leaked weights then sum to n. This weight-sum event is rare for independent uniform shares. Section 6.1 turns this contrast into a two-share lower bound for every modulus N, not just those near powers of two.

Paper organization. Section 2 fixes notation and records the mathematical background. Section 3 develops the Fourier-energy argument for prime-exponent Mersenne moduli, proving Theorem 1. Section 4 extends the argument to near-dyadic moduli using Bernoulli doubling and transference: Theorem 2 is proved in Section 4.3, with the key spectral estimate established in Section 4.4. Section 4.5 proves the sharper bound (Theorem 3). Section 5 proves the corresponding lower bound for an even number of shares (Theorem 4). Section 6 outlines the near-dyadic lower bound (Theorem 5), whose full proof appears in Appendix E, and Section 6.1 proves the two-share obstruction for arbitrary moduli (Theorem 6). Appendix A recalls the bound of Faust et al. [FMM+24] and presents our Chang-enhanced comparison. Appendix B, Appendix C, and Appendix D contain deferred auxiliary proofs; Appendix F presents numerical illustrations.

1.4. AI Disclosure: Use of AI Tools.

We used Google AI Ultra (Gemini) and ChatGPT Pro in preparing this work. Their uses are described below.

(1) ChatGPT suggested the phase $\phi = \exp(-\imath/\sqrt{n})$ in Theorem 5 and the dispersion lemma (Lemma C.4). It also provided a streamlined proof of the technical lemma Claim C.2, which we adapted and independently verified.
(2) ChatGPT supported numerical experiments and data generation (see Figure 2, Figure 3, Figure 4, Figure 5, and Figure 6).
(3) ChatGPT was used for English-language editing, including improvements to grammar, clarity, organization, and presentation.

All AI-generated suggestions were independently examined, revised, and incorporated at the authors' discretion. The authors verified the mathematical arguments and take full responsibility for the correctness and final presentation of the paper.

2. Preliminaries

2.1. Notations and Basic Definitions.

We collect the notation and Fourier conventions used throughout the paper here.

The ambient group and Hamming-weight slices. Retain $G = \mathbb{Z}/N\mathbb{Z}$, its canonical representatives, $\mathrm{wt}(\cdot)$, $W = \mathrm{wt}(G)$, and the nearest-power parameter n from Section 1. Define

$$W_{\mathrm{BIN}} \coloneqq \mathrm{wt}\big(\{0, 1, \ldots, 2^n - 1\}\big) = \{0, 1, \ldots, n\}.$$

Thus, W_{BIN} is the full n-bit weight range.[4] For each $w \in W$, the corresponding *Hamming weight slice* is

$$S_w \coloneqq \{x \in G : \mathrm{wt}(x) = w\}.$$

Bold symbols denote vectors. Thus, if $\mathbf{w} \in W^k$, then $\mathbf{w}_j$ denotes its j-th coordinate, where $j \in \{1, 2, \ldots, k\}$.

[4] In the near-dyadic regime considered here, $W = \{0, 1, \ldots, n-1\}$ when $N < 2^n$, whereas $W = W_{\mathrm{BIN}}$ when $N \geqslant 2^n$.

Fourier conventions. Write $\overline{z}$ for the *complex conjugate* of z and the modulus $|z| = \sqrt{z \cdot \overline{z}}$. A *phase* is a complex number of modulus one.

Define the standard additive character (real x)

$$\mathrm{e}_N(x) \;\coloneqq\; \exp\Big(2\pi\imath\frac{x}{N}\Big).$$

For $f\colon G \to \mathbb{C}$ and $t \in G$, we use the normalized Fourier transform

$$\widehat{f}(t) \;\coloneqq\; \frac{1}{N}\sum_{x\in G} f(x)\,\mathrm{e}_N(-xt).$$

We freely identify a set with its indicator function; thus, for an $S \subseteq G$,

$$\widehat{S}(t) = \frac{1}{N}\sum_{x\in S} \mathrm{e}_N(-xt).$$

The Fourier transform maps a function f to $\widehat{f}$. This map is also invertible, and can be computed via the inversion formula:

$$f(x) = \sum_{t\in G} \widehat{f}(t)\,\mathrm{e}_N(xt) \tag{4}$$

For $1 \leqslant p < \infty$, all L_p-norms are taken with respect to normalized counting measure on G:

$$\|f\|_{L^p} \;\coloneqq\; \left(\frac{1}{N}\sum_{x\in G}|f(x)|^p\right)^{1/p} \tag{5}$$

For $p = \infty$, we set $\|f\|_{L^\infty} \;\coloneqq\; \max_{x\in G}|f(x)|$. With this normalization, Parseval's identity takes the form

$$\frac{1}{N}\sum_{x\in G}|f(x)|^2 = \sum_{t\in G}\Big|\widehat{f}(t)\Big|^2. \tag{6}$$

Local-leakage insecurity. Retain the notation $D(s)$ and $D(U)$ from Section 1. Following the local leakage-resilience metric of [BDIR18, BDIR21, FMM+24], the insecurity associated with a fixed secret $s \in G$ is

$$E_k(s) \;\coloneqq\; \|D(s) - D(U)\|_{\mathrm{TV}} = \frac{1}{2}\sum_{\mathbf{w}\in W^k}\left|\sum_{\substack{t\in G\\ t\neq 0}} \mathrm{e}_N(st)\prod_{j=1}^{k}\widehat{S_{\mathbf{w}_j}}(t)\right|. \tag{7}$$

The second equality follows from Fourier inversion (Equation 4): the zero-frequency term is precisely $D(U)(\mathbf{w})$, and therefore cancels when $D(U)$ is subtracted. The worst-case insecurity is

$$E_k \;\coloneqq\; \max_{s\in G} E_k(s).$$

Asymptotic notation. We use the standard $\ll, \gg, \asymp, \mathcal{O}(\cdot)$, and $\mathrm{o}(\cdot)$ notation. Subscripts indicate the parameters on which the implicit constants may depend; for example, $A \ll_k B$ means $A \leqslant C_k B$ for some constant C_k depending only on k. In our bounds, this dependence is at most exponential: $C_k = C^k$ for an absolute constant C, independent of n, N, and the secret s. Its value may vary from one occurrence to another. The analogous convention applies to $\gg_k$ and $\asymp_k$.

2.2. **Mathematical Background.**

Generating-function formulation. An integer $0 \leqslant x < 2^n$ can be represented as an n-bit binary string $x_{n-1} \dots x_1 x_0$ such that $x = \sum_{j=0}^{n-1} x_j 2^j$ and $x_j \in \{0,1\}$. Then we have $\mathrm{wt}(x) = \sum_{j=0}^{n-1} x_j$ and

$$\prod_{j=0}^{n-1} \big(1 + z \cdot \mathrm{e}_N(-2^j t)\big) = \sum_{0 \leqslant x < 2^n} z^{\mathrm{wt}(x)} \cdot \mathrm{e}_N(-tx) = \sum_{0 \leqslant x < 2^n} z^{\mathrm{wt}(x)} \cdot \left(\prod_{j=0}^{n-1} \mathrm{e}_N(-2^j t x_j) \right)$$
$$= \sum_{w=0}^{n} z^w \sum_{\substack{0 \leqslant x < 2^n \\ \mathrm{wt}(x) = w}} \mathrm{e}_N(-tx).$$

Consequently, the coefficient of z^w is the unnormalized Fourier sum over all n-bit strings of Hamming weight w, i.e., $N \cdot \widehat{B_w}(t)$ (see Equation 18 for definition of the function B_w, which is closely tied to the Hamming slice S_w).

Products indexed by the doubling sequence 2^j also arise in the study of generalized Thue-Morse trigonometric polynomials; see, for example, [FSS23, equation 1.3, 1.4]. Our product has a different phase weighting, and the estimates needed below are proved directly.

Coefficient extraction via Cauchy integration formula. The summary here is based on the treatments of Cauchy coefficient extraction and saddle-point contours in [FS09, Chapter IV and Sections VIII.1–VIII.2]. Let

$$P(z) = \prod_{j=0}^{n-1} (1 + z \cdot \xi_j) \qquad \text{where } |\xi_j| = 1$$

be a complex polynomial. Cauchy's coefficient formula gives, for every $r > 0$,

$$[z^w]P(z) = \frac{1}{2\pi r^w} \int_{-\pi}^{\pi} P\Big(re^{\imath\theta}\Big) \cdot \exp(-\imath w\theta) \mathrm{d}\theta.$$

For $0 < w < n$, we make a standard substitution $p = w/n$ and $r = p/(1-p)$.
Since $1 + re^{\imath\phi} = \big((1-p) + pe^{\imath\phi}\big)/(1-p)$, this yields

$$\Big|[z^w]P(z)\Big| \leqslant \frac{1}{2\pi p^w (1-p)^{n-w}} \int_{-\pi}^{\pi} \prod_{j=0}^{n-1} \Big|(1-p) + p \cdot \xi_j \exp(\imath\theta)\Big| \mathrm{d}\theta.$$

The Bernoulli shift and Mersenne rotations. Let $\mathbb{T} := \mathbb{R}/\mathbb{Z}$ and $T(x) := \{2x\}$. Under binary coding, T is the left shift

$$0.b_0 b_1 b_2 \cdots \xmapsto{T} 0.b_1 b_2 b_3 \cdots,$$

up to the usual ambiguity at dyadic rationals; see [EW11, Chapter 2]. For $N = 2^n - 1$, its restriction to the points a/N recovers cyclic rotation under doubling modulo N: if $a = (b_0 b_1 \cdots b_{n-1})_2$, then

$$T\Big(\frac{a}{N}\Big) = \frac{2a \bmod N}{N} = (0.\overline{b_1 \cdots b_{n-1} b_0})_2. \tag{8}$$

because we would have $a/N = (0.\overline{b_0 b_1 \cdots b_{n-1}})_2$.

Example 2.1. With $n = 5$, $N = 31$, and $a = 19 = (10011)_2$, one has

$$\frac{19}{31} = (0.\overline{10011})_2 \quad \xmapsto{T} \quad (0.\overline{00111})_2 = \frac{7}{31}, \quad 2 \cdot 19 \equiv 7 \pmod{31}.$$

△

For general near-dyadic N, the orbit need not be n-periodic, but $T^j(t/N) = \{2^j t/N\}$ still records the successive binary shifts and serves as a robust substitute for the exact Mersenne orbit structure.

3. Upper Bound: Prime-exponent Mersenne Modulus

In this section, we prove the following theorem.

Theorem 1 (Prime -exponent Mersenne-moduli upper bound)**.** *Let $n \geqslant 3$ be prime and let $N = 2^n - 1$; the modulus N itself need not be prime. For every fixed $k \geqslant 3$ and every $s \in \mathbb{Z}/N\mathbb{Z}$,*

$$E_k(s) \ll n^{-k/2+1} \log n.$$

Consequently, $E_k \ll n^{-k/2+1} \log n$.

We prove Theorem 1 by a Fourier-energy argument that improves the bound of Faust et al. [FMM+24] for the same prime-exponent Mersenne family.

Appendix A.2 records a comparison argument only: replacing the Parseval pointwise estimate in the proof of Faust et al. [FMM+24] by Chang's lemma attains the correct $n^{-1/2}$ scale per additional share, but it requires N to be prime and still does not control $k = 3$.

Proof overview of Theorem 1. For every integer $m \geqslant 1$, define

$$\text{(Energy proxy)} \qquad F_m \coloneqq \sum_{\mathbf{w} \in W^m} \left(\sum_{t \neq 0} \prod_{j=1}^{m} \left| \widehat{S_{\mathbf{w}_j}}(t) \right|^2 \right)^{1/2}. \tag{9}$$

The proof has two layers. Iterating a one-share contraction down to F_1 already proves asymptotic security for every $k \geqslant 3$ and supplies the template used later for near-dyadic moduli. For the sharper Mersenne-type bound of Theorem 1, we stop instead at F_2 and estimate that base case directly.

Step 1: Reduce insecurity to a secret-independent energy. The first step trades one share for the more tractable quantity F_{k-1}; proved in Appendix B.1.

Lemma 3.1 (E_k-to-F_{k-1} reduction)**.** *Let $X_1, \ldots, X_{k-1}, X_k$ be independent and uniform over G, and set*

$$\Delta_{k-1} \coloneqq \Big\| (\mathrm{wt}(X_1), \ldots, \mathrm{wt}(X_{k-1}), X_1 + \cdots + X_{k-1}) - (\mathrm{wt}(X_1), \ldots, \mathrm{wt}(X_{k-1}), X_k) \Big\|_{\mathrm{TV}}$$

Then, for every integer $N \geqslant 2$, every $k \geqslant 2$, and every secret $s \in G = \mathbb{Z}/N\mathbb{Z}$,

$$2E_k(s) \leqslant 2\Delta_{k-1} \leqslant F_{k-1}. \tag{10}$$

The underlying reduction extends verbatim to arbitrary finite abelian groups and arbitrary local leakage functions, with the corresponding energy defined from the leakage fibers. It is purely functional-analytic: it replaces the original secret-dependent L_1 expression by a secret-independent L_2 energy. This loss of one share is decisive because the resulting energies admit an iterative contraction.

Remark 3.2 (Information-theoretic interpretation of F_m)**.** The proxy F_m arose naturally from our Fourier-analytic argument, but it also has a direct information-theoretic interpretation: it is the V_2 correlation measure of Mojahedian et al. [MBG+19, Eq. (9)] between the leakage vector

$$\big(\mathrm{wt}(X_1), \ldots, \mathrm{wt}(X_m)\big) \qquad \text{and} \qquad X_1 + \cdots + X_m.$$

Equivalently, F_m controls, on average over the leakage outcomes, how far the conditional distribution of $X_1 + \cdots + X_m$ is from uniform. Thus, a small value of F_m means that revealing the individual Hamming weights leaves the aggregate sum essentially uniform. Taking $m = (k-1)$, a small value of F_{k-1} implies that the pad $s_1 + \cdots + s_{k-1}$ remains close to uniform despite leakage from the first $(k-1)$ shares. Consequently, the ciphertext $s - \big(s_1 + \cdots + s_{k-1}\big)$ is close to uniform and nearly independent of the secret s, establishing the security of the leaky-encryption construction described

in the introduction. Additive secret sharing reveals only the Hamming weight of the ciphertext, which is a deterministic post-processing step; hence, by data processing, F_{k-1} also controls the insecurity of the original sharing scheme. △

Step 2: Peel off one share. For Mersenne-type moduli, the doubling-orbit structure from Appendix A.1 yields the desired contraction directly at the level of Fourier energies; proved in Appendix B.2.

Lemma 3.3 (F_m-to-F_{m-1} contraction)**.** *Suppose that $N = 2^n - 1$ and n is prime. Then, for every $m \geqslant 2$,*

$$F_m \leqslant n^{-1/2} \cdot F_{m-1}. \tag{11}$$

Iterating Equation 11 in Lemma 3.3 gives

$$F_m \leqslant n^{-(m-1)/2} \cdot F_1. \tag{12}$$

Unlike the pointwise argument of Faust et al. [FMM+24], this step uses the doubling symmetry after the Fourier coefficients have been grouped into an L_2 energy.

Step 3: The preliminary F_1 base case. Parseval's identity reduces F_1 to the Hamming-slice densities, and standard Gaussian estimates for the binomial distribution give the following bound, proved in Appendix B.3.

Lemma 3.4 (The F_1 base case)**.** *For $N \in \{2^n - 1, 2^n\}$,*

$$F_1 \ll n^{1/4}. \tag{13}$$

Combining $(*)$ Equation 10, $(**)$ Equation 12, and $(\dagger)$ Equation 13, we obtain the preliminary estimate

$$2E_k(s) \overset{(*)}{\leqslant} F_{k-1} \overset{(**)}{\leqslant} n^{-(k-2)/2} F_1 \overset{(\dagger)}{\ll} n^{-k/2+5/4}. \tag{14}$$

This already establishes asymptotic security for $k = 3, 4$, and, more generally, for every fixed $k \geqslant 3$. It also gives the basic scaffolding used later in the near-dyadic setting. For $N = 2^n - 1$, however, one can sharpen the base case and recover the stronger bound of Theorem 1.

Step 4: Stop at F_2. For the sharper Mersenne estimate, we avoid the final contraction $F_2 \leqslant n^{-1/2} F_1$ and stop the iteration one step earlier:

$$2E_k(s) \leqslant F_{k-1} \leqslant n^{-(k-3)/2} F_2. \tag{15}$$

The remaining task is a direct estimate of

$$F_2 = \sum_{w_1, w_2 \in W} \left(\sum_{t \neq 0} \left|\widehat{S_{w_1}}(t)\right|^2 \left|\widehat{S_{w_2}}(t)\right|^2 \right)^{1/2}. \tag{16}$$

Lemma B.4 in Appendix B.4 proves

$$F_2 \ll n^{-1/2} \log n, \tag{17}$$

using the following external fourth-moment estimate.

Imported Lemma 3.5 (Technical Input: Energy estimate [BHM+26, Theorem 3])**.** *For $w \in \{0, 1, \ldots, n-1\}$ and S_w denoting the subset of all elements $x \in \mathbb{Z}/(2^n - 1)\mathbb{Z}$ of Hamming weight w, the following bound holds.*

$$\sum_{t \neq 0} \left|\widehat{S_w}(t)\right|^4 \ll n^{-3}.$$

Finally, substituting Equation 17 into Equation 15 gives

$$E_k(s) \ll n^{-(k-3)/2}\, n^{-1/2} \log n = n^{-k/2+1} \log n,$$

as claimed in Theorem 1. □

4. Upper Bound: Moduli close to a Power of Two

We now extend the upper bound beyond Mersenne-type moduli to every modulus in a near-dyadic window. Throughout this section, set $Q := 2^n$ and assume that $|N - Q| \leqslant Q/n$. The main result is the following.

Theorem 2 (Near-dyadic moduli upper bound)**.** *There is an absolute constant $c > 0$ such that the following holds. Let $n \geqslant 3$, let N satisfy $|N - 2^n| \leqslant 2^n/n$, and fix an integer $k \geqslant 3$. Then, for every $s \in \{0, 1, \dots, N-1\}$,*

$$E_k(s) \leqslant (cn)^{-k/2+5/4}.$$

Note that, unlike results in Section 3, n is not required to be prime.

Proof roadmap. Set $Q := 2^n$. We first replace the actual slices in $\{0, \dots, N-1\}$ by the full n-bit range multiplicity functions B_w coming from $\{0, \dots, Q-1\}$, while retaining addition and Fourier analysis modulo N. The reduction to an energy and the F_1 base estimate parallel Section 3. The new input is a one-share contraction for the model energies, obtained from a spectral estimate along the Bernoulli-doubling orbit. Finally, our transference Lemma 4.8 transfers the model bound back to the actual slices, with an error controlled by $|N - Q|/Q$.

4.1. **Full-range Model.** The arithmetic lives on the N-grid $G = \mathbb{Z}/N\mathbb{Z}$, whereas the Hamming-weight geometry is most naturally organized over the full n-bit residue range $R_n := \{0, 1, \dots, Q-1\}$. Our model separates these two roles: the summation variables range over R_n, while all additions and additive characters remain modulo N.

For $w \in W_{\mathrm{BIN}} := \{0, 1, \dots, n\}$ and $x \in G$, define the multiplicity function

$$B_w(x) := \#\Big\{0 \leqslant y < Q : \ y = x \pmod N, \text{ and } \mathrm{wt}(y) = w\Big\}. \tag{18}$$

Each $y \in R_n$ contributes exactly once, namely to $B_{\mathrm{wt}(y)}(y \bmod N)$.

Since $|N - Q| \leqslant Q/2$, every residue class has at most two representatives in $[0, Q)$, and hence $B_w(x) \in \{0, 1, 2\}$. More explicitly:

(i) If $N \geqslant Q$, then

$$B_w(x) = \begin{cases} 1, & 0 \leqslant x < Q \text{ and } \mathrm{wt}(x) = w, \\ 0, & \text{otherwise.} \end{cases}$$

(ii) If $N < Q$, then every $x \in \{0, 1, \dots, N-1\}$ receives the contribution from $y = x$, and the residues $0 \leqslant x < Q - N$ may receive one additional contribution from $y = x + N$.

Example 4.1 (The Mersenne endpoint)**.** Let $N = Q - 1 = 2^n - 1$. Then, for every $0 \leqslant w < n$, one has $B_w = S_w$. The unique element of R_n having Hamming weight n is $Q - 1$, which reduces to $0 \pmod N$; consequently, $B_n = \delta_0$, where δ_0 denotes the indicator of $\{0\} \subseteq G$. △

The Fourier transform of B_w is

$$\widehat{B_w}(t) = \frac{1}{N} \sum_{x \in G} B_w(x)\, \mathrm{e}_N(-xt) = \frac{1}{N} \sum_{\substack{0 \leqslant y < Q \\ \mathrm{wt}(y) = w}} \mathrm{e}_N(-yt). \tag{19}$$

Thus, $\widehat{B_w}(t)$ is the Fourier sum over the full weight-w layer of R_n, sampled using the N-periodic character e_N.

We now introduce the model analogs of the insecurity and its Fourier-energy proxies. For $k \geqslant 2$ and $0 \leqslant s < N$, define

$$E_k^{\mathrm{BIN}}(N, s) := \frac{1}{2} \sum_{\mathbf{w} \in W_{\mathrm{BIN}}^k} \left| \sum_{t \neq 0} \mathrm{e}_N(st) \prod_{j=1}^{k} \widehat{B_{\mathbf{w}_j}}(t) \right|. \tag{20}$$

For every integer $m \geqslant 1$, define

$$F_m^{\mathrm{BIN}}(N) := \sum_{\mathbf{w}\in W_{\mathrm{BIN}}^m} \left(\sum_{t\neq 0} \prod_{j=1}^{m} \left| \widehat{B_{\mathbf{w}_j}}(t) \right|^2 \right)^{1/2}. \tag{21}$$

The distinction from the original energy F_m

$$F_m := \sum_{\mathbf{w}\in W^m} \left(\sum_{t\neq 0} \prod_{j=1}^{m} \left| \widehat{S_{\mathbf{w}_j}}(t) \right|^2 \right)^{1/2}$$

is that the model runs over the complete weight range W_{BIN} of R_n and uses the functions B_w, rather than the actual weight slices $S_w \subseteq G$.

Example 4.2 (Comparing $F_2^{\mathrm{BIN}}(N)$ and F_2 at $N = 2^n - 1$)**.** Retain the setting of Example 4.1. Since $B_w = S_w$ for $w < n$, $B_n = \delta_0$, and $\widehat{B_n}(t) = 1/N$, partitioning the sum according to whether each weight equals n gives

$$F_2^{\mathrm{BIN}}(N) = F_2 + \frac{2}{N} F_1 + \frac{\sqrt{N-1}}{N^2}.$$

Indeed, the two mixed cases each contribute $F_1(N)/N$, while the case $(w_1, w_2) = (n, n)$ contributes $\sqrt{N-1}/N^2$. This identity also illustrates why one should not compare $F_m^{\mathrm{BIN}}(N)$ and F_m by a naïve union bound: the energies are nonlinear expressions in Fourier space, rather than event probabilities in physical space. △

4.2. **Fourier-energy Scaffolding in the Full-range Model.** The model quantities admit the same three-step scaffolding as in the Mersenne case: reduce insecurity to an energy, contract by one share, and finish with a one-share base estimate. A substantive new ingredient is the spectral parameter governing the contraction.

Lemma 4.3 (Model reduction)**.** *For every $k \geqslant 3$ and every $0 \leqslant s < N$,*

$$2E_k^{\mathrm{BIN}}(N, s) \leqslant 2F_{k-1}^{\mathrm{BIN}}(N).$$

The proof in Appendix C.1 is the model analogue of the E_k-to-F_{k-1} reduction. The additional factor of 2, relative to the sharper inequality $2E_k(s) \leqslant F_{k-1}$, comes from the overlap bound $\sum_{w\in W_{\mathrm{BIN}}} B_w(x) \leqslant 2$, for any $x \in G$.

Lemma 4.4 (One-share model contraction)**.** *For every integer $m \geqslant 2$,*

$$F_m^{\mathrm{BIN}}(N) \leqslant \left(\frac{Q}{N}\right)^{1/2} \cdot \alpha^{\mathrm{BIN}}(N)^{1/2} \cdot F_{m-1}^{\mathrm{BIN}}(N),$$

where

$$\alpha^{\mathrm{BIN}}(N) := \max_{t\neq 0} \sum_{w\in W_{\mathrm{BIN}}} \frac{\left|\widehat{B_w}(t)\right|^2}{\rho_w(N)} \quad \text{and} \quad \rho_w(N) := \frac{1}{N}\binom{n}{w}.$$

Appendix C.2 proves this lemma. The factor $(Q/N)^{1/2}$ is harmless throughout the near-dyadic regime. Thus, all of the arithmetic difficulty is concentrated in estimating the single quantity $\alpha^{\mathrm{BIN}}(N)$, proved below in Section 4.4.

Lemma 4.5 (Key spectral estimate)**.** *For $|N - Q| \leqslant Q/n$, $\alpha^{\mathrm{BIN}}(N) \ll n^{-1}$.*

The proof is given in Section 4.4. It follows the orbit $T^j(t/N) = \{2^j t/N\}$ under the Bernoulli map introduced in Section 2 and replaces exact Mersenne periodicity with a quantitative dispersion estimate.

Combining Lemma 4.4 and Lemma 4.5 with $Q/N \leqslant 2$ yields the desired per-share contraction

$$F_m^{\mathrm{BIN}}(N) \ll n^{-1/2} F_{m-1}^{\mathrm{BIN}}(N).$$

At the bottom of the scaffolding, we use the following model analog of the F_1 estimate (proved in Appendix C.3).

Lemma 4.6 (The model F_1 base case)**.** *Uniformly for $|N-Q| \leqslant Q/n$, $F_1^{\mathrm{BIN}}(N) \ll n^{1/4}$.*

We may now finish the model analysis.

Lemma 4.7 (Model bound)**.** *For every $k \geqslant 3$ and every $0 \leqslant s < N$,*

$$E_k^{\mathrm{BIN}}(N,s) \ll_k n^{-k/2+5/4}.$$

Proof of Lemma 4.7. By Lemma 4.3, followed by $(k-2)$ applications of the one-share contraction,

$$E_k^{\mathrm{BIN}}(N,s) \leqslant F_{k-1}^{\mathrm{BIN}}(N) \ll_k n^{-(k-2)/2} \cdot F_1^{\mathrm{BIN}}(N).$$

Invoking Lemma 4.6 gives $E_k^{\mathrm{BIN}}(N,s) \ll_k n^{-(k-2)/2} \cdot n^{1/4} = n^{-k/2+5/4}$. □

4.3. **Transference from the Full-range Model to the N-grid.** It remains to compare the model quantity with the actual insecurity. A direct union bound would charge each share separately for the discrepancy between $[0,Q)$ and $[0,N)$, and would only preserve an $n^{-k/2}$-scale estimate in a much narrower window. The transference theorem instead expands in terms of the number of perturbed factors and controls the surviving factors by the model energies. In this sense, it packages the boundary discrepancy into an economical Fourier-space analog of a union bound.

Lemma 4.8 (Transference)**.** *Suppose that $|N-Q| \leqslant P$. Then, for every $k \geqslant 2$ and every $0 \leqslant s < N$,*

$$E_k(s) \leqslant E_k^{\mathrm{BIN}}(N,s) \;+\; \sum_{\ell=1}^{k-1} \binom{k}{\ell} \cdot \left(\frac{P}{N}\right)^{\ell/2} F_{k-\ell}^{\mathrm{BIN}}(N) \;+\; \left(\frac{P}{N}\right)^{k/2} \sqrt{n}.$$

We use Lemma 4.8 to prove Theorem 2. Lemma 4.8 is proved in Appendix D.1.

Proof of Theorem 2. Proving Theorem 2 needs specializing Lemma 4.8 to $P = |N-Q| \leqslant Q/n$. Since $N \geqslant Q/2$, $P/N \leqslant 2/n$. Moreover, iterating Lemma 4.4 and Lemma 4.5, and applying Lemma 4.6 gives, for every $1 \leqslant \ell < k$,

$$F_{k-\ell}^{\mathrm{BIN}}(N) \ll_k n^{-(k-\ell-1)/2+1/4}.$$

Consequently,

$$\left(\frac{P}{N}\right)^{\ell/2} F_{k-\ell}^{\mathrm{BIN}}(N) \ll_k n^{-\ell/2} \cdot n^{-(k-\ell-1)/2+1/4} = n^{-k/2+3/4}, \tag{22}$$

while

$$\left(\frac{P}{N}\right)^{k/2} \sqrt{n} \ll_k n^{-k/2+1/2}. \tag{23}$$

Both errors are smaller than the model contribution from Lemma 4.7. Hence

$$E_k(s) \ll_k n^{-k/2+5/4},$$

which proves Theorem 2. □

We can extend Lemma 4.8 from the insecurity E_k to the stronger, intermediate quantity Δ_{k-1} (defined in Lemma 3.1) where the adversary learns the Hamming weight of the first $k-1$ shares and the aggregate of them for their k-th leakage. The details are deferred to Appendix D.2 as Corollary 1.

Remark 4.9 (Full disclosure of the aggregate in near-dyadic moduli)**.** The same analysis bounds the stronger quantity Δ_{k-1} of Lemma 3.1, in which the adversary learns the Hamming weights of the first $k-1$ shares together with their exact sum, rather than the Hamming weight of the final share. In particular, Corollary 1 gives, under the same condition as Theorem 2,

$$\Delta_{k-1} \ll_k n^{-k/2+5/4}.$$

Thus, the one-time-pad interpretation in Remark 1.1 (and Remark 3.2 is also justified for near-dyadic moduli. △

4.4. **Proof of Key Spectral Estimate (Lemma 4.5).** We have to prove $\alpha_{\mathrm{BIN}}(N) \ll n^{-1}$. Fix $t \neq 0$. We first discard the non-central weights by a binomial-tail estimate. For the remaining weights, Lemma 4.10 gives a uniform coefficient bound; summing its square against $\rho_w(N)^{-1}$ then yields $\alpha^{\mathrm{bin}}(N) \ll n^{-1}$.

1. From individual coefficients to $\alpha_{\mathrm{BIN}}(N)$. Fix $t \neq 0$ and let

$$W_{\mathrm{cent}} := \left\{ w \in W_{\mathrm{BIN}} : \left| w - \frac{n}{2} \right| \leqslant \sqrt{n \ln n} \right\}.$$

The trivial bound $|\widehat{B_w}(t)| \leqslant \rho_w(N)$ and a binomial tail estimate make the total contribution of $w \notin W_{\mathrm{cent}}$ equal to $\mathcal{O}(n^{-2})$, which is insignificant as compared to n^{-1}.

$$\begin{aligned}
\sum_{w \notin W_{\mathrm{cent}}} \frac{\left|\widehat{B_w}(t)\right|^2}{\rho_w(N)} \leqslant \sum_{w \notin W_{\mathrm{cent}}} \left|\widehat{B_w}(t)\right| &\leqslant \sum_{w \notin W_{\mathrm{cent}}} \rho_w(N) && \text{(trivial bound } |\widehat{B_w}(t)| \leqslant \rho_w(N)) \\
&= \frac{Q}{N} \cdot \Pr\left[|X - n/2| > \sqrt{n \log n} \right] && \text{(where } X \sim \mathrm{Bin}(n, 1/2)) \\
&\leqslant \frac{Q}{N} \cdot 2\exp(-2 \cdot (n \log n)/n) && \text{(by Chernoff bound)} \\
&\ll n^{-2} && \text{(since } Q/N \leqslant 3/2)
\end{aligned}$$

For central weights, we prove:

Lemma 4.10 (Krawtchouk-type Estimate)**.** *For every $w \in W_{\mathrm{cent}}$ and $0 < t < N$, we have:*

$$\left|\widehat{B_w}(t)\right| \ll N^{-1} \binom{n}{w} \cdot \left(\frac{|n-2w|}{n} + \frac{1}{\sqrt{n}} \right).$$

From Lemma 4.10, the cumulative contribution from central weights is:

$$\begin{aligned}
\sum_{w \in W_{\mathrm{cent}}} \frac{\left|\widehat{B_w}(t)\right|^2}{\rho_w(N)} &\ll \sum_{w \in W_{\mathrm{cent}}} \frac{N^{-2}\binom{n}{w}^2}{N^{-1}\binom{n}{w}} \cdot \left(\frac{|n-2w|}{n} + \frac{1}{\sqrt{n}} \right)^2 \\
&\ll \sum_{w \in W_{\mathrm{cent}}} N^{-1} \binom{n}{w} \cdot \left(\frac{(n-2w)^2}{n^2} + \frac{1}{n} \right) \\
&\asymp \sum_{w \in W_{\mathrm{cent}}} Q^{-1} \binom{n}{w} \cdot \left(\frac{(n-2w)^2}{n^2} + \frac{1}{n} \right) \leqslant \frac{4}{n^2} \cdot \mathrm{Var}(X) + \frac{1}{n} \ll \frac{1}{n}. \\
& && \text{(note that } X \sim \mathrm{Bin}(n, 1/2))
\end{aligned}$$

2. From coefficient extraction to individual coefficients. Let us prove Lemma 4.10. For $3 \leqslant n < 9$, the trivial estimate $|\widehat{B_w}(t)| \leqslant \rho_w(N)$ proves Lemma 4.10 after enlarging the absolute constant. Henceforth, assume $n \geqslant 9$.

For $w \in W_{\mathrm{cent}}$, put $p = w/n$ and

$$P_t(z) \;:=\; \prod_{j=0}^{n-1} \bigl(1 + z\,\mathrm{e}_N(-2^j t)\bigr), \qquad \widehat{B}_w(t) = \frac{1}{N}[z^w]P_t(z).$$

Applying the estimate from Section 2 to P_t with $\phi_j := \{2^j t/N\}$ gives

$$\bigl|\widehat{B}_w(t)\bigr| \leqslant \frac{I}{2\pi N p^w (1-p)^{n-w}}, \qquad I \;:=\; \int_{-\pi}^{\pi} \prod_{j=0}^{n-1} L_p(\theta + 2\pi\phi_j)\mathrm{d}\theta,$$

where $L_p(u) := |(1-p) + p\exp(\imath u)|$. By Stirling's formula $p^{-w}(1-p)^{-(n-w)} \asymp \sqrt{n}\binom{n}{w}$, Lemma 4.10 therefore follows from the stated bound on I below.

Lemma 4.11. *For $n \geqslant 9$, $0 < t < N$, $|N - 2^n| \leqslant 2^n/n$, $w \in W_{\mathrm{cent}}$, and $\phi_j = 2^j t/N$, we have:*

$$I = \int_{-\pi}^{\pi} \prod_{j=0}^{n-1} \Bigl|(1-p) + p\exp(\imath\theta + \imath\cdot 2\pi\phi_j)\Bigr| \mathrm{d}\theta \ll \frac{|1-2p|}{\sqrt{n}} + \frac{1}{n},$$

Appendix C.4 proves this lemma; below we present our proof roadmap.

Proof roadmap for Lemma 4.11. We split the factors in the integrand into two groups with different jobs: a short group makes the product small near $\theta = 0$, while the remaining factors give a rapidly decaying upper bound away from 0 Retain $p = w/n$, $\phi_j = \{2^j t/N\}$, and $L_p(u) = |(1-p) + p\exp(\imath u)|$. For $\Omega = \{0, \ldots, n-1\}$ and $J \subseteq \Omega$, write

$$P_J(\theta) \;:=\; \prod_{j \in J} L_p(\theta + 2\pi\phi_j).$$

Thus $I = \int_{-\pi}^{\pi} P_\Omega(\theta)\,\mathrm{d}\theta$.

A few factors make the product small. At angle π, the two terms in L_p point in opposite directions, giving $L_p(\pi) = |1-2p|$. An index with ϕ_j close to $1/2$ therefore gives a small factor when θ is close to 0. When no single index gives a sufficiently strong bound, we use several consecutive factors together. Claim C.1 finds a consecutive group J, independent of θ, with $1 \leqslant \#J \leqslant \ell_0 \;:=\; \bigl\lceil \frac{1}{2}\log_2 n \bigr\rceil$, such that

$$P_J(\theta) \ll |1-2p| + |\theta| + n^{-1/2}, \qquad |\theta| \leqslant \pi.$$

To find this group, replace t by $N - t$ and θ by $-\theta$ if necessary, so that $0 < t/N \leqslant 1/2$. Let r be the first index for which $2^r t/N \in (1/4, 1/2]$. If $r + \ell_0 \leqslant n$, take $J = \{r, r+1, \ldots, r+\ell_0 - 1\}$. Claim C.2 bounds their product by $\mathcal{O}\bigl(|1-2p| + |\theta| + 2^{-\ell_0}\bigr)$, and $2^{-\ell_0} \leqslant n^{-1/2}$. Otherwise, write $t = 2^a b$ with b odd. In this case $b \ll \sqrt{n}$, and the assumption $|N - 2^n| \leqslant 2^n/n$ places ϕ_{n-a-1} within $\mathcal{O}\bigl(n^{-1/2}\bigr)$ of $1/2$, modulo 1. This one factor gives the same bound by Claim C.3.2.

The remaining factors restrict the width. Removing a consecutive group leaves at most two consecutive groups, containing $n - \mathcal{O}(\log n)$ indices in total. The identity $L_p(u)^2 = 1 - 4p(1-p)\sin^2(u/2)$ shows how an angle away from 0 modulo 2π makes a factor smaller. The central-weight condition keeps p away from 0 and 1. Moreover, successive angles satisfy

$$\theta + 2\pi\phi_{j+1} \equiv 2(\theta + 2\pi\phi_j) - \theta \pmod{2\pi}.$$

They therefore cannot all stay close to 0 unless θ is itself close to 0. Our dispersion lemma (Lemma C.4) makes this observation quantitative; together with Claim C.3.1., it gives

$$P_{\Omega\setminus J}(\theta) \leqslant \exp\bigl(-cn\theta^2\bigr), \qquad |\theta| \leqslant \pi,$$

for an absolute constant $c > 0$. The integral of this upper bound is $\mathcal{O}\bigl(n^{-1/2}\bigr)$.

Multiply the bounds, then integrate. The two disjoint groups include every factor. Hence

$$I \ll \int_{-\pi}^{\pi} \bigl(|1-2p| + |\theta| + n^{-1/2}\bigr) \cdot \exp\bigl(-cn\theta^2\bigr)\,\mathrm{d}\theta \ll \frac{|1-2p|}{\sqrt{n}} + \frac{1}{n}.$$

Indeed, the constant terms acquire a factor $n^{-1/2}$ upon integration, while the term containing $|\theta|$ integrates to $\mathcal{O}(n^{-1})$. In particular, when $p = 1/2$, the short group is $\mathcal{O}(n^{-1/2})$ on the scale $|\theta| = \mathcal{O}(n^{-1/2})$, while the width of this region supplies the other factor $n^{-1/2}$.

4.5. **Sharper Bound for Mersenne Moduli.** For $N = 2^n - 1$, the full n-bit range differs from the canonical residue set only at the all-ones string, which reduces to 0 modulo N. Consequently, the model energy $F_2^{\mathrm{BIN}}(N)$ can be compared explicitly with F_2, and Lemma B.4 lets us stop the contraction at two shares. This yields the following refinement of Theorem 2 for every exponent n.

Theorem 3 (Mersenne-moduli upper bound)**.** *There is an absolute constant $c > 0$ such that the following holds. Let $n \geqslant 3$ be an integer and $N = 2^n - 1$, and fix an integer $k \geqslant 3$. Then, for every $s \in \{0, 1, \dots, N-1\}$,*

$$E_k(s) \leqslant (cn)^{-k/2+1} \log n.$$

Equivalently, $E_k(s) \ll_k n^{-k/2+1} \log n$.

Proof of Theorem 3. For $N = Q - 1$, the full n-bit range and the canonical representative set $\{0, 1, \dots, N-1\}$ differ only at the all-ones string $Q - 1$, whose image modulo N is 0 and whose Hamming weight is n. For $N = 2^n - 1$, Example 4.1 gives $B_w = S_w$ for every actual weight $0 \leqslant w < n$. Every summand defining $E_k(s)$ therefore appears unchanged in $E_k^{\mathrm{BIN}}(N, s)$, and the additional summands are nonnegative. Hence

$$E_k(s) \leqslant E_k^{\mathrm{BIN}}(N, s).$$

The same comparison gives $F_m \leqslant F_m^{\mathrm{BIN}}(N)$ for every $m \geqslant 1$. By Lemma 4.3 and the model contraction (Lemma 4.4 and Lemma 4.5),

$$E_k(s) \leqslant F_{k-1}^{\mathrm{BIN}}(N) \ll_k n^{-(k-3)/2} F_2^{\mathrm{BIN}}(N).$$

Using Example 4.2,

$$F_2^{\mathrm{BIN}}(N) = F_2 + \frac{2}{N} F_1 + \frac{\sqrt{N-1}}{N^2}.$$

The estimates established earlier in Lemma 3.4 and Lemma B.4 give

$$F_2 \ll n^{-1/2} \log n \qquad \text{and} \qquad F_1 \ll n^{1/4}.$$

Hence

$$E_k(s) \ll_k n^{-(k-3)/2} \left(n^{-1/2} \log n + \frac{n^{1/4}}{N} + \frac{1}{N^{3/2}} \right).$$

Because $N = 2^n - 1$, the terms containing N^{-1} are exponentially small in n. The dominant contribution is therefore

$$n^{-(k-3)/2} n^{-1/2} \log n = n^{-k/2+1} \log n,$$

as required. □

5. Lower Bound: Mersenne Moduli, Even Number of Shares

Upper bounds show that leakage becomes small; lower bounds identify the remaining correlations, revealing live side-channel threats. For the Mersenne family, the doubling symmetry leads to a particularly transparent obstruction that already gives a $n^{-1/2}$-per-share scale. It will serve as a warm-up for more complex lower-bounding strategies in the next section.

Theorem 4 (Mersenne lower bound)**.** *For any integer $n \geqslant 3$ with $N = 2^n - 1$ and any even integer $k \geqslant 4$,*

$$E_k(0) \geqslant \frac{1}{2} \cdot 10^{-k} \cdot n^{-k/2+1}.$$

Proof idea. Numerical evidence singles out the canonical doubling orbit $\{-2^j : 0 \leqslant j < n\}$ as carrying unusually large aggregate Fourier mass. We choose product phases aligned at the representative frequency $t = -1$. The weights w and $n - w$ pair by complex conjugation, so the resulting contribution is real at every frequency. Because k is even, this sum's k-th power is nonnegative. We may therefore retain only the canonical orbit and obtain

$$2 \cdot E_k(0) \geqslant n \left(\sum_{w \in W} \left| \widehat{S}_w(-1) \right| \right)^k .$$

Lemma 5.1 supplies the required lower bound for the remaining one-frequency mass. The proof below implements this phase choice.

Proof of Theorem 4. Recall first that

$$2 \cdot E_k(0) = \sum_{\mathbf{w} \in W^k} \left| \sum_{t \neq 0} \prod_{j=1}^{k} \widehat{S_{\mathbf{w}_j}}(t) \right| .$$

Step 1. Consider arbitrary phases $\mathbf{\Phi} = \left(\mathbf{\Phi}_{\mathbf{w}} : \mathbf{w} \in W^k \right)$. Then, by duality,

$$2 \cdot E_k(0) = \max_{\mathbf{\Phi}} \left| \sum_{\mathbf{w} \in W^k} \mathbf{\Phi}_{\mathbf{w}} \sum_{t \neq 0} \prod_{j=1}^{k} \widehat{S_{\mathbf{w}_j}}(t) \right| .$$

Step 2. Next, restrict to product phases $\mathbf{\Phi}_{\mathbf{w}} = \prod_{j=1}^{k} \phi_{\mathbf{w}_j}$. This restriction leads to a lower bound.

$$2 \cdot E_k(0) \geqslant \max_{(\phi_w \colon w \in W)} \left| \sum_{\mathbf{w} \in W^k} \sum_{t \neq 0} \prod_{j=1}^{k} \phi_{\mathbf{w}_j} \cdot \widehat{S_{\mathbf{w}_j}}(t) \right| = \max_{(\phi_w \colon w \in W)} \left| \sum_{t \neq 0} \left(\sum_{w \in W} \phi_w \cdot \widehat{S_w}(t) \right)^k \right| . \tag{24}$$

Step 3. Choose specific phases $\phi_0 = 1$ and, for others, ϕ_w defined below:

$$\phi_w = \begin{cases} \overline{\widehat{S_w}(-1)} / |\widehat{S_w}(-1)|, & \text{if } \widehat{S_w}(-1) \neq 0 \\ 1, & \text{otherwise.} \end{cases}$$

For arbitrary t, these specific choices of phases ϕ_w ensure that $\sum_{w \in W} \phi_w \cdot \widehat{S_w}(t)$ is always a real number. This comes from the following reasoning:

(1) For $w = 0$: $\widehat{S_w}(t) \in \mathbb{R}$.
(2) For $w \in W \setminus \{0\}$: Recall, $\widehat{S_{n-w}}(t) = \overline{\widehat{S_w}(t)}$ (because $x \in S_w$ implies $-x \in S_{n-w}$). Furthermore, $\phi_{n-w} = \overline{\phi_w}$ as well. So, $\phi_w \widehat{S_w}(t)$ and $\phi_{n-w} \widehat{S_{n-w}}(t)$ are conjugates of each other, and, hence, their sum is real.

Hence, $\sum_{w \in W} \phi_w \cdot \widehat{S_w}(t) \in \mathbb{R}$ for every t.

Step 4. After that, since k is even, for every t,

$$\left(\sum_{w \in W} \phi_w \cdot \widehat{S_w}(t) \right)^k \geqslant 0$$

As a result, continuing from Equation 24, further restricting the phases to our specific choices, we conclude that

$$\begin{aligned}
2 \cdot E_k(0) &\geqslant \sum_{t \neq 0} \underbrace{\left(\sum_{w \in W} \phi_w \widehat{S_w}(t) \right)^k}_{\geqslant 0} \geqslant \sum_{t \in \{-2^j \,:\, 0 \leqslant j < n\}} \underbrace{\left(\sum_{w \in W} \phi_w \widehat{S_w}(t) \right)^k}_{\geqslant 0} \\
&= n \cdot \left(\sum_{w \in W} \phi_w \widehat{S_w}(-1) \right)^k \\
&\text{(since } \widehat{S_w}(2^j t) = \widehat{S_w}(t) \text{ for } j \in \{0, 1, \ldots, n-1\} \text{ and } \#\{-2^j : 0 \leqslant j < n\} = n\text{)} \\
&= n \cdot \left(\sum_{w \in W} \left| \widehat{S_w}(-1) \right| \right)^k \qquad \text{(due to our specific choice of } \phi_w\text{)} \\
&\geqslant 10^{-k} \cdot n^{-k/2+1}. \qquad \text{(by Lemma 5.1; see Section 5.1)}
\end{aligned}$$

This completes the proof of Theorem 4. □

5.1. **Statement and Proof of Technical Lemma 5.1.**

Lemma 5.1 (Technical: Lower bound)**.** *Let $N = 2^n - 1$ where $n \geqslant 3$ and let $W = \{0, 1, \ldots, n-1\}$. Then, the following inequality holds*

$$\sum_{w \in W} \left| \widehat{S_w}(-1) \right| \geqslant \frac{1}{10} \cdot n^{-1/2}.$$

Proof of Lemma 5.1. Recall from Equation 3 that the slice coefficients have the generating-function representation

$$\widehat{S_w}(t) = \frac{1}{N} \sum_{x \in S_w} \mathrm{e}_N(-xt) = \frac{1}{N} \cdot [z^w] \prod_{j=0}^{n-1} \left(1 + z \cdot \mathrm{e}_N(-2^j t) \right) \tag{25}$$

Define $c_w := N\widehat{S}_w(-1)$, i.e.

$$c_w \;:=\; N\widehat{S}_w(-1) = [z^w] \prod_{j=0}^{n-1} \left(1 + z \cdot \mathrm{e}_N(2^j) \right), \qquad 0 \leqslant w < n.$$

It is hence sufficient to prove the following

$$\sum_{w \in W} |c_w| \geqslant \frac{N}{10} \cdot n^{-1/2}.$$

Other frequencies may carry comparable mass, but they are irrelevant for this one-frequency lower bound.

We first handle $n = 3, 4$ together, and then assume that $n \geqslant 5$.

Case 1: when $n \in \{3, 4\}$**.** Since, $c_0 = 1$, we have

$$\sum_{w \in W} |c_w| \geqslant |c_0| = 1 > \frac{N}{10} \cdot n^{-1/2}, \qquad n \in \{3, 4\}.$$

Case 2: when $n \geqslant 5$. We first introduce the following auxiliary variables:

$$a_w \coloneqq \sum_{\substack{0 \leqslant x < 2^{n-1} \\ \mathrm{wt}(x)=w}} \mathrm{e}_N(x) \qquad \text{and} \qquad a_{-1} := 0.$$

Note that

$$c_w = a_w + \mathrm{e}_N(2^{n-1}) \cdot a_{w-1} = a_w - \mathrm{e}_N\left(\frac{1}{2}\right) \cdot a_{w-1} = (a_w - a_{w-1}) + \left(1 - \mathrm{e}_N\left(\frac{1}{2}\right)\right) \cdot a_{w-1}.$$

So, we get the following lower bound:

$$\sum_{w \in W} |c_w| \geqslant \underbrace{\left(\sum_{w \in W} |a_w - a_{w-1}|\right)}_{\text{Main term}} - \underbrace{\left(\left|1 - \mathrm{e}_N\left(\frac{1}{2}\right)\right| \cdot \sum_{w \in W} |a_{w-1}|\right)}_{\text{Perturbation}} \tag{26}$$

First, we will show that the perturbation is small. Note that

$$\left|1 - \mathrm{e}_N\left(\frac{1}{2}\right)\right| \leqslant \frac{\pi}{N}, \text{ and } |a_{w-1}| \leqslant \binom{n-1}{w-1}$$

From these, the perturbation is upper-bounded as follows:

$$\text{perturbation} \leqslant \frac{\pi}{N} \cdot 2^{n-1} < \pi \tag{27}$$

Next, we lower bound the main term. Let $t \coloneqq \lfloor n/2 \rfloor$. Using the triangle inequality, we have

$$\sum_{w=0}^{t} |a_w - a_{w-1}| \geqslant |a_t - a_{-1}| \geqslant |\mathrm{Im}(a_t - a_{-1})| \geqslant \mathrm{Im}(a_t),$$
$$\sum_{w=t+1}^{n-1} |a_w - a_{w-1}| \geqslant |a_t - a_{n-1}| \geqslant |\mathrm{Im}(a_t - a_{n-1})| \geqslant \mathrm{Im}(a_t) - \mathrm{Im}(a_{n-1}) \geqslant \mathrm{Im}(a_t) - 1.$$

Therefore, the main term will be strictly greater than $2 \cdot \mathrm{Im}(a_t) - 1$. So, it will suffice to lower bound $\mathrm{Im}(a_t)$. Let us denote the j-th octant by the interval

$$\left[\frac{(j-1)\pi}{4}, \frac{j\pi}{4}\right), \qquad j = 1, \ldots, 8.$$

Note that, for every w, a_w is a sum of $\mathrm{e}_N(x)$ with $0 \leqslant x < 2^{n-1}$; since $0 \leqslant 2\pi x/N \leqslant \pi - \pi/N < \pi$, all these terms have nonnegative imaginary parts, and hence the non-selected terms may be discarded when lower-bounding $\mathrm{Im}(a_t)$. Our strategy will be to show that the $\mathrm{Im}(\mathrm{e}_N(x))$ contributions with arguments in octants 2 and 3 already suffice to reach our target lower bound. We use the following two observations.

- If the binary representation of x begins with 010 or 001, then $2^{n-3} \leqslant x \leqslant 3 \cdot 2^{n-3} - 1 < 3N/8$. Hence, $\pi/4 < 2\pi x/N < 3\pi/4$, and therefore we have $\mathrm{Im}(e_N(x)) \geqslant 1/\sqrt{2}$.
- The total number of such x with $\mathrm{wt}(x) = t$ is exactly $2 \cdot \binom{n-3}{t-1}$ because there are $t-1$ many remaining 1's to go into remaining $n-3$ many bit positions, with two prefix options 010 and 001.

Combining these two observations, we get

$$\mathrm{Im}(a_t) \geqslant \frac{1}{\sqrt{2}} \cdot 2 \cdot \binom{n-3}{t-1} = \sqrt{2} \cdot \binom{n-3}{t-1}$$

and hence, by Equation 26 and Equation 27,

$$\sum_{w\in W}|c_w| \geqslant 2\sqrt{2}\cdot\binom{n-3}{t-1} - (\pi+1) > 2\sqrt{2}\cdot\binom{n-3}{t-1} - \frac{29}{7} \qquad \text{(since } \pi < 22/7\text{)}$$

Hence, it remains to prove the following inequality for $n \geqslant 5$.

$$2\sqrt{2}\cdot\binom{n-3}{t-1} - \frac{29}{7} \geqslant \frac{1}{10}\cdot N\cdot n^{-1/2}.$$

For $n = 5$ and $n = 6$, we prove it by computing the exact values of binomial coefficients.

$$\text{If } n=5\colon \sum_{w\in W}|c_w| > 2\sqrt{2}\cdot\binom{2}{1} - \frac{29}{7} = 1.513\ldots > \frac{2^5-1}{10}\cdot 5^{-1/2} = 1.386\ldots$$

$$\text{If } n=6\colon \sum_{w\in W}|c_w| > 2\sqrt{2}\cdot\binom{3}{2} - \frac{29}{7} = 4.342\ldots > \frac{2^6-1}{10}\cdot 6^{-1/2} = 2.571\ldots$$

Now suppose $n \geqslant 7$. The central binomial coefficient estimation [MN23, Section 3.6, Proposition 3.6.2, pp. 96–97] gives us

$$\binom{n-3}{t-1} \geqslant \frac{2^{n-3}}{\sqrt{2(n-2)}}.$$

Therefore,

$$2\sqrt{2}\cdot\binom{n-3}{t-1} - \frac{29}{7} \geqslant \frac{2^{n-2}}{\sqrt{n-2}} - \frac{29}{7} > \frac{N/4}{\sqrt{n}} - \frac{29}{7} \overset{(*)}{>} \frac{N}{10\sqrt{n}},$$

where the last inequality $(*)$ holds for $n \geqslant 7$ and $(2^n-1)/\sqrt{n}$ is increasing for $n \geqslant 1$. □

6. Lower Bound: Moduli close to a Power of Two

We prove a general lower bound for any N that is close to a power of 2.

Theorem 5. *Let $k \geqslant 2$ and $n \geqslant 4$ be integers, let $m \in [0, n-1]$ be real, and let N be an integer satisfying $|N - 2^n| \leqslant 2^m$. Then the following bound holds.*

$$\max_{0\leqslant s<N} E_k(s) \geqslant \sqrt{\frac{1}{N}\sum_{0\leqslant s<N} E_k(s)^2} \geqslant \frac{1}{2}\cdot L^{1/2}\cdot(1/25)^k\cdot n^{-k/2},$$

where $L = \left(n - m - \frac{1}{2}\log_2 n - C\right)_+$, $C > 0$ is an absolute constant, and $(x)_+ = \max\{0, x\}$. In particular, when $m = (1-\varepsilon)n$ and $\varepsilon n \geqslant \log_2 n + 2C$, then (for sufficiently large n, as a function of ε)

$$\max_{0\leqslant s<N} E_k(s) \geqslant \frac{1}{2\sqrt{2}}\cdot\varepsilon^{1/2}\cdot(1/25)^k\cdot n^{-k/2+1/2}.$$

Appendix E proves Theorem 5. Note that Theorem 5 does *not* claim that $E_k(0)$ is large; it proves the existence by demonstrating that the (L^2-)average of $E_k(s)$ is large. Like the proof of Theorem 4, restrict to a specific type of product phases

$$\mathbf{\Phi}_{\mathbf{w}} := \prod_{j=1}^{k}\phi^{\mathbf{w}_j}$$

but this choice is different from that in Theorem 4. For this restriction, the lower bound turns out to be

$$\frac{1}{2}\cdot\sqrt{\sum_{t\neq 0}\underbrace{\left|\sum_{w\in W}\phi^w\cdot\widehat{S_w}(t)\right|^{2k}}_{T(t)^{2k}}}$$

We will identify a set H of candidate frequencies $H\subseteq\{-2^j\,:\,0\leqslant j<n\}$ such that $\#H\geqslant\varepsilon\cdot n/2$ and $T(t)\gg n^{-1/2}$, for every $t\in H$. To this end, because N is very close to 2^n, we argue that instead of the spectral properties of S_w, it suffices to investigate the "idealized set" of all weight-w elements of $\{0,1,\ldots,2^n-1\}$. The analog of $T(t)$ for these idealized sets admits a closed-form expression:

$$\frac{1}{N}\left|\prod_{j=0}^{n-1}\Big(1+\phi\cdot\mathrm{e}_N\big(-2^jt\big)\Big)\right|.$$

We make a careful choice $\phi=\exp(-\imath/\sqrt{n})$ and prove that these idealized quantities are $\gg n^{-1/2}$.

6.1. **Insecurity of Two Shares.** Finally, we present an attack on the $k=2$ case.

Theorem 6. *Let 2^n be a power of two closest to N. For $s=2^n-1 \pmod N$, we have*

$$E_2(s)\geqslant 1-\frac{|2^n-N|}{N}-\mathcal{O}\Big(n^{-1/2}\Big)\geqslant\frac{2}{3}-\mathcal{O}\Big(n^{-1/2}\Big).$$

Proof of Theorem 6. Let $Q\coloneqq 2^n$, $d\coloneqq|N-Q|$, and $\delta\coloneqq d/N$. Consider the event

$$A\coloneqq\Big\{(w_1,w_2)\in W^2\colon w_1+w_2=n\Big\}.$$

We first show that $D(s)(A)\geqslant 1-\delta$, where $s=Q-1 \pmod N$. Suppose first that $N\leqslant Q$, and here $d=Q-N$. For every $x\in\{d,d+1,\ldots,N-1\}$,

$$s-x \pmod N = N+d-1-x=Q-1-x.$$

Thus, x and $s-x \pmod N$ are complementary n-bit strings, and hence $\mathrm{wt}(x)+\mathrm{wt}(s-x \pmod N)=n$. There are $N-d$ such values of x.

If $N>Q$, then for every $x\in\{0,1,\ldots,Q-1\}$.

$$s-x \pmod N = Q-1-x,$$

Their lower n bits are complementary, and both leading bits are zero, so the same weight identity holds. Again, there are $N-d$ such values. Consequently,

$$D(s)(A)\geqslant 1-\frac{d}{N}=1-\delta.$$

For a uniformly random secret, the two shares are independent and uniform over G. Hence

$$D(U)(A)=\frac{1}{N^2}\cdot\#\big\{(x,y)\in\{0,\ldots,N-1\}^2\colon\mathrm{wt}(x)+\mathrm{wt}(y)=n\big\}.$$

If $N\leqslant Q$, then the cardinality of the set above is at most $\binom{2n}{n}$, because the concatenation of the n-bit representations of x and y is a $2n$-bit string of Hamming weight n. So, we have

$$D(U)(A)\leqslant\frac{1}{N^2}\binom{2n}{n}.$$

If $N>Q$, note that $N\leqslant(3/2)\cdot Q<2^{n+1}$. Hence, the cardinality of the set above is at most $\binom{2n+2}{n}$, because the concatenation of $(n+1)$-bit representations of x and y is an n-weight string of $(2n+2)$ bits.

$$D(U)(A)\leqslant\frac{1}{N^2}\binom{2n+2}{n}.$$

Since $Q = 2^n$ is a power of two nearest to N, comparison with $Q/2$ and $2Q$ gives

$$\frac{3}{4}Q \leqslant N \leqslant \frac{3}{2}Q.$$

Thus $d/N \leqslant 1/3$, while $N \asymp 2^n$ and the standard central-binomial estimate imply, in both cases, that

$$D(U)(A) \ll n^{-1/2}.$$

Finally, by the event characterization of total variation distance,

$$E_2(s) \geqslant D(s)(A) - D(U)(A) \geqslant 1 - \frac{d}{N} - \mathcal{O}\Big(n^{-1/2}\Big),$$

as claimed. □

References

[ADN+19] Divesh Aggarwal, Ivan Damgård, Jesper Buus Nielsen, Maciej Obremski, Erick Purwanto, João Ribeiro, and Mark Simkin. Stronger leakage-resilient and non-malleable secret sharing schemes for general access structures. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology – CRYPTO 2019, Part II*, volume 11693 of *Lecture Notes in Computer Science*, pages 510–539, Santa Barbara, CA, USA, August 18–22, 2019. Springer, Cham, Switzerland. doi:10.1007/978-3-030-26951-7_18. ↑6

[AMN+21] Donald Q. Adams, Hemanta K. Maji, Hai H. Nguyen, Minh L. Nguyen, Anat Paskin-Cherniavsky, Tom Suad, and Mingyuan Wang. Lower bounds for leakage-resilient secret-sharing schemes against probing attacks. In *IEEE International Symposium on Information Theory, ISIT 2021, Melbourne, Australia, July 12-20, 2021*, pages 976–981. IEEE, 2021. doi:10.1109/ISIT45174.2021.9518230. ↑3, 6

[BCO04] Eric Brier, Christophe Clavier, and Francis Olivier. Correlation power analysis with a leakage model. In Marc Joye and Jean-Jacques Quisquater, editors, *Cryptographic Hardware and Embedded Systems – CHES 2004*, volume 3156 of *Lecture Notes in Computer Science*, pages 16–29, Cambridge, Massachusetts, USA, August 11–13, 2004. Springer Berlin Heidelberg, Germany. doi:10.1007/978-3-540-28632-5_2. ↑3

[BDIR18] Fabrice Benhamouda, Akshay Degwekar, Yuval Ishai, and Tal Rabin. On the local leakage resilience of linear secret sharing schemes. In Hovav Shacham and Alexandra Boldyreva, editors, *Advances in Cryptology – CRYPTO 2018, Part I*, volume 10991 of *Lecture Notes in Computer Science*, pages 531–561, Santa Barbara, CA, USA, August 19–23, 2018. Springer, Cham, Switzerland. doi:10.1007/978-3-319-96884-1_18. ↑3, 6, 11

[BDIR21] Fabrice Benhamouda, Akshay Degwekar, Yuval Ishai, and Tal Rabin. On the local leakage resilience of linear secret sharing schemes. *Journal of Cryptology*, 34(2):10, April 2021. doi:10.1007/s00145-021-09375-2. ↑3, 6, 11

[BHM+26] Aniruddha Biswas, Jihun Hwang, Hemanta K. Maji, Ilya D. Shkredov, and Xiuyu Ye. Energy estimation of the hamming slice and its applications, 2026. arXiv Preprint arXiv:2609.08056. arXiv:2609.08056. ↑14

[BIS19] Andrej Bogdanov, Yuval Ishai, and Akshayaram Srinivasan. Unconditionally secure computation against low-complexity leakage. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology – CRYPTO 2019, Part II*, volume 11693 of *Lecture Notes in Computer Science*, pages 387–416, Santa Barbara, CA, USA, August 18–22, 2019. Springer, Cham, Switzerland. doi:10.1007/978-3-030-26951-7_14. ↑6

[Bla79] G. R. Blakley. Safeguarding cryptographic keys. *Proceedings of AFIPS 1979 National Computer Conference*, 48:313–317, 1979. ↑3

[BPRW16] Allison Bishop, Valerio Pastro, Rajmohan Rajaraman, and Daniel Wichs. Essentially optimal robust secret sharing with maximal corruptions. In Marc Fischlin and Jean-Sébastien Coron, editors, *Advances in Cryptology – EUROCRYPT 2016, Part I*, volume 9665 of *Lecture Notes in Computer Science*, pages 58–86, Vienna, Austria, May 8–12, 2016. Springer Berlin Heidelberg, Germany. doi:10.1007/978-3-662-49890-3_3. ↑6

[BS19] Saikrishna Badrinarayanan and Akshayaram Srinivasan. Revisiting non-malleable secret sharing. In Yuval Ishai and Vincent Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2019, Part I*, volume 11476 of *Lecture Notes in Computer Science*, pages 593–622, Darmstadt, Germany, May 19–23, 2019. Springer, Cham, Switzerland. doi:10.1007/978-3-030-17653-2_20. ↑6

[CG25] Eshan Chattopadhyay and Jesse Goodman. Leakage-resilient extractors against number-on-forehead protocols. In Michal Koucký and Nikhil Bansal, editors, *57th Annual ACM Symposium on Theory of Computing*, pages 604–614, Prague, Czechia, June 23–27, 2025. ACM Press. doi:10.1145/3717823.3718272. ↑6

[CGG+20] Eshan Chattopadhyay, Jesse Goodman, Vipul Goyal, Ashutosh Kumar, Xin Li, Raghu Meka, and David Zuckerman. Extractors and secret sharing against bounded collusion protocols. In *61st Annual Symposium on Foundations of Computer Science*, pages 1226–1242, Durham, NC, USA, November 16–19, 2020. IEEE Computer Society Press. doi:10.1109/FOCS46700.2020.00117. ↑6

[Cha02] Mei-Chu Chang. A polynomial bound in Freiman's theorem. *Duke Mathematical Journal*, 113(3):399–419, 2002. doi:10.1215/S0012-7094-02-11331-3. ↑6, 33

[CKOS21] Nishanth Chandran, Bhavana Kanukurthi, Sai Lakshmi Bhavana Obbattu, and Sruthi Sekar. Adaptive extractors and their application to leakage resilient secret sharing. In Tal Malkin and Chris Peikert, editors, *Advances in Cryptology – CRYPTO 2021, Part III*, volume 12827 of *Lecture Notes in Computer Science*, pages 595–624, Virtual Event, August 16–20, 2021. Springer, Cham, Switzerland. doi:10.1007/978-3-030-84252-9_20. ↑6

[CKOS22] Nishanth Chandran, Bhavana Kanukurthi, Sai Lakshmi Bhavana Obbattu, and Sruthi Sekar. Short leakage resilient and non-malleable secret sharing schemes. In Yevgeniy Dodis and Thomas Shrimpton, editors, *Advances in Cryptology – CRYPTO 2022, Part I*, volume 13507 of *Lecture Notes in Computer Science*, pages 178–207, Santa Barbara, CA, USA, August 15–18, 2022. Springer, Cham, Switzerland. doi:10.1007/978-3-031-15802-5_7. ↑6

[DDF14] Alexandre Duc, Stefan Dziembowski, and Sebastian Faust. Unifying leakage models: From probing attacks to noisy leakage. In Phong Q. Nguyen and Elisabeth Oswald, editors, *Advances in Cryptology – EUROCRYPT 2014*, volume 8441 of *Lecture Notes in Computer Science*, pages 423–440, Copenhagen, Denmark, May 11–15, 2014. Springer Berlin Heidelberg, Germany. doi:10.1007/978-3-642-55220-5_24. ↑6

[DF09] Persi Diaconis and Jason Fulman. Carries, shuffling, and an amazing matrix. *The American Mathematical Monthly*, 116(9):788–803, 2009. doi:10.4169/193009709X470108. ↑7

[DMMS21] Sébastien Duval, Pierrick Méaux, Charles Momin, and François-Xavier Standaert. Exploring crypto-physical dark matter and learning with physical rounding. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2021(1):373–401, 2021. URL: https://tches.iacr.org/index.php/TCHES/article/view/8738, doi:10.46586/tches.v2021.i1.373-401. ↑7

[DORS08] Yevgeniy Dodis, Rafail Ostrovsky, Leonid Reyzin, and Adam Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. *SIAM Journal on Computing*, 38(1):97–139, 2008. arXiv:https://doi.org/10.1137/060651380, doi:10.1137/060651380. ↑4

[DRS04] Yevgeniy Dodis, Leonid Reyzin, and Adam Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. In Christian Cachin and Jan Camenisch, editors, *Advances in Cryptology – EUROCRYPT 2004*, volume 3027 of *Lecture Notes in Computer Science*, pages 523–540, Interlaken, Switzerland, May 2–6, 2004. Springer Berlin Heidelberg, Germany. doi:10.1007/978-3-540-24676-3_31. ↑4

[EW11] Manfred Einsiedler and Thomas Ward. *Ergodic Theory: with a View Towards Number Theory*, volume 259 of *Graduate Texts in Mathematics*. Springer, London, 2011. doi:10.1007/978-0-85729-021-2. ↑8, 12

[FMM+24] Sebastian Faust, Loïc Masure, Elena Micheli, Maximilian Orlt, and François-Xavier Standaert. Connecting leakage-resilient secret sharing to practice: Scaling trends and physical dependencies of prime field masking. In Marc Joye and Gregor Leander, editors, *Advances in Cryptology – EUROCRYPT 2024, Part IV*, volume 14654 of *Lecture Notes in Computer Science*, pages 316–344, Zurich, Switzerland, May 26–30, 2024. Springer, Cham, Switzerland. doi:10.1007/978-3-031-58737-5_12. ↑3, 4, 6, 7, 10, 11, 13, 14, 31, 33

[FMM+25] Sebastian Faust, Loïc Masure, Elena Michel, Hai Hoang Nguyen, Maximilian Orlt, and François-Xavier Standaert. IP masking with generic security guarantees under minimum assumptions, and applications. In Goichiro Hanaoka and Bo-Yin Yang, editors, *Advances in Cryptology – ASIACRYPT 2025, Part II*, volume 16246 of *Lecture Notes in Computer Science*, pages 544–575, Melbourne, VIC, Australia, December 8–12, 2025. Springer, Singapore, Singapore. doi:10.1007/978-981-95-5096-8_17. ↑4

[FS09] Philippe Flajolet and Robert Sedgewick. *Analytic Combinatorics*. Cambridge University Press, Cambridge, 2009. URL: https://www.cambridge.org/core/books/analytic-combinatorics/7E37474C43E9B95C90BEDE082CF28708. ↑12

[FSS23] Aihua Fan, Jörg Schmeling, and Weixiao Shen. Multifractal analysis of generalized Thue–Morse trigonometric polynomials. *Asian J. Math.*, 27(4):589–620, 2023. ↑12

[FY19] Serge Fehr and Chen Yuan. Towards optimal robust secret sharing with security against a rushing adversary. In Yuval Ishai and Vincent Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2019, Part III*, volume 11478 of *Lecture Notes in Computer Science*, pages 472–499, Darmstadt, Germany, May 19–23, 2019. Springer, Cham, Switzerland. doi:10.1007/978-3-030-17659-4_16. ↑6

[GK18a] Vipul Goyal and Ashutosh Kumar. Non-malleable secret sharing. In Ilias Diakonikolas, David Kempe, and Monika Henzinger, editors, *50th Annual ACM Symposium on Theory of Computing*, pages 685–698, Los Angeles, CA, USA, June 25–29, 2018. ACM Press. doi:10.1145/3188745.3188872. ↑6

[GK18b] Vipul Goyal and Ashutosh Kumar. Non-malleable secret sharing for general access structures. In Hovav Shacham and Alexandra Boldyreva, editors, *Advances in Cryptology – CRYPTO 2018, Part I*, volume 10991 of *Lecture Notes in Computer Science*, pages 501–530, Santa Barbara, CA, USA, August 19–23, 2018. Springer, Cham, Switzerland. doi:10.1007/978-3-319-96884-1_17. ↑6

[HMM+23] Clément Hoffmann, Pierrick Méaux, Charles Momin, Yann Rotella, François-Xavier Standaert, and Balazs Udvarhelyi. Learning with physical rounding for linear and quadratic leakage functions. In Helena Handschuh and Anna Lysyanskaya, editors, *Advances in Cryptology – CRYPTO 2023, Part III*, volume 14083 of *Lecture Notes in Computer Science*, pages 410–439, Santa Barbara, CA, USA, August 20–24, 2023. Springer, Cham, Switzerland. doi:10.1007/978-3-031-38548-3_14. ↑7

[HMNY25] Jihun Hwang, Hemanta K. Maji, Hai H. Nguyen, and Xiuyu Ye. Leakage-resilience of shamir's secret sharing: Identifying secure evaluation places. In Niv Gilboa, editor, *ITC 2025: 6th Conference on Information-Theoretic Cryptography*, volume 343 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 3:1–3:20, Santa Barbara, CA, USA, August 16–17, 2025. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik. doi:10.4230/LIPIcs.ITC.2025.3. ↑6

[HMRS26] Clément Hoffmann, Pierrick Méaux, Mélissa Rossi, and François-Xavier Standaert. Learning with errors with output dependencies: Lwe, lwr, and physical learning problems under the same umbrella. In Magali Bardet and Ruben Niederhagen, editors, *Post-Quantum Cryptography - 17th International Workshop, PQCrypto 2026, Saint-Malo, France, April 14-16, 2026, Proceedings, Part I*, volume 16491 of *Lecture Notes in Computer Science*, pages 287–318. Springer, 2026. doi:10.1007/978-3-032-22695-2_10. ↑7

[HVW20] Carmit Hazay, Muthuramakrishnan Venkitasubramaniam, and Mor Weiss. The price of active security in cryptographic protocols. In Anne Canteaut and Yuval Ishai, editors, *Advances in Cryptology – EUROCRYPT 2020, Part II*, volume 12106 of *Lecture Notes in Computer Science*, pages 184–215, Zagreb, Croatia, May 10–14, 2020. Springer, Cham, Switzerland. doi:10.1007/978-3-030-45724-2_7. ↑6

[IS24] Yuval Ishai and Yifan Song. Leakage-tolerant circuits. In Marc Joye and Gregor Leander, editors, *Advances in Cryptology – EUROCRYPT 2024, Part IV*, volume 14654 of *Lecture Notes in Computer Science*, pages 196–225, Zurich, Switzerland, May 26–30, 2024. Springer, Cham, Switzerland. doi:10.1007/978-3-031-58737-5_8. ↑6

[ISW03] Yuval Ishai, Amit Sahai, and David Wagner. Private circuits: Securing hardware against probing attacks. In Dan Boneh, editor, *Advances in Cryptology – CRYPTO 2003*, volume 2729 of *Lecture Notes in Computer Science*, pages 463–481, Santa Barbara, CA, USA, August 17–21, 2003. Springer Berlin Heidelberg, Germany. doi:10.1007/978-3-540-45146-4_27. ↑3, 6

[KJJ99] Paul C. Kocher, Joshua Jaffe, and Benjamin Jun. Differential power analysis. In Michael J. Wiener, editor, *Advances in Cryptology – CRYPTO'99*, volume 1666 of *Lecture Notes in Computer Science*, pages 388–397, Santa Barbara, CA, USA, August 15–19, 1999. Springer Berlin Heidelberg, Germany. doi:10.1007/3-540-48405-1_25. ↑3

[KMS19] Ashutosh Kumar, Raghu Meka, and Amit Sahai. Leakage-resilient secret sharing against colluding parties. In David Zuckerman, editor, *60th Annual Symposium on Foundations of Computer Science*, pages 636–660, Baltimore, MD, USA, November 9–12, 2019. IEEE Computer Society Press. doi:10.1109/FOCS.2019.00045. ↑6

[Koc96] Paul C. Kocher. Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems. In Neal Koblitz, editor, *Advances in Cryptology – CRYPTO'96*, volume 1109 of *Lecture Notes in Computer Science*, pages 104–113, Santa Barbara, CA, USA, August 18–22, 1996. Springer Berlin Heidelberg, Germany. doi:10.1007/3-540-68697-5_9. ↑3

[KR19] Yael Tauman Kalai and Leonid Reyzin. A survey of leakage-resilient cryptography. In Oded Goldreich, editor, *Providing Sound Foundations for Cryptography: On the Work of Shafi Goldwasser and Silvio Micali*, pages 727–794. Association for Computing Machinery, New York, NY, USA, 2019. doi:10.1145/3335741.3335768. ↑6

[MBG+19] Mohammad Mahdi Mojahedian, Salman Beigi, Amin Gohari, Mohammad Hossein Yassaee, and Mohammad Reza Aref. A correlation measure based on vector-valued lp-norms. *IEEE Trans. Inf. Theory*, 65(12):7985–8004, 2019. doi:10.1109/TIT.2019.2937099. ↑7, 13

[Mes00] Thomas S. Messerges. Using second-order power analysis to attack DPA resistant software. In Çetin Kaya Koç and Christof Paar, editors, *Cryptographic Hardware and Embedded Systems – CHES 2000*, volume 1965 of *Lecture Notes in Computer Science*, pages 238–251, Worcester, Massachusetts, USA, August 17–18, 2000. Springer Berlin Heidelberg, Germany. doi:10.1007/3-540-44499-8_19. ↑3

[MN23] Jiří Matoušek and Jaroslav Nešetřil. *Invitation to Discrete Mathematics*. Clarendon Press / Oxford University Press, Oxford, reprint edition, 2023. URL: https://academic.oup.com/book/54999. ↑24

[MNP+21] Hemanta K. Maji, Hai H. Nguyen, Anat Paskin-Cherniavsky, Tom Suad, and Mingyuan Wang. Leakage-resilience of the Shamir secret-sharing scheme against physical-bit leakages. In Anne Canteaut and François-Xavier Standaert, editors, *Advances in Cryptology – EUROCRYPT 2021, Part II*, volume 12697 of *Lecture Notes in Computer Science*, pages 344–374, Zagreb, Croatia, October 17–21, 2021. Springer, Cham, Switzerland. doi:10.1007/978-3-030-77886-6_12. ↑6

[MNP+22] Hemanta K. Maji, Hai H. Nguyen, Anat Paskin-Cherniavsky, Tom Suad, Mingyuan Wang, Xiuyu Ye, and Albert Yu. Tight estimate of the local leakage resilience of the additive secret-sharing scheme & its consequences. In Dana Dachman-Soled, editor, *ITC 2022: 3rd Conference on Information-Theoretic Cryptography*, volume 230 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 16:1–16:19, Cambridge, MA, USA, July 5–7, 2022. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik. doi:10.4230/LIPIcs.ITC.2022.16. ↑3, 6

[MNPCW22] Hemanta K Maji, Hai H Nguyen, Anat Paskin-Cherniavsky, and Mingyuan Wang. Improved bound on the local leakage-resilience of shamir's secret sharing. In *2022 IEEE International Symposium on Information Theory (ISIT)*, pages 2678–2683. IEEE, 2022. ↑6

[MNPY24] Hemanta K. Maji, Hai H. Nguyen, Anat Paskin-Cherniavsky, and Xiuyu Ye. Constructing leakage-resilient Shamir's secret sharing: Over composite order fields. In Marc Joye and Gregor Leander, editors, *Advances in Cryptology – EUROCRYPT 2024, Part IV*, volume 14654 of *Lecture Notes in Computer Science*, pages 286–315, Zurich, Switzerland, May 26–30, 2024. Springer, Cham, Switzerland. doi:10.1007/978-3-031-58737-5_11. ↑6

[MR10] Christian Mauduit and Joël Rivat. Sur un problème de Gelfond: la somme des chiffres des nombres premiers. *Annals of Mathematics*, 171(3):1591–1646, 2010. doi:10.4007/annals.2010.171.1591. ↑7

[MSV20] Pasin Manurangsi, Akshayaram Srinivasan, and Prashant Nalini Vasudevan. Nearly optimal robust secret sharing against rushing adversaries. In Daniele Micciancio and Thomas Ristenpart, editors, *Advances in Cryptology – CRYPTO 2020, Part III*, volume 12172 of *Lecture Notes in Computer Science*, pages 156–185, Santa Barbara, CA, USA, August 17–21, 2020. Springer, Cham, Switzerland. doi:10.1007/978-3-030-56877-1_6. ↑6

[Ngu25] Hai H. Nguyen. Physical-bit leakage resilience of linear code-based secret sharing. In Serge Fehr and Pierre-Alain Fouque, editors, *Advances in Cryptology – EUROCRYPT 2025, Part VIII*, volume 15608 of *Lecture Notes in Computer Science*, pages 64–93, Madrid, Spain, May 4–8, 2025. Springer, Cham, Switzerland. doi:10.1007/978-3-031-91101-9_3. ↑6

[NRR06] Svetla Nikova, Christian Rechberger, and Vincent Rijmen. Threshold implementations against side-channel attacks and glitches. In Peng Ning, Sihan Qing, and Ninghui Li, editors, *ICICS 06: 8th International Conference on Information and Communication Security*, volume 4307 of *Lecture Notes in Computer Science*, pages 529–545, Raleigh, NC, USA, December 4–7, 2006. Springer Berlin Heidelberg, Germany. doi:10.1007/11935308_38. ↑3

[NRS11] Svetla Nikova, Vincent Rijmen, and Martin Schläffer. Secure hardware implementation of nonlinear functions in the presence of glitches. *Journal of Cryptology*, 24(2):292–321, April 2011. doi:10.1007/s00145-010-9085-7. ↑3, 6

[Pol19] Yury Polyanskiy. Hypercontractivity of spherical averages in Hamming space. *SIAM J. Discrete Math.*, 33(2):731–754, 2019. URL: https://doi.org, doi:10.1137/15M1046575. ↑8

[Sha79] Adi Shamir. How to share a secret. *Communications of the Association for Computing Machinery*, 22(11):612–613, November 1979. doi:10.1145/359168.359176. ↑3

[Shk11] Ilya D. Shkredov. Some applications of w. rudin's inequality to problems of combinatorial number theory. *Uniform Distribution Theory*, 6(2):95–116, 2011. URL: https://arxiv.org/abs/1002.1886. ↑6, 32, 33

[SV19] Akshayaram Srinivasan and Prashant Nalini Vasudevan. Leakage resilient secret sharing and applications. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology – CRYPTO 2019, Part II*, volume 11693 of *Lecture Notes in Computer Science*, pages 480–509, Santa Barbara, CA, USA, August 18–22, 2019. Springer, Cham, Switzerland. doi:10.1007/978-3-030-26951-7_17. ↑6

[SW23] Lukas Spiegelhofer and Michael Wallner. The binary digits of $n + t$. *Ann. Sc. Norm. Super. Pisa Cl. Sci. (5)*, 24(1):1–31, 2023. doi:10.26350/000532_000181. ↑7

Appendix A. Baseline Results

A.1. **Revisiting the Analysis of Faust et al.** This subsection recalls the Fourier argument of Faust et al. [FMM+24]. Throughout, we assume

$$N = 2^n - 1 \qquad \text{with } n \text{ prime;}$$

the modulus N itself need not be prime. For completeness, we review how the Mersenne rotation identity (Equation 8) gives nonzero doubling orbits of length n, and how their pointwise Fourier bound yields the insecurity estimate $E_k \ll_k n^{-k/4+1}$.

Proposition A.1 (Faust et al. [FMM+24, Theorem 6])**.** *Let $N = 2^n - 1$ where n is a prime number.* [5] *Then, for any $k \geqslant 3$ and $s \in G$,*

$$E_k(s) \ll_k n^{-k/4+1}$$

and this gives asymptotic security for $k \geqslant 5$.

Doubling-orbit structure. By the Mersenne rotation identity (Equation 8), multiplication by 2 modulo N acts as a one-position cyclic rotation of the n-bit representation and therefore preserves each slice S_w. The orbit length remains to be determined. For $x \neq 0$, consider

$$x \longmapsto 2x \longmapsto \cdots \longmapsto 2^{n-1}x \longmapsto 2^n x = x \pmod N.$$

This orbit has length exactly n. Let R denote cyclic rotation of an n-bit string by one position. Since R^n is the identity, every orbit under R has size dividing n. As n is prime, each orbit therefore has size either 1 or n. An orbit has size 1 only if the bit string is fixed by a one-position rotation, in which case all of its bits are equal. Its Hamming weight is then either 0 or n. Consequently, no element of S_w, for $1 \leqslant w \leqslant n-1$, can lie in an orbit of size 1. Thus, the doubling map partitions every nontrivial Hamming-weight slice S_w into orbits of length exactly n.

A pointwise Fourier bound. Let $\rho_w := |S_w|/N = \frac{1}{N}\binom{n}{w}$ denote the density of S_w. Since doubling permutes S_w,

$$\widehat{S_w}(t) = \widehat{S_w}(2t) = \cdots = \widehat{S_w}(2^{n-1}t).$$

For every nonzero frequency t, the frequencies $t, 2t, \ldots, 2^{n-1}t \pmod N$ are distinct by the same orbit argument. Parseval's identity therefore gives

$$\rho_w = \sum_{u \in G} \left|\widehat{S_w}(u)\right|^2 \geqslant \sum_{j=0}^{n-1} \left|\widehat{S_w}(2^j t)\right|^2 = n\left|\widehat{S_w}(t)\right|^2.$$

Consequently,

$$\max_{t \neq 0} \left|\widehat{S_w}(t)\right| \leqslant \sqrt{\frac{\rho_w}{n}}. \tag{28}$$

For the endpoint slice $S_0 = \{0\}$, the same estimate follows directly from $|\widehat{S_0}(t)| = 1/N$ and $N \geqslant n$.

[5] Faust et al. [FMM+24, Theorem 6] state their result for $p = 2^n - 1$ with n merely odd, but it rests on [FMM+24, Theorem 7] whose proof partitions the nonzero frequencies into classes closed under $\alpha \mapsto 2\alpha$ and asserts each class has at least n elements. This holds when n is prime, but fails for n odd composite: for example, take $n = 9$ and $\alpha = (001)^3 = 73$, its orbit has size 3 not $n = 9$. Their notations ($\mathbb{F}_p$, etc.) suggests they intended p to be a Mersenne *prime*, which forces n prime (as they also observed in [FMM+24, Remark 1]). We therefore make the condition explicit. Primality of p was also not used in their proof.

The resulting insecurity bound. We now insert Equation 28 into the Fourier expression for $E_k(s)$. For the first two factors, we retain the full Fourier sum and apply Cauchy–Schwarz; for the remaining $k-2$ factors, we use the pointwise estimate above. Thus,

$$\begin{aligned}
2 \cdot E_k &= \max_{s \in G} 2 \cdot E_k(s) \\
&= \max_{s \in G} \sum_{\mathbf{w} \in W^k} \left| \sum_{t \neq 0} \mathrm{e}_N(st) \prod_{j=1}^{k} \widehat{S_{\mathbf{w}_j}}(t) \right| && \text{(by Equation 7)} \\
&\leqslant \max_{s \in G} \sum_{\mathbf{w} \in W^k} \sum_{t \neq 0} \prod_{j=1}^{k} \left| \widehat{S_{\mathbf{w}_j}}(t) \right| && \text{(triangle inequality)} \\
&= \sum_{\mathbf{w} \in W^k} \sum_{t \neq 0} \prod_{j=1}^{k} \left| \widehat{S_{\mathbf{w}_j}}(t) \right| \\
&\leqslant \left(\sum_{w \in W} \sqrt{\rho_w / n} \right)^{k-2} \sum_{w_1, w_2 \in W} \sum_{t \neq 0} \left| \widehat{S_{w_1}}(t) \right| \left| \widehat{S_{w_2}}(t) \right| && \text{(by Equation 28)} \\
&\leqslant n^{-(k-2)/2} \cdot \left(\sum_{w \in W} \sqrt{\rho_w} \right)^{k-2} \sum_{w_1, w_2 \in W} \sqrt{\sum_{t \neq 0} \left| \widehat{S_{w_1}}(t) \right|^2} \cdot \sqrt{\sum_{t \neq 0} \left| \widehat{S_{w_2}}(t) \right|^2} \\
& && \text{(by Cauchy–Schwarz)} \\
&< n^{-(k-2)/2} \cdot \left(\sum_{w \in W} \sqrt{\rho_w} \right)^{k-2} \sum_{w_1, w_2 \in W} \sqrt{\rho_{w_1} \cdot \rho_{w_2}} && \text{(by Parseval's identity)} \\
&= n^{-(k-2)/2} \left(\sum_{w \in W} \sqrt{\rho_w} \right)^{k} \\
&\ll_k n^{-(k-2)/2} \cdot n^{k/4} = n^{-k/4+1} && \text{(by Appendix B.3, a standard sum)}
\end{aligned}$$

Here, the penultimate line uses the standard estimate

$$\sum_{w \in W} \sqrt{\rho_w} \ll n^{1/4}$$

Accordingly, this analysis establishes asymptotic security only for $k \geqslant 5$. Our upper-bound argument replaces this pointwise treatment with an iterative Fourier-energy contraction, allowing each additional share to contribute a net factor of order $n^{-1/2}$.

A.2. **A Sharper Baseline Bound via Chang's Lemma.** In this subsection only, assume that

$$N = 2^n - 1 \text{ is a Mersenne prime}$$

We keep the triangle-inequality framework of Faust et al. unchanged and improve only the pointwise estimate inserted into it.

Proposition A.2 (Chang-enhanced Faust et al. bound)**.** *If $N = 2^n - 1$ is a Mersenne prime, then, for every fixed $k \geqslant 3$ and every $s \in G$,*

$$E_k(s) \ll_k n^{-k/2+3/2} \cdot (\log n)^{(k-2)/2}. \tag{29}$$

In particular, this refinement proves asymptotic security for $k \geqslant 4$.

We use the following version of Chang's lemma stated in [Shk11, Theorem 1.2].

Imported Lemma A.3 (Chang's Lemma [Cha02, Lemma 3.1] and [Shk11, Theorem 1.2])**.** *Let G be a finite abelian group, $\Lambda \subseteq \widehat{G}$ be a dissociated set: a subset $\Lambda \subseteq G$ is said to be* dissociated *if*

$$\sum_{\lambda\in\Lambda} \varepsilon_\lambda \cdot \lambda = 0 \text{ with } \varepsilon_\lambda \in \{0, \pm 1\}$$

implies $\varepsilon_\lambda = 0$ for all $\lambda \in \Lambda$. Let $S \subseteq G$ be a subset such that $\#S = \delta \cdot \#G$ for some $\delta \in (0, 1]$. Then, for some absolute constant $C > 0$,

$$\sum_{\lambda\in\Lambda} \left|\widehat{S}(\lambda)\right|^2 \leqslant C \cdot \delta^2 \log(1/\delta)$$

Proof of Proposition A.2. Fix a nonzero frequency t. Since N is prime, t is invertible, and we claim that

$$\Lambda_t := \{t, 2t, \ldots, 2^{n-2}t\}$$

is dissociated. To see this, consider any linear combination evaluating to 0. Indeed, after multiplying this $\{0, \pm 1\}$-relation by t^{-1}, the relation becomes one among the distinct powers $1, 2, \ldots, 2^{n-2}$. Its absolute value is strictly smaller than N, so congruence modulo N is equality over the integers; uniqueness of binary expansion then makes the relation trivial.

By the doubling established in Appendix A.1, $\widehat{S_w}$ is constant on Λ_t. Chang's large-spectrum lemma (Imported Lemma A.3) therefore gives, for every $w \in W$,

$$\left|\widehat{S_w}(t)\right| \ll \rho_w \sqrt{\frac{\log(1/\rho_w)}{n-1}}. \tag{30}$$

Recall that Parseval's identity previously gave only $\leqslant \sqrt{\rho_w/n}$.

We now repeat only the final aggregation in the argument of Faust et al. [FMM+24]: retain two Fourier factors in ℓ_2, and apply (30) to the remaining $(k-2)$ factors. Parseval's and Cauchy-Schwarz yield

$$2 \cdot E_k(s) \ll_k \left(\sum_{w\in W} \sqrt{\rho_w}\right)^2 \cdot \left(\frac{1}{\sqrt{n-1}} \sum_{w\in W} \rho_w \sqrt{\log\left(\frac{1}{\rho_w}\right)}\right)^{k-2}. \tag{31}$$

The first factor is $O(n^{1/2})$, by the estimate in Appendix B.3 (also used in Appendix A.1). For the second, Jensen's inequality and the entropy bound give

$$\sum_{w\in W} \rho_w \sqrt{\log\left(\frac{1}{\rho_w}\right)} \leqslant \left(\sum_{w\in W} \rho_w \log\left(\frac{1}{\rho_w}\right)\right)^{1/2} \leqslant \sqrt{\log(\#W)} \leqslant \sqrt{\log n}.$$

Substitution into Equation 31 proves Equation 29. □

What this comparison establishes. The comparison isolates two independent limitations: dividing by the frequency t requires every nonzero frequency to be a unit, while the coefficient-wise triangle inequality is too costly in the three-share case. The energy argument in Section 3 removes both losses.

Appendix B. Proofs for Analytical Proxies

B.1. Reduction to F_{k-1}: Proof of Lemma 3.1.

Set

$$D' := (\mathrm{wt}(X_1), \ldots, \mathrm{wt}(X_{k-1})) \qquad \text{and} \qquad Y := X_1 + \cdots + X_{k-1}.$$

For a fixed secret $s \in G$, consider the deterministic map

$$\Psi_s \colon W^{k-1} \times G \longrightarrow W^k, \qquad \Psi_s(\mathbf{w}, x) := (\mathbf{w}, \mathrm{wt}(s - x)).$$

Then $\Psi_s(D', Y)$ has distribution $D(s)$. Moreover, $\Psi_s(D', X_k)$ has distribution $D(U)$. Indeed, after averaging the secret uniformly over G, the final share is uniform and independent of the first $(k-1)$

shares; and $s - X_k$ is likewise uniform over G and independent of D'. Hence, by the data processing inequality for the total variation distance,

$$E_k(s) = \big\|\Psi_s(D', Y) - \Psi_s(D', X_k)\big\|_{\mathrm{TV}} \leqslant \big\|(D', Y) - (D', X_k)\big\|_{\mathrm{TV}} = \Delta_{k-1}.$$

Thus, it remains to prove that $2\Delta_{k-1} \leqslant F_{k-1}$.

For $\mathbf{w} = (w_1, \dots, w_{k-1}) \in W^{k-1}$ and $x \in G$, Fourier inversion formula (Equation 4) gives

$$\Pr[D' = \mathbf{w},\, Y = x] = \frac{1}{N} \sum_{t \in G} \mathrm{e}_N(xt) \prod_{j=1}^{k-1} \widehat{S_{\mathbf{w}_j}}(t),$$

whereas

$$\Pr[D' = \mathbf{w},\, X_k = x] = \frac{1}{N} \prod_{j=1}^{k-1} \rho_{\mathbf{w}_j} \quad \text{and} \quad \rho_w := \frac{|S_w|}{N},$$

is precisely the zero-frequency term. Therefore

$$2\Delta_{k-1} = \sum_{\mathbf{w} \in W^{k-1}} \sum_{x \in G} \left| \frac{1}{N} \sum_{t \neq 0} \mathrm{e}_N(xt) \prod_{j=1}^{k-1} \widehat{S_{\mathbf{w}_j}}(t) \right|.$$

To bound the inner sum for a fixed $\mathbf{w} \in W^{k-1}$, instantiate Claim B.1 below with $\Omega = G$,

$$f_\omega \;\coloneqq\; \mathbb{1}_{\{-\omega\}} \qquad \text{and} \qquad z_t \;\coloneqq\; \prod_{j=1}^{k-1} \widehat{S_{\mathbf{w}_j}}(t).$$

Under our Fourier convention,

$$\widehat{f_\omega}(t) = \frac{1}{N} \mathrm{e}_N(\omega t).$$

Since $(f_\omega)_{\omega \in G}$ is a partition of G, we have $\big\|\sum_{\omega \in \Omega} |f_\omega|\big\|_{L^\infty} = 1$. Claim B.1 consequently yields

$$\sum_{x \in G} \left| \frac{1}{N} \sum_{t \neq 0} e_N(xt) \prod_{j=1}^{k-1} \widehat{S_{w_j}}(t) \right| \leqslant \left(\sum_{t \neq 0} \prod_{j=1}^{k-1} \left| \widehat{S_{w_j}}(t) \right|^2 \right)^{1/2}.$$

Summing over $\mathbf{w} \in W^{k-1}$ gives $2\Delta_{k-1} \leqslant F_{k-1}$. Together with the data-processing inequality above, this proves $2E_k(s) \leqslant 2\Delta_{k-1} \leqslant F_{k-1}$. □

Claim B.1. Let Ω be a finite set, and for each $\omega \in \Omega$ let $f_\omega \colon \mathbb{Z}/N\mathbb{Z} \to \mathbb{C}$. For any complex numbers $(z_t \colon 0 < t < N)$, we have

$$\sum_{\omega \in \Omega} \left| \sum_{t \neq 0} z_t \cdot \widehat{f_\omega}(t) \right| \quad \leqslant \left\| \sum_{\omega \in \Omega} |f_\omega| \right\|_{L^2} \cdot \sqrt{\sum_{t \neq 0} |z_t|^2} \quad \leqslant \left\| \sum_{\omega \in \Omega} |f_\omega| \right\|_{L^\infty} \cdot \sqrt{\sum_{t \neq 0} |z_t|^2}.$$

A useful way to interpret the claim is as follows. Suppose that, for every $x \in \mathbb{Z}/N\mathbb{Z}$, we have $\sum_{\omega \in \Omega} |f_\omega(x)| \leqslant B$, where $B > 0$ is an absolute constant. For instance, if the f_ω are the indicator functions of a partition of $\mathbb{Z}/N\mathbb{Z}$, then $B = 1$. Later we will also encounter function families for which $B = 2$. Under this pointwise upper bound, the rightmost expression simplifies to $\ll \sqrt{\sum_{t \neq 0} |z_t|^2}$.

Proof of Claim B.1. Consider the dual characterization. We remind the reader that phases are unimodular complex numbers. Over phases $\mathbf{\Phi} = (\mathbf{\Phi}_\omega \colon \omega \in \Omega)$ we have the following guarantee:

$$\sum_{\omega \in \Omega} \left| \sum_{t \neq 0} z_t \cdot \widehat{f_\omega}(t) \right| = \sup_{\mathbf{\Phi}} \left| \sum_{\omega \in \Omega} \mathbf{\Phi}_\omega \sum_{t \neq 0} z_t \cdot \widehat{f_\omega}(t) \right|$$

$$
\begin{aligned}
&= \sup_{\mathbf{\Phi}} \left| \sum_{t\neq 0} z_t \cdot \left(\sum_{\omega\in\Omega} \mathbf{\Phi}_\omega \widehat{f_\omega}(t) \right) \right| \\
&= \sup_{\mathbf{\Phi}} \left| \sum_{t\neq 0} z_t \cdot \widehat{f_{\mathbf{\Phi}}}(t) \right| && \text{(define } f_{\mathbf{\Phi}} \coloneqq \sum_{\omega\in\Omega} \mathbf{\Phi}_\omega \cdot f_\omega\text{)} \\
&\leqslant \sup_{\mathbf{\Phi}} \sqrt{\sum_{t\neq 0} |z_t|^2} \cdot \|f_{\mathbf{\Phi}}\|_{L^2}. && \text{(Cauchy-Schwarz and Parseval's)}
\end{aligned}
$$

Pointwise, $|f_\Phi(x)| \leqslant \sum_{\omega\in\Omega} |f_\omega(x)|$, and hence

$$\|f_\Phi\|_{L^2} \leqslant \left\| \sum_{\omega\in\Omega} |f_\omega| \right\|_{L^2}$$

This proves the first inequality. The second inequality follows from the monotonicity of norms. □

B.2. **Reduction from F_m to F_{m-1}: Proof of Lemma 3.3.** We begin with two technical claims.

Claim B.2. For any $t \neq 0$,

$$\sum_{w\in W} \frac{\left|\widehat{S_w}(t)\right|^2}{\rho_w} \leqslant n^{-1}$$

where $\rho_w = |S_w|/N$.

Proof of Claim B.2. Consider the function $f_t \colon G \to \mathbb{C}$ defined by

$$f_t = \sum_{w\in W} z_{w,t} \cdot S_w \quad \text{where } z_{w,t} \coloneqq \frac{\overline{\widehat{S_w}(t)}}{\rho_w}$$

Note that $\widehat{f_t} = \sum_{w\in W} z_{w,t} \cdot \widehat{S_w}$. In particular,

$$\widehat{f_t}(t) = \sum_{w\in W} \frac{\left|\widehat{S_w}(t)\right|^2}{\rho_w} \in \mathbb{R}_{\geqslant 0}$$

By rotation symmetry and n being prime, we have $\#\{t, 2t, \ldots, 2^{n-1}t\} = n$ and $\widehat{f_t}(t) = \widehat{f_t}(2t) = \cdots = \widehat{f_t}(2^{n-1}t)$. By Parseval's theorem, we get

$$
\begin{aligned}
n\left|\widehat{f_t}(t)\right|^2 = \sum_{0\leqslant j<n} \left|\widehat{f_t}(2^j t)\right|^2 &\leqslant \sum_u \left|\widehat{f_t}(u)\right|^2 \\
&= \frac{1}{N} \sum_x |f_t(x)|^2 && \text{(by Parseval's)} \\
&= \frac{1}{N} \sum_{w\in W} \rho_w N \cdot |z_{w,t}|^2 \\
&= \sum_{w\in W} \frac{\left|\widehat{S_w}(t)\right|^2}{\rho_w} = \widehat{f_t}(t).
\end{aligned}
$$

From the extreme LHS and RHS expressions, we conclude that $\widehat{f_t}(t) \leqslant n^{-1}$. □

Claim B.3. Let Ω be a finite set of indices. Let $a_t, b_{w,t}, \mu_w$ be nonnegative real numbers, and assume that $\mu_w = 0$ implies $b_{w,t} = 0$ for every $t \neq 0$. Put $\Omega_+ = \{w \in \Omega \colon \mu_w > 0\}$. Suppose that $\alpha \geqslant 0$ satisfies

$$\sum_{w\in\Omega_+} \frac{b_{w,t}}{\mu_w} \leqslant \alpha \qquad \text{for every } t \neq 0.$$

Then,

$$\sum_{w\in\Omega} \sqrt{\sum_{t\neq 0} a_t \cdot b_{w,t}} \leqslant \alpha^{1/2} \cdot \sqrt{\sum_{w\in\Omega} \mu_w} \cdot \sqrt{\sum_{t\neq 0} a_t}.$$

Proof of Claim B.3. The terms with $w \notin \Omega_+$ vanish. If $\Omega_+ = \emptyset$, both sides are zero; otherwise, apply the weighted Cauchy–Schwarz argument below over Ω_+.

$$\begin{aligned}
\sum_{w\in\Omega_+} \sqrt{\sum_{t\neq 0} a_t \cdot b_{w,t}} &= \sum_{w\in\Omega_+} \sqrt{\mu_w} \cdot \sqrt{\sum_{0<t<N} a_t \cdot b_{w,t}/\mu_w} \\
&\leqslant \sqrt{\sum_{w\in\Omega_+} \mu_w} \cdot \sqrt{\sum_{0<t<N} a_t \sum_{w\in\Omega_+} b_{w,t}/\mu_w} \qquad \text{(by Cauchy-Schwarz)} \\
&\leqslant \left(\max_{0<t<N} \sum_{w\in\Omega_+} b_{w,t}/\mu_w\right)^{1/2} \cdot \sqrt{\sum_{w\in\Omega_+} \mu_w} \cdot \sqrt{\sum_{t\neq 0} a_t}.
\end{aligned}$$

The claim follows since

$$\max_{0<t<N} \sum_{w\in\Omega_+} \frac{b_{w,t}}{\mu_w} \leqslant \alpha \quad \text{and} \quad \sum_{w\in\Omega_+} \mu_w = \sum_{w\in\Omega} \mu_w.$$

□

Now, to prove Lemma 3.3, we use these two technical claims.

$$\begin{aligned}
F_m &= \sum_{\mathbf{w}\in W^m} \sqrt{\sum_{t\neq 0} \prod_{j=1}^{m} \left|\widehat{S_{w_j}}(t)\right|^2} \\
&= \sum_{\mathbf{w}\in W^{m-1}} \sum_{w\in W} \sqrt{\sum_{t\neq 0} \underbrace{\prod_{j=1}^{m-1} \left|\widehat{S_{w_j}}(t)\right|^2}_{a_t :=} \cdot \underbrace{\left|\widehat{S_w}(t)\right|^2}_{b_{w,t}}} \\
&\leqslant \sum_{\mathbf{w}\in W^{m-1}} n^{-1/2} \cdot \sqrt{1} \cdot \sqrt{\sum_{t\neq 0} \prod_{j=1}^{m-1} \left|\widehat{S_{w_j}}(t)\right|^2} \\
&= n^{-1/2} \cdot F_{m-1}. \qquad \text{(by Claim B.3 and } \mu_w = \#S_w/N \text{ and } \alpha = n^{-1} \text{ from Claim B.2)}
\end{aligned}$$

This completes the proof of Lemma 3.3.

B.3. **Estimating F_1: Proof of Lemma 3.4.** We present a self-contained elementary upper bound. Let us begin by first upper-bounding the following expression.

$$\sum_{0\leqslant w\leqslant n} \sqrt{\binom{n}{w} \cdot 2^{-n}}$$

Set $p_w \coloneqq \binom{n}{w} \cdot 2^{-n}$ and $a_w = 1 + (w - n/2)^2/n$. Then we have

$$\begin{aligned}
\sum_{0 \leqslant w \leqslant n} \sqrt{p_w} &\leqslant \left(\sum_{0 \leqslant w \leqslant n} p_w a_w \right)^{1/2} \cdot \left(\sum_{0 \leqslant w \leqslant n} a_w^{-1} \right)^{1/2} \\
&= \left(1 + \frac{\mathrm{Var}(X)}{n} \right)^{1/2} \cdot \left(\sum_{0 \leqslant w \leqslant n} \frac{1}{1 + (w - n/2)^2/n} \right)^{1/2} && (X \sim \mathrm{Bin}(n, 1/2)) \\
&\leqslant \frac{\sqrt{5}}{2} \cdot 2 \left(1 + \int_0^{\infty} \frac{1}{1 + t^2/n} \mathrm{d}t \right)^{1/2} && (\text{because } \mathrm{Var}(X) = n/4) \\
&= \sqrt{5} \cdot \left(1 + \sqrt{n} \int_0^{\pi/2} \mathrm{d}\theta \right)^{1/2} && (\text{substituting } t = \sqrt{n} \cdot \tan\theta) \\
&\leqslant \sqrt{5} + \sqrt{5\pi/2} \cdot n^{1/4} \ll n^{1/4}. && (\text{Jensen's inequality on } \sqrt{\cdot})
\end{aligned}$$

Using this estimate, we can upper bound F_1.

$$F_1 = \sum_{w \in W} \sqrt{\sum_{t \neq 0} \left| \widehat{S_w}(t) \right|^2} \asymp \sum_{0 \leqslant w \leqslant n} \sqrt{\binom{n}{w} \cdot 2^{-n}} \ll n^{1/4}.$$

after using Parseval's identity and accommodating some edge terms. □

B.4. **Tighter F_2 Upper Bound (Lemma B.4).** When n is prime, note that by our one-share contraction and F_1 estimation, we already have the upper bound $F_2 \leqslant n^{-1/2} F_1 \ll n^{-1/4}$. The objective here is to improve the upper bound to $n^{-1/2} \cdot \mathsf{poly}(\log n)$, roughly a further $n^{-1/4}$ improvement.

Lemma B.4. *Let $N = 2^n - 1$ for any integer $n \geqslant 3$.*

$$A \coloneqq \sum_{w \in W} \left(\sum_{t \neq 0} \left| \widehat{S_w}(t) \right|^4 \right)^{1/4} \ll n^{-1/4} \cdot \sqrt{\log n}.$$

In particular, as a consequence, $F_2 \ll n^{-1/2} \cdot \log n$.

Proof of Lemma B.4. From the upper bound on A, the bound $F_2 \leqslant A^2 \ll n^{-1/2} \cdot \log n$ follows from the Cauchy-Schwarz inequality.

For the first bound, we will use Imported Lemma 3.5. From this technical lemma, the proof of the theorem follows directly. Consider the contribution to A by the central weights $W_{\mathrm{cent}} \coloneqq \left\{ w \in W : |n/2 - w| \leqslant \sqrt{n \log n} \right\}$.

$$Q_1 = \sum_{w \in W_{\mathrm{cent}}} \left(\sum_{t \neq 0} \left| \widehat{S_w}(t) \right|^4 \right)^{1/4} \ll \#W_{\mathrm{cent}} \cdot n^{-3/4} \ll n^{-1/4} \cdot \sqrt{\log n}.$$

The contribution from the remaining weights is

$$\begin{aligned}
Q_2 &= \sum_{w \in W \setminus W_{\mathrm{cent}}} \left(\sum_{t \neq 0} \left| \widehat{S_w}(t) \right|^4 \right)^{1/4} \\
&\leqslant \sum_{w \in W \setminus W_{\mathrm{cent}}} \left(\sum_{t \neq 0} \left| \widehat{S_w}(t) \right|^2 \right)^{1/2} && (\text{for } a, b \geqslant 0,\ a^2 + b^2 \leqslant (a+b)^2)
\end{aligned}$$

$$\begin{aligned}
&\leqslant \sum_{w\in W\setminus W_{\mathrm{cent}}} \sqrt{\rho_w} && \text{(by Parseval's identity)}\\
&\leqslant \sqrt{|W\setminus W_{\mathrm{cent}}|}\cdot\sqrt{\sum_{w\in W\setminus W_{\mathrm{cent}}}\rho_w} && \text{(by Cauchy–Schwarz)}\\
&\leqslant \sqrt{n}\cdot\sqrt{2^n/N}\cdot\sqrt{\Pr\left[|\mathrm{Bin}(n,1/2)-n/2|>\sqrt{n\log n}\right]}\\
&< \sqrt{n}\cdot\sqrt{2^n/N}\cdot\sqrt{2\exp(-2\log n)} && \text{(by Chernoff bound)}\\
&\leqslant \sqrt{2^{n+1}/N}\cdot n^{-1/2} < 2n^{-1/2}.
\end{aligned}$$

After this, $A = Q_1 + Q_2 \ll n^{-1/4}\cdot\sqrt{\log n}$ upper bound is immediate. □

Appendix C. Full Range Model

Throughout this appendix, we write $W = W_{\mathrm{BIN}}$.

C.1. **Proof of Lemma 4.3.** We will use Claim B.1 with $f_w = B_w$, and in that claim, we will use the fact that $\left\| \sum_{w\in W} |B_w| \right\|_{L^\infty} \leqslant 2$. Then Lemma 4.3 is immediate at this point:

$$\begin{aligned}
2\cdot E_k^{\mathrm{BIN}}(N,s) &= \sum_{\mathbf{w}\in W^k}\left|\sum_{t\neq 0}\mathrm{e}_N(st)\prod_{j=1}^{k}\widehat{B_{\mathbf{w}_j}}(t)\right|\\
&= \sum_{\mathbf{w}'\in W^{k-1}}\sum_{\mathbf{w}_k\in W}\left|\sum_{t\neq 0}\underbrace{\mathrm{e}_N(st)\prod_{j=1}^{k-1}\widehat{B_{\mathbf{w}'_j}}(t)}_{z_t\,:=}\cdot\widehat{B_{\mathbf{w}_k}}(t)\right| && \text{(rearranging)}\\
&\leqslant \sum_{\mathbf{w}'\in W^{k-1}} 2\cdot\left(\sum_{t\neq 0}\prod_{j=1}^{k-1}\left|\widehat{B_{\mathbf{w}'_j}}(t)\right|^2\right)^{1/2} = 2\cdot F_{k-1}^{\mathrm{BIN}}(N). && \text{(by Claim B.1)}
\end{aligned}$$

as desired. □

C.2. **Proof of Lemma 4.4.** We apply Claim B.3 with $\mu_w = \rho_w(N)$, $\alpha = \alpha^{\mathrm{BIN}}(N)$, and $\sum_{w\in W}\rho_w(N) = Q/N$.

$$\begin{aligned}
F_m^{\mathrm{BIN}}(N) &= \sum_{\mathbf{w}'\in W^{m-1}}\sum_{\mathbf{w}_m\in W}\left(\sum_{t\neq 0}\underbrace{\prod_{j=1}^{m-1}\left|\widehat{B_{\mathbf{w}'_j}}(t)\right|^2}_{a_t}\cdot\underbrace{\left|\widehat{B_{\mathbf{w}_m}}(t)\right|^2}_{b_{\mathbf{w}_m,t}}\right)^{1/2}\\
&\leqslant \sum_{\mathbf{w}'\in W^{m-1}}\alpha^{\mathrm{BIN}}(N)^{1/2}\cdot(Q/N)^{1/2}\cdot\left(\sum_{t\neq 0}\prod_{j=1}^{m-1}\left|\widehat{B_{\mathbf{w}'_j}}(t)\right|^2\right)^{1/2} && \text{(By Claim B.3)}\\
&= (Q/N)^{1/2}\cdot\alpha^{\mathrm{BIN}}(N)^{1/2}\cdot F_{m-1}^{\mathrm{BIN}}(N).
\end{aligned}$$

□

C.3. **Proof of Lemma 4.6.**

$$\begin{aligned}
F_1^{\mathrm{BIN}}(N) &= \sum_{w\in W}\left(\sum_{t\neq 0}\left|\widehat{B_w}(t)\right|^2\right)^{1/2} \\
&< \sum_{w\in W}\left(\frac{1}{N}\sum_{0\leqslant x<N} B_w(x)^2\right)^{1/2} && \text{(by Parseval's)} \\
&\leqslant \sum_{w\in W}\left(\frac{2}{N}\cdot\binom{n}{w}\right)^{1/2} && \text{(because } \sum_{0\leqslant x<N} B_w(x)=\tbinom{n}{w} \text{ and } B_w(x)\leqslant 2\text{)} \\
&\ll \frac{1}{N^{1/2}}\cdot Q^{1/2}n^{1/4} \asymp n^{1/4}. && \text{(see Lemma 3.4 and since } Q/N\leqslant 2\text{)}
\end{aligned}$$

□

C.4. **Integral Estimation: Proof of Lemma 4.11.** We let again $\phi_j = \{2^j t/N\}$, $p = w/n$, and $L_p(u) = |(1-p)+p\cdot\exp(\imath u)|$. For $|w-n/2|\leqslant\sqrt{n\log n}$, our target is to prove the following integral

$$\begin{aligned}
I &= \int_{-\pi}^{\pi}\left|\prod_{j=0}^{n-1}(1-p)+p\cdot\exp\left(\imath\cdot\left(\theta-\frac{2\pi\cdot 2^j t}{N}\right)\right)\right|\mathrm{d}\theta \\
&= \int_{-\pi}^{\pi}\left|\prod_{j=0}^{n-1}(1-p)+p\cdot\exp\left(\imath\cdot\left(\theta+\frac{2\pi\cdot 2^j t}{N}\right)\right)\right|\mathrm{d}\theta = \int_{-\pi}^{\pi}\prod_{j=0}^{n-1}L_p(\theta+2\pi\phi_j)\,\mathrm{d}\theta.
\end{aligned}$$

Thus, to prove Lemma 4.11, we need to prove the following inequality:

$$\text{Target:}\qquad \int_{-\pi}^{\pi}\prod_{j=0}^{n-1}L_p(\theta+2\pi\phi_j)\mathrm{d}\theta \ll \frac{|1-2p|}{\sqrt{n}}+\frac{1}{n}. \tag{32}$$

Let $\Omega=\{0,1,\ldots,n-1\}$ and, for $J\subseteq\Omega$, define

$$P_J(\theta) \coloneqq \prod_{j\in J}L_p(\theta+2\pi\phi_j)$$

Claim C.1. Assume that $|N-2^n|\leqslant 2^n/n$ and $p\in[0,1]$. For $n\geqslant 2$ and every $1\leqslant t<N$, there is an interval $J\subseteq\Omega$ such that

(1) $1\leqslant \#J\leqslant\left\lceil\frac{1}{2}\log_2 n\right\rceil$, and
(2) $P_J(\theta)\ll|1-2p|+|\theta|+n^{-1/2}$, for all $|\theta|\leqslant\pi$.

We prove Claim C.1 below after this proof. For now, we use Claim C.1 to prove our target Equation 32 which proves our goal Lemma 4.11. Note that $\Omega\setminus J$ splits into (at most) two intervals, say of size s and t.

Recall that $w\in W_{\mathrm{cent}}$. Now, in W_{cent}, consider the smallest relative weight

$$\tau(n) := \frac{n/2-\sqrt{n\ln n}}{n}$$

$\tau(n)>0$ and increasing for $n\geqslant n_0:=9$. We choose $\varepsilon_0:=\tau(n_0)$. Then, we will have $\varepsilon_0\leqslant w/n\leqslant(1-\varepsilon_0)$ for all $n\geqslant n_0$ and $w\in W_{\mathrm{cent}}$.[6]

[6]In the current proof, the emphasis is not on optimizing the constants, but on getting the asymptotics correct. If one wants to specialize the analysis for larger n_0, then a larger ε_0 can be chosen. However, the objective was to obtain results for as many values of n as possible, so we chose the smallest n_0.

We use Claim C.3.1. and Lemma C.4 to upper bound

$$P_{\Omega\setminus J}(\theta) \leqslant \exp\left(-\frac{2\varepsilon_0(1-\varepsilon_0)}{10}\cdot(s-1)_+\sin^2\left(\frac{\theta}{2}\right)\right)\cdot\exp\left(-\frac{2\varepsilon_0(1-\varepsilon_0)}{10}\cdot(t-1)_+\sin^2\left(\frac{\theta}{2}\right)\right).$$

Since, $\#J \leqslant \lceil \frac{1}{2}\log_2 n\rceil$, we have

$$(s-1)_+ + (t-1)_+ \geqslant (s+t) - 2 \geqslant n - \left\lceil\frac{1}{2}\log_2 n\right\rceil - 2 \geqslant n/2. \qquad \text{(for } n \geqslant 8\text{)}$$

Using the fact that $\sin(|\theta|/2) \geqslant |\theta|/\pi$, when $|\theta| \leqslant \pi$, we get the bound:

$$P_{\Omega\setminus J}(\theta) \leqslant \exp(-Cn\theta^2),$$

where $C > 0$ is an absolute constant.

After that, using Claim C.1, we have

$$P_\Omega(\theta) \ll \Big(\,|1-2p| + |\theta| + n^{-1/2}\,\Big)\cdot\exp(-Cn\theta^2).$$

Therefore, using the standard Gaussian integral, we get:

$$\int_{-\pi}^{\pi} P_\Omega(\theta)\mathrm{d}\theta \ll \frac{|1-2p|}{\sqrt{n}} + \frac{1}{n},$$

completing the proof of Lemma 4.11. □

Proof of Claim C.1. Without loss of generality assume that $0 < t/N \leqslant 1/2$, because (if needed) one can consider $N - t$ instead of t to make it $< 1/2$, which leads to the same answer because of evenness and 2π periodicity of $L_p(\cdot)$ function.

Define

$$h := \lfloor\log_2(N/t)\rfloor,\ h' := n - h,\ \text{and}\ \ell_0 := \left\lceil\frac{1}{2}\log_2 n\right\rceil.$$

Then, $2^{-(h+1)} < t/N \leqslant 2^{-h}$, and $1 \leqslant h \leqslant n$ (roughly speaking, the first non-zero bit appears at h-th position).

Case: $h' < \ell_0$**.** (a huge run of 0's in the head). Write $t = 2^a b$, where b is odd. Consider $J := \{j^*\}$, where $j^* = n - a - 1$. Note that $t < 2^n$ because $t/N \leqslant 1/2$ and $N < 2^{n+1}$. Therefore, $J \subseteq \Omega$.

After that, we want to prove that b is very small. Note:

$$b \leqslant t \leqslant N\cdot 2^{-h} < 2^{n+1-h} = 2^{h'+1} \leqslant 2^{\ell_0} \leqslant 2\cdot\sqrt{n}.$$

Below, we will use the notation $\|x\|_{\mathbb{T}}$ to denote the distance of x from the nearest integer; here $\mathbb{T} = \mathbb{R}/\mathbb{Z}$. Our objective, next, is to show that ϕ_{j^*} is close to $1/2$.

$$\begin{aligned}
\left\|\{\phi_{j^*}\} - \frac{1}{2}\right\|_{\mathbb{T}} &= \left\|\phi_{j^*} - \frac{1}{2}\right\|_{\mathbb{T}} && \text{(because } \|\{a\}+b\|_{\mathbb{T}} = \|a+b\|_{\mathbb{T}}\text{)}\\
&= \left\|2^{n-j^*-1}\cdot\frac{2^a b}{N} - \frac{1}{2}\right\|_{\mathbb{T}} = \left\|\frac{b}{2}\cdot\frac{2^n}{N} - \frac{1}{2}\right\|_{\mathbb{T}}\\
&= \left\|\frac{b}{2}\cdot\frac{2^n}{N} - \frac{b}{2}\right\|_{\mathbb{T}} && \text{(because } b \text{ is odd)}\\
&\leqslant \frac{b}{2}\cdot\frac{|N-2^n|}{N} && \text{(because } \|ab\|_{\mathbb{T}} \leqslant |a|\cdot|b|\text{)}\\
&\leqslant \sqrt{n}\cdot\frac{2^n}{Nn} \leqslant 2n^{-1/2}
\end{aligned}$$

Using this fact, the distance between $\theta + 2\pi\phi_{j^*}$ and $\pi + 2\pi\mathbb{Z}$ is

$$\leqslant |\theta| + 2\pi \left\| \{\phi_{j^*}\} - \frac{1}{2} \right\|_{\mathbb{T}} \ll |\theta| + n^{-1/2}.$$

Since these two angles are close, by using Claim C.3.2., we have:

$$P_J(\theta) \ll |1-2p| + |\theta| + n^{-1/2}.$$

Case: $h' \geqslant \ell_0$**.** Define $J := \{h-1, h, \ldots, h+\ell_0-2\}$. It is clear that $J \subseteq \Omega$ and $1/4 < 2^{h-1}t/N \leqslant 1/2$. Suppose we can prove that

$$\text{Leftover target:} \qquad P_J(\theta) \ll |1-2p| + |\theta| + 2^{-\#J},$$

then, we are done. Because,

$$P_J(\theta) \ll |1-2p| + |\theta| + 2^{-\#J} = |1-2p| + |\theta| + 2^{-\ell_0} \leqslant |1-2p| + |\theta| + n^{-1/2},$$

The leftover target is a special case of Claim C.2 below (with $v = \ell_0$ and $x = 2^{h-1}t/N$). □

Claim C.2. For every $p \in [0,1]$, integer $v \geqslant 1$, $x \in [1/4, 3/4]$, and $\theta \in \mathbb{R}$

$$\prod_{u=0}^{v-1} L_p(\theta + 2\pi\{2^u x\}) \ll |1-2p| + |\theta| + 2^{-v}.$$

Proof of Claim C.2. We begin with a few identities and inequalities:

(1) Let $(1-\alpha) = (1-2p)^2$. Then,

$$L_p(u)^2 = (1-\alpha) + \alpha \cdot \cos^2(u/2) \tag{33}$$

This is because,

$$\begin{aligned} L_p(u)^2 &= \Big((1-p) + p\exp(\imath u)\Big) \cdot \Big((1-p) + p\exp(-\imath u)\Big) \\ &= (1-p)^2 + p^2 + 2p(1-p)\cos(u) \\ &= (1-2p)^2 + 2p(1-p)\Big(\cos(u) + 1\Big) \\ &= (1-2p)^2 + 4p(1-p)\cos^2(u/2) \\ &= (1-\alpha) + \alpha \cdot \cos^2(u/2) \qquad (\text{let } (1-\alpha) = (1-2p)^2) \end{aligned}$$

(2) For $\alpha, x, y \in [0,1]$, we have:

$$\Big((1-\alpha) + \alpha \cdot x\Big) \cdot \Big((1-\alpha) + \alpha \cdot y\Big) \leqslant \Big((1-\alpha) + \alpha \cdot xy\Big) \tag{34}$$

(3) For unimodular ζ, ξ and $k \geqslant 0$, the following bound holds.

$$\left|(1-\xi)\prod_{j=0}^{k-1}\Big(1+\zeta\xi^{2^j}\Big)\right| \leqslant 2 + 2^{k+1}|1-\zeta|. \tag{35}$$

To show this, we proceed by induction on $k \in \{0, 1, \ldots\}$. For the base case $k = 0$, the LHS is $|1-\xi| \leqslant 2$, which is obviously smaller than the RHS.

Now, we go to the inductive step. By the induction hypothesis (with ζ and ξ^2 and $(k-1)$), we have

$$\left|(1-\xi^2)\prod_{j'=0}^{k-2}\Big(1+\zeta\big(\xi^2\big)^{2^{j'}}\Big)\right| \leqslant 2 + 2^k|1-\zeta|$$

Next, manipulate the LHS of the inductive hypothesis:

$$\begin{aligned}
&\left|(1-\xi)\prod_{j=0}^{k-1}\left(1+\zeta\xi^{2^j}\right)\right| \\
&= \left|(1-\xi)(1+\zeta\xi)\prod_{j'=0}^{k-2}\left(1+\zeta\left(\xi^2\right)^{2^{j'}}\right)\right| \\
&\leqslant \left|(1-\xi)(1+\xi)\prod_{j'=0}^{k-2}\left(1+\zeta\left(\xi^2\right)^{2^{j'}}\right)\right| + \left|(1-\xi)(\zeta\xi-\xi)\prod_{j'=0}^{k-2}\left(1+\zeta\left(\xi^2\right)^{2^{j'}}\right)\right| \\
&\qquad\qquad \text{(by triangle inequality)} \\
&\leqslant 2+2^k|1-\zeta| + \left|(1-\xi)(\zeta\xi-\xi)\prod_{j'=0}^{k-2}\left(1+\zeta\left(\xi^2\right)^{2^{j'}}\right)\right| \qquad \text{(by the inductive hypothesis)} \\
&= 2+2^k|1-\zeta| + \left|(1-\xi)(1-\zeta)\prod_{j'=0}^{k-2}\left(1+\zeta\left(\xi^2\right)^{2^{j'}}\right)\right| \qquad \text{(because } \xi \text{ is unimodular)} \\
&\leqslant 2+2^k|1-\zeta| + 2^k|1-\zeta| \qquad \text{(all factors are at most 2 in magnitude)} \\
&= 2+2^{k+1}|1-\zeta|,
\end{aligned}$$

which completes the induction.

Now, to prove the target statement (Claim C.2), we upper bound:

$$\begin{aligned}
\prod_{u=0}^{v-1} L_p(\theta+2\pi\{2^u x\}) &= \sqrt{\prod_{u=0}^{v-1}\Big((1-\alpha)+\alpha\cdot\cos^2(\theta/2+\pi\{2^u x\})\Big)} && \text{(by Equation 33)} \\
&\leqslant \sqrt{(1-\alpha)+\alpha\prod_{u=0}^{v-1}\cos^2(\theta/2+\pi\{2^u x\})} && \text{(apply Equation 34 repetitively)} \\
&\leqslant \sqrt{(1-\alpha)+\prod_{u=0}^{v-1}\cos^2(\theta/2+\pi\{2^u x\})} && \text{(because } \alpha\in[0,1]) \\
&\leqslant |1-2p| + \left|\prod_{u=0}^{v-1}\cos(\theta/2+\pi\{2^u x\})\right| && \text{(because } \sqrt{1-\alpha}=|1-2p|)
\end{aligned}$$

So, it would suffice to prove that

$$\left|\prod_{u=0}^{v-1}\cos(\theta/2+\pi\{2^u x\})\right| \ll |\theta|+2^{-v}$$

Let us rewrite the LHS of the inequality above:

$$\begin{aligned}
\left|\prod_{u=0}^{v-1}\cos(\theta/2+\pi\{2^u x\})\right| &= \prod_{u=0}^{v-1}\frac{1}{2}\cdot\left|1+\exp(\imath(\theta+2\pi\{2^u x\}))\right| \\
&= 2^{-v}\prod_{u=0}^{v-1}\left|1+\exp(\imath\theta)\cdot\exp(\imath 2\pi x)^{2^u}\right| \\
&\leqslant 2^{-v}\cdot\left(2+2^{v+1}|1-\exp(\imath\theta)|\right)\cdot|1-\exp(2\pi\imath x)|^{-1} \qquad \text{(by Equation 35)}
\end{aligned}$$

$$\leqslant \sqrt{2}\cdot 2^{-v} + \sqrt{2}\cdot 2|\sin(\theta/2)| \qquad (\text{as } |1-\exp(2\pi\imath x)| \geqslant \sqrt{2} \text{ for } x \in [1/4, 3/4])$$
$$\leqslant \sqrt{2}\cdot\big(|\theta| + 2^{-v}\big).$$

This completes the proof of Claim C.2. □

C.5. **Technical Results for Integral Estimation.** Recall from Section 4.4 (and in the proof of Lemma 4.11 in Appendix C.4) that

$$L_p(\varphi) \;\coloneqq\; \Big|(1-p) + p\cdot\exp(\imath\varphi)\Big|. \tag{36}$$

Claim C.3. For $\varepsilon_0 \leqslant p \leqslant (1-\varepsilon_0)$ and every $\varphi \in [-\pi, \pi]$ the following bounds hold:

(1) $L_p(\varphi) \leqslant \exp\big(-2\varepsilon_0(1-\varepsilon_0)\cdot\sin^2(\varphi/2)\big)$. In particular, when $\varepsilon_0 = 1/3$,

$$L_p(\varphi) \leqslant \exp\left(-\frac{4}{9}\cdot\sin^2\left(\frac{\varphi}{2}\right)\right).$$

(2) If the distance between π and $\varphi + 2\pi\mathbb{Z}$ is $\leqslant \varepsilon$, then for every $p \in [0,1]$,

$$L_p(\varphi) \leqslant |1-2p| + \varepsilon.$$

Proof of Claim C.3. We have:

$$\begin{aligned} L_p(\varphi)^2 &= (1-p)^2 + p^2 + 2(1-p)p\cdot\cos\varphi \\ &= 1 - 4p(1-p)\cdot\sin^2(\varphi/2) \\ &\leqslant \exp\big(-4p(1-p)\cdot\sin^2(\varphi/2)\big) && (\text{because } 1-t \leqslant \exp(-t)) \\ &\leqslant \exp\big(-4\varepsilon_0(1-\varepsilon_0)\cdot\sin^2(\varphi/2)\big). && (\text{because } \varepsilon_0 \leqslant p \leqslant (1-\varepsilon_0)) \end{aligned}$$

We rewrite

$$\begin{aligned} L_p(\varphi)^2 &= \cos^2(\varphi/2) + (1-2p)^2\cdot\sin^2(\varphi/2) \\ &\leqslant \cos^2(\varphi/2) + |1-2p|^2 \\ &\leqslant (\,\cos(\varphi/2) + |1-2p|\,)^2. && (\text{By Minkowski}) \end{aligned}$$

For the second bound, note that $\cos^2(\varphi/2) = \sin^2((\pi-\varphi)/2) \leqslant \varepsilon^2$. □

Lemma C.4 (Dispersion Lemma)**.** *The following bound holds for any n, α, β:*

$$S_n(\alpha,\beta) \;\coloneqq\; \sum_{j=0}^{n-1} \sin^2\big(\alpha + 2^j\beta\big) \geqslant \frac{n-1}{10}\sin^2\alpha.$$

Proof of Lemma C.4. Define $\theta_j = \alpha + 2^j\beta$. Note that $\theta_{j+1} = 2\theta_j - \alpha \pmod{\pi}$. So,

$$|\sin(\alpha)| = |\sin(2\theta_j - \theta_{j+1})|.$$

Next, note that

$$|\sin(2\theta_j - \theta_{j+1})| \leqslant |\sin(2\theta_j)| + |\sin\theta_{j+1}| \leqslant 2|\sin\theta_j| + |\sin\theta_{j+1}|.$$

From this, we conclude that

$$\begin{aligned} \sin^2(2\theta_j - \theta_{j+1}) &\leqslant 4\sin^2\theta_j + \sin^2\theta_{j+1} + 4\cdot|\sin\theta_j\sin\theta_{j+1}| \\ &\leqslant 5(\sin^2\theta_j + \sin^2\theta_{j+1}). && (\text{because } 4ab \leqslant a^2 + 4b^2) \end{aligned}$$

We conclude:

$$\sin^2(\alpha) \leqslant 5\big(\sin^2\theta_j + \sin^2\theta_{j+1}\big)$$

Adding over all $j \in \{0, 1, \ldots, n-2\}$, we get:

$$(n-1)\sin^2\alpha \leqslant 5\sum_{j=0}^{n-2}\left(\sin^2\theta_j + \sin^2\theta_{j+1}\right) \leqslant 10\sum_{j=0}^{n-1}\sin^2\theta_j.$$

and this concludes the proof of Lemma C.4. □

Appendix D. Transference Lemma

D.1. **Proof of Transference Lemma (Lemma 4.8).** Throughout this appendix, we write $W = W_{\text{BIN}} = \{0, 1, \ldots, n\}$ and set $S_w = 0$ whenever the actual weight-w slice is empty. This zero extension does not change $E_k(s)$.

Our aim is to express S_w, the actual Hamming slice restricted to $\{0, \ldots, N-1\}$, as a mild perturbation of B_w, a function that samples the N-grid at the Hamming slice restricted to the full n-bit range $\{0, \ldots, Q-1\}$.

(1) First, suppose $N > Q$.
Consider the set $P_w := \{Q \leqslant y < N : \text{wt}(y) = w\}$. $B_w = S_w \cap \{0, 1, \ldots, Q-1\}$. Therefore, $S_w = B_w + P_w$.
(2) Next, suppose $N \leqslant Q$.
Consider the set $P_w := \{N \leqslant y < Q : \text{wt}(y) = w\}$. $y \in P_w$ gets accounted at $x = y - N$. Therefore, $S_w = B_w - P_w$.

We formally define $P_w \colon G \to \{0, 1\}$, for $w \in W_{\text{BIN}}$ as follows:

$$P_w(x) := \begin{cases} \mathbb{1}(Q \leqslant x < N) \cdot \mathbb{1}(\text{wt}(x) = w), & N > Q \\ \mathbb{1}(0 \leqslant x < Q - N) \cdot \mathbb{1}(\text{wt}(x+N) = w), & N < Q \\ 0, & N = Q. \end{cases}$$

and $\sigma := (-1)^{\mathbb{1}(N<Q)}$. As a result, in general, we have $S_w = B_w + \sigma \cdot P_w$, where $\sigma \in \{\pm 1\}$ depending on $N > Q$ or $N \leqslant Q$. And, the sets $(P_w : w \in W)$ are pairwise disjoint and their cumulative density is

$$\sum_{w \in W} \frac{\#P_w}{N} \leqslant \frac{P}{N}.$$

Now, for $\mathbf{w} \in W^k$, let us compute the product of the Fourier coefficients that appears in our target quantity $E_k(s)$. The terms in the summand can be expanded as:

$$\prod_{j=1}^{k} \widehat{S_{\mathbf{w}_j}}(t) = \prod_{j=1}^{k}\left(\widehat{B_{\mathbf{w}_j}}(t) + \sigma \cdot \widehat{P_{\mathbf{w}_j}}(t)\right) = \sum_{J \subseteq \{1,2,\ldots,k\}} \sigma^{\#J}\left(\prod_{j \notin J}\widehat{B_{\mathbf{w}_j}}(t)\right)\left(\prod_{j \in J}\widehat{P_{\mathbf{w}_j}}(t)\right)$$

For $J \subseteq \{1, 2, \ldots, k\}$, define the restriction:

$$2 \cdot E_k(s)|_J := \sum_{\mathbf{w} \in W^k}\left|\sum_{t \neq 0} \mathrm{e}_N(st) \cdot \left(\prod_{j \notin J}\widehat{B_{\mathbf{w}_j}}(t)\right) \cdot \left(\prod_{j \in J}\widehat{P_{\mathbf{w}_j}}(t)\right)\right|$$

By the triangle inequality, we have

$$E_k(s) \leqslant \sum_J E_k(s)|_J$$

So, it suffices to upper bound the contributions of individual $E_k(s)|_J$'s.

Case $J = \emptyset$. Note that $E_k(s)|_J = E_k^{\text{BIN}}(s)$

Case $J = \{1, 2, \ldots, k\}$. In this case, we have

$$2 \cdot E_k(s)|_J = \sum_{\mathbf{w} \in W^k} \left| \sum_{t \neq 0} \mathrm{e}_N(st) \prod_{j=1}^{k} \widehat{P_{\mathbf{w}_j}}(t) \right|$$

We will use the L^2 version of the Claim B.1's upper bound. Fix $\mathbf{w} \in W^{k-1}$ (the first $(k-1)$ indices) and we will perform summation over the last index $\mathbf{w}_k \in W$ with the following setup:

(1) $f_w = P_w$
(2) $z_t = \mathrm{e}_N(st) \cdot \prod_{1 \leqslant j < k} \widehat{P_{\mathbf{w}_j}}(t)$

After this, we will have:

$$\begin{aligned} 2 \cdot E_k(s)|_J &\leqslant \sum_{\mathbf{w} \in W^{k-1}} \left\| \sum_{\mathbf{w}_k \in W} P_{\mathbf{w}_k} \right\|_{L^2} \cdot \sqrt{\sum_{t \neq 0} \prod_{1 \leqslant j < k} \left| \widehat{P_{\mathbf{w}_j}}(t) \right|^2} \\ &\leqslant \sqrt{\frac{P}{N}} \sum_{\mathbf{w} \in W^{k-1}} \sqrt{\sum_{t \neq 0} \prod_{1 \leqslant j < k} \left| \widehat{P_{\mathbf{w}_j}}(t) \right|^2} \\ &\qquad \text{(because } P_w \text{ are pairwise disjoint and } \textstyle\sum_{w \in W} P_w \text{ has density } \leqslant P/N\text{)} \end{aligned}$$

Roughly speaking, this is the density-sensitive version of Lemma 3.1.

When $k = 2$, the interpolation argument below requires zero applications: the preceding estimate already leaves a single P-index. For $k \geqslant 3$, we apply Claim B.3. Fix $\mathbf{w} \in W^{k-2}$, the first $(k-2)$ indices, and we will perform summation over the last index $\mathbf{w}_{k-1} \in W$ with the following setup:

(1) $a_t = \prod_{1 \leqslant j < k-1} \left| \widehat{P_{\mathbf{w}_j}}(t) \right|^2$
(2) $b_{w,t} = \left| \widehat{P_w}(t) \right|^2$
(3) $\mu_w = \#P_w / N$, i.e., P_w's density
(4) $\alpha = 1$

After using Claim B.3, because $\sum_{w \in W} \mu_w \leqslant P/N$, we have the inequality:

$$E_k(s)|_J \leqslant \left(\frac{P}{N}\right)^{1/2} \cdot \left(\frac{P}{N}\right)^{1/2} \sum_{\mathbf{w} \in W^{k-2}} \sqrt{\sum_{t \neq 0} \prod_{1 \leqslant j < k-1} \left| \widehat{P_{\mathbf{w}_j}}(t) \right|^2}$$

Repeatedly applying this step eliminates one index from the end. We stop when we have one index left over.

$$\begin{aligned} E_k(s)|_J &\leqslant \left(\frac{P}{N}\right)^{(k-1)/2} \sum_{w \in W} \sqrt{\sum_{t \neq 0} \left| \widehat{P_w}(t) \right|^2} \\ &\leqslant \left(\frac{P}{N}\right)^{(k-1)/2} \cdot \sqrt{n} \cdot \sqrt{\frac{P}{N}} && \text{(by Parseval and Cauchy-Schwarz)} \\ &= \left(\frac{P}{N}\right)^{k/2} \cdot \sqrt{n}. \end{aligned}$$

Case $\emptyset \neq J \subsetneq \{1, 2, \ldots, k\}$. This case covers all remaining cases. In this case, we have

$$2 \cdot E_k(s)|_J \coloneqq \sum_{\mathbf{w} \in W^k} \left| \sum_{t \neq 0} \mathrm{e}_N(st) \cdot \left(\prod_{j \notin J} \widehat{B_{\mathbf{w}_j}}(t) \right) \cdot \left(\prod_{j \in J} \widehat{P_{\mathbf{w}_j}}(t) \right) \right|.$$

Next, we use the upper-bounding strategy we used above for $J = \{1, 2, \ldots, k\}$. Denote $\overline{J} := \{1, 2, \ldots, k\} \setminus J$. We will get the inequality:

$$E_k(s)|_J \leqslant \left(\frac{P}{N}\right)^{\#J/2} \sum_{\mathbf{w}_{\overline{J}} \in W^{k-\#J}} \sqrt{\sum_{t\neq 0}\left(\prod_{j\in \overline{J}} \left|\widehat{B_{\mathbf{w}_j}}(t)\right|^2\right)}$$
$$= \left(\frac{P}{N}\right)^{\#J/2} \cdot F^{\mathrm{BIN}}_{k-\#J}(N).$$

After this, the Transference lemma (Lemma 4.8) is immediate:

$$\begin{aligned} E_k(s) &\leqslant \sum_{J\subseteq\{1,2,\ldots,k\}} E_k(s)|_J \\ &\leqslant E_k(s)|_\emptyset + E_k(s)|_{\{1,2,\ldots,k\}} + \sum_{\emptyset\neq J\subsetneq\{1,2,\ldots,k\}} E_k(s)|_J \\ &\leqslant E^{\mathrm{BIN}}_k(N,s) + \left(\frac{P}{N}\right)^{k/2}\sqrt{n} + \sum_{0<\ell<k}\binom{k}{\ell}\cdot\left(\frac{P}{N}\right)^{\ell/2} F^{\mathrm{BIN}}_{k-\ell}(N) \end{aligned}$$

□

D.2. **Statement and Proof of Corollary 1.** Recall that for every integer $m \geqslant 1$,

$$2\cdot\Delta_m = \sum_{\mathbf{w}\in W^m}\frac{1}{N}\sum_{s\in G}\left|\sum_{t\neq 0}\mathrm{e}_N(st)\prod_{j=1}^m \widehat{S_{\mathbf{w}_j}}(t)\right| \tag{37}$$

by the Fourier inversion formula (Equation 4). We can define the corresponding full-range model by replacing $\widehat{S_w}$ by $\widehat{B_w}$, as we did for E_k:

$$2\cdot\Delta^{\mathrm{BIN}}_m(N) := \sum_{\mathbf{w}\in W^m}\frac{1}{N}\sum_{s\in G}\left|\sum_{t\neq 0}\mathrm{e}_N(st)\prod_{j=1}^m \widehat{B_{\mathbf{w}_j}}(t)\right| \tag{38}$$

Corollary 1. *Suppose that $|N-Q| \leqslant P$. Then, for every $m \geqslant 2$,*

$$\Delta_m \leqslant \Delta^{\mathrm{BIN}}_m(N) + \left(\frac{P}{N}\right)^{m/2}\sqrt{n} + \sum_{0<\ell<m}\binom{m}{\ell}\cdot\left(\frac{P}{N}\right)^{\ell/2} F^{\mathrm{BIN}}_{m-\ell}(N)$$

If moreover $P = |N-Q| \leqslant Q/n$, then $\Delta_m \ll_m n^{-(m-1)/2+1/4}$; equivalently, for all $k \geqslant 3$,

$$\Delta_{k-1} \ll_k n^{-k/2+5/4}.$$

Proof of Corollary 1. Comparing Δ_m (Equation 37) with $E_m(s)$ (Equation 7), and $\Delta^{\mathrm{BIN}}_m(N)$ (Equation 38) with $E^{\mathrm{BIN}}_m(N,s)$ (Equation 20), respectively, yields the following relationships:

$$\Delta_m = \frac{1}{N}\sum_{s\in G}\frac{1}{2}\sum_{\mathbf{w}\in W^m}\left|\sum_{t\neq 0}\mathrm{e}_N(st)\prod_{j=1}^m \widehat{S_{\mathbf{w}_j}}(t)\right| = \frac{1}{N}\sum_{s\in G}E_m(s) \tag{39}$$

$$\Delta^{\mathrm{BIN}}_m(N) = \frac{1}{N}\sum_{s\in G}\frac{1}{2}\sum_{\mathbf{w}\in W^m}\left|\sum_{t\neq 0}\mathrm{e}_N(st)\prod_{j=1}^m \widehat{B_{\mathbf{w}_j}}(t)\right| = \frac{1}{N}\sum_{s\in G}E^{\mathrm{BIN}}_m(N,s) \tag{40}$$

By Lemma 4.8,

$$\begin{aligned}\Delta_m &= \frac{1}{N}\sum_{s\in G} E_m(s) && \text{(by Equation 39)}\\ &\leqslant \frac{1}{N}\sum_{s\in G}\left(E_m^{\mathrm{BIN}}(N,s) + \left(\frac{P}{N}\right)^{m/2}\sqrt{n} + \sum_{0<\ell<m}\binom{m}{\ell}\cdot\left(\frac{P}{N}\right)^{\ell/2} F_{m-\ell}^{\mathrm{BIN}}(N)\right)\\ &\qquad\qquad \text{(averaging Lemma 4.8 with } k=m \text{ over } s\in G)\\ &= \Delta_m^{\mathrm{BIN}}(N) + \left(\frac{P}{N}\right)^{m/2}\sqrt{n} + \sum_{0<\ell<m}\binom{m}{\ell}\cdot\left(\frac{P}{N}\right)^{\ell/2} F_{m-\ell}^{\mathrm{BIN}}(N).\\ &\qquad\qquad \text{(by Equation 40: second and third summands do not depend on } s)\end{aligned}$$

For the second bound, we first use Claim B.1.

$$\begin{aligned}2\Delta_m^{\mathrm{BIN}}(N) &= \sum_{\mathbf{w}\in W^m}\frac{1}{N}\sum_{x\in G}\left|\sum_{t\neq 0}\mathrm{e}_N(xt)\prod_{j=1}^{m}\widehat{B_{\mathbf{w}_j}}(t)\right| && \text{(by definition: Equation 38)}\\ &\leqslant \sum_{\mathbf{w}\in W^m}\sqrt{\sum_{t\neq 0}\prod_{j=1}^{m}\left|\widehat{B_{\mathbf{w}_j}}(t)\right|^2} && \text{(using Claim B.1 with } f_x=\mathbb{1}_{\{-x\}})\\ &= F_m^{\mathrm{BIN}}(N) && (41)\end{aligned}$$

Take $P=|N-Q|\leqslant Q/n$. By Equation 41 and Lemmas 4.4 to 4.6,

$$\Delta_m^{\mathrm{BIN}}(N) \ll F_m^{\mathrm{BIN}}(N) \ll_m n^{-(m-1)/2}\cdot F_1^{\mathrm{BIN}}(N) \ll_m n^{-(m-1)/2+1/4}.$$

Equation 22 and Equation 23, with $k=m$, give the same bound for the remaining terms. Hence,

$$\Delta_m \ll_m n^{-(m-1)/2+1/4}.$$

Taking $m=k-1$ gives the final claim. □

Appendix E. Lower Bound: Proof of Theorem 5

Just for this section, we consider $W=\{0,1,\ldots,n\}$; any empty S_w slice is identified by the 0 function. Since $m\leqslant(n-1)$, this will capture all possible weights. Let $\delta := |2^n-N|/N \leqslant 2^m/N$ denote the distortion from the ideal behavior. In our proof, in fact, we will use the quantity

$$L_* \;:=\; \#\left\{0\leqslant j<n-1 \;:\; 2^j\delta \leqslant \frac{1}{10\pi}\cdot n^{-1/2}\right\}$$

If $L_*=0$, the desired lower bound is trivial. Henceforth assume $L_*>0$.

Claim E.1. Let C be an absolute constant with $C\geqslant 1+\log_2(10\pi)$. Then,

$$L_* \geqslant \left(n-m-\frac{1}{2}\log_2 n - C\right)_+ = L$$

We defer the proof of Claim E.1 to the end of this section, after the proof of Theorem 5. For phases

$$\mathbf{\Phi} := \left(\,\mathbf{\Phi}_{\mathbf{w}} \;:\; \mathbf{w}\in W^k\,\right)$$

we know that

$$2\cdot E_k(s) = \max_{\mathbf{\Phi}}\left|\sum_{\mathbf{w}\in W^k}\mathbf{\Phi}_{\mathbf{w}}\sum_{t\neq 0}\mathrm{e}_N(st)\prod_{j=1}^{k}\widehat{S_{\mathbf{w}_j}}(t)\right|$$

We will restrict to

$$\mathbf{\Phi}_{\mathbf{w}} \coloneqq \prod_{j=1}^{k} \phi^{\mathbf{w}_j}$$

for some appropriately chosen phase ϕ. Then, for this restriction, we have

$$2 \cdot E_k(s) \geqslant \left| \sum_{\mathbf{w} \in W^k} \sum_{t \neq 0} \mathrm{e}_N(st) \prod_{j=1}^{k} \phi^{\mathbf{w}_j} \widehat{S_{\mathbf{w}_j}}(t) \right| = \underbrace{\left| \sum_{t \neq 0} \left(\sum_{w \in W} \phi^w \widehat{S_w}(t) \right)^k \cdot \mathrm{e}_N(st) \right|}_{=: f(s)} =: |f(s)| \tag{42}$$

Note that $\widehat{f}(0) = 0$, and for non-zero t

$$\widehat{f}(t) = \left(\sum_{w \in W} \phi^w \widehat{S_w}(t) \right)^k.$$

Define the set $H := \left\{ -2^j : 0 \leqslant j < L_* \right\}$. Since $L_* \leqslant n-1$, we have $2^{L_*-1} \leqslant 2^{n-2} < N$; hence, all elements in H remain distinct in $\mathbb{Z}/N\mathbb{Z}$, i.e. $|H| = L_*$.[7] So,

$$\begin{aligned}
2 \cdot E_k = \max_{0 \leqslant s < N} 2 \cdot E_k(s) &\geqslant \sqrt{\frac{1}{N} \sum_{0 \leqslant s < N} \left(2E_k(s)\right)^2} && \text{(monotonicity of norms)} \\
&\geqslant \sqrt{\frac{1}{N} \sum_{0 \leqslant s < N} |f(s)|^2} && \text{(by Equation 42 above)} \\
&= \sqrt{\sum_{t} \left| \widehat{f}(t) \right|^2} && \text{(Parseval's identity)} \\
&\geqslant \sqrt{\sum_{t \in H} \left| \widehat{f}(t) \right|^2} \\
&\geqslant L_*^{1/2} \cdot \min_{t \in H} \left| \widehat{f}(t) \right| = L_*^{1/2} \min_{t \in H} \left| \sum_{w} \phi^w \widehat{S_w}(t) \right|^k. && \text{(since } |H| = L_*)
\end{aligned}$$

So, to prove Theorem 5, it will suffice to show that: There is a choice of phase ϕ such that for all $0 \leqslant j < L_*$, the following estimate is satisfied.

$$\left| \sum_{w \in W} \phi^w \widehat{S_w}(-2^j) \right| \geqslant \left(\frac{1}{25} \right) \cdot n^{-1/2}. \tag{43}$$

Let us begin towards lower-bounding this target expression. We begin by observing that our target expression (inside the $|\cdot|$) is the following polynomial.

$$T(t) \coloneqq \sum_{w \in W} \phi^w \widehat{S_w}(-t) = \frac{1}{N} \sum_{w \in W} \phi^w \sum_{x \in S_w} \mathrm{e}_N(xt) = \frac{1}{N} \sum_{0 \leqslant x < N} \phi^{\mathrm{wt}(x)} \mathrm{e}_N(xt).$$

We will compare this expression against an ideal baseline:

$$I(t) \coloneqq \frac{1}{N} \sum_{0 \leqslant x < 2^n} \phi^{\mathrm{wt}(x)} \mathrm{e}_N(xt) = \frac{1}{N} \prod_{\ell=0}^{n-1} \left(1 + \phi \cdot \mathrm{e}_N\left(2^\ell t\right) \right).$$

[7] Since $|N - 2^n| \leqslant 2^m$ with $m \leqslant n-1$, we have $N \geqslant 2^n - 2^{n-1} = 2^{n-1}$. Then, $2^{L_*-1} \leqslant 2^{n-2} < N$ because $L_* \leqslant n-1$. So for $0 \leqslant i < j < L_*$, we have $0 < 2^j - 2^i < N$, and also $-2^i \not\equiv -2^j \pmod{N}$. Hence, $-2^0, \ldots, -2^{L_*-1}$ must be distinct in $\mathbb{Z}/N\mathbb{Z}$.

By the union bound, we have

$$|T(t)| \geqslant |I(t)| - \frac{|2^n - N|}{N} = |I(t)| - \delta.$$

Thus, it suffices to choose ϕ such that

$$|I(t)| \geqslant \frac{21}{275} \cdot n^{-1/2} \quad \text{for all } t \in \left\{2^j : 0 \leqslant j < L_*\right\}$$

and it can absorb δ into it, which has magnitude at most $1/(10\pi) \cdot n^{-1/2}$.

Let us substitute $\phi = \exp(-\imath\theta)$, where $\theta = n^{-1/2}$ is a positive angle. Fix j and let us begin on our new target of proving the lower bound:

$$\left|I(2^j)\right| = \frac{1}{N} \prod_{\ell=0}^{n-1} \left|1 + \exp\left(-\imath\theta + 2\pi\imath \cdot \frac{2^{\ell+j}}{N}\right)\right| \geqslant \frac{21}{275} \cdot n^{-1/2}. \tag{44}$$

Our aim is to lower-bound the RHS terms separately based on which of the three conditions it satisfies below.

Case 1: $\ell + j \geqslant n$**.** Define the gap $g = \ell + j - n$. Note that $0 \leqslant g < j$. For this analysis, we will need the signed version of δ defined as $\Delta \coloneqq (2^n - N)/N$. Using this, we can write:

$$\exp\left(2\pi\imath \cdot \frac{2^{\ell+j}}{N}\right) = \exp\left(2\pi\imath \cdot \frac{2^{n+g}}{N}\right) = \exp\Big(2\pi\imath \cdot 2^g \cdot (1 + \Delta)\Big) = \exp\Big(2\pi\imath \cdot 2^g\Delta\Big)$$

Note that the magnitude of this angle

$$|2\pi \cdot 2^g\Delta| \leqslant \pi \cdot 2^j\delta \leqslant \frac{1}{10} \cdot n^{-1/2},$$

because $g < j$ and $2^j\delta \leqslant (10\pi)^{-1} \cdot n^{-1/2}$ by the definition of the quantity L_*. These angles will be significantly small in magnitude compared to $\theta = n^{-1/2}$. As a consequence, we will have:

$$\left|1 + \exp\left(-\imath\theta + 2\pi\imath \cdot \frac{2^{\ell+j}}{N}\right)\right| \geqslant 2\cos\left(\frac{n^{-1/2} + \frac{1}{10} \cdot n^{-1/2}}{2}\right) = 2\cos\left(\frac{11 \cdot n^{-1/2}}{20}\right) \tag{45}$$

Now, for fixed j, we estimate the cumulative contribution of all ℓ satisfying $\ell + j \geqslant n$.

$$\begin{aligned}
\prod_{\substack{0 \leqslant \ell < n: \\ \ell + j \geqslant n}} \left|1 + \exp\left(-\imath\theta + 2\pi\imath \cdot \frac{2^{\ell+j}}{N}\right)\right| &= \prod_{g=0}^{j-1} \left|1 + \exp\left(-\imath\theta + 2\pi\imath \cdot \frac{2^{n+g}}{N}\right)\right| \\
&\geqslant \left(2\cos\left(\frac{11 \cdot n^{-1/2}}{20}\right)\right)^j && \text{(by Equation 45 above)} \\
&\geqslant 2^j \cdot \exp\left(-\frac{2}{\pi} \cdot \left(\frac{11}{20}\right)^2 \cdot n^{-1} j\right) && \text{(by Lemma E.2)} \\
&\geqslant 2^j \cdot \frac{4}{5}. && \text{(because } j \leqslant n \text{ and } \exp(-1/5) \geqslant 1 - 1/5.\text{)}
\end{aligned}$$

Case 2: $\ell + j = n - 1$**.** In this case, we have

$$\exp\left(2\pi\imath \cdot \frac{2^{\ell+j}}{N}\right) = \exp\left(2\pi\imath \cdot \frac{2^{n-1}}{N}\right) = \exp\left(\pi\imath \cdot \frac{2^n}{N}\right) = -\exp(\pi\imath \cdot \Delta)$$

Our concern is that the magnitude of the term $\left|1+\phi\cdot \mathrm{e}_N(2^{\ell+j})\right|$ may become too small. However, the magnitude of Δ being much smaller than $\theta = n^{-1/2}$ will save us. Indeed,

$$\begin{aligned}\left|1+\exp\left(-\imath\theta+2\pi\imath\frac{2^{\ell+j}}{N}\right)\right| &= \left|1-\exp(-\imath\theta+\pi\imath\cdot\Delta)\right| = 2\sin\left(\frac{|\theta-\pi\Delta|}{2}\right)\\ &\geqslant 2\sin\left(\frac{9\theta}{20}\right) && \text{(because } \pi|\Delta|\leqslant\pi\delta\leqslant\theta/10\text{)}\\ &\geqslant \frac{9}{5\pi}\cdot n^{-1/2}. && \text{(because } \sin t\geqslant(2/\pi)\cdot t\text{)}\end{aligned}$$

Case 3: $\ell+j<n-1$**.** Again, we write the gap $g=n-(\ell+j)$. Note that $2\leqslant g\leqslant n-j$. As before, we rewrite:

$$\exp\left(2\pi\imath\cdot\frac{2^{\ell+j}}{N}\right)=\exp\left(2\pi\imath\cdot\frac{2^{n-g}}{N}\right)=\exp\big(2\pi\imath\cdot 2^{-g}\cdot(1+\Delta)\big)$$

Now, we lower-bound the magnitude of the term

$$\left|1+\exp\left(-\imath\theta+2\pi\imath\cdot\frac{2^{\ell+j}}{N}\right)\right| = 2\cos\left(-\frac{\theta}{2}+\pi 2^{-g}(1+\Delta)\right).$$

Let us consider the cumulative contributions of these terms:

$$\begin{aligned}\prod_{\substack{0\leqslant\ell<n:\\ \ell+j<n-1}}\left|1+\exp\left(-\imath\theta+2\pi\imath\frac{2^{\ell+j}}{N}\right)\right| &= \prod_{g=2}^{n-j}2\cos\left(-\frac{\theta}{2}+\pi 2^{-g}+\pi 2^{-g}\Delta\right)\\ &= 2^{n-j-1}\prod_{g'=0}^{n-j-2}\cos\left(\frac{\pi}{4}\cdot 2^{-g'}-\left(\frac{\theta}{2}-\frac{\pi}{4}\Delta\cdot 2^{-g'}\right)\right).\end{aligned}$$

Note that $\pi 2^{-g}\in(0,\pi/4]$. The angle $\pi 2^{-g}\Delta=(\pi/4)\Delta\cdot 2^{-g'}$ will have magnitude $\leqslant\pi\delta/4$, which will be much smaller than $\theta/2$. The argument of the cosine is always $\leqslant\pi/4$ (in magnitude) and is the difference of two positive angles.

We continue simplifying the RHS. We will use the fact that for positive u,v, we have $(u-v)^2\leqslant u^2+v^2$, and Lemma E.2.

$$\begin{aligned}&\prod_{\substack{0\leqslant\ell<n:\\ \ell+j<n-1}}\left|1+\exp\left(-\imath\theta+2\pi\imath\cdot\frac{2^{\ell+j}}{N}\right)\right|\\ &\geqslant 2^{n-j-1}\prod_{g'=0}^{n-j-2}\exp\left(-\frac{2}{\pi}\cdot\underbrace{\left(\frac{\pi}{4}\cdot 2^{-g'}\right)^2}_{u^2}-\frac{2}{\pi}\cdot\underbrace{\left(\frac{\theta}{2}-\frac{\pi}{4}\Delta 2^{-g'}\right)^2}_{v^2}\right) && \text{(by Lemma E.2)}\\ &\geqslant 2^{n-j-1}\prod_{g'=0}^{n-j-2}\exp\left(-\frac{2}{\pi}\cdot\left(\frac{\pi}{4}\cdot 2^{-g'}\right)^2-\frac{2}{\pi}\cdot\left(\frac{\theta}{2}+\frac{\pi}{4}\delta 2^{-g'}\right)^2\right)\\ &\geqslant 2^{n-j-1}\exp\left(-\frac{\pi}{6}-\frac{1}{2\pi}\underbrace{-\frac{\theta\delta}{2}\cdot 2-\frac{\pi}{8}\delta^2\cdot\frac{4}{3}}_{\text{small}\ll n^{-1}}\right)\\ &\overset{(*)}{\geqslant} 2^{n-j}\cdot\frac{1}{4}.\end{aligned}$$

The $(*)$ step uses explicit numerical calculation: for all $n \geqslant 4$ with $\theta = n^{-1/2}$ and $\delta \leqslant 1/(10\pi)\cdot n^{-1/2}$,

$$\frac{\pi}{6} + \frac{1}{2\pi} + \theta\delta + \frac{\pi}{6}\delta^2 \leqslant \frac{\pi}{6} + \frac{1}{2\pi} + \frac{1}{10\pi n} + \frac{1}{600\pi n} \leqslant \frac{\pi}{6} + \frac{1}{2\pi} + \frac{1}{40\pi} + \frac{1}{2400\pi} < \log 2$$

Putting together the three pieces. We can substitute the three lower bounds to obtain a lower bound on our ideal polynomial.

$$\begin{aligned}
\left|I(2^j)\right| &= \frac{1}{N}\prod_{\ell=0}^{n-1}\left|1 + \exp\left(-\imath\theta + 2\pi\imath\cdot\frac{2^{\ell+j}}{N}\right)\right| \\
&\geqslant \frac{2^n}{N}\cdot\frac{4}{5}\cdot\frac{9}{5\pi}\cdot\frac{1}{4}\cdot n^{-1/2} \\
&\geqslant \frac{2}{3}\cdot\frac{9}{25\pi}\cdot n^{-1/2} = \frac{6}{25\pi}\cdot n^{-1/2} && \text{(because } 2^n/N \geqslant 2/3 \text{ and } \pi < 22/7)
\end{aligned}$$

This is what we set ourselves out to achieve in Equation 44, and this implies the following lower bound for our target polynomial:

$$\left|T(2^j)\right| \geqslant \left(\frac{6}{25\pi} - \frac{1}{10\pi}\right)\cdot n^{-1/2} = \frac{7}{50\pi}\cdot n^{-1/2} > \frac{1}{25}\cdot n^{-1/2}. \qquad (\pi < 3.5)$$

which was our target from Equation 43. So, all in one:

$$E_k = \max_{0\leqslant s<N} E_k(s) \geqslant \sqrt{\frac{1}{N}\sum_{0\leqslant s<N} E_k(s)^2} \geqslant \frac{1}{2}\cdot L_*^{1/2}\cdot\left(\frac{1}{25}\right)^k\cdot n^{-k/2} \geqslant \frac{1}{2}\cdot L^{1/2}\cdot\left(\frac{1}{25}\right)^k\cdot n^{-k/2}$$

which gives us Theorem 5. □

Lemma E.2 (Gaussian Lower bound on Cosine)**.** *For all $|x| \leqslant \pi/4$, we have*

$$\cos x \geqslant \exp\left(-\frac{2}{\pi}x^2\right).$$

Proof of Lemma E.2. We prove that the following function is non-negative for $x \in [0, \pi/4]$.

$$f(x) = \ln\cos x + \frac{2}{\pi}\cdot x^2.$$

This follows from the fact that $f'(x) = -\tan x + (4/\pi)\cdot x \geqslant 0$ when $x \in [0, \pi/4]$. The equality holds at $x = 0$ and $x = \pi/4$, and between those extremes we have $\tan x \leqslant (4/\pi)\cdot x$. □

We conclude this section with the proof of Claim E.1, which was deferred from the proof of Theorem 5.

Proof of Claim E.1. Write $A := n - m - (1/2)\log_2 n - C$ so that $L = (A)_+$. Since $|N - 2^n| \leqslant 2^m$ with $m \leqslant n-1$, we have $N \geqslant 2^n - 2^m \geqslant 2^{n-1}$, and so

$$\delta = \frac{|2^n - N|}{N} \leqslant \frac{2^m}{N} \leqslant 2^{m-n+1}.$$

Consequently, for every integer j with $0 \leqslant j \leqslant A$, since $C \geqslant 1 + \log_2(10\pi)$,

$$2^j\delta \leqslant 2^A\cdot 2^{m-n+1} = 2^{1-C}n^{-1/2} \leqslant \frac{1}{10\pi}\cdot n^{-1/2},$$

so every such j is counted by L_*. Moreover, $m \geqslant 0$ and $C > 5$ give $A < n-1$, so all these j lie in the range $\{0, 1, \ldots, n-2\}$. If $A > 0$, the number of such integers is $\lfloor A\rfloor + 1 \geqslant A$, and therefore $L_* \geqslant A = L$; otherwise, we would just have $L = 0 \leqslant L_*$ trivially.

Appendix F. Collection of Useful Figures

F.1. **Comparison of our Proxy (Figure 3).** The triangle inequality proxy is defined as

$$T_k := \sum_{\mathbf{w}\in W^k} \sum_{t\neq 0} \prod_{j=1}^{k} \left|\widehat{S_{\mathbf{w}_j}}(t)\right|.$$

We know that $2 \cdot E_k(s) \leqslant T_k$. However, it is unclear how our proxy F_{k-1} compares to T_k. Figure 3 illustrates this comparison.

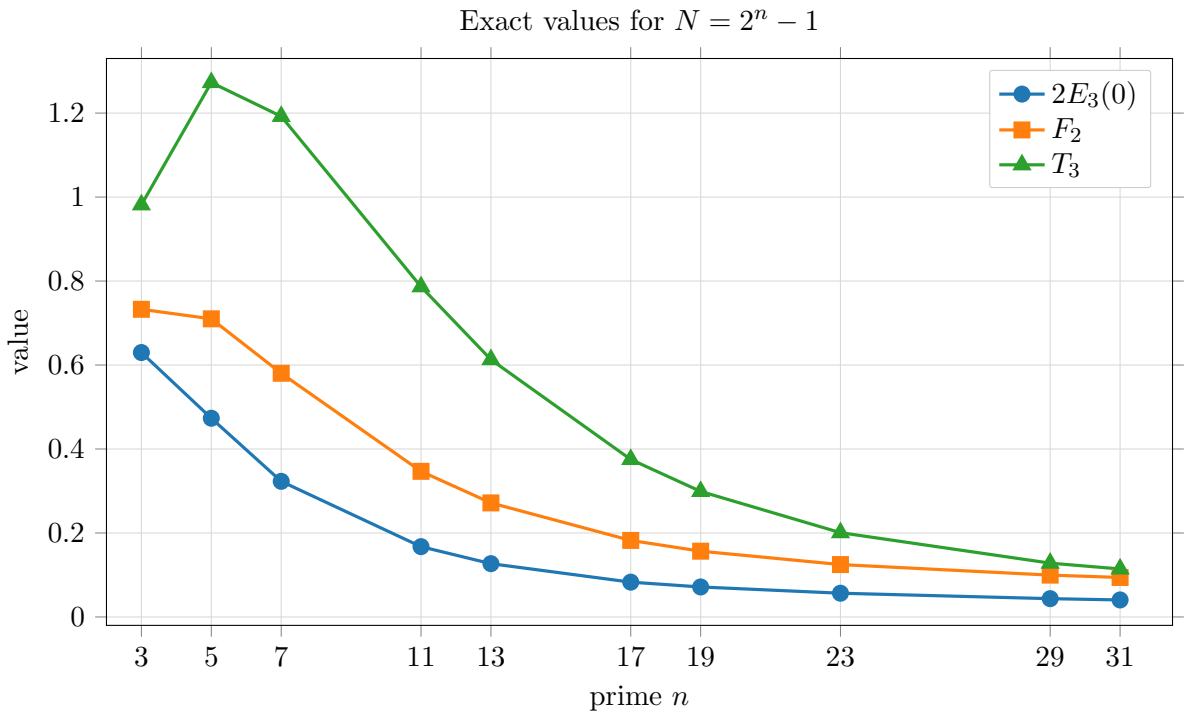

Figure 3. Comparison of $2 \cdot E_3(0)$, F_2, and T_3 for primes $3 \leqslant n < 32$ and $N = 2^n - 1$.

F.2. **Comparison of $\sum_{w\in W}\left|\widehat{S_w}(1)\right|$ and $\sum_{w\in W}\left|\widehat{S_w}(3)\right|$ (Figure 4).** Define

$$A := \sum_{w\in W}\left|\widehat{S_w}(1)\right| \quad \text{and} \quad B := \sum_{w\in W}\left|\widehat{S_w}(3)\right|.$$

Figure 4 plots the ratio of B/A for moduli $N = 2^n - 1$, and prime $n \leqslant 31$. Over the computed range, the aggregate Fourier masses remain comparable. The data suggest a positive limiting ratio; convergence is not proved here.

F.3. **Plot of $E_3(0)$ (Figure 5).** Experimentally, it seems that $E_3(0)$ behaves like $a \cdot n^{-1}$, where

$$a = \mathrm{E}\Big[\ |\ \Psi(X,Y,Z)\ |\ \Big]$$

where X, Y, Z are independent standard normal random variables, and

$$\Psi(x,y,z) = \sum_{\mathrm{cyc}} (x^2 - 1)yz.$$

Figure 5 indicates that $n \cdot E_3(0)$ is approximately affine in $1/n$ over the computed range. For $N = 2^n - 1$ and primes $41 \leqslant n \leqslant 257$, we have the plot for $n \cdot E_3(0)$ versus $1/n$.

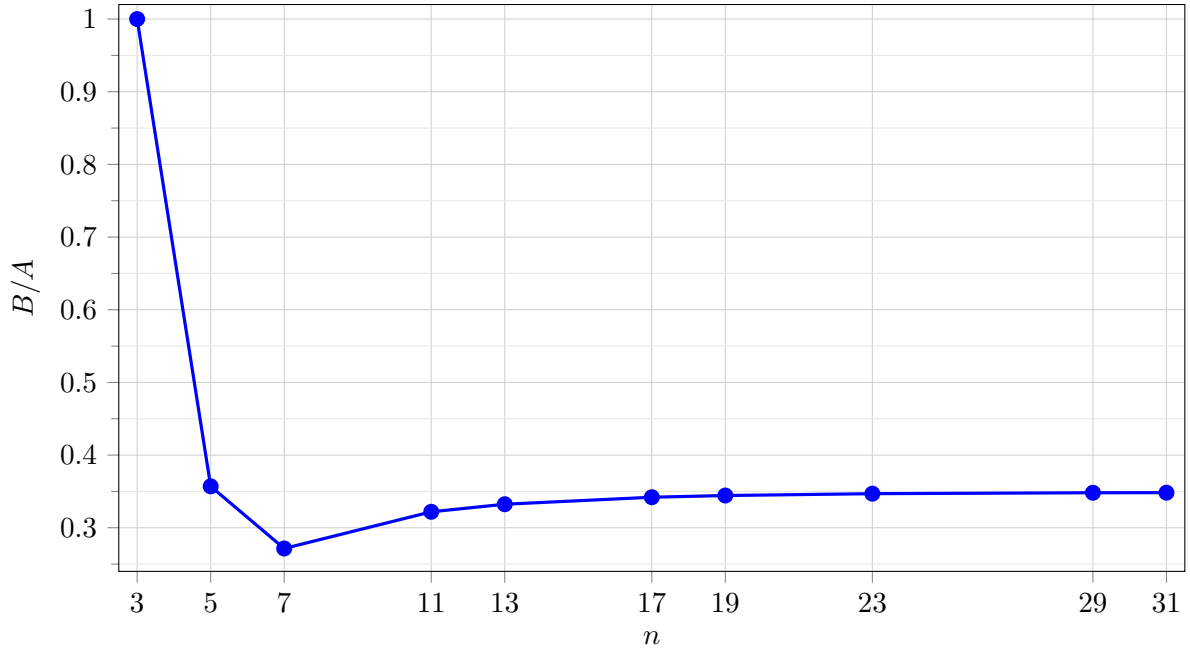

FIGURE 4. For $N = 2^n - 1$ and prime $3 \leqslant n \leqslant 31$, plot of $\sum_{w \in W} \left|\widehat{S_w}(3)\right| / \sum_{w \in W} \left|\widehat{S_w}(1)\right|$.

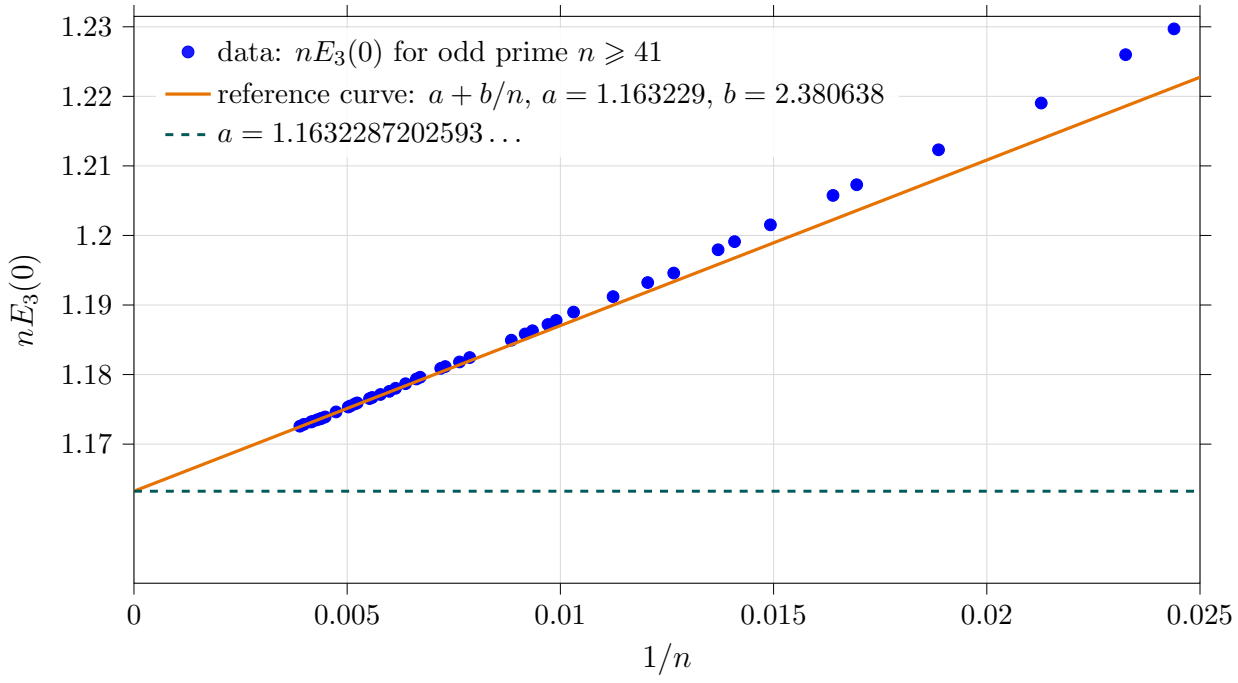

FIGURE 5. For $41 \leqslant n \leqslant 257$ and $N = 2^n - 1$, plots $n \cdot E_3(0)$ versus $1/n$ plot, along with an affine reference curve; no lower-bound claim is made.

F.4. **Plot of $|\widehat{S_w}(3)|/|\widehat{S_w}(1)|$ (Figure 6).** Figure 4 compared the frequencies $t = 1$ and $t = 3$ after aggregating over all weights. Figure 6 instead fixes one central slice ($w = \lfloor n/2 \rfloor$ for $N = 2^n - 1$)

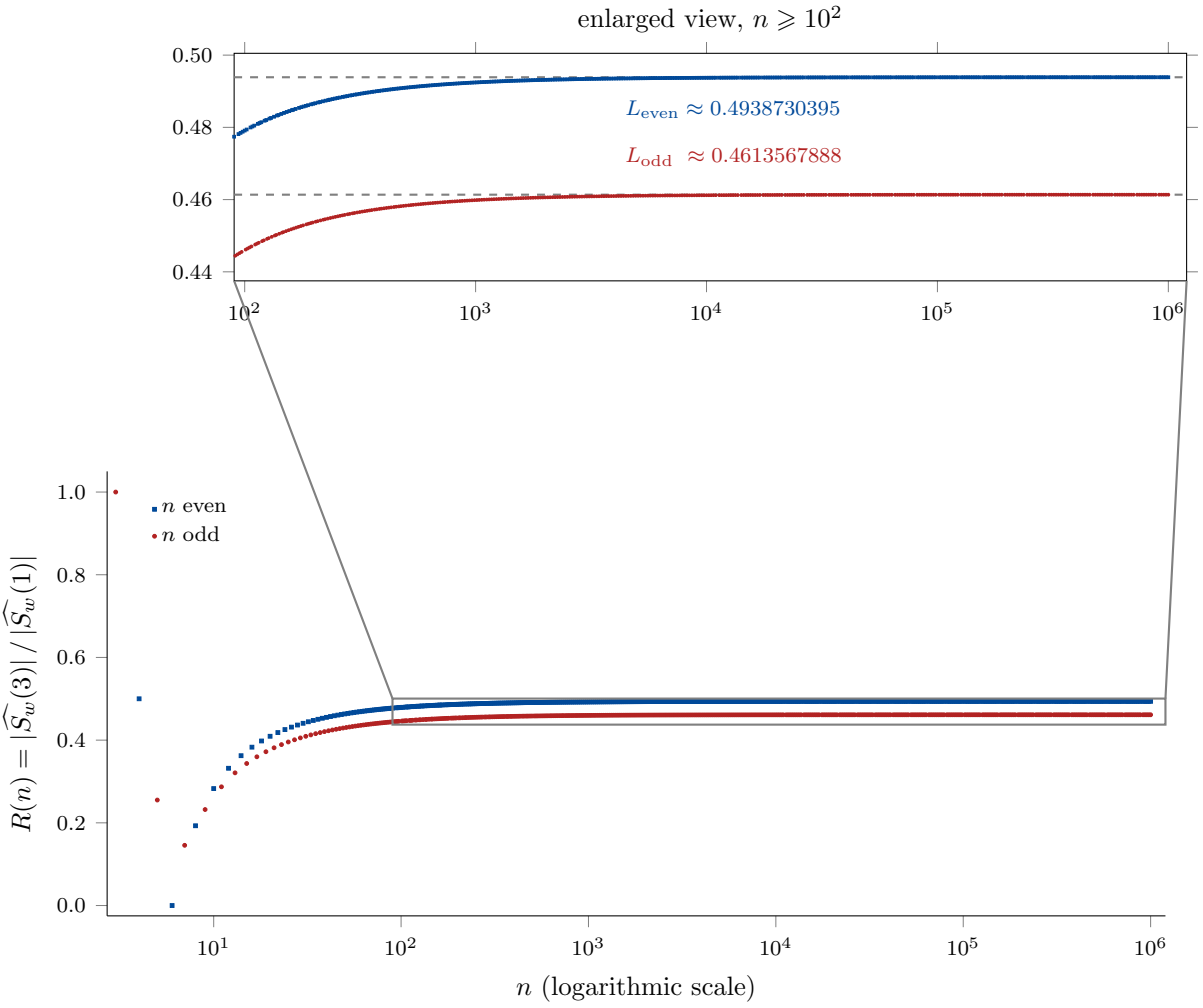

FIGURE 6. $R(n) = |\widehat{S_w}(3)|/|\widehat{S_w}(1)|$ for $w = \lfloor n/2 \rfloor$ and $N = 2^n - 1$. The dashed levels in the enlargement are the numerically extrapolated limits L_{even} and L_{odd}; they are not proved limiting constants.

and plots the ratio

$$R(n) = \frac{\left|\widehat{S_w}(3)\right|}{\left|\widehat{S_w}(1)\right|} \quad \text{for } n \leqslant 10^6$$

The even and odd subsequences appear to approach different positive limits. The dashed levels are numerical extrapolations; neither convergence nor the limiting constants are proved here.