A Spectral Theory of Local-to-Global Leakage

Rei Ueno[1][0000−0002−9754−6792]

Kyoto University, Japan. ueno.rei.2e@kyoto-u.ac.jp

Abstract. This paper studies how local leakages characterize the global leakage in additive secret sharing over a finite abelian group, which is a special case of linear secret sharing and side-channel masking. Our main finding is structural: local-to-global leakage is governed by two disjoint components—the magnitude of local leakage and its spectral alignment across shares, while existing literature attempted to characterize the global leakage only by local mutual information (MI). We first give an exact decomposition of leakage in additive secret sharing. The decomposition reveals a structural parameter invisible to MI, while previous literature conjectured that the local MI alone suffices to characterize global leakage. We then show that this missing parameter is indispensable: two leakage channels with identical local MI can induce exponentially different global leakage. Our exact decomposition identifies the missing structural quantity and completely characterizes this phenomenon.

We then show separation theorems with its realizability—local MI, despite governing the classical leakage bounds, does not even determine the exponential rate at which leakage vanishes under homomorphic composition. More strongly, we establish a universal separation theorem: even when all permutation-invariant scalar measures of local leakage agree, the global leakage can exhibit exponentially different decay rates.

We further provide the characterization of the spectral alignment factor and derive its lower bound through character-order multiplicity and Möbius inversion. We then extend the framework to automorphism-weighted sharing. We establish that automorphism-weighted sharing can achieve perfect leakage cancellation with finitely many shares, even for non-trivial full-support local leakage channels. In sharp contrast, such perfect cancellation is impossible for ordinary additive sharing with non-trivial i.i.d local leakage channels and finitely many shares.

Consequently, local leakage magnitudes correctly measure the average amount of local leakage, but discard the compatibility structure needed for composition. Magnitude-only bounds must optimize over structural alignments that the actual leakage channel does not realize.

Table of Contents

1 Introduction . . . 3
1.1 The problem and existing approaches . . . 3
1.2 Main contribution . . . 4
1.3 Paper organization . . . 4
2 Local-to-Global Leakage: A First Separation . . . 5
2.1 Problem setting . . . 5
2.2 A motivating counterexample: equal local MI, different global leakage . . . 5
2.3 What information existing bounds lose . . . 6
3 The Exact Local-to-Global Decomposition . . . 7
3.1 Spectral representation of leakage . . . 7
3.2 The exact decomposition theorem . . . 8
3.3 Shannon equivalence . . . 8
4 Beyond Mutual Information: Exponential Separation and Structural Amplification . . . 10
4.1 Realizability of spectral leakage profile . . . 10
4.2 Exact security exponent by structure . . . 12
4.3 Exponential separation theorems . . . 13
4.4 When are magnitude-only bounds tight? . . . 15
5 Quotient Spectral Theory . . . 16
5.1 Cyclic group . . . 16
5.2 Non-cyclic finite abelian groups . . . 18
5.3 Automorphisms . . . 19
6 Canonical Leakage Models . . . 22
6.1 Deterministic/Noise-free leakages . . . 22
6.2 HW leakage with Gaussian noise . . . 23
7 Secret sharing and masking . . . 24
7.1 Additive secret sharing . . . 24
7.2 Side-channel masking . . . 25
8 Limitations and open problems . . . 25
A Equivalence between KL and chi-square divergences . . . 26
B Proofs . . . 27
B.1 Proof of Theorem 1 . . . 27
B.2 Proof of Lemma 1 . . . 28
B.3 Proof of Theorem 2 . . . 29

1 Introduction

1.1 The problem and existing approaches

Local-to-Global information leakage. Let Z be a random variable on finite abelian group $\mathbb{G}$ as a secret information. To protect Z against leakages, it can be homomorphically decomposed into $d+1$ shares $S_0, S_1, \dots S_d$ over $\mathbb{G}$ to store them locally. Each share S_i follows a uniform distribution of $\mathbb{G}$ and its information leaks through a channel τ_i as $\tau_i(S_i) \to \boldsymbol{L}_i$. We may omit the subscript i, as we basically assume that the leakage is independent and identically distributed (i.i.d). The problem is how the global leakage $I(Z;\boldsymbol{X})$ is characterized by the local leakages, where $I(A;B)$ is a mutual information (MI) between random variables A and B, and $\boldsymbol{X} = (\boldsymbol{L}_0, \boldsymbol{L}_1, \dots, \boldsymbol{L}_d)$. This *local-to-global leakage in additive secret sharing* has been a major and significant problem mainly studied by two independent research fields: linear secret sharing (LSS) [BOGW88, Bea91, SV19] and side-channel masking [KJJ99, CJRR99].

The conventional pictures. In the research field of LSS, the global leakage bounds have been developed considering the local leakage as a deterministic function $\tau_i : \mathbb{G} \to \mathbb{Z}_n$, where $\mathbb{Z}_n = \mathbb{Z}/n\mathbb{Z}$ and n is a positive integer ($n \geq 2$). The analysis has been mainly devoted to bounding the global leakage for a given leakage function. In the research field of *side-channel masking*, contrarily, an average-case local leakage has been a major focus. The leakage function is stochastic and represented using a conditional probability distribution $p_{S_i|\boldsymbol{L}_i}$ of S_i given its continuous noisy leakage $\boldsymbol{L}_i \in \mathbb{R}$, and the local leakage is characterized by $I(S_i;\boldsymbol{L}_i)$. These previous studies conjectured that the global MI is bounded using a constant C_d as

$$I(Z;\boldsymbol{X}) \lesssim C_d I(S;\boldsymbol{L})^{d+1},$$

as by Duc, Faust, and Standaert [DFS15]. This indicates an exponential security amplification by the decomposition order d—the smaller local MI yields an exponentially smaller global MI. It might imply that the global leakage can be asymptotically characterized by the local MI; namely, a function F exists such that

$$I(Z;\boldsymbol{X}) \approx F(I(S_0;\boldsymbol{L}_0), I(S_1;\boldsymbol{L}_1), \dots, I(S_d;\boldsymbol{L}_d)) \text{ as } d \to \infty,$$

as in [BCG+23]. This conjecture may sound natural because it simply means that the global MI is bounded solely by the average-case local leakages.

The surprise. In this paper, we demonstrate an exponential separation of global MI from local MI. Let τ and τ' be different leakage channels such that $I(S;\boldsymbol{L}^\tau) = I(S;\boldsymbol{L}^{\tau'})$, where $\boldsymbol{L}^\tau$ and $\boldsymbol{L}^{\tau'}$ are the local leakage through τ and τ', respectively. Let $I(Z;\boldsymbol{X}^\tau)$ and $I(Z;\boldsymbol{X}^{\tau'})$ be the global leakages through local leakage channels τ and τ', respectively. In this setting, the local leakages are identical for τ and τ'; hence, $I(Z;\boldsymbol{X}^\tau) \asymp I(Z;\boldsymbol{X}^{\tau'})$ would hold if the above conjecture was true. However, counterexamples exist such that

$$\frac{I(Z;\boldsymbol{X}^\tau)}{I(Z;\boldsymbol{X}^{\tau'})} = \exp\left(\Omega(d)\right),$$

as shown in Section 2.2. This means that *local MI only does not even determine the exponential rate of global leakage.*

1.2 Main contribution

We present a novel spectral theory to characterize the local-to-global leakage. Our main finding is structural: local-to-global leakage is governed by two independent components—the magnitude of local leakage and its spectral alignment probability across shares. Existing scalar measures capture the former but not the latter. This implies an impossibility of existing approaches—local scalar leakage strength (*e.g.*, MI and total variation distance (TV)) does not determine even the exponential rate of global leakage.

Exact local-to-global decomposition and characterization. We present the first exact decomposition of leakages through χ^2-information theory, which is represented by

$$I_{\chi^2}(Z;\boldsymbol{X}) = \kappa_d \prod_{i=0}^{d} I_{\chi^2}(S_i;\boldsymbol{L}_i),$$

where $I_{\chi^2}(Z;\boldsymbol{X})$ and $I_{\chi^2}(S_i;\boldsymbol{L}_i)$ are the global and local leakages represented by χ^2-MI [NN14, AA19], respectively; and κ_d is the structural decomposition factor, defined as multi-collision probability of Fourier modes under local leakages (*i.e.*, *spectral alignment probability)*. We then show its Shannon equivalence—under a strict weak worst-case condition, the global MI is characterized by

$$I(Z;\boldsymbol{X}) = \frac{1+o(1)}{2\ln 2}\kappa_d \prod_{i=0}^{d} \alpha_i I(S_i;\boldsymbol{L}_i),$$

where α_i is the *leakage sharpness*, which is a conversion factor between Shannon and χ^2-MI, controlled by $\operatorname{ess\,sup}_{s,\boldsymbol{l}} p_{S_i|\boldsymbol{L}_i}(s \mid \boldsymbol{l})$.

Exponential MI separation. We demonstrate the exponential separation of global MI with its realizability proof. We construct infinitely many pairs of leakage channels with identical local χ^2-MI but distinct exponential rates of global leakage. More strongly, we establish a universal separation theorem: even when all permutation-invariant scalar measures of local leakage agree, the global leakage can have exponentially different decay rates.

Quotients, automorphisms, and finite-share perfect leakage cancellation. We characterize the spectral alignment probability κ_d through character-order multiplicity and Möbius inversion. Then, the problem is group-theoretically generalized to automorphism-weighted sharing. We establish that automorphism-weighted sharing can achieve perfect leakage cancellation with finitely many shares, even for non-trivial full-support local leakage channels. In contrast, perfect cancellation is impossible for ordinary additive sharing with non-trivial i.i.d channels and finitely many shares.

1.3 Paper organization

Section 2 describes our formulation of the problem with a counterexample and the limitations of MI and existing approaches. Section 3 presents our exact local-to-global decomposition with a general formalization. Section 4 presents the exponential separation of MI, its realizability, and impossibility of magnitude-only characterization of leakage, and discusses exact exponential rate of global MI. Section 5 presents the quotient spectral theory for a fine-grained analysis of the spectral alignment and automorphism-weighted sharing. Section 6 evaluates bounds for canonical leakage models to validate our theory for non-trivial settings not as a proof artifact. Section 7 compares our results with existing theories and bounds in the literature of LSS and masking. Finally, Section 8 concludes this paper with some discussions. Appendices include proofs of some theorems and lemmas.

2 Local-to-Global Leakage: A First Separation

2.1 Problem setting

In this paper, we consider $\mathbb{G}$ as general finite abelian group $\mathbb{Z}_{m_1} \times \mathbb{Z}_{m_2} \times \cdots \times \mathbb{Z}_{m_k}$ unless especially specified. Let m be the cardinality of $\mathbb{G}$. If $\mathbb{G}$ is cyclic, we simply write $\mathbb{G} = \mathbb{Z}_m$. A secret variable Z is decomposed into $d+1$ shares $S_0, S_1, \dots, S_d$ as $Z = S_0 + S_1 + \cdots + S_d$. Each share leaks independently through the channel τ_i as $\tau_i(S_i) \to \boldsymbol{L}_i$. The problem is how the local leakages $I(S; \boldsymbol{L})$ compose into the global leakage $I(Z; \boldsymbol{X})$, where $\boldsymbol{X} = (\boldsymbol{L}_0, \boldsymbol{L}_1, \dots, \boldsymbol{L}_d)$. We may omit the subscript i, assuming that the leakages are i.i.d. As aforementioned, it was conjectured that the global MI $I(Z; \boldsymbol{X})$ can be characterized by local MI $I(S; \boldsymbol{L})$; that is, $I(Z; \boldsymbol{X}) \asymp I(S; \boldsymbol{L})^{d+1}$, even though some studies mentioned the MI insufficiency with regard to the strong worst-case leakage [DFS15, BBM26]. Throughout this paper, we assume that S and $\boldsymbol{L}$ are not independent, which means that $p_{S|\boldsymbol{L}}$ is not a uniform distribution and $I(S; \boldsymbol{L}) > 0$, to exclude a trivial case that $I(Z; \boldsymbol{X}) = 0$ for any d.

2.2 A motivating counterexample: equal local MI, different global leakage

Let $\mathbb{G} = \mathbb{Z}_{2^l}$ with $l \geq 2$. Consider two deterministic leakage channels $\boldsymbol{L}^\tau = \mathrm{LSB}(S)$ and $\boldsymbol{L}^{\tau'} = \mathrm{MSB}(S)$, where $\mathrm{LSB}(S)$ and $\mathrm{MSB}(S)$ are the least and most significant bit of S, respectively. Since both LSB and MSB are balanced functions and $|\mathrm{Im}\,\tau| = |\mathrm{Im}\,\tau'| = 2$, the local leakages are equal as $I(S; \boldsymbol{L}^\tau) = I(S; \boldsymbol{L}^{\tau'}) = 1$. On one hand, since $\boldsymbol{L}_\tau = S \bmod 2$ holds, the local leakages $\boldsymbol{L}_0^\tau, \boldsymbol{L}_1^\tau, \dots, \boldsymbol{L}_d^\tau$ additively composes into

$$\tau(Z) = \tau(S_0) + \tau(S_1) + \cdots + \tau(S_d),$$

which implies that the global leakage is given by $I(Z; \boldsymbol{X}^\tau) = 1$. In other words, the LSB leakage divides shares by subgroup of $2\mathbb{Z}_{2^l}$, and the shared value Z can be also divided by the subgroups from the homomorphic aggregation.

On the other hand, such an exact homomorphic aggregation does not exist for the channel τ'. Instead, the MSB leakage limits the share into a half interval of either $[0, 2^{l-1})$ or $[2^{l-1}, 2^l)$. Hence, $p_{S|\boldsymbol{L}}$ is considered as a half-interval distribution of $\mathbb{Z}_{2^l}$. It is not uniform over $\mathbb{Z}_{2^l}$; hence, $\max_{w \neq 0} |\hat{p}_{S|\boldsymbol{L}}(w \mid \boldsymbol{L})| = \rho^* < 1$ holds, where $\hat{f}$ is a discrete Fourier transform (DFT) of f (see Section 4.1 for a rigorous proof). Thus, as in [UIM+26, Lemma 1], Jensen and Parseval state that

$$I(Z; \boldsymbol{X}^{\tau'}) \leq \log \left(1 + \sum_{w \neq 0} \prod_{i=0}^{d} \mathbb{E}_{\boldsymbol{L}} \left| \hat{p}_{S_i|\boldsymbol{L}_i}(w \mid \boldsymbol{L}) \right|^2 \right) = O\left((m-1)(\rho^*)^{2(d+1)} \right), \tag{1}$$

by adding random variables (*i.e.*, convoluting probability distributions).

Equation (1) implies $I(Z; \boldsymbol{X}^{\tau'}) = \exp(-\Omega(d))$, while $I(Z; \boldsymbol{X}^\tau) = 1$ holds. These facts constitute a counterexample of $I(Z; \boldsymbol{X}^\tau)/I(Z; \boldsymbol{X}^{\tau'}) = \exp(\Omega(d))$ with $I(S; \boldsymbol{L}^\tau) = I(S; \boldsymbol{L}^{\tau'})$. This phenomenon (*i.e.*, exponential separation) cannot be explained even by the strong worst-case leakage [DFS15, BBM26]. Our exact decomposition framework presented in Section 3 reveals that the spectral alignment essentially explains this phenomenon in general and enables to derive an asymptotic expansion and exact characterization.

Spectral interpretation. The LSB leakage concentrates to the order-2 Fourier character (*i.e.*, -1), but has Fourier coefficient of zero at other non-trivial characters, meaning that $\rho^* = 1$. Hence, the global leakage does not decay even if the posterior distributions are convoluted, which yields $I(Z; \boldsymbol{X}^\tau) = 1$. In contrast, the MSB leakage is dispersed to order-2^l characters, yielding $\rho^* < 1$, while the number of non-zero Fourier characters is 2^{l-1}. Thus, the concentration or dispersion of leakages among Fourier characters plays an essential role in determining the exponential rate. However, local MI does not capture such spectral phenomenon in principle, as it represents an average amount of leakage through a summation of all non-trivial Fourier characters given leakage.

2.3 What information existing bounds lose

Bounds of global MI by local MIs have been mainly studied in the side-channel masking community [DFS15, IUH22, MRS22, BCG+23, UIM+26]. These studies have mainly featured the total variation distance between the posterior distribution $p_{S|\boldsymbol{L}}$ and uniform distribution $p_S = U(\mathbb{G})$ to provide bounds by local MI through Pinsker's inequality [PF64, FHT03], where $U(\mathbb{G})$ is the uniform distribution over $\mathbb{G}$. In CRYPTO 2026 [UIM+26], Ueno, Inoue, Minematsu, Ito, and Homma presented an upper-bound for general finite abelian group as

$$I(Z; \boldsymbol{X}) \le \log\left(1 + (m-1)(2\ln 2)^{d+1} I(S; \boldsymbol{L})^{d+1}\right)$$

We review their proofs to identify the fundamental limitation of this approach—the above spectral information is lost in the bound. Note that no lower bound is known due to this limitation.

Their proof relies on the triangle and Pinsker's inequalities as

$$\left(\sum_{w \neq 0} \mathbb{E}_{\boldsymbol{L}}\left[\left|\hat{p}_{S|\boldsymbol{L}}(w \mid \boldsymbol{L})\right|^2\right]\right)^2 \le 2\ln 2 \sum_{w \neq 0} I(S; \boldsymbol{L}) = 2\ln(2)(m-1) I(S; \boldsymbol{L}).$$

The left-hand term represents the Fourier energy of the posterior distribution $p_{S|\boldsymbol{L}}$ and captures the spectral information. However, in the right-hand sides, the dependence on Fourier mode disappears, because they evaluated Fourier energy as accumulation of mode-wise leakages greater than/equal to the worst-case amount in terms of $w \neq 0$, which loses the spectral information entirely.

Using data processing inequality, we show that this approach has an essential difficulty in exploiting the spectral information. Let $Y^{(w)} = S + \ker \chi_w$, where χ_w is the Fourier character over $\mathbb{G}$. Using the relation, the above bound can be strengthened as

$$\left(\sum_{w \neq 0} \mathbb{E}_{\boldsymbol{L}}\left[\left|\hat{p}_{S|\boldsymbol{L}}(w \mid \boldsymbol{L})\right|^2\right]\right)^2 \le 2\ln 2 \sum_{w \neq 0} I(Y^{(w)}; \boldsymbol{L}).$$

This is a Fourier-mode aware bound. Meanwhile, $I(Y^{(w)}; \boldsymbol{L}) \le I(S; \boldsymbol{L})$ holds owing to the data processing inequality on Markov chain $Y^{(w)} \leftrightarrow S \leftrightarrow \boldsymbol{L}$, because $Y^{(w)}$ is a function of S. For example of $\mathbb{G} = \mathbb{Z}_2^l$, $I(Y^{(i)}; \boldsymbol{L})$ is one bit at maximum because Fourier character over $\mathbb{Z}_2^l$ is equal to Walsh character $\chi_w(s) = (-1)^{\langle s|w \rangle}$. In contrast, $I(S; \boldsymbol{L})$ is an l-bit information at maximum.

Meanwhile, as in the counterexample of Section 2.2, the data processing inequality can hold with equality (*i.e.*, $I(S; \boldsymbol{L}^{(\tau)}) = I(Y^{(2^{l-1})}; \boldsymbol{L}^\tau)$, implying that we cannot strengthen it unless assumption about leakage. To make matters worse, unless the assumption about leakage, we cannot characterize $I(Y^{(w)}; \boldsymbol{L})$ anymore, while for some w and channel, it can be significantly smaller than $I(S; \boldsymbol{L})$, as

in the counterexample of MSB. Hence, the data processing inequality explains that we are forced to evaluate $I(\boldsymbol{Y}^{(w)}; \boldsymbol{L})$ from above as its potential maximum value as large as $I(S; \boldsymbol{L})$, which means that *the bound by local MI (Section 2.3) essentially loses the spectral information at all and corresponds to a worst-case evaluation in terms of w and τ*. This feature weakens their bounds and makes it difficult to derive lower bounds of global MI. Our theory presented in Section 3.2 further explains and characterizes these limitations in existing approaches. Notably, [UIM+26] identified mode-wise leakage $I(Y^{(w)}; \boldsymbol{L})$ as determining asymptotic behavior of global MI (*i.e.*, whether $I(Z; \boldsymbol{X}) = o(1)$ or $O(1)$). Still, its worst-mode analysis did not characterize the leakage.

The bounds from the research community of LSS considered deterministic leakage functions without focus on MI; hence, their tailored bounds would be stronger than the bound in masking context as their analyses utilize the spectral phenomena of the target leakage function (*e.g.*, [BDIR18, MNPC+22,BHMY26,KA26]). However, what parameter characterizes the global leakage is unclear, as their bounds are given by the order d and the cardinality/order of $\mathbb{G}$, but not local MI.

3 The Exact Local-to-Global Decomposition

3.1 Spectral representation of leakage

Fourier analysis is an important tool to analyze the global leakage as in counterexample in Section 2.2 and used in the previous literature [BDIR18,BHMY26,UIM+26]. We first define the Fourier mode of leakage, which plays an important role in characterizing the leakage channel.

Definition 1 (Fourier mode of leakage). *Let $\hat{p}_{S_i|\boldsymbol{L}_i}(w \mid \boldsymbol{l}) = \mathbb{E}_S[\chi_w(S) \mid \boldsymbol{L} = \boldsymbol{l}]$ be the discrete Fourier transform (DFT) of $p_{S_i|\boldsymbol{L}_i}(\cdot \mid \boldsymbol{l})$ given $\boldsymbol{l}$, where χ_w is the Fourier character over $\mathbb{G}$. For given w, the Fourier mode of leakage is defined as*

$$\lambda_i(w) = \mathbb{E}_{\boldsymbol{L}} \left[\left| \hat{p}_{S_i|\boldsymbol{L}_i}(w \mid \boldsymbol{L}) \right|^2 \right].$$

In this paper, we primarily focus on the χ^2-MI as a major leakage quantity, which is considered as a variant of MI by replacing the KL divergence to the chi-square divergence. Our exact local-to-global decomposition is first established in the χ^2-MI, and then we prove its Shannon equivalence.

Definition 2 (χ^2-mutual information (χ^2-MI)). *Let $\chi^2(p \parallel q)$ be the chi-square divergence between discrete probability distributions p and q over $\mathcal{X}$ defined as*

$$\chi^2(p \parallel q) = \sum_{x \in \mathcal{X}} \frac{(p(x) - q(x))^2}{q(x)}.$$

The χ^2-MI between random variables X and Y is defined as

$$I_{\chi^2}(X; Y) = \mathbb{E}_Y \left[\chi^2 \left(p_{X|Y}(\cdot \mid Y) \,\middle\|\, p_X \right) \right].$$

Note that $I_{\chi^2}(S; \boldsymbol{L}) > 0$ if $p_{S|\boldsymbol{L}}$ is not a uniform distribution. One important property of χ^2-MI is that, when we consider a uniform distribution (p_S here), it can be decomposed into the sum of non-trivial Fourier modes as

$$I_{\chi^2}(S; \boldsymbol{L}) = \sum_{w \neq 0} \lambda(w),$$

according to Parseval's theorem. This also implies that, in our problem setting, $I_{\chi^2}(S; \boldsymbol{L})$ is equivalent to the quantity named leakage energy in [UIM+26, Definition 8].

3.2 The exact decomposition theorem

We here present our exact local-to-global decomposition theorem in the χ^2-MI notion to explain its underlying mathematical mechanism, which is followed by its compatibility interpretation and Shannon equivalence.

Theorem 1 (Exact spectral decomposition). *For each i, let $q_i(w) = \lambda_i(w)/\sum_{w\neq 0}\lambda_i(w)$ be a probability distribution of the normalized Fourier leakage spectrum over $\mathcal{W}\setminus\{0\}$ conditioned by leakage. The global leakage is exactly decomposed in χ^2-MI as*

$$I_{\chi^2}(Z;\boldsymbol{X}) = \kappa_d \prod_{i=0}^{d} I_{\chi^2}(S_i;\boldsymbol{L}_i),$$

where $\kappa_d = \Pr[W_0 = W_1 = \cdots = W_d]$ ($W_i \sim q_i$) is the spectral alignment factor. If leakages are independent and identically distributed (i.i.d) for all shares (i.e., $q_0 = q_1 = \cdots = q_d$), the spectral alignment probability is equivalent to

$$\kappa_d = 2^{-dH_{d+1}(q)},$$

where $H_{d+1}(q)$ is the Rényi-$(d+1)$ entropy of the normalized Fourier leakage spectrum.

Proof. This is proven by tightening an inequality between the diagonal sum and all combination sum, represented by

$$I_{\chi^2}(Z;\boldsymbol{X}) = \sum_{w\neq 0}\prod_{i=0}^{d}\lambda_i(w) \le \prod_{i=0}^{d}\sum_{w\neq 0}\lambda_i(w) = \prod_{i=0}^{d} I_{\chi^2}(S_i;\boldsymbol{L}_i).$$

We then interpret the conversion factor κ_d as a diagonal sum of probability distributions $q_i(w) = \lambda_i(w)/I_{\chi^2}(S_i;\boldsymbol{L}_i)$. See Appendix B.1. □

The compatibility interpretation. The exact decomposition constitutes an interchange of sum and product—representing the relation between diagonal sum and all combination sum. This indicates that only the diagonal Fourier modes of local leakages contribute to the global leakage under the algebraic reconstruction constraint, while cross-mode combinations cannot simultaneously contribute. Hence, the exact decomposition needs accumulating only the compatible Fourier modes, which are aligned with probability $\kappa_d = \Pr[W_0 = W_1 = \cdots = W_d]$ ($W_i \sim q_i$), equivalently $2^{-dH_{d+1}(q)}$ under the i.i.d assumption ($q_i = q$ for all i). In contrast, all incompatible combinations should be eliminated. The number of such combinations is exponential in d and their occurrence can be with an exponentially overwhelming probability. Thus, the spectral alignment probability κ_d is an essential and indispensable factor to characterize the global MI, which represents the MI insufficiency. Existing MI-only bounds might evaluate all Fourier-mode combinations including non-compatible alignments, potentially incurring a multiplicative slack of $\exp(\Theta(d))$.

3.3 Shannon equivalence

We first show the conversion between Shannon MI and χ^2-MI. With this, we prove an upper bound on $I(Z;\boldsymbol{X})$, then provide its exact asymptotic expansion by second-order approximation. The Shannon bounds imply that, if we recover the essential information discarded by local MI bounds, the Rényi-2 masking leakage exactly factorizes the spectral alignment and leakage sharpness factors.

Lemma 1 (Conversion between Shannon MI and χ^2-MI). *Let $f(r)$ and $g(r)$ be measurable functions $\mathbb{R}^+ \to \mathbb{R}$ defined as*

$$f(r) = r \ln r - r + 1, \qquad g(r) = \frac{(r-1)^2}{f(r)},$$

respectively ($g(1) \coloneqq 2$). For all i, for given s and $\boldsymbol{l}$, let $r(s, \boldsymbol{l})$ be a likelihood ratio defined using a Radon–Nikodim derivative as

$$r(s, \boldsymbol{l}) = \frac{\mathrm{d}p_{S,\boldsymbol{L}}}{\mathrm{d}(p_S \otimes p_{\boldsymbol{L}})}(s, \boldsymbol{l}) = \frac{\mathrm{d}p_{\boldsymbol{L}|S=s}}{\mathrm{d}p_{\boldsymbol{L}}}(\boldsymbol{l}). \tag{2}$$

Let $\mu_i(s, \mathrm{d}\boldsymbol{l})$ be a D_{KL}-weighted product probability measure of p_S and $p_{\boldsymbol{L}}$, defined as

$$\mu_i(s, \mathrm{d}\boldsymbol{l}) = \frac{f(r_i)}{\ln(2) I(S; \boldsymbol{L})} p_{S_i}(s) p_{\boldsymbol{L}_i}(\mathrm{d}\boldsymbol{l}).$$

It holds that

$$I_{\chi^2}(S_i; \boldsymbol{L}_i) = \alpha_i I(S_i; \boldsymbol{L}_i)$$

where $\alpha_i = \ln(2) \mathbb{E}_{\mu_i} g(r_i)$ and $\mathbb{E}_{\mu_i}$ is an expectation in terms of the probability measure μ_i.

Proof. Functions f and g represent KL and χ^2 divergences using likelihood ratio, respectively. We derive the conversion factor as a KL-weighted expectation using f and g. See Appendix B.2. □

The conversion factor α represents the leakage sharpness. The χ^2-MI is tail-sensitive more than Shannon MI. Hence, α acts as a weight for rare but strong leakage (*i.e.*, large $\max_s p_{S|\boldsymbol{L}}(s \mid \boldsymbol{l})$ but small $p_{\boldsymbol{L}}(\boldsymbol{l}) > 0$). Recall that $r(s, \boldsymbol{l})$ represents a likelihood ratio $p_{S|\boldsymbol{L}=\boldsymbol{l}}/p_S$, which implies that $I(S; \boldsymbol{L}) = \mathbb{E}_{p_S \otimes p_{\boldsymbol{L}}} f(r)$ and $I_{\chi^2}(S; \boldsymbol{L}) = \mathbb{E}_{p_S \otimes p_{\boldsymbol{L}}}[(r-1)^2]$. Considering rare but strong leakage (*i.e.*, $r \to \infty$, while r is bounded for given finite groups), it holds that

$$\frac{(r-1)^2}{f(r)} \approx \frac{r^2}{r \log r} \to \infty. \tag{3}$$

Equation (3) implies that a rare leakage contributes to χ^2-MI more greatly than Shannon MI, if the leakage has a much larger posterior probability than its prior. The leakage sharpness factor α captures this phenomenon. Thus, $\alpha I(S; \boldsymbol{L})$ is an effective leakage amount in our problem setting.

Furthermore, our exact local-to-global decomposition enables a more interesting interpretation of α, beyond the tail-sensitivity of χ^2-MI. Theorem 1 states that the local leakage is decomposed into spectral coefficients and then propagates to global leakage through the aligned Fourier modes (*i.e.*, squared power spectral density for each $w \neq 0$), whose sum equals χ^2-MI. As $\sum_w \lambda_w = I_{\chi^2}(S; \boldsymbol{L})$ holds, χ^2-MI quantifies the local-to-global propagation of information more directly than Shannon MI. Hence, $I_{\chi^2}(S; \boldsymbol{L}) = \alpha I(S; \boldsymbol{L})$ can be interpreted as a new finding that the quantity should be the product of the local-to-global propagation factor and the average information.

Using Theorem 1 and Lemma 1 we derive sharp bounds in Shannon as Theorem 2, which states that the global leakage is exactly characterized by three factors—spectral alignment probability, leakage sharpness, and leakage magnitude (*i.e.*, local MI).

Theorem 2 (Sharp bound on global leakage in Shannon). *If strict weak worst-case leakage condition, which is defined as* $\rho^* = \sup_i \max_{w\neq 0} \operatorname{ess\,sup}_{\boldsymbol{l}} \left|\hat{p}_{S_i|\boldsymbol{L}_i}(w \mid \boldsymbol{l})\right| < 1$*, holds, the global Shannon MI is asymptotically expanded as*

$$I(Z;\boldsymbol{X}) = \frac{1+o(1)}{2\ln 2}\kappa_d \prod_{i=0}^{d} \alpha_i I(S_i;\boldsymbol{L}_i) = \Theta\left(\kappa_d \prod_{i=0}^{d} \alpha_i I(S_i;\boldsymbol{L}_i)\right) \quad \text{as } d\to\infty.$$

Proof. This is proven using the second-order Taylor approximation of D_{KL} by χ^2 divergence in Appendix A. See Appendix B.3. □

The convergence radius of Taylor expansion means the weak worst-case global leakage regime such that $\operatorname{ess\,sup}_{z,\boldsymbol{x}} p_{Z|\boldsymbol{X}}(z \mid \boldsymbol{x}) < 2/m$. Under the strict weak worst-case leakage condition, the additive secret masking itself pushes the global leakage into the weak worst-case global leakage regime. The error term $O((m-1)(\rho^*)^{d+1})$ includes the group order, which have been considered in a proof artifact in existing upper bounds [DFS15, BBM26] and actually vanishes as $d\to\infty$. In addition, Lemma 3 implies that the leakage sharpness factor α converges to $2\ln 2$ at the limit of weak worst-case leakage. When the worst-case local leakage is in the weak regime of the convergence radius, we can apply the Taylor expansion to $I(S_i;\boldsymbol{L}_i)$, yielding Corollary 1.

Corollary 1 (Sharp bound of global MI in weak worst-case local leakage regime). *If* $\operatorname{ess\,sup}_{s,\boldsymbol{l}} p_{S_i|\boldsymbol{L}_i}(s \mid \boldsymbol{l}) < 2/m$ *holds, the global MI is asymptotically expanded as*

$$I(Z;\boldsymbol{X}) = \left(\kappa_d (2\ln 2)^d \prod_{i=0}^{d} I(S_i;\boldsymbol{L}_i)\right)(1+o(1))^{d+1} \quad \text{as } d\to\infty,\ \operatorname{ess\,sup}_{s,\boldsymbol{l}} \left|\frac{p_{S_i|\boldsymbol{L}_i}(s \mid \boldsymbol{l})}{p_S(s)} - 1\right| \to 0.$$

Proof. It is obvious from Theorem 2, Lemma 3 for $I(S_i;\boldsymbol{L}_i)$, and $\alpha_i \to 2\ln 2$. □

Corollary 1 implies that $2\ln 2$ represents an essential factor for Shannon–χ^2 conversion and local-to-global propagation. While the local MI requires $d+1$ conversions from nat to bit for the local-to-global propagation, it is finally aggregated as the global MI, which is inversely converted from bit to nat. Hence, the exponent of $2\ln 2$ is not $d+1$ but d.

4 Beyond Mutual Information: Exponential Separation and Structural Amplification

4.1 Realizability of spectral leakage profile

As in Section 3, we present the separation in χ^2-MI to explain its mathematical mechanism, which is followed by Shannon through Shannon–χ^2 conversion. We first prove the separation in a case that the leakage is extremely concentrated in low order character and dispersed to high order characters. In this section, we assume $\mathbb{G} \neq \mathbb{Z}_2$ nor $\mathbb{Z}_3$, where the separation does not exist due to the impossibility of the spectral information among Fourier modes.

We first show a rigorous proof of the counterexample of Section 2.2 as canonical leakage, which is followed by the realizability proof for general case.

Proposition 1 (Spectral concentration vs. dispersion (canonical example)). *Let* $\tau = \mathrm{LSB}$ *and* $\tau' = \mathrm{MSB}$ *be local leakage channel over* $\mathbb{G} = \mathbb{Z}_{2^l}$ *($l \geq 2$). It holds that*

$$I_{\chi^2}(S;\boldsymbol{L}^{\tau}) = I_{\chi^2}(S;\boldsymbol{L}^{\tau'}), \qquad \frac{I_{\chi^2}(Z;\boldsymbol{X}^{\tau})}{I_{\chi^2}(Z;\boldsymbol{X}^{\tau'})} = \exp(\Omega(d)).$$

Proof. From Theorem 1 and the assumption that $I_{\chi^2}(S;\boldsymbol{L}^\tau) = I_{\chi^2}(S;\boldsymbol{L}^{\tau'})$, it holds

$$\frac{I_{\chi^2}(Z;\boldsymbol{X}^\tau)}{I_{\chi^2}(Z;\boldsymbol{X}^{\tau'})} = \frac{\kappa_d^\tau}{\kappa_d^{\tau'}},$$

where κ_d^τ and $\kappa_d^{\tau'}$ are the spectral alignment probabilities for τ and τ', respectively. It suffices to prove that κ_d^τ and $\kappa_d^{\tau'}$ can be exponentially different even if $I_{\chi^2}(S;\boldsymbol{L}^\tau) = I_{\chi^2}(S;\boldsymbol{L}^{\tau'})$.

Recall that $I_{\chi^2}(S;\boldsymbol{L}^\tau) = \sum_{w\neq 0}\lambda^\tau(w)$, where $\lambda^\tau(w)$ is the power spectral density for τ. On one hand, the LSB leakage is concentrated to the order-2 character (*e.g.*, $\boldsymbol{L}^\tau = S \bmod 2$ over $\mathbb{Z}_{2^l}$ with $l>1$), yielding $\lambda^\tau(w)=1$ at the order-2 character and $\lambda^\tau(w)=0$ at other characters with $w\neq 0$. On the other hand, the MSB leakage is dispersed to higher-order order characters with odd w, as the power spectral density is given by

$$\lambda^{\tau'}(w) = \begin{cases} \frac{1}{2^{2l}\sin^2\left(\frac{\pi w}{2^l}\right)} < 1 & \text{for odd } w, \\ 0 & \text{for even } w. \end{cases}$$

Hence, we conclude

$$\frac{\kappa_d^\tau}{\kappa_d^{\tau'}} = \frac{\sum_{w\neq 0}\lambda^\tau(w)^{d+1}}{\sum_{w\neq 0}\lambda^{\tau'}(w)^{d+1}} = \frac{1}{\sum_{w\neq 0}\lambda^{\tau'}(w)^{d+1}} = \exp(\Omega(d)),$$

as required. □

Proposition 1 demonstrates that the exponential rate depends on how the spectral leakage information is concentrated to a Fourier mode or dispersed, as Section 4.2 exactly shows that the exponential rate is determined by $\|q\|_\infty$. In contrast, it only shows its existence of separation in a case. Theorem 3 generalizes it to any local χ^2-MI by establishing its realizability; namely, pairs of channels separating global MI exist for any sufficiently small local χ^2-MI value for some $\mathbb{G}$, where the spectrum can vary, not limited to deterministic leakages over $\mathbb{Z}_{2^l}$.

Theorem 3 (Realizability of spectrally distinct leakage channels). *Let $\mathbb{G} \neq \mathbb{Z}_2$ nor $\mathbb{Z}_3$ be a finite abelian group. For any conjugate-symmetric probability distribution q over non-trivial characters, full-support stochastic channels exist with sufficiently small local χ^2-MI. Then, we can construct infinitely many pairs of stochastic channels τ and τ' such that $I_{\chi^2}(S;\boldsymbol{L}^\tau) = I_{\chi^2}(S;\boldsymbol{L}^{\tau'})$ but $I_{\chi^2}(Z;\boldsymbol{X}^\tau)$ and $I_{\chi^2}(Z;\boldsymbol{X}^\tau)$ have strictly different exponential rates. Moreover, both channels satisfy the strict weak worst-case leakage condition:* $\max_{w\neq 0}\operatorname{ess\,sup}_{\boldsymbol{l}}\hat{p}_{S|\boldsymbol{L}}(w\mid \boldsymbol{l}) < 1$.

Proof. Define

$$g_q(s) = \sum_{w\neq 0}\sqrt{q(w)}\chi_w(s), \qquad G_q = \max_s |g_q(s)|, \qquad h_q(s) = \frac{g_q(s)}{G_q},$$

where g_q is the inverse DFT corresponding to the power density spectrum (*i.e.*, $|\hat{g}_q(w)|^2 = q(w)$ holds for any $w\neq 0$), G_q is a normalization term, and h_q is a bounded function of realization to construct a stochastic channel. Here, $g_q(s)$ is real for any s because of the conjugate symmetry of q. In addition, $\mathbb{E}_S g_q(S) = 0$ and $\mathbb{E}_S g_q(S)^2 = 1$ hold owing to the conjugate symmetry and Parseval,

respectively. In addition, $-1 \le h_q(s) \le 1$ holds for any s. Define a binary leakage $\boldsymbol{l} \in \{-1, 1\}$. For any $0 < \epsilon < 1$, define the leakage channel as

$$p_{\boldsymbol{L}|S}(\boldsymbol{l} \mid s) = \frac{1 + \boldsymbol{l}\epsilon h_q(s)}{2}.$$

This is a full-support probability distribution because it is in the range of $(0, 1)$ as $|h_q(s)| \le 1$ and summation over $\boldsymbol{l} \in \{-1, 1\}$ equals one for any s.

Suppose that $S \sim U(\mathbb{G})$, which is followed by $p_{\boldsymbol{L}}(1) = p_{\boldsymbol{L}}(-1) = 1/2$ because

$$p_{\boldsymbol{L}}(\boldsymbol{l}) = \mathbb{E}_S p_{\boldsymbol{L}}(\boldsymbol{l} \mid S) = \frac{1 + \boldsymbol{l}\epsilon \mathbb{E}_S h_q(S)}{2} = \frac{1}{2}.$$

Note that $\mathbb{E}_S h_q(S) = 0$ holds because $\sum_s g_q(s) = 0$. Hence, according to Bayes' theorem,

$$p_{S|\boldsymbol{L}}(s \mid \boldsymbol{l}) = \frac{1 + \boldsymbol{l}\epsilon h_q(s)}{m}, \qquad \frac{p_{S|\boldsymbol{L}}(s \mid \boldsymbol{l})}{p_S(s)} - 1 = \boldsymbol{l}\epsilon h_q(s).$$

Hence, for $w \neq 0$, we have

$$\hat{p}_{S|\boldsymbol{L}}(w \mid \boldsymbol{l}) = \boldsymbol{l}\epsilon \hat{h}_q(w), \qquad \lambda(w) = \frac{\epsilon^2 q(w)}{G_q^2},$$

which is followed by

$$I_{\chi^2}(S; \boldsymbol{L}) = \frac{\epsilon^2}{G_q^2}.$$

Here, ϵ can take any real value in $(0, 1)$ and hence $I_{\chi^2}(S; \boldsymbol{L})$ is continuous in $(0, 1/G_q^2)$. Thus, for spectral profiles q^τ and q^τ, there exist infinitely many pairs of channels τ and τ' such that $I_{\chi^2}(S; \boldsymbol{L}^\tau) = I_{\chi^2}(S; \boldsymbol{L}^{\tau'})$ with appropriate ϵ. This proves the claimed realization. □

4.2 Exact security exponent by structure

In literature, the exponential decay of global MI has been believed by the *noise amplification*, which means that the exponential rate is determined by the product of local MI or other related leakage magnitude [CJRR99, DFS16]. However, our exact decomposition theorem implies that the global leakage amount is characterized not only by the noise, but also the algebraic structure of sharing itself contributes to the exponential rate through the spectral alignment probability. We identify another exponent exactly as stated in Theorem 4.

Theorem 4 (Exponential rate of global χ^2-MI). *Let $I_{\chi^2}(S; \boldsymbol{L}) = M$ denote the local leakage amount. The exponential rate of global χ^2-MI is given by*

$$-\lim_{d\to\infty} \frac{1}{d} \log I_{\chi^2}(Z; \boldsymbol{X}) = -\log M + H_\infty(q).$$

Proof. Recall that $\kappa_d = \sum_{w\neq 0} q(w)^{d+1}$, which is followed by

$$-\lim_{d\to\infty} \frac{1}{d} \log \kappa_d = -\log \|q\|_\infty = H_\infty(q), \tag{4}$$

where H_∞ is the min-entropy. From Theorem 1 and Equation (4), we conclude

$$-\lim_{d\to\infty}\frac{1}{d}\log I_{\chi^2}(Z;\boldsymbol{X}) = -\lim_{d\to\infty}\left(\log I_{\chi^2}(S;\boldsymbol{L})^{1+\frac{1}{d}} + \frac{1}{d}\log\kappa_d\right) = -\log M + H_\infty(q),$$

as required. □

The full spectrum determines the exact leakage at every d whereas its largest spectral component determines the asymptotic exponent. We then derive the global MI exponent in Shannon.

Corollary 2 (Exponential rate of global Shannon MI). *Under the strict weak worst-case leakage assumption, the exponential rate of global Shannon MI is given by*

$$-\lim_{d\to\infty}\frac{1}{d}\log I(Z;\boldsymbol{X}) = -\log\left(\alpha I(S;\boldsymbol{L})\right) + H_\infty(q).$$

Proof. It is obvious from Theorem 4 and the Shannon–χ^2 equivalence. □

4.3 Exponential separation theorems

We first prove the exponential separation in the χ^2-MI, followed by Shannon.

Theorem 5 (Exponential separation of global leakage in χ^2-MI). *Let τ and τ' be leakage channels such that $I_{\chi^2}(S;\boldsymbol{L}^\tau) = I_{\chi^2}(S;\boldsymbol{L}^{\tau'}) = M$. For $\mathbb{G}\neq\mathbb{Z}_2$ nor $\mathbb{Z}_3$, pairs of τ and τ' exist such that $I_{\chi^2}(Z;\boldsymbol{X}^\tau)$ and $I_{\chi^2}(Z;\boldsymbol{X}^{\tau'})$ have different exponential rates; namely, it holds*

$$-\lim_{d\to\infty}\frac{1}{d}\log\frac{I_{\chi^2}(Z;\boldsymbol{X}^\tau)}{I_{\chi^2}(Z;\boldsymbol{X}^{\tau'})} > 0.$$

Proof. When $\|q^\tau\|_\infty < \|q^{\tau'}\|_\infty$ without loss of generality, according to Theorem 4, it holds

$$-\lim_{d\to\infty}\frac{1}{d}\log\frac{I_{\chi^2}(Z;\boldsymbol{X}^\tau)}{I_{\chi^2}(Z;\boldsymbol{X}^{\tau'})} = H_\infty(q^\tau) - H_\infty(q^{\tau'}) > 0.$$

Thus, it suffices to prove that pairs of τ and τ' exist such that $I_{\chi^2}(S;\boldsymbol{L}^\tau) = I_{\chi^2}(S;\boldsymbol{L}^{\tau'}) = M$ but $\|q^\tau\|_\infty < \|q^{\tau'}\|_\infty$. Theorem 3 states that, pairs of τ and τ' exist such that $I_{\chi^2}(S;\boldsymbol{L}^\tau) = I_{\chi^2}(S;\boldsymbol{L}^\tau)$ with $0 < M < \min(G_{q^\tau}^{-2}, G_{q^{\tau'}}^{-2})$ by appropriately choosing q^τ and $q^{\tau'}$ with $\|q^\tau\|_\infty < \|q^{\tau'}\|_\infty$. □

The local leakage magnitude $I_{\chi^2}(S;\boldsymbol{L})$ and spectral alignment κ_d are independent: fixing the local χ^2-MI does not determine the alignment exponent—local MI and spectral alignment probability are independent degrees of freedom. Then, Corollary 3 states the equivalence of separation in χ^2-MI to Shannon.

Corollary 3 (Separation theorem in Shannon). *Let τ and τ' be leakage channels such that $I(S;\boldsymbol{L}^\tau) = I(S;\boldsymbol{L}^{\tau'}) = M$ for sufficiently small M. For $\mathbb{G}\neq\mathbb{Z}_2$ nor $\mathbb{Z}_3$, pairs of τ and τ' exist such that $I(Z;\boldsymbol{X}^\tau)$ and $I(Z;\boldsymbol{X}^{\tau'})$ can have different exponential rates. Under the strict weak worst-case leakage assumption, it holds*

$$-\lim_{d\to\infty}\frac{1}{d}\log\frac{I(Z;\boldsymbol{X}^\tau)}{I(Z;\boldsymbol{X}^{\tau'})} > 0.$$

Proof. According to Lemma 1 with regard additive error (see Lemma 3), for τ and q^τ in the proof of Theorem 3, it holds

$$I(S;\boldsymbol{L}^\tau) = \frac{\epsilon^2}{2\ln(2)G_{q^\tau}^2} + O(\epsilon^4),$$

likewise $I(S;\boldsymbol{L}^{\tau'})$. By continuity of the local Shannon MI with respect to ϵ, we can independently choose the channel parameters so that both channels have the same sufficiently small local Shannon MI M. In this case, Lemma 1 and Theorem 5 imply

$$-\lim_{d\to\infty}\frac{1}{d}\log\frac{I(Z;\boldsymbol{X}^\tau)}{I(Z;\boldsymbol{X}^{\tau'})} = -\log\frac{\alpha^\tau}{\alpha^{\tau'}} + H_\infty(q^\tau) - H_\infty(q^{\tau'}).$$

Here, $\alpha^\tau/\alpha^{\tau'} \to 1$ at the limit of weak worst-case leakage, followed by $\log(\alpha^\tau/\alpha^{\tau'}) = o(1)$ as $M \to 0$. This completes the proof. □

This theorem strictly separates the Shannon MI for weak worst-case leakage regime of sufficiently small M, where $H_\infty(q^\tau) - H_\infty(q^{\tau'})$ is not canceled out. Furthermore, for some $\mathbb{G}$, we can construct pairs of leakage channels such that $I(S;\boldsymbol{L}^\tau) = I(S;\boldsymbol{L}^{\tau'})$ exactly holds while $I_{\chi^2}(S;\boldsymbol{L}^\tau) = I_{\chi^2}(S;\boldsymbol{L}^{\tau'})$. The pair of LSB and MSB leakages is such an example. For permutation-related leakage channels, every *permutation-invariant* local leakage magnitude, including Shannon MI and χ^2-MI, takes the same value. For example, define I_{L_θ}-MI and $I_{L_\theta^\psi}$-MI as

$$I_{L_\theta}(S;\boldsymbol{L}) = \mathbb{E}_{\boldsymbol{L}}\left\|p_{S|\boldsymbol{L}}(\cdot\mid\boldsymbol{L}) - p_S\right\|_\theta, \qquad I_{L_\theta^\psi}(S;\boldsymbol{L}) = \mathbb{E}_{\boldsymbol{L}}\left\|p_{S|\boldsymbol{L}}(\cdot\mid\boldsymbol{L}) - p_S\right\|_\theta^\psi,$$

for any $\theta \geq 1$, which are MI induced by L_θ distance and its θ-th power, respectively. This class of leakage magnitudes includes the total variation distance (TV), statistical distance (SD), and *bias* [PR13], which has been analyzed in the cryptologic community and used for bounding the global MI [DFS15, UIM+26, BBM26]. Theorem 6 states an impossibility of characterizing the global MI by any permutation-invariant local leakage magnitude.

Theorem 6 (Universal separation of permutation-invariant local leakage magnitudes). *Let $J(S;\boldsymbol{L})$ be a permutation-invariant local leakage magnitude, such as χ^2-MI, Shannon MI, L_θ-MI, and L_θ^ψ-MI. For some $\mathbb{G}$, pairs of (full-support) leakage channels τ and τ' exist such that global Shannon MI exponents are different while their local magnitudes are identical; namely, even if $J(S;\boldsymbol{L}^\tau) = J(S;\boldsymbol{L}^\tau)$, it can hold*

$$-\lim_{d\to\infty}\frac{1}{d}\log\frac{I(Z;\boldsymbol{X}^\tau)}{I(Z;\boldsymbol{X}^{\tau'})} = H_\infty(q^\tau) - H_\infty(q^{\tau'}) > 0.$$

Proof. Define the channels τ and τ' such that $p_{S|\boldsymbol{L}^\tau}(s\mid\boldsymbol{l}) = p_{S|\boldsymbol{L}^{\tau'}}(\pi(s)\mid\boldsymbol{l})$ for any s and $\boldsymbol{l}$, where π is a permutation on $\mathbb{G}$. For any permutation-invariant magnitudes, $J(S;\boldsymbol{L}^\tau) = J(S;\boldsymbol{L}^\tau)$ holds including χ^2-MI. In other words, since the two channels are related by an input permutation, every permutation-invariant local leakage magnitude takes the same value for both channels. Meanwhile, $\|q^\tau\|_\infty$ and $\|q^{\tau'}\|_\infty$ can be different because the permutation does not generally preserve the Fourier mode alignment. Thus, by using Theorem 5, we prove the separation if such channels are realizable. For $\mathbb{G} = \mathbb{Z}_{2^l}$ with $l \geq 2$, by letting $h^\tau(s) = (-1)^{\mathrm{LSB}(s)}$ and $h^{\tau'}(s) = (-1)^{\mathrm{MSB}(s)}$, which yields $\|q^\tau\|_\infty \neq \|q^{\tau'}\|_\infty$, Theorem 3 provides realization of full-support channels satisfying the strict weak worst-case condition. □

In Theorem 6, we only showed its realizability for $\mathbb{Z}_{2^l}$ ($l \geq 2$). For even-order groups (excluding $\mathbb{Z}_2$ and $\mathbb{Z}_2^2$), Theorem 3 may be a realization with balanced functions $h^\tau(s) = (-1)^{\mathbb{1}_{\mathcal{S}}(s)}$ and $\boldsymbol{L}^{\tau'}(s) = (-1)^{\mathbb{1}_{\mathcal{S}'}(s)}$ such that $\|q^\tau\|_\infty \neq \|q^{\tau'}\|_\infty$, where $\mathbb{1}_{\mathcal{S}}$ is the indicator function of set $\mathcal{S} \subset \mathbb{G}$. This strategy cannot prove the case of $\mathbb{Z}_2^2$ because all balanced functions over $\mathbb{Z}_2^2$ are affine, implying $\|q^\tau\|_\infty = \|q^{\tau'}\|_\infty$. For odd-order groups, it may be proven using subsets of $\mathbb{G}$ such that $|\mathcal{S}| = |\mathcal{S}'| \geq 3$ that makes the maximum of power spectrum different, which we need to check for each $\mathbb{G}$. In summary, $\mathbb{Z}_2$ and $\mathbb{Z}_3$ are not separable, while separation in $\mathbb{Z}_5$ and $\mathbb{Z}_2^2$ cannot be proven by using the above realization. It can be extended to other even-order groups, and odd-order groups with $m \geq 7$ may be proven by the above strategy.

4.4 When are magnitude-only bounds tight?

Theorem 4 and Corollary 2 reveal when the magnitude-only bounds suffer from exponential slack, as stated by Theorem 7.

Theorem 7 (Necessary and sufficient condition for exponential tightness). *A necessary and sufficient condition for spectral-alignment relaxation to be exponentially tight is $H_\infty(q) = 0$.*

Proof. It is obvious from Theorem 4 and Corollary 2. □

The condition $H_\infty(q) = 0$ is equivalent to that q is a one-hot distribution. The LSB leakage has such a profile; namely, the leakage is completely concentrated to the order-2 character. Hence, $I(Z; \boldsymbol{X}^{\mathrm{LSB}}) = \prod_{i=0}^{d} I(S_i; \boldsymbol{L}_i^{\mathrm{LSB}})$ holds. Corollary 4 states that the condition is restrictive—only the LSB-related leakage can make spectral-alignment relaxation exponentially tight.

Corollary 4 (Only local leakages concentrating to order-2 character can make spectral-alignment relaxation exponentially tight). *Spectral-alignment relaxation can be exponentially tight if and only if the leakage depends exclusively on a non-trivial parity of the LSBs of even-order invariant factors and is exclusively concentrated to a single non-trivial order-2 character. If the group order is odd, the local-MI-only bound of global MI must have an exponential slack incurred by the spectral alignment probability.*

Proof. Recall that a necessary and sufficient condition is $H_\infty(q) = 0$, which implies that q is a one-hot distribution. Also recall that q is the probability distribution of normalized power spectral density over the frequency $\mathcal{W} \setminus \{0\}$, equivalently to the set of non-trivial Fourier characters $\{\, \chi_w \mid w \neq 0 \,\}$. Because of the conjugate symmetry, $q(\chi_w) = q(\bar{\chi}_w)$ always holds for any w. Unless $\chi_w = \bar{\chi}_w$, $|\mathrm{supp}(q)| \geq 2$ holds, which cannot be a one-hot distribution. Hence, q can be one-hot if and only if it has non-zero value at a non-trivial character satisfying $\chi_w = \bar{\chi}_w$. Order-2 characters are only such non-trivial characters, because $\chi_w = \bar{\chi}_w$ holds if and only if their imaginary parts are zero and hence χ_w takes values in $\{-1, 1\}$.

Let $\mathbb{G} = \bigoplus_{j=1}^{k} \mathbb{Z}_{m_j}$ and let $w = (w_1, w_2, \ldots, w_k)$ be the frequency over $\mathbb{G}$. The j-th invariant has no non-trivial order-2 character when m_j is odd; if m_j is even, the non-trivial order-2 characters are at $w_j = m_j/2$. Let $\mathcal{J}$ be the set of indices of invariant factors with even order and let $s = (s_1, s_2, \ldots, s_k) \in \mathbb{G}$. All order-2 characters of $\mathbb{G}$ are represented by

$$\chi_w(s) = (-1)^{\sum_{j' \in \mathcal{J}'} \mathrm{LSB}(s_{j'})} = (-1)^{\sum_{j' \in \mathcal{J}'} \mathrm{LSB}(s_{j'}) \bmod 2},$$

where $\mathcal{J}'$ is a non-trivial subset of $\mathcal{J}$. This character represents the claimed LSB-XOR leakage.

If the group order is odd, no invariant factor is even and has non-trivial order-2 characters. Hence, $H_\infty(q) > 0$ always holds for odd-order group. This completes the proof. □

Thus, MI-only bounds can be exponentially tight only if the leakage power is entirely concentrated on a parity character—non-zero XOR of the LSBs of the even-order invariant factors. In contrast, odd-order group always has a strictly positive exponent of spectral alignment probability, that makes spectral-alignment relaxation exponentially loose.

The LSB leakage is especially dangerous in terms of global leakage amount. Although this has been claimed by Faust, Masure, Micheli, Nguyen, Orlt, and Standaert in Eurocrypto 2024 [FMM+24], our theory explains what, how, and when leakage is dangerous. Leakage exclusively through LSB-XOR is sufficient for exponential tightness. However, leakage that reveals an LSB-XOR together with additional information need not satisfy this condition, as illustrated by Hamming-weight leakage over $\mathbb{Z}_2^l$.

5 Quotient Spectral Theory

In Section 3, we derived the *character* spectral theory $I_{\chi^2}(Z;\boldsymbol{X}) = \kappa_d I_{\chi^2}(S;\boldsymbol{L})^{d+1}$, where $\kappa_d = \sum_{w\neq 0} q(w)^{d+1}$. For arbitrary non-trivial character $\chi_w : \mathbb{G} \to \mathbb{C}^\times$, it holds that $\mathbb{G}/\ker\chi_w \simeq \operatorname{Im}\chi_w$ according to the first homomorphism theorem. Since $\mathbb{G}$ is finite, $\operatorname{Im}\chi_w \subset \mathbb{C}^\times$ is also finite, and hence it is cyclic and isomorphic to $\mathbb{Z}_{\gamma_w} = \mathbb{Z}/\gamma_w\mathbb{Z}$. This implies that the Fourier mode $\lambda(w)$ represents how much information about primitive Fourier mode specific to quotient $\mathbb{G}/\ker\chi_w$ leaks through $\boldsymbol{L}$. As a result, we characterize and derive a lower bound of κ_d using the Euler totient. Then, we discuss automorphism-weighted sharing, which can de-align the spectrum by structure beyond the lower bound.

5.1 Cyclic group

In this subsection, we analyze the properties of the spectral alignment structure through the quotient spectra for cyclic group, followed by generalization to non-cyclic finite abelian group in Section 5.2.

Main theorem. Define the order-γ primitive spectral mass as

$$Q(\gamma) = \sum_{w\in\mathcal{W}_\gamma} q(w),$$

where $\mathcal{W}_\gamma = \{\, w \mid \operatorname{ord}\chi_w = \gamma \,\}$. Here, Q_γ defines a probability distribution over $\mathcal{R} = \{\, \operatorname{ord}\gamma_w \mid \gamma_w \text{ is non-trivial character} \,\}$ and represents the quotient leakage spectrum. Then, the spectral alignment probability is separated into *what leaks* (*i.e.*, $Q(\gamma)$) and *how much it leaks* (*i.e.*, $2^{-H_{d+1}(q_\gamma)}$) in a quotient-level manner, stated by Theorem 8.

Theorem 8 (Exact quotient alignment decomposition). *Let $q_\gamma(w) = q(w)/Q(\gamma)$ be the conditional probability of normalized Fourier leakage spectrum given $\gamma \in \mathcal{R}$ for γ with $Q(\gamma) > 0$. The spectral alignment probability is given by*

$$\kappa_d = \sum_{\gamma\mid m} Q(\gamma)^{d+1} \sum_{w\in\mathcal{W}_\gamma} q_\gamma(w)^{d+1} = \sum_{\gamma\mid m} 2^{-dH_{d+1}(q_\gamma)} Q(\gamma)^{d+1}.$$

Proof. Recall that $\kappa_d = \Pr[W_0 = W_1 = \cdots = W_d]$. By marginalizing κ_d in terms of $W_0, W_1, \ldots, W_d$ and accumulating the conditional probabilities in quotient-wise events that $\operatorname{ord}\chi_{W_i} = \gamma$ holds for

all i, we conclude

$$\kappa_d = \sum_{\gamma|m} \Pr[\operatorname{ord}\chi_W = \gamma]^{d+1} \Pr[W_0 = W_1 = \cdots = W_d \mid \forall i, \operatorname{ord}\chi_{W_i} = \gamma] = \sum_{\gamma|m} Q(\gamma)^{d+1} \sum_{w\in\mathcal{W}_\gamma} q_\gamma(w)^{d+1},$$

as required. □

Here, $Q(\gamma)$ and $2^{-dH_{d+1}(q_\gamma)}$ represent the quotient-order probability mass and within-quotient alignment, respectively. Theorem 8 states a separation of κ_d into these two components: $Q(\gamma)$ and $\kappa_{d,\gamma} = \sum_{w\in\mathcal{W}_\gamma} q_\gamma(w)^{d+1}$.

Spectral multiplicity. The number of characters with γ is $\phi(\gamma)$, where ϕ is Euler's totient function. This implies that κ_d can be characterized using the Euler totient. Theorem 9 states that $\kappa_{d,\gamma} = \sum_{w\in\mathcal{W}_\gamma} q_\gamma(w)^{d+1}$ represents the *primitive-order multiplicity* and characterizes the spectral alignment probability through the Euler totient.

Theorem 9 (Euler multiplicity bound). *It holds that*

$$\kappa_{d,\gamma} = 2^{-dH_{d+1}(q_\gamma)} \geq \phi(\gamma)^{-d}, \qquad \kappa_d \geq \sum_\gamma \frac{Q(\gamma)^{d+1}}{\phi(\gamma)^d} \tag{5}$$

where the equality holds if the spectra are uniform among each primitive order class.

Proof. Since the Rényi entropy takes the maximum for uniform distribution, it suffices to prove the equality of Equation (5) assuming that q_γ is uniform (*i.e.*, the leakage is uniform among the primitive order class). Because the number of characters with order γ equals $\phi(\gamma)$, with the uniformity assumption, $q_\gamma(w) = \phi(\gamma)^{-1}$ holds for any $w \in \mathcal{W}_\gamma$. According to the definition of $\kappa_{d,\gamma}$ and $|\mathcal{W}_\gamma| = \phi(\gamma)$, we conclude

$$\kappa_{d,\gamma} = \phi(\gamma)\left(\frac{1}{\phi(\gamma)^{d+1}}\right) = \phi(\gamma)^{-d},$$

as required. The lower bound of κ_d directly follows this. □

This bound explains the separation in counterexample of LSB and MSB leakages with $m = 2^l$ in Section 2.2. The LSB and MSB leakages correspond to the order-2 character and characters with odd w, respectively. Even if the leakage mass is identical, the LSB leakage character with $\gamma = 2$ cannot be de-aligned because $\phi(2) = 1$, while an alignment suppression of $\phi(2^l)^{-d} = 2^{-(l-1)d}$ at most is possible for the MSB leakage characters with $\gamma = 2^l$. In contrast, if m is a prime, all Fourier characters have no non-trivial lower-order character but the order is m for all characters. As $\phi(\gamma) = m - 1$, the alignment suppression may be maximized in terms of Euler totient. Prime-order cyclic groups allow potentially stronger spectral de-alignment because all non-trivial characters belong to a single order class of size $m - 1$.

Primitive–Quotient leakage conversion. $Q(\gamma)$ exactly represents the aggregate leakage power of order-γ character, which may be artificial and hard to observe directly. Meanwhile, operational leakage is frequently given through a quotient, which is cumulative over all subquotients of $\mathbb{Z}_\gamma$, as an upper bound of global leakage is characterized by $I(Y^{(\gamma)}; \boldsymbol{L})$ in CRYPTO 2026 [UIM+26], where

$Y^{(w)} = S + \ker \chi_w$ is the quotient variable. In our exact decomposition framework, the quotient leakage is naturally defined as $I_{\chi^2}(Y^{(w)}; \boldsymbol{L})$. However, this quantity does not directly characterize the leakage concentration. For example of LSB leakage in $\mathbb{Z}_{2^l}$ in Section 2.2, the leakage is clearly concentrated to order-2 character as $I_{\chi^2}(Y^{(2)}; \boldsymbol{L}) = 1$, while order-$2^j$ characters for all $1 < j < l$ also have $I_{\chi^2}(Y^{(2^j)}; \boldsymbol{L}) = 1$. In contrast, $Q(2) = 1$ and $Q(2^j) = 0$ holds for this example, which means that $Q(\gamma)$ appropriately measures the leakage concentration. Since $I_{\chi^2}(Y^{(\gamma)}; \boldsymbol{L})$ is cumulative, we need to evaluate $I_{\chi^2}(Y^{(\gamma)}; \boldsymbol{L})$ for all non-trivial divisors of γ to exactly identify which algebraic resolution carries the leakage, as stated by Theorem 10.

Theorem 10 (Primitive–Quotient composition/decomposition). *Consider a cyclic group. Let $R(\gamma) = I_{\chi^2}(Y^{(w)}; \boldsymbol{L})/I_{\chi^2}(S; \boldsymbol{L})$ be the normalized quotient leakage with $\gamma = \operatorname{ord} \chi_w$. Let μ_{M} denote the Möbius function. It holds*

$$R(\gamma) = \sum_{g|\gamma} Q(g), \qquad Q(\gamma) = \sum_{g|\gamma} \mu_{\mathrm{M}}\left(\frac{\gamma}{g}\right) R(g).$$

Proof. By definition, $I_{\chi^2}(Y^{(w)}; \boldsymbol{L}) = \sum_{g|\gamma} \sum_{u \in \mathcal{W}_g} \lambda(u)$ holds owing to the orthogonality of χ^2-MI. By normalizing this, we have the expression of $R(\gamma)$. The inverse expression is derived from it according to the Möbius inversion formula. □

By Theorem 10, individual Fourier coefficients need not be inspected directly. The entire primitive leakage spectrum can be recovered from information leaked about natural quotient variables.

Role of each quantity and spectrum. In summary, q, Q, and R represent the normalized Fourier-mode-wise, order-wise, and quotient-wise leakages as probability distributions or cumulative mass, respectively, and they are fine-to-coarse in this order. Our quotient spectral theory composes/decomposes these probability distributions with regard to the spectral alignment probability. The spectral alignment is separated into the quotient-level alignment and within-quotient primitive alignment as shown in Theorems 8 and 10, which can be characterized by Euler totient. What each quantity and spectrum describes is summarized as follows: Local MI $I(S; \boldsymbol{L})$ tells us how much information leaks. Leakage sharpness α tells us how concentrated that information is. The quotient spectrum R tells us what algebraic information leaks. Spectral alignment κ_d tells us whether the same algebraic information survives across shares.

5.2 Non-cyclic finite abelian groups

In Section 5, we focused on cyclic group for simple and intuitive explanation. Our quotient spectral theory may be extended to non-cyclic finite abelian group. The number-theoretic Möbius inversion may be replaced with quotient-poset inversion [Rot64] for the generalization. Let $\gamma_{w_j,j} = \gcd(w_j, m_j)$. The order of character $w = (w_1, w_2, \dots, w_k)$ is given by

$$\gamma_w = \operatorname{lcm}\left(\frac{m_1}{\gamma_{w_1,1}}, \frac{m_2}{\gamma_{w_2,2}}, \dots, \frac{m_k}{\gamma_{w_k,k}}\right),$$

as discussed in Section 2.3. Here, $\mathbb{G}/\ker \chi_w$ is cyclic. In addition, we need the number of exact order-γ characters, which is a generalization of the Euler totient, as stated by Lemma 2

Lemma 2 (Number of order-γ characters). *Let $\mathbb{G} = \bigoplus_{j=1}^{k} \mathbb{Z}_{m_j}$. Let $\phi_{\mathbb{G}}(\gamma)$ be the number of order-γ characters in $\mathbb{G}$. It holds*

$$\phi_{\mathbb{G}}(\gamma) = \sum_{g|\gamma} \mu_{\mathrm{M}}\left(\frac{\gamma}{g}\right) \prod_{j=1}^{k} \gcd(m_j, g). \tag{6}$$

Proof. Let $\zeta(\gamma)$ be the number of elements of $\{\, w \mid \gamma \text{ is divisible by } \operatorname{ord} \chi_w \,\}$. According to [BS03, Lemma 1], it holds

$$\zeta(\gamma) = \prod_{j=1}^{k} \zeta_j(\gamma), \tag{7}$$

where $\zeta_j(\gamma) = \gcd(m_j, \gamma)$ is the number of elements that divides γ. Meanwhile, by definition, it holds

$$\zeta(\gamma) = \sum_{g|\gamma} \phi_{\mathbb{G}}(g). \tag{8}$$

Combining Equations (7) and (8) with the Möbius inversion formula, we conclude Equation (6). □

Note that $\phi_{\mathbb{G}}(\gamma) = \phi(\gamma)$ holds when $\mathbb{G}$ is cyclic. Hence, the lower bound on κ_d for non-cyclic group by the quotient spectral theory is represented by

$$\kappa_d \geq \sum_{\gamma} \frac{Q(\gamma)^{d+1}}{\phi_{\mathbb{G}}(\gamma)^d}.$$

5.3 Automorphisms

The problem is group-theoretically generalized before presenting the exact decomposition theorem. Let φ_i be an invertible automorphism over $\mathbb{G}$ for the i-th share. The secret variable Z is decomposed as $Z = \sum_i \varphi_i(S_i)$, where S_i is sampled from uniform distribution over $\mathbb{G}$. Due to definition, φ_i must be a bijection. This is a group-theoretic generalization, which covers additive decomposition discussed in the LSS and masking communities (*e.g.*, automorphism-weighted sharing and inner-product masking). The exact decomposition admits a finer characterization through the quotient structure of finite abelian groups with this extension, developed in this section. As a main result of this section, we show that there exist non-trivial (full-support) leakage channels and automorphisms which completely cancel the global MI as $I(Z; \boldsymbol{X}) = 0$ with *finite* shares, while $I(Z; \boldsymbol{X}) > 0$ always holds for all finite d for non-automorphism-weighted sharing with the i.i.d leakage.

Our quotient spectral theory through Euler totient and Möbius inversion is particularly useful to explain the leakage from the automorphism-weighted secret sharing. Consider automorphisms in the Fourier domain as $\varphi_i^* \in \mathrm{Aut}(\mathbb{G}) : w \to A_i^* w$, where A_i^* is a representation of φ_i in the Fourier domain. Since an automorphism is an order-preserving permutation over $\mathbb{G}$, it always holds $\operatorname{ord} \chi_w = \operatorname{ord} \chi_{A_i^* w}$ for any w, while the primitive Fourier modes are permuted. This means that each automorphism preserves the quotient-order mass $Q(\gamma)$ of its corresponding local spectrum, while it de-aligns the within-quotient primitives.

Recall that the spectral alignment probability is $\kappa_d = \sum_{w\neq 0}\prod_{i=0}^{d} q_i(w)$, when the automorphism-weights are absent (*i.e.*, φ_i is an identity for all i). The quotient spectral theory demonstrated its separation into quotient-level alignment and within-quotient primitive alignment. In an automorphism-weighted case, the spectral alignment probability is given by

$$\kappa_d(\varphi_0,\varphi_1,\dots,\varphi_d) = \sum_{w\neq 0}\prod_{i=0}^{d} q_i(A_i^* w) = \Pr[(A_0^*)^{-1}W_0 = (A_1^*)^{-1}W_1 = \cdots = (A_d^*)^{-1}W_d]. \tag{9}$$

Equation (9) represents the automorphism-induced spectral de-alignment. If the automorphism-weights are absent, the spectral alignment probability is dominated by

$$\frac{\kappa_d}{m-1} \leq \max_{w\neq 0}\prod_{i=0}^{d} q_i(w) = \max_{w} q(w)^{d+1} \leq \kappa_d,$$

under the leakage-i.i.d assumption, which implies that q_i takes the maximum at an identical w for all i. In contrast, owing to the automorphism-induced de-alignment, we cannot interchange the max and product and the spectral alignment probability is given by

$$\frac{\kappa_d(\varphi_0,\varphi_1,\dots,\varphi_d)}{m-1} \leq \max_{w\neq 0}\prod_{i=0}^{d} q_i(A_i^* w) \leq \kappa_d(\varphi_0,\varphi_1,\dots,\varphi_d).$$

If not all automorphisms are identical, $q_i(A_i^* w)$ can take the maximum at different w for each i. In this case, the spectral alignment probability can be smaller than that without automorphism-weights, making the resulting global MI exponent larger as

$$-\frac{1}{d}\log\kappa_d(\varphi_0,\varphi_1,\dots,\varphi_d) = -\frac{1}{d}\log\max_{w\neq 0}\prod_{i=0}^{d} q_i(A_i^* w) + O\left(\frac{\log(m-1)}{d}\right),$$

provided that $\kappa_d > 0$. However, even if not all automorphisms are identical, the spectrum may remain aligned; in such a case, the automorphism does not reduce the global MI. In the following, we discuss its relation to the quotient de-alignment and character orders.

Primitive order and multiplicity. Consider a cyclic group $\mathbb{G} = \mathbb{Z}_m$ with m having prime factors $\{\, p > 1 \mid m \text{ is divisible by } p \,\}$. Let p be a prime factor of m. In $\mathbb{G}$, the order-p characters correspond to the frequency of $w_\ell = m\ell/p$ for $\ell = 1, 2, \dots, p-1$. For any $s_i \in \mathbb{Z}_m$ and automorphism φ_i, an element in $a \in \mathbb{Z}_m^\times$ exists such that

$$\varphi_i : s_i \mapsto a s_i \bmod m, \qquad \varphi_i^* : w \mapsto a w \bmod m, \tag{10}$$

in the spatial and Fourier domains, respectively. Equation (10) is followed by

$$aw \bmod m = (a\ell \bmod p)\frac{m}{p}.$$

Hence, any automorphism is order-preserving in Fourier domain; it maps order-γ characters to the same order character. This also holds for non-cyclic finite abelian groups. For order-γ characters, the number of available spectral modes is $\phi(\gamma)$ (for non-cyclic finite abelian groups, $\phi_{\mathbb{G}}(\gamma)$).

Global leakage cancellation by de-alignment. Recall that $\kappa_d \geq \sum_\gamma Q(\gamma)^{d+1}/\phi(\gamma)^d > 0$ for non-automorphism-weighted sharing, implying that $I(Z;\boldsymbol{X}) > 0$ holds for any finite d. In contrast, because automorphism moves the leakage concentration from characters to others with same order, the automorphism-weighted sharing may achieve $I(Z;\boldsymbol{X}) = 0$ with finitely many shares for some non-trivial channels. We first introduce Proposition 2, which shows a realizability of such channels with an example.

Proposition 2 (Example of finite-share perfect leakage cancellation and its realizability). *Let $\mathbb{G} = \mathbb{Z}_5$. Let $\varphi_0 = \mathrm{id}$ and $\varphi_1(S_1) = 2S_1$. Consider an automorphism-weighted sharing $Z = \varphi_0(S_0) + \varphi_1(S_1) = S_0 + 2S_1$. Full-support leakage channels exist such that $I(Z;\boldsymbol{X}) = 0$.*

Proof. Cyclic group $\mathbb{Z}_5$ has order-5 characters at $w = 1$, 2, 3, and 4. As q is conjugate symmetry, $q(w) = q(-w)$ always holds. Consider a normalized spectrum such that

$$q(1) = q(4) = \frac{1}{2}, \qquad q(2) = q(3) = 0,$$

where $\operatorname{supp} q = \{1, 4\}$. When applying the automorphism $\varphi_1(S_1) = 2S_1$, the transformed support is $\operatorname{supp}(q \circ \varphi_1^*) = \{2, 3\}$. Hence, $\operatorname{supp}(q \circ \varphi_0^*) \cap \operatorname{supp}(q \circ \varphi_1^*) = \varnothing$ holds—the supports are disjoint. This implies that

$$\kappa_1(\varphi_0, \varphi_1) = \sum_{w \neq 0} q(w)q(2w) = 0,$$

followed by $I(Z;\boldsymbol{X}) = 0$. A family of channels to realize q can be constructed in the same manner as Theorem 3 using a binary leakage $\boldsymbol{l} \in \{-1, 1\}$ as

$$p_{\boldsymbol{L}|S}(\boldsymbol{l} \mid s) = \frac{1 + \boldsymbol{l}\epsilon \cos(2\pi s/5)}{2}, \tag{11}$$

for any $0 < \epsilon < 1$. This completes the proof. □

Proposition 2 clearly demonstrates an example of spectral de-alignment and shows that Euler multiplicity bound does not hold for automorphism-weighted sharing. Character multiplicity $|\mathcal{W}_\gamma| = \phi_G(\gamma)$ determines the number of available modes for de-alignment and the spectral support $\operatorname{supp} q_\gamma \subset \mathcal{W}_\gamma$ represents how the leakage is distributed to the modes. Then, the automorphism orbit determines an actual alignment. Automorphism maps characters to others with same order, while it cannot necessarily map multiple characters to others simultaneously in an arbitrary manner. Theorem 11 characterizes the essential difference of automorphism-weighted sharing from non-weighted one. For automorphism-weighted sharing, the spectral support plays an essential role for characterizing global leakage (perfect cancellation) in addition to the maximum spectral power (exponential rate).

Theorem 11 (Spectral de-alignment and finite-share perfect leakage cancellation). *Let $\mathbb{G}$ be a finite abelian group. Let τ be a non-trivial local leakage channel with Fourier power spectrum λ. For automorphism $\varphi_0, \varphi_1, \ldots, \varphi_d$, $I(Z;\boldsymbol{X}) = 0$ holds if and only if*

$$\bigcap_{i=0}^{d} (A_i^*)^{-1}(\operatorname{supp} \lambda) = \varnothing. \tag{12}$$

Furthermore, there exist finite abelian groups and non-trivial leakage channels for which this condition is achievable with finitely many shares, whereas it is impossible for non-weighted sharing under i.i.d leakage assumption.

Proof. Equation (12) is necessary and sufficient condition for $\kappa_d = 0$ by definition. Its realizability can be shown in the same manner as Proposition 2 using the binary leakage and Equation (11). □

Low-order obstruction of de-alignment. For cyclic group $\mathbb{Z}_m$ with even order, only one order-2 character exists. Since $\phi(2) = 1$, the order-2 character is the fixed point for any automorphism. Thus, the LSB leakage, which exactly corresponds to the order-2 character, is canonical and can never be de-aligned by any automorphism in $\mathbb{Z}_m$ with even m. This explains why and how the LSB leakage is dangerous by structure, as claimed in [FMM+24]. Note that this does not generally apply to non-cyclic groups, which may have multiple order-2 characters.

6 Canonical Leakage Models

This section evaluates the exponential rate of global MI. The goal of this section is to concretely demonstrate our exact decomposition and separation theorems. Hence, in this section, we utilize some canonical leakages over $\mathbb{Z}_2^l$, $\mathbb{Z}_{2^l}$, and $\mathbb{Z}_p$, where p is an l-bit prime, as typical and simple examples.

AI disclosure. Although the core idea, methodology, and how to evaluate the rates are completely driven by the authors, some concrete calculations in this section were carried out by ChatGPT. The authors checked and verified the correctness of ChatGPT outputs.

6.1 Deterministic/Noise-free leakages

In this subsection, we evaluate the LSB, MSB, and Hamming weight (HW) leakages. Table 1 summarizes the global MI exponents for each condition, which are defined by $-\lim_{d\to\infty} \frac{1}{d} \log I(Z; \boldsymbol{X})$. Over cyclic groups, the coefficient $\lambda_*^{(m)}$ is defined as

$$\lambda_* = \max_{w\neq 0} \lambda(q) = \max_{w\neq 0} \frac{1}{m} \sum_{h=0}^{l} n_h^{-1} \left| \sum_{\mathrm{HW}(s)=h} \exp\left(\frac{2\pi \mathrm{i} w s}{m}\right) \right|^2 ,$$

where $n_h = |\{\mathrm{HW}(s) = h\}|$ for given h. When $m = 2^l$, it is given as $n_h = \binom{l}{h}$. The expressions of $\lambda_*^{(m)}$ are exact finite characterizations, while closed-form evaluations of these maxima for general l remain unproved. Yet, the results in Table 1 are consistent with previous literature [BHMY26].

For $\mathbb{Z}_2^l$, we suppose that the element is given in an l-bit string; hence, the LSB and MSB leakages mean the value at the first and l-th invariant factor, respectively. In this case, MSB equals LSB at the l-th invariant factor, meaning that $H_\infty(q) = 0$ and the exponent by spectral alignment is exactly zero. In addition, the exponent is zero also for HW leakage. The parity of the HW equals the XOR of all coordinate bits, which is a nontrivial linear character over $\mathbb{Z}_2^l$, as in [IUH22, UIM+26]. For $\mathbb{G} = \mathbb{Z}_2^l$, the global leakage exponent vanishes if and only if the local leakage perfectly determines a non-trivial linear parity of the share.

In contrast, for cyclic groups, we find the exponential separation. For cyclic groups, LSB and MSB leakages have identical local χ^2-MI whenever both are non-constant binary deterministic functions, regardless of whether the group order is even or odd. These examples of canonical leakage clearly demonstrate that the spectral information is indispensable as the global MI exponent contains the Fourier character information.

Table 1: Exponential rates of global Shannon mutual information under noiseless local leakage. The group and bit length l are fixed as $d \to \infty$. All expressions are exact, except that closed-form evaluations of the maximum spectral power for HW leakage over $\mathbb{Z}_{2^l}$ and $\mathbb{Z}_p$ remain undeveloped.

Group	Leakage	$I_{\chi^2}(S;\boldsymbol{L})$	Global MI exponent
$\mathbb{Z}_2^l$	LSB	1	0
	MSB	1	0
	HW	l	0
$\mathbb{Z}_{2^l}$	LSB	1	0
	MSB	1	$\log \frac{2^{2l}\sin^2(\pi/2^l)}{4}$
	HW†	l	$-\log \lambda_*^{(2^l)}$
$\mathbb{Z}_p$	LSB	1	$\log\left((p^2-1)\sin^2\frac{\pi}{2p}\right)$
	MSB	1	$\log \frac{2^{l-1}(p-2^{l-1})\sin^2(\pi/p)}{\sin^2(\pi 2^{l-1}/p)}$
	HW†	$\left\lvert \mathrm{Im}(\mathrm{HW}\lvert_{\mathbb{Z}_p})\right\rvert - 1$	$-\log \lambda_*^{(p)}$

† implicitly exact, yet closed form is not derived.

6.2 HW leakage with Gaussian noise

In this subsection, we analyze a stochastic leakage over $\mathbb{Z}_2^l$, $\mathbb{Z}_{2^l}$, and $\mathbb{Z}_p$. We utilize an HW leakage with additive Gaussian noise, defined as

$$\boldsymbol{L} = \mathrm{HW}(S) + N, \qquad N \sim \mathcal{N}(0, \sigma^2).$$

This canonical leakage model has been used and analyzed for a long time in the side-channel masking community [MOP07]. For ease of analysis and presentation, we here suppose the low signal-to-noise ratio (SNR). For given $\mathbb{G}$, let

$$v = \mathrm{Var}_S\, \mathrm{HW}(S), \qquad \rho = \frac{v}{\sigma^2},$$

be the signal component and SNR, respectively.

Table 2 summarizes the global MI exponent under these conditions. The values follow from the low-SNR first-order approximation as

$$I(S;\boldsymbol{L}) = \frac{\rho}{2\ln 2} + O(\rho^2), \qquad I_{\chi^2}(S;\boldsymbol{L}) = \rho + O(\rho^2), \tag{13}$$

according to the I-MMSE (information-theoretic minimum mean square error) formula [GSV05] in the low-SNR regime (corresponding to the Shannon–Hartley upper bound). Hence, in Table 2, all local MIs are given by Equation (13).

In this situation, the global MI exponent is given by

$$-\lim_{d\to\infty} \frac{1}{d} \log I(Z;\boldsymbol{X}) = -\log(\rho\omega_m) + O(\rho),$$

Table 2: Asymptotic security exponents under Gaussian Hamming-weight leakage in the low-SNR regime.

Group	$I_{\chi^2}(S;\boldsymbol{L})$	Global MI exponent
$\mathbb{Z}_2^l$ and $\mathbb{Z}_{2^l}$	$\rho + O(\rho^2)$	$-\log\frac{\rho}{l} + O(\rho)$
$\mathbb{Z}_p$	$\rho + O(\rho^2)$	$-\log(\rho\omega_p) + O(\rho)$

where the coefficient ω_m is defined as

$$\omega_m = \frac{1}{v}\max_{w\neq 0}\left|\widehat{\mathrm{HW}}(x)\right|^2.$$

Here, ω_m is exactly $\omega_m = 1/l$ when $\mathbb{G} = \mathbb{Z}_2^l$ or $\mathbb{Z}_{2^l}$, because the DFT can be exactly calculated over them. The coefficient ω_m represents the ratio of Fourier characters with strongest leakages among entire Fourier energy in the HW leakage model at the low-SNR limit.

In the low-SNR regime, the dominant spectral power is determined by the Fourier coefficients of the HW function, leading to the same normalized maximum spectral power for these two groups. In contrast, ω_m has a non-trivial value for $\mathbb{Z}_p$. This is because each bit is neither independent nor uniformly distributed in the binary representation of $\mathbb{Z}_p$ element, implying that the Fourier support of each bit function is *not* disjoint to 2-adic Fourier characters and Fourier coefficients for different bits interfere with one another. This property is completely different from $\mathbb{Z}_2^l$ and $\mathbb{Z}_{2^l}$, where each bit function is different Fourier character and the Fourier support of each bit function is confined to characters with a distinct 2-adic valuation, respectively. The interference depends on p; hence, ω_p explicitly appears in the exponent. The enhanced leakage resilience over prime-order groups [DFS16] is explained by a smaller maximum normalized spectral power such that $\omega_m < 1/l$, when observed, rather than by the prime-order group structure alone.

7 Secret sharing and masking

7.1 Additive secret sharing

Deterministic leakage functions have been studied in the LSS community. Because the MI $I(S;\boldsymbol{L})$, χ^2-MI $I_{\chi^2}(S;\boldsymbol{L})$, and λ are easily computed for deterministic leakage functions by a tailored calculus, our exact bound can evaluate the leakage resilience of LSS against them as a generalization of existing bounds, as in Section 6.

In [BHMY26], Biswas, Hwang, Maji, and Ye also presented tight bounds for the HW leakage *via* a tailored Krawtchouk polynomial analysis. In one sense, our theory unifies the existing approaches in LSS and side-channel masking, as it exactly characterizes the global leakage by the spectral alignment factor (which can be analyzed by, for example, Krawtchouk polynomial) and leakage magnitudes (*e.g.*, MI and TV). Both factors can be discussed solely by the posterior probability; hence, their combination enables to handle both deterministic and stochastic leakages. Notably, our global MI rate is consistent with [BHMY26], while a direct comparison would be difficult as the focused metrics are different.

7.2 Side-channel masking

In the side-channel community, the upper bound of global MI or total variation distance by local leakage magnitudes have been presented. Using our exact decomposition theorem, we can also derive an upper bound of $I(Z;\boldsymbol{X})$ as follows:

Corollary 5 (Upper bound on Shannon MI). *Global leakage $I(Z;\boldsymbol{X})$ is bounded by*

$$I(Z;\boldsymbol{X}) \le \log\left(1 + I_{\chi^2}(Z;\boldsymbol{X})\right) = \log\left(1 + \kappa_d \prod_{i=0}^{d} I_{\chi^2}(S_i;\boldsymbol{L}_i)\right)$$
$$= \log\left(1 + \kappa_d \prod_{i=0}^{d} \alpha_i I(S_i;\boldsymbol{L}_i)\right),$$

Proof. The equality is obvious from Theorem 1. The inequality is derived by [UIM+26, Lemma 1] through Jensen and Parseval.

Corollary 5 strengthens conventional magnitude-only bounds by explicitly incorporating spectral alignment for general finite abelian group [DFS15, IUH22, MRS22, BCG+23, UIM+26], where the group-order loss $m-1$, which corresponds to the error term $O((m-1)(\rho^*)^{d+1})$ in Theorem 2, is eliminated. Existing bounds were magnitude-only, including local MI, TV (*i.e.*, L_1-MI) [DFS15], and bias (*i.e.*, L_2-MI) [PR13], which can be exponentially loose except for LSB(-XOR) leakages as shown by Section 4.4. Note that the upper bounds rely on major inequality such as Chernoff, Jensen, and Pinsker; their approach would be difficult to derive lower bounds and asymptotic expansion, which are enabled by our exact local-to-global decomposition without these inequalities.

Recently, in CRYPTO 2026 [BBM26], Béguinot, Brian, and Masure developed another bound claiming to remove/mitigate the *field-size loss* in the masking security bounds. Their methodology focuses on the worst-case leakage strength, which would be the same scope of our leakage sharpness. They showed that a separation of average local leakage (*i.e.*, local total variation distance, corresponding to the local MI) and the worst-case leakage concentration. This approach and its follow-ups would be going toward removing the field-size slack from upper bounds. Thus, our spectral theory would be orthogonal to their approaches to derive stronger upper bounds by magnitudes.

This applies to IP masking, which corresponds to the automorphism-weighted sharing. Our quotient spectral theory demonstrates that the global MI is suppressed when the number of low-order characters is smaller. This suggests that prime-order cyclic groups can offer improved leakage resilience in certain leakage models, consistent with observations in previous studies [DFS16, MMMS23].

8 Limitations and open problems

This paper showed an exact characterization of local-to-global leakage by the leakage spectrum and magnitude. However, the leakage spectrum (even quotient spectral profile) is not always operational. Hence, bounding and characterizing bounds by intermediate spectral–magnitude quantities, which are more operational than (quotient) spectral leakage, would be an important direction for future work for further analysis and understanding of local-to-global leakage. Our theory may be orthogonal to studies on strengthening magnitude-based bounds like Béguinot, Brian, and Masure [BBM26].

Another important open problem is application of our theory to moderately-stochastic leakages beyond the canonical leakages (*i.e.*, deterministic and low-SNR regime) analyzed in this paper. By this, we may be able to further characterize the local-to-global leakage by noise and SNR from an information-theoretic perspective.

AI methodology

We used ChatGPT to assist with drafting, refining, and verifying the proofs. The tool materially affected Section 6. More details can be found in Section 6. The authors verified the correctness and originality of all content, including references.

A Equivalence between KL and chi-square divergences

Lemma 3 (Second-order approximation of D_{KL} by χ^2). *Let p and q be discrete probability distributions over $\mathcal{X}$ such that* $\mathrm{supp}(p) \subset \mathrm{supp}(q)$. *$D_{\mathrm{KL}}(p \parallel q)$ is approximated as*

$$D_{\mathrm{KL}}(p \parallel q) = \begin{cases} \dfrac{\chi^2(p \parallel q)}{2 \ln 2} + O(\|r\|_\infty^3) & \text{w.r.t. additive error,} \\ \dfrac{\chi^2(p \parallel q)}{2 \ln 2}(1 + O(\|r\|_\infty)) & \text{w.r.t. relative error,} \end{cases}$$

as $p \to q$, where $r(x) = p(x)/q(x) - 1$ for any x and $\|\cdot\|_\infty$ is the L_∞ norm. The convergence radius is $\|r\|_\infty < 1$

Proof. Let $r(x) = p(x)/q(x) - 1$. Recall that

$$\begin{aligned} D_{\mathrm{KL}}(p \parallel q) &= \sum_x p(x) \log \left(\frac{p(x)}{q(x)} \right) \\ &= \sum_x q(x)(1 + r(x)) \log (1 + r(x)) \,. \end{aligned}$$

Using the third-order Taylor expansion of $(1+x)\log(1+x)$ at $x = 0$, we have $(1+x)\log(1+x) = (\ln 2)^{-1}(x + x^2/2 - x^3/6 + O(x^4))$. Hence, D_{KL} is expressed by

$$D_{\mathrm{KL}}(p \parallel q) = \frac{1}{\ln 2} \sum_x q(x) \left(r(x) + \frac{r(x)^2}{2} - \frac{r(x)}{6} + O(r(x)^4) \right).$$

Here, $\sum_x q(x) r(x) = 0$ and $\sum_x q(x) r(x)^2 = \chi^2(p \parallel q)$ due to its definition. Thus, the second-order approximation is expressed as

$$D_{\mathrm{KL}}(p \parallel q) \approx \frac{\chi^2(p \parallel q)}{2 \ln 2}.$$

The additive error term is given by

$$\frac{\sum_x q(x) r(x)^3}{6} + \sum_x O(q(x) r(x)^4) = O(\|r\|_\infty^3),$$

because the third-order term is bounded as $\sum_x q(x)\,|r(x)|^3 \le \|r\|_\infty^3$. In addition, the relative error is given by $1 + O(\|r\|_\infty)$, since the third-order term is also bounded as

$$\frac{\sum_x q(x)\,|r(x)|^3}{6} \le \|r\|_\infty \chi^2(p \parallel q).$$

This completes the proof.

B Proofs

B.1 Proof of Theorem 1

Proof. By definition, the global χ^2-MI is given by

$$I_{\chi^2}(Z; \boldsymbol{X}) = \sum_{w \neq 0} \mathbb{E}_{\boldsymbol{X}} \left[\left| \hat{p}_{Z|\boldsymbol{X}}(w \mid \boldsymbol{X}) \right|^2 \right]. \tag{14}$$

Since Z is additive composition of $S_0, S_1, \ldots, S_d$, it holds that

$$\begin{aligned}
\hat{p}_{Z|\boldsymbol{X}}(w \mid \boldsymbol{X}) &= \mathbb{E}_Z \left[\chi_w(Z) \mid \boldsymbol{X} \right] \\
&= \mathbb{E}_Z \left[\chi_w(S_0 + S_1 + \cdots + S_d) \mid \boldsymbol{X} \right] \\
&= \mathbb{E}_Z \left[\prod_{i=0}^{d} \chi_w(S_i) \,\middle|\, \boldsymbol{X} \right] \\
&= \prod_{i=0}^{d} \mathbb{E}_S \left[\chi_w(S_i) \mid \boldsymbol{L}_i \right] \\
&= \prod_{i=0}^{d} \hat{p}_{S_i|\boldsymbol{L}_i}(w \mid \boldsymbol{L}_i),
\end{aligned} \tag{15}$$

for any $w \neq 0$, as $S_0, S_1, \ldots, S_d$ and $\boldsymbol{L}_0, \boldsymbol{L}_1, \ldots, \boldsymbol{L}_d$ are independent except for $\tau_i(S_i) \to \boldsymbol{L}_i$ owing to the independence assumption. By combining Equations (14) and (15), we have

$$I_{\chi^2}(Z; \boldsymbol{X}) = \sum_{w \neq 0} \prod_{i=0}^{d} \lambda_i(w),$$

which is followed by

$$\begin{aligned}
\sum_{w \neq 0} \prod_{i=0}^{d} \lambda_i(w) &= \frac{\prod_{i=0}^{d} \sum_{w \neq 0} \lambda_i(w)}{\prod_{i=0}^{d} \sum_{w \neq 0} \lambda_i(w)} \sum_{w \neq 0} \prod_{i=0}^{d} \lambda_i(w) \\
&= \frac{\sum_{w \neq 0} \prod_{i=0}^{d} \lambda_i(w)}{\prod_{i=0}^{d} \sum_{w \neq 0} \lambda_i(w)} \prod_{i=0}^{d} \sum_{w \neq 0} \lambda_i(w) \\
&= \left(\sum_{w \neq 0} \prod_{i=0}^{d} \frac{\lambda_i(w)}{\sum_{w' \neq 0} \lambda_i(w')} \right) \left(\prod_{i=0}^{d} \sum_{w \neq 0} \lambda_i(w) \right) \\
&= \left(\sum_{w \neq 0} \prod_{i=0}^{d} q_i(w) \right) \left(\prod_{i=0}^{d} I_{\chi^2}(S_i; \boldsymbol{L}_i) \right).
\end{aligned}$$

Here, $\prod_{i=0}^{d} q_i(w)$ represents the probability of an event that w is selected from the probability distribution q_i for all i. Hence, $\sum_{w\neq 0}\prod_{i=0}^{d} q_i(w)$ represents the sum of probabilities for all $w \neq 0$ that all W_i collide at $W_i = w$, which means that $\sum_{w\neq 0}\prod_{i=0}^{d} q_i(w) = \Pr[W_0 = W_1 = \cdots = W_d] = \kappa_d$ holds, where $W_i \sim q_i$.

When local leakages are i.i.d, $q_i = q$ holds for any i. Hence, κ_d is equivalent to the collision probability of $d+1$ elements sampled from q, which is represented by Rényi-$(d+1)$ entropy as

$$\Pr[W_0 = W_1 = \cdots = W_d] = 2^{-dH_{d+1}(q)} \quad (W_i \sim q \text{ for all } i).$$

This completes the proof. □

B.2 Proof of Lemma 1

Proof. For given $\boldsymbol{l}$, denote KL and χ^2 divergences in base 2 as

$$D_{\boldsymbol{l}} = D_{\mathrm{KL}}(p_{S|\boldsymbol{L}}(\cdot \mid \boldsymbol{l}) \parallel p_S) = \frac{1}{\ln 2}\sum_s p_S(s) f(r(s,\boldsymbol{l})),$$
$$\chi^2_{\boldsymbol{l}} = \chi^2(p_{S|\boldsymbol{L}}(\cdot \mid \boldsymbol{l}) \parallel p_S) = \sum_s p_S(s)(r(s,\boldsymbol{l})-1)^2.$$

By definition, for given $\boldsymbol{l}$, it holds

$$\chi^2_{\boldsymbol{l}} = \sum_s p_S(s) f(r) g(r),$$

which is followed by

$$\begin{aligned}\chi^2_{\boldsymbol{l}} &= \ln(2) D_{\boldsymbol{l}} \frac{\sum_s p_S(s) f(r) g(r)}{\sum_s p_S(s) f(r)} \\ &= \ln(2) D_{\boldsymbol{l}} \sum_s \nu_{\boldsymbol{l}}(s) g(r) \\ &= \ln(2) D_{\boldsymbol{l}} \mathbb{E}_{\nu_{\boldsymbol{l}}} g(r),\end{aligned}$$

where $\nu_{\boldsymbol{l}} = \frac{f(r)}{\ln(2) D_{\boldsymbol{l}}} p_S(s)$ is a D_{KL}-weighted probability measure on s as $\sum_s \nu_{\boldsymbol{l}}(s) = 1$. By taking the expectation over $\boldsymbol{L}$, we have

$$I_{\chi^2}(S;\boldsymbol{L}) = \ln(2) I(S;\boldsymbol{L}) \mathbb{E}_{\eta} \mathbb{E}_{\nu} g(r),$$

where $\eta(\mathrm{d}\boldsymbol{l}) = \frac{D_{\boldsymbol{l}}}{\ln(2) I(S;\boldsymbol{L})} p_{\boldsymbol{L}}(\mathrm{d}\boldsymbol{l})$ is a D_{KL}-weighted probability measure on $\boldsymbol{l}$ as $\int \eta(\mathrm{d}\boldsymbol{l}) = 1$. Using η and ν_l, we have the D_{KL}-weighted probability measure $\mathrm{d}\mu$ and conversion factor $\mathbb{E}_{\mu} g(r)$ as

$$\mu(s, \mathrm{d}\boldsymbol{l}) = \eta(\mathrm{d}\boldsymbol{l}) \nu_{\boldsymbol{l}}(s) = \frac{f(r)}{\ln(2) I(S;\boldsymbol{L})} p_S(s) p_{\boldsymbol{L}}(\mathrm{d}\boldsymbol{l}),$$
$$\mathbb{E}_{\mu} g(r) = \mathbb{E}_{\eta} \mathbb{E}_{\nu_{\boldsymbol{l}}} g(r),$$

respectively. This completes the proof. □

B.3 Proof of Theorem 2

Proof. If $\rho^* = \operatorname{ess\,sup}_{\boldsymbol{l}} |\hat{p}_{S_i|\boldsymbol{L}_i}(w \mid \boldsymbol{l})| < 1$ holds for any $w \neq 0$, $p_{Z|\boldsymbol{X}}$ uniformly converges to the uniform distribution over $\mathbb{G}$ with an exponential rate of ρ^*, followed by a uniform convergence of $\operatorname{ess\,sup}_{z,\boldsymbol{x}} r(z, \boldsymbol{x}) \to 1$ where $r(z, \boldsymbol{x}) = p_{Z|\boldsymbol{X}}(z \mid \boldsymbol{x})/p_Z(z)$ (precisely, r is defined as the Radon–Nikodim derivative in Equation (2)). According to the triangle inequality and inverse DFT, it holds

$$\left| \frac{p_{Z|\boldsymbol{X}}(z \mid \boldsymbol{x})}{p_Z(z)} - 1 \right| \leq \sum_{w \neq 0} \left| \hat{p}_{Z|\boldsymbol{X}}(w \mid \boldsymbol{x}) \right| \leq (m-1)(\rho^*)^{d+1},$$

implying that it holds

$$\left\| \frac{p_{Z|\boldsymbol{X}}(z \mid \boldsymbol{x})}{p_Z(z)} - 1 \right\|_\infty \leq (m-1)(\rho^*)^{d+1} = o(1), \tag{16}$$

as $d \to \infty$, because $\rho^* < 1$ owing to the assumption.

Instantiate Lemma 3 with $p = p_{Z|\boldsymbol{X}}$ and $q = p_Z$. According to Equation (16), it uniformly holds

$$D_{\mathrm{KL}}(p_{Z|\boldsymbol{X}}(\cdot \parallel \boldsymbol{x}) \parallel p_Z) = \frac{1 + O((m-1)(\rho^*)^{d+1})}{2 \ln 2} \chi^2(p_{Z|\boldsymbol{X}}(\cdot \parallel \boldsymbol{x}) \parallel p_Z).$$

By taking expectation over $\boldsymbol{X}$, we conclude

$$I(Z; \boldsymbol{X}) = \frac{(1 + O((m-1)(\rho^*)^{d+1})}{2 \ln 2} I_{\chi^2}(Z; \boldsymbol{X}),$$

as $d \to \infty$ for given m. This completes the proof. □

References

AA19. Emmanuel Abbe and Enric Boix Adserà. Subadditivity beyond trees and the chi-squared mutual information. In *2019 IEEE International Symposium on Information Theory (ISIT)*, pages 697–701, 2019.

BBM26. Julien Béguinot, Gianluca Brian, and Loïc Masure. Bypassing the random-probing model in masking security proofs. In *Advances in Cryptology—CRYPTO 2026*, pages 231–261, 2026.

BCG+23. Julien Béguinot, Wei Cheng, Sylvain Guilley, Yi Liu, Loïc Masure, Olivier Rioul, and François-Xavier Standaert. Removing the field size loss from Duc et al.'s conjectured bound for masked encodings. In *International Workshop on Constructive Side-Channel Analysis and Secure Design (COSADE)*, pages 86–104, 2023.

BDIR18. Fabrice Benhamouda, Akshay Degwekar, Yuval Ishai, and Tal Rabin. On the local leakage resilience of linear secret sharing schemes. In *Advances in Cryptology—CRYPTO 2018*, pages 531–561. Springer International Publishing, 2018.

Bea91. Donald Beaver. Efficient multiparty protocols using circuit randomization. In *Advances in Cryptology—CRYPTO '91*, pages 420–432, 1991.

BHMY26. Aniruddha Biswas, Jihun Hwang, Hemanta K. Maji, and Xiuyu Ye. Resilience of inner-product masking scheme against hamming weight leakage. In *Information-Theoretic Cryptography Conference (ITC)*, 2026.

BOGW88. Michael Ben-Or, Shafi Goldwasser, and Avi Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation. In *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing*, pages 1–10, 1988.

BS03. Simon R. Blackburn and Andrew J. Spencer. Products of subsets in an abelian group. *Journal of Combinatorial Theory, Series A*, 103(1):53–68, 2003.

CJRR99. Suresh Chari, Charanjit S. Jutla, Josyula R. Rao, and Pankaj Rohatgi. Towards sound approaches to counteract power-analysis attacks. In *Advances in Cryptology—CRYPTO' 99*, pages 398–412, 1999.

DFS15. Alexandre Duc, Sebastian Faust, and François-Xavier Standaert. Making masking security proofs concrete: Or how to evaluate the security of any leakage device. In *Advances in Cryptology—EUROCRYPT 2015*, volume 9056 of *Lecture Notes in Computer Science*, pages 401–429. Springer, 2015.

DFS16. Stefan Dziembowski, Sebastian Faust, and Maciej Skórski. Optimal amplification of noisy leakages. In *Theory of Cryptography*, pages 291–318, 2016.

FHT03. Alexei A Fedotov, Peter Harremoës, and Flemming Topsoe. Refinements of Pinsker's inequality. *IEEE Transactions on Information Theory*, 49(6):1491–1498, 2003.

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*, pages 316–344, 2024.

GSV05. Dongning Guo, Shlomo Shamai, and Sergio Verdú. Mutual information and minimum mean-square error in Gaussian channels. *IEEE Transactions on Information Theory*, 51(4):1261–1282, 2005.

IUH22. Akira Ito, Rei Ueno, and Naofumi Homma. On the success rate of side-channel attacks on masked implementations: Information-theoretical bounds and their practical usage. In *ACM SIGSAC Conference on Computer and Communications Security (CCS 2022)*, pages 1521–1535, 2022.

KA26. Hiroki Koga and Hiroto Abe. New tight bounds on the local leakage resilience of the additive (n, n)-threshold scheme determined by the eigenvalues of circulant matrices. In *Advances in Cryptology—ASIACRYPT 2025*, pages 34–66, 2026.

KJJ99. Paul Kocher, Joshua Jaffe, and Benjamin Jun. Differential power analysis. In *Advances in Cryptology—CRYPTO' 99*, volume 1666 of *Lecture Notes in Computer Science*, pages 388–397, 1999.

MMMS23. Loïc Masure, Pierrick Méaux, Thorben Moos, and François-Xavier Standaert. Effective and efficient masking with low noise using small-Mersenne-prime ciphers. In *Advances in Cryptology—EUROCRYPT 2023*, pages 596–627, 2023.

MNPC+22. Hamanta 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 *Information-Theoretic Cryptography (ITC)*, 2022.

MOP07. Stefan Mangard, Elisabeth Oswald, and Thomas Popp. *Power Analysis Attacks: Revealing the Secrets of Smart Cards*. Springer New York, 2007.

MRS22. Loïc Masure, Olivier Rioul, and François-Xavier Standaert. A nearly tight proof of Duc et al.'s conjectured security bound for masked implementations. In *Smart Card Research and Advanced Applications: 21st International Conference, CARDIS 2022*, pages 69–81, 2022.

NN14. Frank Nielsen and Richard Nock. On the chi square and higher-order chi distances for approximating f-divergences. *IEEE Signal Processing Letters*, 21(1):10–13, 2014.

PF64. Mark S Pinsker and Amiel Feinstein. *Information and information stability of random variables and processes*, volume 12. Holden-Day San Francisco, 1964.

PR13. Emmanuel Prouff and Matthieu Rivain. Masking against side-channel attacks: A formal security proof. In *Advances in Cryptology—EUROCRYPT 2013*, pages 142–159, 2013.

Rot64. Gian-Carlo Rota. *On the foundations of combinatorialtheory I. Theory of Möbius functions*, volume 2. Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete, 1964.

SV19. Akshayaram Srinivasan and Prashant Nalini Vasudevan. Leakage resilient secret sharing and applications. In *Advances in Cryptology—CRYPTO 2019*, page 480–509, 2019.

UIM+26. Rei Ueno, Akiko Inoue, Kazuhiko Minematsu, Akira Ito, and Naofumi Homma. A formal security proof of masking: Reduction from strong noisy leakage to probing model without random probing and application to LR primitive. In *Advances in Cryptology—CRYPTO 2026*, pages 326–358, 2026.