Indistinguishability of Sum of Permutations

A Fourier Analytic Route to Classical and Quantum Security

Ritam Bhaumik[1], Chun Guo[2,3,4], Xiaoning Guo[2,3,4], and Ashwin Jha[5]

[1] CRC, TII, Abu Dhabi, UAE
bhaumik.ritam@gmail.com
[2] School of Cyber Science and Technology;
[3] State Key Laboratory of Cryptography and Digital Economy Security;
[4] Key Laboratory of Cryptologic Technology and Information Security of Ministry of Education, Shandong University, Qingdao, China
chun.guo.sc@gmail.com,xxiaoningguo@gmail.com
[5] University of Wuppertal, Wuppertal, Germany
ashwin.jha@outlook.de

Abstract. We study classical and quantum indistinguishability of sums of independent random permutations and related transformations from permutations to functions. Let G be a finite abelian group of order N, and let $\pi_+^k(x) = \pi_1(x) + \cdots + \pi_k(x)$ for $k \geq 2$ independent uniform random permutations of G. We give a unified Fourier analytic treatment in which the construction is represented by its probability density and a distinguisher by its acceptance function, with the classical and quantum query models imposing different restrictions on the Fourier support of the latter.

Classically, we obtain the bound $O_k(q/N^{k-1/2})$ for every $q < N$, and refine it below the birthday threshold to $O_k(q^2/N^k)$. In the quantum model, a simulation argument gives $O_k(N^{-(k-3/2)})$ for $q \leq (N-1)/2$, while Fourier interpolation gives concrete finite bounds up to $q \leq 4N/15$ and the query-dependent bounds

$$O\left(\min\left\{N^{-1/2}, \frac{q^3}{N^2} + \frac{1}{N}\right\}\right), \qquad O_k\left(\min\left\{\frac{q^3}{N^k}, N^{-(k-3/2)}\right\}\right),$$

for $k = 2$ and $k \geq 3$, respectively, throughout $1 \leq q \leq (N-1)/2$. For $q = 1$, the first bound sharpens to $O(N^{-2})$. Over $G = \mathbb{F}_2^n$, a one-query Fourier attack matches the order of our one-query bound, while an $N/2$-query parity attack with advantage $1/2$ shows that our bounds reach the constant-advantage query threshold.

We further study two variants of sum of permutations over binary vector spaces. First, we allow arbitrary surjective linear postprocessing, which includes truncation, and obtain classical and quantum bounds that retain the output-size dependence. Second, we analyse Dinur's variable-output single-permutation construction, LXoP, for every fixed output width, and derive its classical and quantum security bounds; for one- and two-block outputs, we give concrete quantum security bounds.

Keywords: sum of permutations, LXoP, Fourier analysis, quantum security, PRP-to-PRF conversion

1 Introduction

Block ciphers, or pseudorandom permutation (PRP) candidates, are among the most thoroughly studied primitives in symmetric-key cryptography, while many applications [RS06, Che22, Nan24] require a pseudorandom function (PRF). Apart from a few examples like SURF [Ber97], SipHash [AB12], and some AES-based constructions [MN17b, BIL+21, FGL+24], relatively few dedicated PRF candidates have been proposed or analysed. This makes permutation-to-function transformation a natural route to pseudorandom function designs. It is well-known [Rao66, Fre77, BKR94, CN08] that a random permutation itself can be distinguished from a random function by looking for collisions, with constant advantage after about $\sqrt{N}$ classical queries, where N denotes the ambient group size. A natural question is therefore how to combine a small number of random permutation calls so that the resulting function remains indistinguishable from random well beyond the birthday bound.

In this line of research, three simple constructions are the sum of permutations [BKR98], encrypted Davies-Meyer [CS16] and its dual [MN17a]. In particular, the sum of permutations construction has received a lot of attention from the community. Let G be a finite abelian group of order N, written additively, and let $\pi_1, \ldots, \pi_k$ be independent uniform random permutations of G. For $k \geq 2$, define the sum of k permutations by

$$\pi_+^k(x) \coloneqq \pi_1(x) + \cdots + \pi_k(x).$$

When $G = \mathbb{F}_2^n$, this is the usual XOR-of-permutations construction. The two-permutation case goes back to Bellare et al. [BKR98], and the k-permutation generalisation [Luc00] is due to Lucks. The construction is simple to evaluate: the permutation calls are independent and can be made in parallel, followed by a group addition. Its security analysis, however, has proved considerably more involved, and has occupied the community for more than two decades.

Classical security of sum of permutations. A long line of work [BI99, Luc00, Pat08, CLP15, MP15, Ebe17, DHT17, DNS22, CP20, Din24] has investigated how much security this simple transformation provides. Adding independent permutation outputs allows collisions, so the absence of collisions no longer distinguishes the construction from a random function. The outputs nevertheless remain dependent, because each constituent permutation samples without replacement. An analysis must bound the effect of these dependencies on the joint distribution of all queried outputs, including when the number of queries is close to N. Over the years, different proof techniques have been used to prove this result when $G = \mathbb{F}_2^n$, with various degrees of success. The first verifiable proof [Luc00] was given by Lucks who employed the game-playing technique to prove PRF advantage of $O(q^3/N^2)$, where q is the number of queries. Bellare and Impagliazzo, and later, Dai et al. devised new statistical techniques [BI99, DHT17] to show bounds of $O_n(q/N)$ and $O(\sqrt{q^3/N^3})$, respectively. Patarin's mirror theory [Pat08, Pat10, Pat16] has been the main tool to study the problem from a combinatorial viewpoint, and a verified proof [DNS22] of the mirror theory result yields a bound of $O(q^2/N^2)$ for $q \le N/17$ (see also [CP20] for the single-permutation variant). On the other hand, for $\sqrt{N} \le q < N$, the best-known attack [Pat13] due to Patarin only achieved an advantage of $\Omega(\sqrt{q^2/N^3})$, leaving a tightness gap. This gap was recently filled [Din24] by Dinur who established a bound of $O(\sqrt{q^2/N^3})$ for $q < N/2$ and $N \ge 1000$, which is tight for $q \ge \sqrt{N}$. Shortly after, it came to light that Eberhard had already proved [Ebe17] the same bound for any finite abelian group and for every $q < N$, in a work [Ebe17] that was largely unknown in the cryptography community. Both Eberhard and Dinur applied Fourier analytic techniques to study the problem. Dinur also studied the sum of k permutations construction, and obtained a bound of $q/N^{k-1/2}$ for $q < N/2$ queries.

Variants of sum of permutations. Several variants of the sum of permutations have been proposed to obtain shorter outputs or to generate several output blocks from a single permutation. For shorter outputs, summation can be combined with the *truncation* technique [HWKS98]. Let $G = \mathbb{F}_2^n$, with $|G| = N$, and let $\tau : G \to H$ be a surjective linear map, where $|H| = M \le N$. The corresponding construction is

$$x \longmapsto \tau\big(\pi_1(x) + \pi_2(x)\big),$$

where π_1 and π_2 are independent random permutations of G. When τ is a coordinate projection, Choi et al. [CKLL22] proved a classical distinguishing bound of $O(q\sqrt{qM}/N^2)$ for $q \le N/4$.

For outputs in the larger range G^w, Iwata's $\mathsf{XORP}[w]$ construction [Iwa06] uses a single random permutation $\pi : G \to G$ and is given by

$$x \longmapsto \big(\pi(0 \,\|\, x) + \pi(1 \,\|\, x), \ldots, \pi(0 \,\|\, x) + \pi(w \,\|\, x)\big).$$

Here the prefixes are encoded in $s = \lceil \log_2(w+1) \rceil$ bits and $x \in \mathbb{F}_2^{n-s}$. The construction produces w output blocks using $w+1$ permutation calls. For fixed w, its classical distinguishing advantage is $O_w(q/N)$ [Pat16, IMV16, BN18]. Dinur [Din25] introduced the LXoP family, which also produces w output blocks from $w+1$ calls, but combines consecutive permutation outputs using a linear automorphism $\sigma : G \to G$:

$$x \longmapsto \big(\pi(0 \,\|\, x) + \sigma(\pi(1 \,\|\, x)), \ldots, \pi((w-1) \,\|\, x) + \sigma(\pi(w \,\|\, x))\big).$$

For $w \in \{1, 2\}$, he proved a classical distinguishing advantage bound of $O(q/N^{3/2})$ for $q \le N/16$ and $q \le N/32$, respectively, and $N \ge 1024$.

Quantum security of sum of permutations. The quantum setting raises a different difficulty. For these input-symmetric constructions, a classical distinguisher can be reduced to observing the outputs on a fixed set of queried inputs. This remains true even when the queries are adaptive, as the labels of fresh inputs do not affect the distribution of the next answer. A quantum distinguisher may instead query a superposition of inputs and combine information from different positions by interference. Its behaviour is therefore not determined by the output distribution on any one fixed set of q inputs.

Quantum security of XOR-of-PRP constructions has been studied [MS17] by Mennink and Szepieniec, including key-recovery attacks and a security analysis. There, the construction oracle is queried classically, while the adversary may use offline quantum computation. This is the so-called Q1 security model. In this paper, we study information-theoretic security in the Q2 model, where the adversary can make quantum superposition queries to the construction. The adversary is computationally unbounded, and we count only queries to the construction. Thus the question concerns what can be learned from the oracle itself, independently of the time complexity of processing its answers.

The sum of permutations has a global constraint that makes this question particularly interesting. Every complete truth table satisfies

$$\sum_{x\in G} \pi_+^k(x) = k \sum_{y\in G} y.$$

For a uniform random function, the sum on the left is uniform on G, so the same equality holds with probability exactly $1/N$. The complete-table distribution of π_+^k therefore has statistical distance at least $1-1/N$ from uniform. At the same time, the classical bounds show that every proper restriction is close to uniform. How much of this dependence can a quantum algorithm detect with few queries? At what query count does the checksum become visible?

Our bounds give a sharp answer to the second question in the binary setting. For every fixed $k \geq 2$, the advantage tends to zero for all $q \leq N/2-1$, whereas an $N/2$-query parity algorithm has advantage $1/2$. For $q < N/2$, we exploit the well-known fact [BBC+98, Zha12b] that the acceptance probability of a q-query quantum adversary depends only on the joint distributions of oracle outputs on all sets of at most $2q$ inputs. This allows us to focus on dependencies among small collections of outputs, even though each quantum query may involve a superposition over the entire domain.

1.1 Our Results

We develop a common analytic framework for classical and quantum distinguishing games, and apply it to sums of independent permutations, linear-output variants, and a variable-output construction based on a single permutation. In what follows, all bounds are information-theoretic, and the number k of independent permutations, or the output width w, is fixed in the corresponding asymptotic statements.

Sum of independent permutations. For every finite abelian group G of order N, we prove the classical bounds

$$\left\|\pi_+^k - \mathsf{g}\right\|_{(q)} \leq \begin{cases} 0, & 0 \leq q \leq 1, \\ O_k(q^2/N^k), & 2 \leq q \leq \sqrt{N}, \\ O_k(q/N^{k-1/2}), & \sqrt{N} < q < N, \end{cases}$$

where $\mathsf{g}\colon G \to G$ is a uniform random function. This gives a single proof over arbitrary finite abelian groups, and refines the bound when $q < \sqrt{N}$. For $k = 2$, the bound above that threshold is Eberhard's bound; our proof extends it to every fixed $k \geq 2$.

For quantum queries, we obtain

$$\left\||\pi_+^2\rangle - |\mathsf{g}\rangle\right\|_{(q)} \leq O\left(\min\left\{N^{-1/2}, \frac{q^3}{N^2} + \frac{1}{N}\right\}\right),$$

$$\left\||\pi_+^k\rangle - |\mathsf{g}\rangle\right\|_{(q)} \leq O_k\left(\min\left\{\frac{q^3}{N^k}, N^{-(k-3/2)}\right\}\right) \qquad (k \geq 3),$$

throughout $1 \leq q \leq (N-1)/2$. For $k = 2$ and $q = 1$, the bound improves to $O(N^{-2})$. The bounds involving q^3 describe the behaviour in the range $q < \sqrt{N}$, while the bounds independent of q remain small throughout the stated range. We also give concrete bounds for $N \geq 1024$ and $q \leq 4N/15$.

Over $G = \mathbb{F}_2^n$, a one-query algorithm has advantage $1/(2(N-1)^k)$, matching the order of our upper bound at one query. For $n \geq 2$, the $N/2$-query parity attack, together with the uniform bound of $O_k(N^{-(k-3/2)})$ throughout $q \leq N/2-1$, establishes the exact number of queries required to achieve a constant advantage.

Linear-output-postprocessing. For $G = \mathbb{F}_2^n$, let $\tau\colon G \to H = \mathbb{F}_2^m$ be surjective and linear, and put $M = 2^m$. We prove a classical bound $O_k(q\sqrt{M}/N^k)$ for $\tau\pi_+^k$ throughout $q < N$. In particular, for the truncated construction considered [CKLL22] by Choi et al., our bound improves their single-user bound $O(q\sqrt{qM}/N^2)$ by a factor $\sqrt{q}$. We also obtain quantum bounds that retain the same dependence on M, both for small query counts and throughout $q \leq (N-1)/2$.

Variable output from a single permutation. We analyse Dinur's LXoP family for every fixed width w, under the admissibility condition stated in Section 6. The construction produces w output blocks using $w+1$ permutation calls. We obtain classical advantage $O_w(q/N^{3/2})$ and quantum advantage $O_w(N^{-1/2})$ for query counts up to a sufficiently small constant fraction of the construction's input-domain size. We also give concrete quantum bounds for $w = 1, 2$.

1.2 Technical Overview and Organisation

Densities and query restrictions. In Section 2, we represent the construction by its probability density φ and a distinguisher by its acceptance probability function a. Relative to the reference product measure, the distinguishing advantage is then given by the inner product $|\langle a, \varphi - 1\rangle|$. To analyse this inner product, we first choose an orthonormal basis for functions on the output set, with respect to the reference output distribution, containing the constant function 1. We use this same basis for each truth-table entry $f(x)$. Taking products of these basis functions, with one factor for each input, gives an orthonormal basis for functions of the complete truth-table. We expand a and φ in this product basis and group the terms according to the set S of inputs at which the factor is nonconstant. These groups give the orthogonal *pure-support* components. A component indexed by S depends only on the outputs at inputs in S, and averaging over any one of these outputs, with the others fixed, gives zero. Orthogonality then expresses the distinguishing advantage as the absolute value of a sum of inner products between a- and φ-components with matching supports.

This decomposition is defined for an arbitrary product measure, and the grouped components do not depend on the particular basis chosen. Under the uniform measure on a finite abelian group, we may choose the group characters as our basis. The components are then precisely the terms of the Fourier expansion grouped by their support. For the constructions studied in this paper, input symmetry lets us restrict a classical distinguisher to a fixed set Q of q inputs. Only supports contained in Q can contribute to its advantage. For a quantum distinguisher, Lemma 2, shows that only components supported on at most $2q$ inputs can contribute, although different components may involve different inputs. Since $0 \le a \le 1$, Cauchy-Schwarz bounds the remaining inner product by one half of the L_2 norm of the corresponding density projection. This yields the Fourier interpolation bound of Theorem 1.

Injection moments and the first nonzero level. Section 3 applies this framework to sums of independent permutations. On a fixed set of distinct inputs, the outputs of one permutation form a uniformly random injection. Summing independent permutation outputs convolves their densities, so the Fourier coefficients of the sum are powers of those of one injection. Eberhard's fourth-moment bound [Ebe17], together with a bound on the largest nonconstant coefficient, gives the global classical bound for every $k \ge 2$. We also derive a higher-moment bound for the Fourier coefficients of the injection density, which controls the components involving three or more inputs. In the classical analysis, we apply it to the joint output distribution on q fixed inputs. In the quantum analysis, we apply it to the corresponding components of the distribution of the complete truth table, retaining only those involving at most $2q$ inputs. The detailed coefficient bounds and moment calculations are given in Appendices A and B.

We call the collection of components supported on exactly t inputs *level* t. Since the one-point marginals are uniform, level one vanishes. To obtain a more precise bound at small query counts, we treat level two separately before applying Cauchy-Schwarz. This component is a scalar multiple of the difference between the number of output collisions and its expected value under a uniform random function. The classical proof bounds its contribution in L_1 by summing over the queried pairs. For quantum queries, Appendix C bounds the correlation of this collision count with the acceptance function. We obtain an $O(q^3)$ bound using Zhandry's small-range method [Zha12a]: the centred collision count is the first-order deviation from a uniform function when outputs are sampled through a large but finite intermediate range. The common moment estimate bounds the remaining levels.

Simulation from the checksum. To reach every quantum query count below $N/2$, in Section 4, we use Eberhard's bound along with a checksum-based simulation argument. Any $N-1$ outputs of π_+^k determine the last output. Start instead with independent uniform values on $N-1$ inputs and complete the table using the same checksum. The resulting function is uniform among all functions with that checksum and is $(N-1)$-wise independent. It is therefore indistinguishable from a uniform random function whenever $2q \le N-1$. It remains to compare the sum of permutations with this completed reference function. Both complete tables are determined by their restrictions to the chosen $N-1$ inputs, so their statistical distance is bounded by the classical estimate on those inputs. This gives the quantum bound $O_k(N^{-(k-3/2)})$ throughout the required range.

Linear-output-postprocessing and shared permutation calls. Section 5 considers applying a surjective linear map τ to each output. A Fourier character of the projected output $\tau(y)$ can also be viewed as a character of the original output y: its index η is replaced by $\tau^\top\eta$. Thus only indices in the subspace $\mathrm{Im}(\tau^\top)$, which has size M, enter the analysis. To bound their contribution, we use the fact that applying the same invertible linear map to every output of a random injection preserves its distribution. This symmetry lets us average over subspaces of size M and compare the Fourier moments restricted to such a subspace with the unrestricted moments. This is where the output size M enters the bounds. For the refined quantum estimate, we also distinguish tuples of character indices according to the dimension of their span. The probability that all indices lie in a random subspace of size M depends on this dimension. Keeping this dependence gives a sharper bound for the higher levels.

For the LXoP family studied in Section 6, each construction output consists of w blocks obtained from $w+1$ permutation calls. Expanding a character of the construction output gives a product of characters of the underlying permutation outputs, with one index for each call. These permutation outputs form a random injection, so we can again use its Fourier coefficients. The shared calls force the indices associated with each construction input to satisfy a linear relation. We exploit this relation when applying a Dinur-style recursive identity [Din25] that replaces an injection coefficient by a sum of coefficients involving fewer permutation outputs. To obtain a bound on the sum of squared coefficients, we must control how many original index tuples can give the same tuple after a reduction. By choosing the reductions in a suitable order, the linear relation and the admissibility condition let us recover the removed index. This controls the number of times each reduced coefficient is counted and gives the required second-moment bound. Summing over the supports permitted in each query model then gives the classical and quantum bounds.

2 Fourier Analytic Framework

This section first fixes the notation used throughout the paper and then develops a unified analytic framework for classical and quantum distinguishing games. The framework is inspired by the Fourier analytic methods employed by Eberhard [Ebe17] and Dinur [Din24, Din25].

Notation. For a positive integer n, write $[n] \coloneqq \{1,\dots,n\}$ and $[0] \coloneqq \varnothing$. For $0 \le k \le n$, let $(n)_k \coloneqq n(n-1)\cdots(n-k+1)$, with $(n)_0 = 1$. Empty sums and products are understood as zero and one, respectively. The constants in $O_k(\cdot)$ and $O_w(\cdot)$ may depend only on the indicated fixed parameter.

Let Q be a probability measure on a non-empty finite set Ω. We write $x \sim \mathsf{Q}$ to mean that x is drawn according to Q, and write $\mathrm{supp}(\mathsf{Q})$ for its support. We denote the uniform measure on Ω by U_Ω, abbreviated to U when Ω is clear. A *density* relative to Q is a function $\varphi\colon \mathrm{supp}(\mathsf{Q}) \to \mathbb{R}_{\ge 0}$ satisfying $\mathbf{E}_{x\sim\mathsf{Q}}[\varphi(x)] = 1$. We sometimes identify φ with the probability distribution that assigns mass $\varphi(z)\mathsf{Q}(z)$ to each $z \in \mathrm{supp}(\mathsf{Q})$, when the reference measure is clear. The reference measure Q itself has density 1.

Let $L_2(\mathsf{Q})$ be the space of complex-valued functions on $\mathrm{supp}(\mathsf{Q})$. For $h, g \in L_2(\mathsf{Q})$ and $1 \le p < \infty$, write

$$\langle h, g\rangle_{\mathsf{Q}} \coloneqq \mathbf{E}_{x\sim\mathsf{Q}}[h(x)\overline{g(x)}], \qquad \|h\|_{p,\mathsf{Q}} \coloneqq \left(\mathbf{E}_{x\sim\mathsf{Q}}[|h(x)|^p]\right)^{1/p}.$$

For densities φ and ψ relative to Q, their statistical distance is given by

$$\mathrm{TVD}_{\mathsf{Q}}(\varphi,\psi) \coloneqq \frac{1}{2}\|\varphi-\psi\|_{1,\mathsf{Q}},$$

whence by Cauchy-Schwarz,

$$\mathrm{TVD}_{\mathsf{Q}}(\varphi,\psi) \le \frac{1}{2}\|\varphi-\psi\|_{2,\mathsf{Q}}. \tag{1}$$

2.1 Product Spaces and Support Decomposition

We recall the standard support decomposition on a finite product space (see [O'D14, Sections 8.1 and 8.3]), which requires no algebraic structure on the output set.

Let $\mathcal{X}$ and $\mathcal{Y}$ be finite sets, with $\mathcal{Y} \ne \varnothing$. Fix a probability measure P on $\mathcal{Y}$ and equip $\mathcal{Y}^{\mathcal{X}}$ with the product measure $\mathsf{P}^{\mathcal{X}}$. Unless stated otherwise, densities, expectations, inner products, and norms on this space are taken with respect to $\mathsf{P}^{\mathcal{X}}$, whose subscript is omitted.

Let $L_2^0(\mathsf{P}) \coloneqq \left\{g \in L_2(\mathsf{P}) : \mathbf{E}_{y\sim\mathsf{P}}[g(y)] = 0\right\}$. Choose an orthonormal basis B_0 of $L_2^0(\mathsf{P})$ and put $B = \{1\} \cup B_0$. For $\lambda = (\lambda_x)_{x\in\mathcal{X}} \in B^{\mathcal{X}}$, define

$$\Phi_\lambda(f) \coloneqq \prod_{x\in\mathcal{X}} \lambda_x(f(x)), \qquad \mathrm{supp}(\lambda) \coloneqq \{x \in \mathcal{X} : \lambda_x \neq 1\}.$$

We call $\mathrm{supp}(\lambda)$ the *pure support* of Φ_λ. The functions Φ_λ form an orthonormal basis of $L_2(\mathsf{P}^{\mathcal{X}})$.

For $S \subseteq \mathcal{X}$, let Π_S be the orthogonal projection onto the subspace

$$\mathcal{H}_S \coloneqq \mathrm{Span}\{\Phi_\lambda : \mathrm{supp}(\lambda) = S\}.$$

Then

$$L_2(\mathsf{P}^{\mathcal{X}}) = \bigoplus_{S\subseteq\mathcal{X}} \mathcal{H}_S, \qquad h = \sum_{S\subseteq\mathcal{X}} \Pi_S h, \qquad \|h\|_2^2 = \sum_{S\subseteq\mathcal{X}} \|\Pi_S h\|_2^2. \tag{2}$$

The spaces $\mathcal{H}_S$ depend only on the orthogonal decomposition $L_2(\mathsf{P}) = \mathrm{Span}\{1\} \oplus L_2^0(\mathsf{P})$, and not on the particular basis chosen for $L_2^0(\mathsf{P})$. The space $\mathcal{H}_\varnothing$ consists of the constant functions, so $\Pi_\varnothing h = \mathbf{E}[h] \cdot 1$. In particular, $\Pi_\varnothing \varphi = 1$ for every density φ. If the distribution represented by φ has marginal P at x, then

$$\Pi_{\{x\}}\varphi = \mathbf{E}[\varphi \mid f(x)] - 1 = 0. \tag{3}$$

For integers $t, d \geq 0$, write

$$h_{=t} \coloneqq \sum_{|S|=t} \Pi_S h, \qquad \Pi_{\leq d} h \coloneqq \sum_{\substack{S\subseteq\mathcal{X} \\ |S|\leq d}} \Pi_S h.$$

We shall repeatedly use the following coordinate dependence property of the product basis. A function $h \in L_2(\mathsf{P}^{\mathcal{X}})$ depends only on the coordinates in $Q \subseteq \mathcal{X}$ if and only if

$$\Pi_S h = 0 \qquad \text{for every } S \not\subseteq Q. \tag{4}$$

If these projections vanish, the expansion of h contains only basis functions supported in Q, which depend only on Q. Conversely, if h depends only on Q, every coefficient involving a mean zero basis function at a coordinate outside Q vanishes after averaging over that coordinate.

2.2 Distinguishing Advantage and Query Support Cutoffs

For any function $f \in \mathcal{Y}^{\mathcal{X}}$, we denote the classical oracle corresponding to f by the function itself, while the quantum oracle is denoted by $|f\rangle$.

Let O_0 and O_1 be oracles sampled according to densities φ and ψ relative to $\mathsf{P}^{\mathcal{X}}$, respectively, and accessed in the same model. For an oracle algorithm $\mathcal{A}$ and a fixed function f, let $a(f) \in [0,1]$ be the probability that $\mathcal{A}$ outputs 1 given oracle access to f. Define

$$\|O_0 - O_1\|_{\mathcal{A}} \coloneqq \left|\Pr[\mathcal{A}^{O_0} = 1] - \Pr[\mathcal{A}^{O_1} = 1]\right|.$$

We write $\|O_0 - O_1\|_{(q)}$ for the supremum over all algorithms making at most q queries in the indicated access model. We work in the information-theoretic setting: algorithms may be computationally unbounded, and only oracle queries are counted. Then,

$$\|O_0 - O_1\|_{\mathcal{A}} = \left|\mathbf{E}[a\varphi] - \mathbf{E}[a\psi]\right| = \left|\langle a, \varphi - \psi\rangle\right| \leq \mathrm{TVD}_{\mathsf{P}^{\mathcal{X}}}(\varphi, \psi), \tag{5}$$

where the final inequality follows by writing $\langle a, \varphi - \psi\rangle = \mathbf{E}[(a - \frac{1}{2})(\varphi - \psi)]$, since $\mathbf{E}[\varphi - \psi] = 0$ and $|a - \frac{1}{2}| \leq \frac{1}{2}$, so that the inner product is at most $\frac{1}{2}\|\varphi - \psi\|_1$. For the remainder of this paper, we assume $\psi = 1$.

We next apply the support decomposition (2) to the inner product in (5). The constant component of $\varphi - 1$ is zero, and its projection onto every non-empty support S is $\Pi_S\varphi$. Orthogonality therefore gives

$$\|O_0 - O_1\|_{\mathcal{A}} = \left|\sum_{\varnothing\neq S\subseteq\mathcal{X}} \langle \Pi_S a, \Pi_S \varphi\rangle\right|. \tag{6}$$

For $j \geq 1$, put

$$\Delta_j(a, \varphi) \coloneqq \sum_{\substack{S\subseteq\mathcal{X} \\ |S|=j}} \langle \Pi_S a, \Pi_S \varphi\rangle.$$

The first two levels will often be treated directly, while the remaining levels are bounded in L_2.

Lemma 1. *Let $O_0 \sim \varphi$ and $O_1 \sim 1$ be function oracles in the same access model. Suppose $\Pi_S a = 0$ whenever $|S| > d$. Then, for every integer $0 \le r \le d$,*

$$\|O_0 - O_1\|_{\mathcal{A}} \le \sum_{j=1}^{r} |\Delta_j(a,\varphi)| + \frac{1}{2}\left(\sum_{\substack{S\subseteq\mathcal{X}\\ r+1\le|S|\le d}} \|\Pi_S\varphi\|_2^2\right)^{1/2}.$$

If a depends only on the coordinates in $Q \subseteq \mathcal{X}$, the sum in the remainder may be restricted to $S \subseteq Q$.

Proof. By (6), only non-empty supports contribute to the advantage. Since $\Pi_S a = 0$ for $|S| > d$, we have

$$\langle a, \varphi - 1\rangle = \sum_{j=1}^{r} \Delta_j(a,\varphi) + \sum_{\substack{S\subseteq\mathcal{X}\\ r<|S|\le d}} \langle \Pi_S a, \Pi_S\varphi\rangle.$$

Put $m = \mathbf{E}[a]$. Since $0 \le a \le 1$,

$$\sum_{\varnothing\ne S\subseteq\mathcal{X}} \|\Pi_S a\|_2^2 = \|a-m\|_2^2 = \mathbf{E}[a^2] - m^2 \le m - m^2 \le \frac{1}{4}.$$

By Cauchy-Schwarz,

$$\left|\sum_{\substack{S\subseteq\mathcal{X}\\ r<|S|\le d}} \langle \Pi_S a, \Pi_S\varphi\rangle\right| \le \frac{1}{2}\left(\sum_{\substack{S\subseteq\mathcal{X}\\ r<|S|\le d}} \|\Pi_S\varphi\|_2^2\right)^{1/2}.$$

The triangle inequality gives the claim. If a depends only on Q, then (4) gives $\Pi_S a = 0$ whenever $S \not\subseteq Q$, so these supports may also be omitted. □

Remark 1. The constructions considered below are invariant in distribution under input permutations, as is the product reference measure. Thus, for $q \le |\mathcal{X}|$, a classical q-query distinguisher may be restricted to any fixed set Q of size q: answer its j-th fresh query using the j-th point of Q, and replay answers to repeated queries. Input symmetry preserves the transcript distribution under both oracle distributions. By (4), the resulting acceptance function satisfies $\Pi_S a_Q = 0$ whenever $S \not\subseteq Q$.

For a q-query quantum distinguisher, every nonzero pure-support component of the acceptance probability involves at most $2q$ input coordinates. We use the standard quantum function-oracle model, with a fixed injective encoding of $\mathcal{Y}$ into a finite abelian group and translation by the encoded output; all applications below instantiate it by the usual addition oracle on a finite abelian output group.

Lemma 2. *Let $\mathcal{A}$ be a quantum algorithm making at most q queries to a function $f : \mathcal{X} \to \mathcal{Y}$, and let $a(f)$ be its acceptance probability. Then*

$$\Pi_S a = 0 \qquad \text{whenever } |S| > 2q.$$

Proof. Write the function oracle as

$$U_f\,|x,r,z\rangle = |x\rangle \otimes V_{f(x)}\,|r\rangle \otimes |z\rangle\,,$$

where $V_y\,|r\rangle = |r + \mathrm{enc}(y)\rangle$ translates the answer register by the fixed encoding of $y \in \mathcal{Y}$, and z denotes the internal registers of the algorithm. For a basis label $u = (x,r,z)$, every matrix entry $\langle u|U_f|u'\rangle$ therefore depends on f only through the coordinate $f(x)$.

Let $\rho_t(f)$ denote the density matrix of the algorithm after the t-th query and the subsequent oracle-independent operation, with ρ_0 the initial state. We prove inductively that every matrix entry of $\rho_t(f)$, viewed as a function of f, is a linear combination of functions each depending on at most $2t$ input coordinates.

The claim is immediate for $t = 0$, since ρ_0 is independent of f. Suppose it holds after $t-1$ queries, and consider the state immediately after the next oracle call,

$$\rho_t'(f) = U_f\,\rho_{t-1}(f)\,U_f^{\dagger}.$$

For basis labels $u = (x, r, z)$ and $v = (x', r', z')$, we have

$$\langle u|\rho_t'(f)|v\rangle = \sum_{u',v'} \langle u|U_f|u'\rangle\,\overline{\langle v|U_f|v'\rangle} \cdot \langle u'|\rho_{t-1}(f)|v'\rangle.$$

The first two factors depend only on $f(x)$ and $f(x')$. By the induction hypothesis, the last factor is a linear combination of functions each depending on a set $Q \subseteq \mathcal{X}$ with $|Q| \le 2(t-1)$. Multiplying any such function by the first two factors produces a function depending only on

$$Q \cup \{x, x'\},$$

which has size at most $2t$. Hence every entry of $\rho_t'(f)$ has the required form.

If Φ_t denotes the oracle-independent operation following the query, then every entry of $\rho_t(f) = \Phi_t(\rho_t'(f))$ is a fixed linear combination of entries of $\rho_t'(f)$, with coefficients independent of f. We emphasise that the density-matrix formulation also allows mixed initial states, classical randomness, and intermediate measurements, with their outcomes retained in the internal registers. The same coordinate-dependence bound therefore holds for $\rho_t(f)$. This completes the induction.

Finally, for a fixed acceptance measurement operator M,

$$a(f) = \mathrm{Tr}(M\rho_q(f))$$

is a fixed linear combination of entries of $\rho_q(f)$. Thus a is a linear combination of functions each depending on at most $2q$ input coordinates. For any such function h, depending only on Q, and any S with $|S| > 2q$, we have $S \not\subseteq Q$. By (4), $\Pi_S h = 0$. Linearity of Π_S now gives

$$\Pi_S a = 0 \qquad \text{whenever } |S| > 2q. \qquad \square$$

Lemmas 1 and 2, with $r = 0$ and $d = 2q$, yield the following Fourier interpolation result, which we use repeatedly.

Theorem 1 (Fourier interpolation). *Let $O_0 \sim \varphi$ and $O_1 \sim 1$ be quantum function oracles. For every $q \ge 0$,*

$$\|O_0 - O_1\|_{(q)} \le \frac{1}{2}\|\Pi_{\le 2q}(\varphi - 1)\|_2 = \frac{1}{2}\left(\sum_{\substack{S\subseteq\mathcal{X}\\ 1\le|S|\le 2q}} \|\Pi_S\varphi\|_2^2\right)^{1/2}.$$

2.3 Character Bases over Finite Abelian Groups

We now instantiate the support decomposition of Subsection 2.1. Let G be a finite abelian group of order $N \ge 2$, written additively, and set $\mathcal{Y} = G$ and $\mathsf{P} = \mathsf{U}_G$. Unless stated otherwise, all densities, expectations, inner products, and norms below are relative to the corresponding uniform measures.

Let $\widehat{G}$ denote the character group of G. A character $\psi\colon G \to \mathbb{C}$ satisfies $\psi(x+y) = \psi(x)\psi(y)$ and $|\psi(x)| = 1$, with $\psi^{-1}(x) = \overline{\psi(x)}$. The principal character $\hat{\mathbf{1}}$ is the constant function 1. The characters form a canonical orthonormal basis of $L_2(\mathsf{U}_G)$, abbreviated to $L_2(G)$; the nonprincipal characters form a basis of its mean zero subspace. Thus, in Section 2.1, we may take $B_0 = \widehat{G} \setminus \{\hat{\mathbf{1}}\}$ and $B = \widehat{G}$. We refer to [Ter99, Bab02] for the standard background in Fourier analysis.

For $\chi = (\chi_x)_{x\in\mathcal{X}} \in \widehat{G}^{\mathcal{X}}$, the corresponding product basis function from Section 2.1 is

$$\chi(f) := \Phi_\chi(f) = \prod_{x\in\mathcal{X}} \chi_x(f(x)), \quad \mathrm{supp}(\chi) = \{x \in \mathcal{X} : \chi_x \ne \hat{\mathbf{1}}\}, \quad |\chi| := |\mathrm{supp}(\chi)|.$$

For $h\colon G^{\mathcal{X}} \to \mathbb{C}$, define its Fourier coefficient at χ by

$$\widehat{h}(\chi) := \langle h, \chi\rangle = \mathbf{E}_{f\sim\mathsf{U}_{G^{\mathcal{X}}}}[h(f)\overline{\chi(f)}].$$

Fourier inversion and Parseval's identity give

$$h = \sum_{\chi\in\widehat{G}^{\mathcal{X}}} \widehat{h}(\chi)\chi, \qquad \|h\|_2^2 = \sum_{\chi\in\widehat{G}^{\mathcal{X}}} |\widehat{h}(\chi)|^2. \tag{7}$$

The support projections from Section 2.1 become

$$\Pi_S h = \sum_{\mathrm{supp}(\chi)=S} \widehat{h}(\chi)\chi, \qquad h_{=t} = \sum_{|\chi|=t} \widehat{h}(\chi)\chi.$$

Thus $h_{=t}$ is the level t Fourier component. We call $\{(\psi,\dots,\psi) : \psi \in \widehat{G}\}$ the diagonal subgroup of $\widehat{G}^{\mathcal{X}}$ and write $\hat{\mathbf{1}}^{(\mathcal{X})}$ for the principal character on $G^{\mathcal{X}}$. For $\ell > 0$,

$$\sum_{|\chi|=t} |\widehat{h}(\chi)|^{\ell}$$

is the ℓ-th Fourier moment of h at level t. For functions h_1, h_2 on $G^{\mathcal{X}}$, define their convolution by

$$(h_1 * h_2)(f) \coloneqq \mathbf{E}_{g\sim \mathsf{U}_{G^{\mathcal{X}}}}[h_1(g)h_2(f-g)].$$

Then

$$\widehat{h_1 * h_2}(\chi) = \widehat{h}_1(\chi)\widehat{h}_2(\chi), \qquad \widehat{h^{*k}}(\chi) = \widehat{h}(\chi)^k.$$

Walsh masks. When $G = \mathbb{F}_2^n$, we identify G with its character group through

$$\chi_w(x) = (-1)^{\langle w,x\rangle}, \qquad \chi_u\chi_v = \chi_{u+v}.$$

Here $\langle\cdot,\cdot\rangle$ is the binary dot product. We call w a *mask* and use the same term for tuples of masks. The principal character corresponds to 0. For a linear map L between binary vector spaces,

$$\chi_\eta(Lx) = \chi_{L^\top\eta}(x),$$

so pulling a character back through L corresponds to applying $L^\top$ to its mask.

2.4 Injection Densities and Fourier Coefficients

Let Ω be a finite ambient set and let $\varnothing \neq A \subseteq \Omega$. Write 1_A for its indicator and

$$\mu_A = \frac{|\Omega|}{|A|}1_A$$

for the density of U_A relative to the ambient measure U_Ω.

Uniform Density on Injections. For $0 \le d \le N$, let $\mathcal{I}_d$ denote the set of injections $[d] \to G$. Its density relative to U_{G^d} is given by

$$\mu_{\mathcal{I}_d} = \frac{N^d}{(N)_d}1_{\mathcal{I}_d}.$$

We also refer to this as the injection density over G^d. Fixing an ordering of G identifies $\mathcal{I}_N$ with the set $\mathcal{S}$ of permutations of G, and G^N with the complete function space G^G. In particular,

$$\mu_{\mathcal{S}} = \frac{N^N}{N!}1_{\mathcal{S}}.$$

Restricting a uniform element of $\mathcal{I}_d$ to any t-subset $S \subseteq [d]$ gives a uniform element of $\mathcal{I}_t$. Indeed, each injective assignment of t distinct outputs to S has $(N-t)_{d-t}$ extensions to an injection on $[d]$, and $(N)_d = (N)_t(N-t)_{d-t}$. Accordingly, if $\chi \in \widehat{G}^d$ has support $S = \{i_1,\dots,i_t\}$, with $i_1 < \cdots < i_t$, then

$$\widehat{\mu_{\mathcal{I}_d}}(\chi) = \widehat{\mu_{\mathcal{I}_t}}(\chi_{i_1},\dots,\chi_{i_t}). \tag{8}$$

Stated differently, inserting principal characters does not change the corresponding Fourier coefficient of the injection density. We also use (8) when $d = N$, with $\mu_{\mathcal{I}_N}$ identified with $\mu_{\mathcal{S}}$.

Two further symmetries will be useful. Negating every output preserves the injection distribution, so its Fourier coefficients are real. Translating every output by a fixed $z \in G$ also preserves this distribution. Consequently,

$$\widehat{\mu_{\mathcal{I}_d}}(\chi) = 0 \quad \text{unless} \quad \prod_{i=1}^{d}\chi_i = \hat{\mathbf{1}}.$$

Indeed, translation multiplies the coefficient by $\overline{\prod_i \chi_i(z)}$, which must be one for every z if the coefficient is nonzero.

Since permuting the input coordinates preserves the uniform injection distribution, the fixed-support moments of the injection density depend only on the size of the support. For $1 \le t \le N$ and $\ell > 0$, we write

$$\Theta_{t,\ell} := \sum_{\alpha \in (\widehat{G}\setminus\{\hat{\mathbf{1}}\})^t} \left|\widehat{\mu_{I_t}}(\alpha)\right|^\ell, \qquad \Lambda_t := \max_{\alpha \in (\widehat{G}\setminus\{\hat{\mathbf{1}}\})^t} \left|\widehat{\mu_{I_t}}(\alpha)\right|,$$

and abbreviate $\Theta_t := \Theta_{t,2}$. Thus, $\Theta_{t,\ell}$ is the ℓ-th Fourier moment of μ_{I_d} on one fixed support of size t, while Λ_t is the corresponding maximum Fourier coefficient. Accordingly, for every $1 \le t \le d \le N$,

$$\sum_{\substack{\chi \in \widehat{G}^d \\ |\chi| = t}} \left|\widehat{\mu_{I_d}}(\chi)\right|^\ell = \binom{d}{t}\Theta_{t,\ell}, \qquad \text{and} \qquad \sum_{\substack{\chi \in \widehat{G}^d \\ \chi \neq \hat{\mathbf{1}}^{(d)}}} \left|\widehat{\mu_{I_d}}(\chi)\right|^\ell = \sum_{t=1}^{d} \binom{d}{t}\Theta_{t,\ell},$$

where the first quantity is the ℓ-th Fourier moment of μ_{I_d} at level t. The quantities Θ_t and $\binom{d}{t}\Theta_t$ are referred to as the Fourier weights for μ_{I_t} and μ_{I_d} at level t in [Din24], respectively. For $k \ge 2$, bounding all but two factors by the maximum gives

$$\Theta_{t,2k} \le \Lambda_t^{2k-2}\Theta_t.$$

We shall also use Eberhard's coefficient-merging identity [Ebe17, Section 4]. For $2 \le t \le N$ and nonprincipal characters $\chi_1, \ldots, \chi_t \in \widehat{G}$,

$$\widehat{\mu_{I_t}}(\chi_1, \ldots, \chi_t) = -\frac{1}{N-t+1}\sum_{i=1}^{t-1} \widehat{\mu_{I_{t-1}}}(\chi_1, \ldots, \chi_i\chi_t, \ldots, \chi_{t-1}). \tag{9}$$

Condition on the first $t-1$ distinct outputs $y_1, \ldots, y_{t-1}$. Since χ_t is nonprincipal,

$$\mathbf{E}\left[\overline{\chi_t(y_t)} \mid y_1, \ldots, y_{t-1}\right] = -\frac{1}{N-t+1}\sum_{i=1}^{t-1} \overline{\chi_t(y_i)}.$$

Multiplying by $\prod_{j=1}^{t-1}\overline{\chi_j(y_j)}$ and taking expectation over the uniform injection on the first $t-1$ coordinates gives the identity.

3 Sum of Independent Permutations

Given $k \ge 2$ independent uniform random permutations $\pi_1, \ldots, \pi_k$ of G, define

$$\pi_+^k(x) := \pi_1(x) + \cdots + \pi_k(x), \qquad x \in G.$$

For $G = \mathbb{F}_2^n$, this is the usual k-XOR-of-permutations construction. We use k for the number of permutations and reserve t for a Fourier level. The density on any fixed set of $d < N$ distinct inputs is

$$\varphi_{d,k} = \mu_{I_d}^{*k},$$

so its Fourier coefficients are $\widehat{\mu_{I_d}}(\chi)^k$. We first record the two global estimates that will be used throughout the section, starting with Eberhard's result for $k = 2$.

Proposition 1 (Theorem 1.5 in [Ebe17]). *Let G be a finite abelian group of order N. For every $1 \le d < N$,*

$$\sum_{t=1}^{d}\binom{d}{t}\Theta_{t,4} = \sum_{\chi \in \widehat{G}^d \setminus \{\hat{\mathbf{1}}^{(d)}\}} |\widehat{\mu_{I_d}}(\chi)|^4 = \|\mu_{I_d}^{*2} - 1\|_2^2 \le O\left(\frac{d^2}{N^3}\right).$$

Proof. The inequality follows directly from [Ebe17, Theorem 1.5]. The two equalities follow from Parseval, convolution-multiplication duality, and a grouping of the characters by their Fourier level. □

We next record a slight generalisation of the smoothing argument used by Dinur that isolates the only additional step needed for $k > 2$.

Proposition 2. *For every $1 \le d < N$, every $k \ge 2$, and every set $\Gamma \subseteq \widehat{G}^d \setminus \{\hat{\mathbf{1}}^{(d)}\}$,*

$$\sum_{\chi\in\Gamma} |\widehat{\mu_{I_d}}(\chi)|^{2k} \le \lambda_N^{2k-4} \sum_{\chi\in\Gamma} |\widehat{\mu_{I_d}}(\chi)|^4, \tag{10}$$

where $\lambda_N := \binom{N}{2}^{-1/2}$. In particular,

$$\sum_{t=1}^{d} \binom{d}{t} \Theta_{t,2k} \le O_k\left(\frac{d^2}{N^{2k-1}}\right). \tag{11}$$

Proof. Fix $\chi \neq \hat{\mathbf{1}}^{(d)}$, and pad it with principal characters to a character $\chi' \in \widehat{G}^N$. By (8),

$$\widehat{\mu_{I_d}}(\chi) = \widehat{\mu_S}(\chi').$$

Since $d < N$ and χ is nonprincipal, the padded tuple has both a principal and a nonprincipal coordinate, and is not diagonal. Eberhard's non-diagonal coefficient bound [Ebe17, Proof of Theorem 1.4, p. 18], after multiplication by $N^N/N!$ to pass to densities, gives $|\widehat{\mu_{I_d}}(\chi)| \le \lambda_N$. Hence

$$|\widehat{\mu_{I_d}}(\chi)|^{2k} \le \lambda_N^{2k-4} |\widehat{\mu_{I_d}}(\chi)|^4.$$

Summing over Γ proves (10). Taking all nonprincipal characters and applying Proposition 1 gives (11). □

Remark 2. The same bound also applies to the complete permutation density μ_S for every character $\chi \in \widehat{G}^G$ with $1 \le |\chi| < N$. Indeed, if $S = \mathrm{supp}(\chi)$ and $|S| = t < N$, then the restriction of a uniform permutation to S is a uniform injection, and hence $\widehat{\mu_S}(\chi) = \widehat{\mu_{I_t}}(\chi|_S)$.

The next estimate controls the levels above two in both query models. The parameter m counts the input positions over which these levels are summed: we use $m = q$ for the classical case and $m = N$ for the quantum case.

Lemma 3. *Let G be a finite abelian group of order $N \ge 1024$. For every integer $3 \le T \le 8N/15$,*

$$\sum_{t=3}^{T} \binom{N}{t} \Theta_{t,4} < \frac{13}{4N^2}, \tag{12}$$

and

$$\max_{3\le t\le T} \Lambda_t \le \binom{N}{4}^{-1/2}. \tag{13}$$

Consequently, for integers $k \ge 2$ and $T \le m \le N$,

$$\sum_{t=3}^{T} \binom{m}{t} \Theta_{t,2k} \le \frac{\binom{m}{3}}{\binom{N}{3}} \frac{13}{4N^2} \binom{N}{4}^{-(k-2)}. \tag{14}$$

The proof is given in Appendix B.

3.1 Refined Classical Security

The $O_k(q/N^{k-1/2})$ bound for every $q < N$ follows from the preceding smoothing argument. To improve the bound below the birthday threshold, we first compute the two-point distribution and then apply the moment bound to the levels above two.

Lemma 4. *Let $x \neq x'$ and let $F = \pi_+^k$. The density $\varphi_{2,k}$ of $(F(x), F(x'))$ on G^2 satisfies*

$$\varphi_{2,k}(y, y') - 1 = \eta_k\big(N\mathbf{1}_{\{y\}}(y') - 1\big), \qquad \eta_k = \left(-\frac{1}{N-1}\right)^k.$$

If $\tau\colon G \to H$ is a surjective homomorphism and $|H| = M$, the two-point density $\nu_{2,k}$ of τF satisfies

$$\nu_{2,k}(y, y') - 1 = \eta_k\big(M\mathbf{1}_{\{y\}}(y') - 1\big). \tag{15}$$

Proof. For each permutation, put $Z_i = \pi_i(x') - \pi_i(x)$. The variable $\pi_i(x)$ is uniform on G, Z_i is uniform on $G \setminus \{0\}$, and the two are independent. The pairs are independent across i. Hence $F(x)$ is uniform and independent of $S_k = Z_1 + \cdots + Z_k$, while $F(x') = F(x) + S_k$. If $p_j(z) = \Pr[S_j = z]$, then convolution with the uniform distribution on $G \setminus \{0\}$ gives

$$p_{j+1}(z) = \frac{1 - p_j(z)}{N - 1}.$$

Subtracting $1/N$ and iterating from $p_1(0) = 0$, $p_1(z) = 1/(N-1)$ for $z \neq 0$, yields

$$p_k(z) = \frac{1}{N} + \eta_k \left(1_{\{0\}}(z) - \frac{1}{N} \right).$$

The claimed density identity follows from $\varphi_{2,k}(y, y') = N^2 \Pr[F(x) = y, F(x') = y'] = N p_k(y' - y)$. Pushing forward through τ sends uniform measure on G to uniform measure on H, yielding (15). □

For a fixed pair $\{x, x'\}$, the corresponding support-two density projection therefore has

$$\|\Pi_{\{x,x'\}} \varphi\|_2^2 = (N-1)|\eta_k|^2,$$

and its total-variation contribution is

$$|\eta_k| \left(1 - \frac{1}{N} \right) = \frac{1}{N(N-1)^{k-1}}.$$

Consequently, after restricting to a fixed set Q of q inputs (by Remark 1),

$$|\Delta_2| \leq \min \left\{ \frac{\binom{q}{2}}{N(N-1)^{k-1}}, \frac{|\eta_k|}{2} \sqrt{\binom{q}{2}(N-1)} \right\}. \tag{16}$$

The first estimate uses the pairwise total variation and the triangle inequality: since $\Pi_{\{x,x'\}} \varphi$ has mean zero, $|\langle \Pi_{\{x,x'\}} a, \Pi_{\{x,x'\}} \varphi \rangle| = |\langle a - \frac{1}{2}, \Pi_{\{x,x'\}} \varphi \rangle| \leq \frac{1}{2} \|\Pi_{\{x,x'\}} \varphi\|_1$, as in (5). The second is Cauchy-Schwarz on level two.

Theorem 2. *Fix $k \geq 2$, and let $\mathsf{g} \colon G \to G$ be uniform. For every $q < N$,*

$$\left\| \pi_+^k - \mathsf{g} \right\|_{(q)} \leq \begin{cases} 0, & 0 \leq q \leq 1, \\ O_k(q^2/N^k), & 2 \leq q \leq \sqrt{N}, \\ O_k(q/N^{k-1/2}), & \sqrt{N} < q < N. \end{cases}$$

In particular, $\left\| \pi_+^k - \mathsf{g} \right\|_{(q)} \leq O_k(q/N^{k-1/2})$ throughout $q < N$.

Proof. The case $q = 0$ is immediate. For the global estimate with $q \geq 1$, Remark 1 reduces the analysis to a fixed set of q inputs. The density there is $\varphi_{q,k} = \mu_{I_q}^{*k}$, and Proposition 1, followed by Proposition 2, gives

$$\|\varphi_{q,k} - 1\|_2 \leq O_k \left(\frac{q}{N^{k-1/2}} \right).$$

The statistical-distance bound follows from (1).

For $q \leq 1$, the advantage is zero because every one-point marginal is uniform. Assume now $2 \leq q \leq \sqrt{N}$. By Remark 1, fix a set Q of q queried inputs. Equation (3) removes level one, while (16) gives $O_k(q^2/N^k)$ at level two. For $N \geq 1024$ and $q \geq 3$, apply (14) with $m = T = q$. Lemma 1 then bounds the remainder by

$$O_k \left(\frac{q^{3/2}}{N^{2k-3/2}} \right) = O_k \left(\frac{q^2}{N^k} \right).$$

For $q = 2$, the remainder is empty. The finitely many smaller group orders are absorbed into the implicit constant. □

3.2 Quantum Security

We first state the bounds for the full range $2q < N$. They combine a direct analysis of the low-support Fourier components with the simulation argument of Section 4.

Theorem 3. *Let G be a finite abelian group of order N, and let $\mathsf{g}\colon G \to G$ be uniform. For $1 \le q \le (N-1)/2$,*

$$\left\| |\pi_+^2\rangle - |\mathsf{g}\rangle \right\|_{(q)} \le O\left(\min\left\{N^{-1/2}, \frac{q^3}{N^2} + \frac{1}{N}\right\}\right).$$

For every fixed $k \ge 3$,

$$\left\| |\pi_+^k\rangle - |\mathsf{g}\rangle \right\|_{(q)} \le O_k\left(\min\left\{\frac{q^3}{N^k}, \frac{1}{\sqrt{N^{2k-3}}}\right\}\right).$$

For $q = 1$, the $k = 2$ bound sharpens to $O(N^{-2})$.

The bounds independent of q follow from Corollary 2. Appendix C treats level two using Zhandry's small-range distributions and proves the query-dependent bounds, including the one-query refinement.

The direct Fourier calculation also gives explicit constants. Lemma 3 applies up to support size $8N/15$, which allows $q \le 4N/15$ quantum queries. This extends the range $q \le N/4$ obtained from moment estimates restricted to support size at most $N/2$.

Corollary 1. *Let G be a finite abelian group of order $N \ge 1024$, fix $k \ge 2$, and let $\mathsf{g}\colon G \to G$ be uniform. For $0 \le q \le 4N/15$, we have*

$$\left\| |\pi_+^k\rangle - |\mathsf{g}\rangle \right\|_{(q)} \le \sqrt{\frac{N}{8(N-1)^{2k-2}} + \frac{13}{16N^2}\binom{N}{4}^{-(k-2)}} < 0.355\,\frac{\sqrt{N}}{(N-1)^{k-1}}.$$

Proof. The case $q = 0$ is immediate, so assume $q \ge 1$. Regrouping the complete-table Fourier coefficients by support, and using Remark 2 and $\Theta_1 = 0$, gives

$$\sum_{1 \le |\chi| \le 2q} |\widehat{\mu_S}(\chi)|^{2k} = \sum_{t=2}^{2q} \binom{N}{t} \Theta_{t,2k}.$$

The exact support-two contribution is $N/(2(N-1)^{2k-2})$, by Lemma 4. Equation (14), with $m = N$ and $T = 2q$, bounds the remaining contribution; this remainder is zero when $q = 1$. Consequently,

$$\sum_{1 \le |\chi| \le 2q} |\widehat{\mu_S}(\chi)|^{2k} \le \frac{N}{2(N-1)^{2k-2}} + \frac{13}{4N^2}\binom{N}{4}^{-(k-2)}.$$

Theorem 1 gives the first inequality. For the second, $(N-1)^2 \le \binom{N}{4}$ when $N \ge 1024$, so the square of the bound, after division by $N/(N-1)^{2k-2}$, is at most

$$\frac{1}{8} + \frac{13(N-1)^2}{16N^3}\left(\frac{(N-1)^2}{\binom{N}{4}}\right)^{k-2} \le \frac{1}{8} + \frac{13}{16 \cdot 1024} < 0.355^2. \qquad \square$$

3.3 Quantum Attacks over $\mathbb{F}_2^n$

We now take $G = \mathbb{F}_2^n$ and $N = 2^n$. For distinct $x, x' \in G$ and nonzero $\lambda \in G$, there exists a one-query quantum algorithm that returns $b_\lambda(x, x') = \langle \lambda, f(x) \oplus f(x') \rangle$ with certainty. This is the usual two-point phase-query algorithm, written with an arbitrary output mask; see also Bonnetain et al. [BLNPS21]. The algorithm and its proof are given in Appendix D.

Proposition 3. *For every $k \ge 2$, there exists a one-query quantum algorithm that distinguishes $|\pi_+^k\rangle$ from $|\mathsf{g}\rangle$ with advantage $1/(2(N-1)^k)$.*

Proof. Fix distinct $x, x' \in G$ and nonzero $\lambda \in G$. The one-query algorithm computes $b_\lambda(x, x')$. Use the acceptance function $a(f) = \frac{1}{2}(1 + (-1)^{b_\lambda(x,x')})$, which accepts exactly when $b_\lambda(x, x') = 0$. This bit is uniform under g. By Lemma 4, $\mathbf{E}[(-1)^{b_\lambda(x,x')}] = (-1/(N-1))^k$ under π_+^k. Hence the distinguishing advantage is $1/(2(N-1)^k)$. $\square$

Proposition 4. *For $n \geq 2$ and $k \geq 2$, there exists an $N/2$-query quantum algorithm that distinguishes $|\pi_+^k\rangle$ from $|\mathsf{g}\rangle$ with advantage $1/2$.*

Proof. Fix nonzero $\lambda \in G$ and partition G into $N/2$ disjoint pairs. Applying the one-query algorithm to each pair and XORing the outputs computes $\bigoplus_{x \in G} \langle \lambda, f(x) \rangle$ using $N/2$ queries. For $f = \pi_+^k$, this bit is zero because every permutation ranges over G and $\bigoplus_{y \in G} y = 0$ for $n \geq 2$. Under g, it is uniform. Accepting when it is zero gives advantage $1/2$. □

For $G = \mathbb{F}_2^n$ with $n \geq 2$, the one-query attack matches the asymptotic order of our quantum upper bound at $q = 1$. The matching half-domain threshold is discussed in Section 4.

4 Quantum Security by Completion

We next use the checksum invariance property to extend quantum security to every query count below half the domain size. The argument applies more generally when any prescribed values on a fixed number of inputs have a unique completion within a given class of functions. Completing independent uniform values gives a reference function with limited independence, to which the quantum support cutoff applies.

Theorem 4. *Let $\mathcal{X}$ be a finite set of size D, let G be a finite abelian group, and fix $1 \leq d \leq D$. Let $\mathcal{C} \subseteq G^{\mathcal{X}}$ be such that, for every $S \subseteq \mathcal{X}$ of size d, the restriction map*

$$\mathcal{C} \longrightarrow G^S, \qquad f \longmapsto f|_S,$$

is a bijection. Let $F \colon \mathcal{X} \to G$ be a random function taking values in $\mathcal{C}$, and let $\mathsf{g} \colon \mathcal{X} \to G$ be uniform. For every q with $2q \leq d$,

$$\| |F\rangle - |\mathsf{g}\rangle \|_{(q)} \leq \|F - \mathsf{g}\|_{(d)} .$$

Proof. Fix a set S of size d, and let $E_S \colon G^S \to \mathcal{C}$ be the inverse restriction map. Then $E_S(F|_S) = F$, whereas $F_{\mathcal{C}} := E_S(\mathsf{g}|_S)$ is uniform on $\mathcal{C}$. The restriction hypothesis for every d-point set implies that $F_{\mathcal{C}}$ is d-wise independent. Lemma 2 therefore gives

$$\| |F_{\mathcal{C}}\rangle - |\mathsf{g}\rangle \|_{(q)} = 0 \qquad (2q \leq d),$$

equivalently by Zhandry's theorem [Zha12b, Theorem 3.1].

For a fixed q-query distinguisher $\mathcal{A}$, let $a_S(T)$ be its acceptance probability when given the complete oracle $E_S(T)$. A computationally unbounded classical distinguisher can query all points of S and accept with probability $a_S(T)$. Its acceptance probabilities against F and g are exactly those of $\mathcal{A}$ against F and $F_{\mathcal{C}}$. Thus

$$\| |F\rangle - |\mathsf{g}\rangle \|_{\mathcal{A}} = \| |F\rangle - |F_{\mathcal{C}}\rangle \|_{\mathcal{A}} \leq \|F - \mathsf{g}\|_{(d)} .$$

Taking the supremum over $\mathcal{A}$ proves the claim. □

The sum of permutations construction falls in the $d = D - 1$ case. Indeed, we can use its checksum to reconstruct the final output.

Corollary 2. *Let G be a finite abelian group of order N. Fix some $k \geq 2$, and $q \leq (N-1)/2$. Let $\mathsf{g} \colon G \to G$ denote a uniform random function. Then,*

$$\left\| |\pi_+^k\rangle - |\mathsf{g}\rangle \right\|_{(q)} \leq O_k\left(N^{-(k-3/2)}\right).$$

Proof. Let $c_k := k \sum_{y \in G} y$ and define the set

$$\mathcal{C}_k := \left\{ f \in G^G : \sum_{x \in G} f(x) = c_k \right\}.$$

Then, for any $r \in G$ and $S = G \setminus \{r\}$, the restriction map

$$\mathcal{C}_k \longrightarrow G^S, \qquad f \longmapsto f|_S,$$

is a bijection. Indeed, given any $f' \in G^S$, we get a unique $f \in C_k$ defined as follows:

$$f(x) \coloneqq \begin{cases} f'(x), & x \neq r, \\ c_k - \sum_{x' \neq r} f'(x'), & x = r. \end{cases}$$

Every realisation of π_+^k satisfies

$$\sum_{x \in G} \pi_+^k(x) = k \sum_{y \in G} y = c_k,$$

and hence $\pi_+^k \in C_k$. Theorem 4 with $d = N-1$, followed by Theorem 2 at $N-1$ classical queries, gives

$$\left\| |\pi_+^k\rangle - |\mathsf{g}\rangle \right\|_{(q)} \leq \left\| \pi_+^k - \mathsf{g} \right\|_{(N-1)} \leq O_k\left(\frac{N-1}{N^{k-1/2}}\right).$$

This proves the claimed bound for $2q \leq N-1$. □

Sharp query threshold for the binary setting. When $G = \mathbb{F}_2^n$ with $n \geq 2$, Corollary 2 and Proposition 4 meet at consecutive query counts. For every fixed $k \geq 2$, the advantage is $O_k(N^{-(k-3/2)})$ for every $q \leq N/2 - 1$, while an $N/2$-query algorithm has advantage at least $1/2$. Thus $N/2$ is the asymptotic constant-advantage quantum query threshold for the binary construction.

5 Linear-Output-Postprocessing Variants

Throughout this section, $G = \mathbb{F}_2^n$, $N = 2^n$, and we identify characters with Walsh masks. For a mask tuple $\chi \in G^{\mathcal{X}}$, write

$$r(\chi) \coloneqq \dim \operatorname{Span}\{\chi_x : x \in \mathcal{X}\}.$$

For $L \in \mathrm{GL}(n,2)$ and $x = (x_1, \ldots, x_d)$, write $Lx = (Lx_1, \ldots, Lx_d)$. The injection density is invariant under this diagonal action, since L preserves pairwise distinctness. Averaging over the general linear group therefore controls the Fourier mass inside a fixed subspace.

Proposition 5. *Suppose $f : G^d \to \mathbb{R}$ is invariant under the diagonal action of $\mathrm{GL}(n,2)$, and let $\ell > 0$. For every m-dimensional subspace $W \leq G$,*

$$\sum_{\substack{\chi \in W^d \\ \chi \neq 0}} |\widehat{f}(\chi)|^\ell \leq \frac{|W|-1}{N-1} \sum_{\substack{\chi \in G^d \\ \chi \neq 0}} |\widehat{f}(\chi)|^\ell.$$

Proof. For $S \subseteq G^d$, write

$$M_\ell(S) \coloneqq \sum_{\substack{\chi \in S \\ \chi \neq 0}} |\widehat{f}(\chi)|^\ell.$$

Using the invariance of f, we have $\widehat{f}(L^\top \chi) = \widehat{f}(\chi)$ for all $L \in \mathrm{GL}(n,2)$. It follows that $M_\ell(W^d) = M_\ell((L^\top W)^d)$. Averaging over $L \sim \mathsf{U}_{\mathrm{GL}(n,2)}$ gives

$$M_\ell(W^d) = \sum_{\substack{\chi \in G^d \\ \chi \neq 0}} |\widehat{f}(\chi)|^\ell \Pr_L\left[\operatorname{Span}\{\chi_1, \ldots, \chi_d\} \subseteq L^\top W\right]. \tag{17}$$

We claim that

$$\Pr\left(\operatorname{Span}\{\chi_1, \ldots, \chi_d\} \subseteq L^\top W\right) \leq \frac{|W|-1}{|G|-1}. \tag{18}$$

Let $r = r(\chi)$. The subspace $L^\top W$ is uniform among the m-dimensional subspaces of G. For $r \leq m$, its probability of containing a fixed r-dimensional subspace is

$$p_{n,m}(r) = \frac{\binom{n-r}{m-r}_2}{\binom{n}{m}_2} = \prod_{i=0}^{r-1} \frac{2^m - 2^i}{2^n - 2^i}, \tag{19}$$

and $p_{n,m}(r) = 0$ for $r > m$. Here $\binom{}{}_2$ denotes the Gaussian binomial coefficient. Since $p_{n,m}(r) \leq p_{n,m}(1) = (|W|-1)/(N-1)$ for $r \geq 1$, claim (18) follows. Substituting it in (17) proves the result. □

The same argument applies after restricting to any set of masks invariant under the diagonal action; in particular, support size and rank may be fixed simultaneously.

5.1 Generalised and Truncated Sums of Permutations

Fix $0 \le m \le n$, let $H = \mathbb{F}_2^m$, $M = 2^m$, and let $\tau\colon G \to H$ be a surjective linear map. Define

$$\tau\pi_+^k(x) \coloneqq \tau(\pi_+^k(x)) = \tau(\pi_1(x)) + \cdots + \tau(\pi_k(x)).$$

Projection onto m fixed coordinates recovers the truncated sum of permutations studied in [CKLL22]. Proposition 5 and (11) give the following improved classical indistinguishability bound for $\tau\pi_+^k$.

Theorem 5. *Fix $k \ge 2$ and $q < N$, and let $\mathsf{g}\colon G \to H$ be uniform. Then*

$$\left\|\tau\pi_+^k - \mathsf{g}\right\|_{(q)} \le O_k\left(\frac{q\sqrt{M}}{N^k}\right).$$

Proof. The cases $q = 0$ and $M = 1$ are immediate. By Remark 1, fix a set of q distinct inputs, and let ν_k be the density of the construction on H^q. Character pullback gives

$$\widehat{\nu_k}(\eta) = \widehat{\mu_{\mathcal{I}_q}}(\tau^\top\eta)^k. \tag{20}$$

Since τ is surjective, $\tau^\top$ identifies H with the subspace $W = \tau^\top(H) \le G$, of size M. Parseval's identity, Proposition 5, and (11) give

$$\|\nu_k - 1\|_2^2 = \sum_{\chi \in W^q \setminus \{0\}} |\widehat{\mu_{\mathcal{I}_q}}(\chi)|^{2k} \le \frac{M-1}{N-1} O_k\left(\frac{q^2}{N^{2k-1}}\right) = O_k\left(\frac{q^2 M}{N^{2k}}\right).$$

The statistical-distance bound proves the claim. □

For $k = 2$, this improves the $O(q\sqrt{qM}/N^2)$ single-user bound of [CKLL22] by a factor $\sqrt{q}$ in their common parameter range.

5.2 Quantum Security

The quantum bounds retain the dependence on the output size M. As in the unprojected case, we combine a bound from completion with a more precise treatment of level two. Keeping the rank of the pulled-back masks improves the bound for the remaining levels.

Theorem 6. *Let $\mathsf{g}\colon G \to H$ be uniform. For $1 \le q \le (N-1)/2$,*

$$\left\||\tau\pi_+^2\rangle - |\mathsf{g}\rangle\right\|_{(q)} \le O\left(\min\left\{\frac{\sqrt{M}}{N}, \frac{q^3+M}{N^2}\right\}\right). \tag{21}$$

For every fixed $k \ge 3$,

$$\left\||\tau\pi_+^k\rangle - |\mathsf{g}\rangle\right\|_{(q)} \le O_k\left(\min\left\{\frac{q^3}{N^k}, \frac{\sqrt{M}}{N^{k-1}}\right\}\right).$$

For $q = 1$, the $k = 2$ bound sharpens to $O(N^{-2})$. For $M = 1$, the advantage is zero.

The proof is given in Appendix C.3. We record the simulation bound and the concrete Fourier bound below.

The projected checksum determines the final output from the other $N-1$ outputs, so Theorem 4 applies.

Corollary 3. *Let $G = \mathbb{F}_2^n$, $H = \mathbb{F}_2^m$, $N = 2^n$, $M = 2^m$, and let $\tau\colon G \to H$ be surjective and linear. Fix $k \ge 2$, and let $\mathsf{g}\colon G \to H$ be uniform. For every $q \le (N-1)/2$,*

$$\left\||\tau\pi_+^k\rangle - |\mathsf{g}\rangle\right\|_{(q)} \le O_k\left(\frac{\sqrt{M}}{N^{k-1}}\right).$$

Proof. Every realisation $f = \tau\pi_+^k$ satisfies $\sum_{x\in G} f(x) = \tau(k\sum_{y\in G} y)$. Any $N-1$ values determine the remaining value uniquely. Theorem 4, with output group H and $d = N-1$, followed by Theorem 5, gives

$$\left\||\tau\pi_+^k\rangle - |\mathsf{g}\rangle\right\|_{(q)} \le \left\|\tau\pi_+^k - \mathsf{g}\right\|_{(N-1)} \le O_k\left(\frac{(N-1)\sqrt{M}}{N^k}\right).$$

□

For explicit constants, we restrict the Fourier mass in Corollary 1 to masks in the image of $\tau^\top$.

Corollary 4. *Assume $N \geq 1024$, fix $k \geq 2$, and let $\mathsf{g}\colon G \to H$ be uniform. For $0 \leq q \leq 4N/15$, we have*

$$\left\| |\tau\pi_+^k\rangle - |\mathsf{g}\rangle \right\|_{(q)} \leq \sqrt{\frac{M-1}{N-1}\left(\frac{N}{8(N-1)^{2k-2}} + \frac{13}{16N^2}\binom{N}{4}^{-(k-2)}\right)}.$$

Proof. For the complete function table, the pullback identity (20) holds with $\mathcal{I}_q$ replaced by $\mathcal{S}$. The map $\tau^\top$ preserves support. Proposition 5, restricted to levels at most $2q$, therefore multiplies the Fourier mass used in Corollary 1 by at most $(M-1)/(N-1)$. Theorem 1 proves the claim. □

6 The Variable-Output LXoP Family

Throughout this section, $G = \mathbb{F}_2^n$ and $N = 2^n$. We identify characters with Walsh masks. Iwata's XORP$[w]$ construction [Iwa06] produces w output blocks from $w+1$ domain-separated permutation calls. Dinur defined [Din25] the following LXoP family and proposed the admissibility condition below, leaving its general-w security analysis to future work. We use our notation for this family and give classical and quantum bounds for every fixed w.

Fix $w \geq 1$, put $s = \lceil \log_2(w+1) \rceil$, and assume $s < n$. The construction has input domain $\mathbb{F}_2^{n-s}$, of size $D = N/2^s$, and produces w elements of G on each input. For $j \in \{0, \ldots, w\}$, let

$$\iota_j\colon \mathbb{F}_2^{n-s} \longrightarrow G, \qquad \iota_j(x) = \langle j \rangle_s \parallel x,$$

where $\langle j \rangle_s$ denotes the s-bit binary representation of j. We call $\sigma \in \mathrm{GL}(n,2)$ *w-admissible* if

$$I + \sigma^r \in \mathrm{GL}(n,2) \qquad \text{for every } 1 \leq r \leq w. \tag{22}$$

For $w = 1$, this is the linear orthomorphism condition on σ. The same condition suffices for $w = 2$, since $I + \sigma^2 = (I+\sigma)^2$ over $\mathbb{F}_2$.

Given a uniform random permutation π of G, define $\pi_+^{\sigma,w}\colon \mathbb{F}_2^{n-s} \to G^w$ by

$$\pi_+^{\sigma,w}(x) \coloneqq \big(\pi(\iota_0(x)) + \sigma(\pi(\iota_1(x))), \ldots, \pi(\iota_{w-1}(x)) + \sigma(\pi(\iota_w(x)))\big). \tag{23}$$

In particular, $\pi_+^{\sigma,1}$ and $\pi_+^{\sigma,2}$ are respectively the constructions LXoP$[\sigma, n]$ and LXoP$[\sigma, 2, n]$ of Dinur.

The condition (22) is non-vacuous for every fixed w. For example, after identifying G with $\mathbb{F}_{2^n}$, multiplication by a primitive field element gives a w-admissible linear map whenever $w < 2^n - 1$.

The domain separators ensure that all $(w+1)q$ permutation inputs used by q distinct construction inputs are distinct. Thus their outputs form a uniform injection, even though the w output blocks at one input share permutation calls. We must account for these shared calls when bounding the Fourier moments. The resulting classical bound is as follows.

Theorem 7. *Fix $w \geq 1$. There exists a constant $c_w > 0$, depending only on w, such that the following holds. Let σ be w-admissible, let $\mathsf{g}_w\colon \mathbb{F}_2^{n-s} \to G^w$ be a uniform random function, and let $1 \leq q \leq c_w D$. Then*

$$\left\| \pi_+^{\sigma,w} - \mathsf{g}_w \right\|_{(q)} \leq O_w\left(\frac{q}{N^{3/2}}\right).$$

We prove the theorem using a fixed-support moment bound that will also be used for quantum queries. The full recursive calculation is given in Appendix E; here we describe the masks to which it applies.

Pulling an output character back through the linear combining map gives a character on the underlying permutation outputs. Write $A = \sigma^\top$ and define

$$T_{\sigma,w}(\alpha_1, \ldots, \alpha_w) \coloneqq \big(\alpha_1, A\alpha_1 + \alpha_2, \ldots, A\alpha_{w-1} + \alpha_w, A\alpha_w\big) \in G^{w+1}. \tag{24}$$

For example, for $w = 2$ the output character (α_1, α_2) pulls back to $(\alpha_1, A\alpha_1 + \alpha_2, A\alpha_2)$. Extending the notation, given a mask $\alpha = (\alpha^{(1)}, \ldots, \alpha^{(q)}) \in (G^w)^q$, where $\alpha^{(i)} = (\alpha_{i,1}, \ldots, \alpha_{i,w}) \in G^w$, we apply $T_{\sigma,w}$ separately to each tuple $\alpha^{(i)}$ and use the same notation for the resulting map:

$$T_{\sigma,w}(\alpha) \coloneqq \big(T_{\sigma,w}(\alpha^{(1)}), \ldots, T_{\sigma,w}(\alpha^{(q)})\big) \in (G^{w+1})^q.$$

We distinguish between construction masks and injection masks. A construction mask $\alpha = (\alpha^{(1)}, \dots, \alpha^{(q)}) \in (G^w)^q$ assigns one tuple of labels to each construction input. Its *construction level* is the number of inputs for which $\alpha^{(i)} \neq 0$. Here a nonzero tuple may still have some zero entries. Its pullback $\beta = T_{\sigma,w}(\alpha) \in (G^{w+1})^q$ is an injection mask, with one label for each underlying permutation-output position. Its *injection level* is the number of nonzero entries of β. Writing

$$\beta = (\beta^{(1)}, \dots, \beta^{(q)}), \qquad \beta^{(i)} = (\beta_{i,0}, \dots, \beta_{i,w}) \in G^{w+1},$$

we call $\beta^{(i)}$ the i-th *block* of the injection mask.

Let $\xi_{n,q}^{(\sigma,w)}$ denote the density on any fixed set of q distinct construction inputs. Expanding each output character through the linear combining map gives

$$\widehat{\xi_{n,q}^{(\sigma,w)}}(\alpha) = \widehat{\mu_{\mathcal{I}_{(w+1)q}}}\big(T_{\sigma,w}(\alpha)\big). \tag{25}$$

Indeed, $\langle \alpha, \sigma(y) \rangle = \langle A\alpha, y \rangle$, so the character on one block has labels $(\alpha_1, A\alpha_1 + \alpha_2, \dots, A\alpha_w)$ on the corresponding permutation outputs. Those outputs form a uniform injection.

For $1 \le t \le D$, write

$$L_t := \sum_{\substack{\alpha \in (G^w)^t \\ \alpha^{(i)} \neq 0 \text{ for all } i}} |\widehat{\xi_{n,t}^{(\sigma,w)}}(\alpha)|^2 = \sum_{\substack{\alpha \in (G^w)^t \\ \alpha^{(i)} \neq 0 \text{ for all } i}} |\widehat{\mu_{\mathcal{I}_{(w+1)t}}}\big(T_{\sigma,w}(\alpha)\big)|^2. \tag{26}$$

Thus L_t is the second Fourier moment on one fixed support of t construction inputs. The corresponding injection coefficients need not all lie at the same level, because some coordinates can vanish after applying $T_{\sigma,w}$. We record the structure of the pullback before stating the moment bound. For one block, write

$$\beta = T_{\sigma,w}(\alpha) = (\beta_0, \dots, \beta_w), \qquad C_j = A^{w-j}.$$

By (24)

$$\sum_{j=0}^{w} C_j \beta_j = \sum_{j=0}^{w} A^{w-j} \beta_j = 0. \tag{27}$$

Indeed, after substitution, each α_j occurs twice and cancels. Conversely, the original labels are recovered by $\alpha_1 = \beta_0$ and $\alpha_{j+1} = \beta_j + A\alpha_j$ for $1 \le j < w$. Hence $T_{\sigma,w}$ is injective. Since every C_j is invertible, a nonzero construction block $\alpha^{(i)}$ cannot have a pullback $\beta^{(i)}$ with exactly one nonzero entry: otherwise (27) would force that entry to be zero. Hence each nonzero construction block contributes between 2 and $w+1$ nonzero entries to the injection mask. Consequently, construction level t pulls back to injection levels between $2t$ and $(w+1)t$.

To bound L_t, we partition the pullbacks by their exact injection support and iterate (9). At each step we choose a coordinate in an untouched block. The block relation and w-admissibility let us recover the removed label, so distinct masks remain distinct within each branch. Each merge touches at most two blocks, allowing $\lceil t/2 \rceil$ steps. Appendix E gives the recovery argument and the resulting bounds.

Lemma 5. *Fix $w \ge 1$ and let σ be w-admissible. The fixed-support moments satisfy*

$$L_1 = O_w(N^{-3}). \tag{28}$$

There is a constant C_w, depending only on w, such that for $t \ge 2$ with $(w+1)t \le N/8$,

$$L_t \le \left(C_w \frac{t^{3/2}}{N^{3/2}} \right)^t. \tag{29}$$

The proof is given in Appendix E.

Proof of Theorem 7. Let $\xi_q = \xi_{n,q}^{(\sigma,w)}$, and recall L_t from (26). Each L_t is the contribution of one fixed set of t construction inputs. There are $\binom{q}{t}$ such supports among the q queried inputs. Thus input symmetry and Parseval give

$$\|\xi_q - 1\|_2^2 = \sum_{t=1}^{q} \binom{q}{t} L_t. \tag{30}$$

Put $a = w+1$. Lemma 5 bounds level one and, whenever $aq \le N/8$, every level up to q. Using $\binom{q}{t} \le (eq/t)^t$, we obtain

$$\binom{q}{t} L_t \le \left(C'_w \frac{q\sqrt{t}}{N^{3/2}} \right)^t .$$

The ratio of consecutive terms on the right is at most $C''_w q\sqrt{q+1}/N^{3/2}$, where $C''_w = \sqrt{e}C'_w$ absorbs the factor $((t+1)/t)^{t/2} \le \sqrt{e}$. Choose $c_w > 0$ sufficiently small so that $c_w \le (8a)^{-1}$ and this ratio is at most $1/2$ whenever $q \le c_w D \le c_w N$. Such a choice is possible since $q\sqrt{q+1}/N^{3/2} \le \sqrt{2}(q/N)^{3/2}$ for $q \ge 1$. The same choice ensures $at \le aq \le N/8$ for every $t \le q$. The contribution of all $t \ge 2$ in (30) is then $O_w(q^2/N^3)$. Together with (28),

$$\|\xi_q - 1\|_2^2 \le O_w\left(\frac{q}{N^3} + \frac{q^2}{N^3}\right) = O_w\left(\frac{q^2}{N^3}\right),$$

since $q \ge 1$. Taking square roots and using statistical distance at most one half of the L_2 distance proves the result. □

6.1 Quantum Security

The same fixed-support moments control quantum distinguishing advantage as well. We now count supports in the complete construction domain and sum only the levels up to $2q$, as required by Theorem 1.

Corollary 5. *Fix $w \ge 1$. There exists a sufficiently small constant $c_w > 0$, depending only on w, such that the following holds. Let $s = \lceil \log_2(w+1) \rceil$, let σ be w-admissible, and let $\mathsf{g}_w \colon \mathbb{F}_2^{n-s} \to G^w$ be uniform. For every $0 \le q \le c_w D$,*

$$\left\| |\pi_+^{\sigma,w}\rangle - |\mathsf{g}_w\rangle \right\|_{(q)} \le O_w\left(\frac{1}{\sqrt{N}}\right).$$

Proof. The case $q = 0$ is immediate; assume $q \ge 1$. Let $D = N/2^s$ be the size of the actual domain, let ξ denote the density of the complete $\pi_+^{\sigma,w}$ function, and let L_t denote the fixed-support second Fourier moment from (26). By input symmetry,

$$\sum_{|\alpha|=t} |\widehat{\xi}(\alpha)|^2 = \binom{D}{t} L_t.$$

By Lemma 5, $L_1 = O_w(N^{-3})$ and $L_t \le (C_w t^{3/2}/N^{3/2})^t$ when $(w+1)t \le N/8$. Hence, for $t \ge 2$,

$$\binom{D}{t} L_t \le \left(C'_w \sqrt{\frac{t}{N}} \right)^t =: A_t, \qquad \frac{A_{t+1}}{A_t} \le C''_w \sqrt{\frac{t+1}{N}}.$$

Shrink c_w from Theorem 7, if necessary, so that

$$c_w \le \min\left\{ 2^{-s-1}, \frac{1}{16(w+1)}, \frac{1}{8(C''_w)^2} \right\}.$$

Since $q \le c_w D \le c_w N$, we have $2q \le \min\{D, N/(8(w+1))\}$, so Lemma 5 applies at every level $t \le 2q$, and the ratio is at most $1/2$ for $2 \le t < 2q$. Hence

$$\sum_{t=2}^{2q} \binom{D}{t} L_t \le O_w(A_2) = O_w(1/N),$$

while the level-one contribution is $DL_1 = O_w(N^{-2})$. Therefore

$$\sum_{1 \le |\alpha| \le 2q} |\widehat{\xi}(\alpha)|^2 \le O_w(1/N).$$

The result follows from Theorem 1. □

Concrete bounds for one- and two-block output. For $w \in \{1, 2\}$, Dinur gives sharper fixed-support moments. Applying Theorem 1 to these estimates gives the following concrete bounds. The proof is given in Appendix F.

Corollary 6. *Let $N = 2^n \geq 2^{10}$, and let σ be 1-admissible (equivalently, 2-admissible). For a uniform function $\mathsf{g}\colon \mathbb{F}_2^{n-1} \to G$,*

$$\left\| |\pi_+^{\sigma,1}\rangle - |\mathsf{g}\rangle \right\|_{(q)} \leq \sqrt{\frac{1887}{N}} \qquad (0 \leq q \leq N/64).$$

For a uniform function $\mathsf{g}_2\colon \mathbb{F}_2^{n-2} \to G^2$,

$$\left\| |\pi_+^{\sigma,2}\rangle - |\mathsf{g}_2\rangle \right\|_{(q)} < \sqrt{\frac{37}{2N}} \qquad (0 \leq q \leq N/384).$$

The two constants are not comparable, as they come from different estimates: for $w = 1$ we bound the level sums by a geometric series with a crude ratio, whereas for $w = 2$ we evaluate the first few levels numerically and bound only the tail by a geometric series. The same numerical treatment would also reduce the constant $w = 1$.

Acknowledgements

The authors would like to thank Benoît Cogliati and the GAPS 2025 discussion group — Wonseok Choi, Chandranan Dhar, Ravindra Jejurikar, Jannis Leuther, Jordan Naccache, Abishanka Saha, André Schrottenloher, and Rentaro Shiba — for valuable discussions on the topic. A part of this result was conceived while A.J. was supported (in parts) by the German Research Foundation (DFG) within the framework of the Excellence Strategy of the Federal Government and the States – EXC 2092 CaSa – 390781972, and the Chair of Symmetric Cryptography at Ruhr University Bochum.

Responsible Disclosure of AI Usage. We used ChatGPT 6, GPT-5.6, and Google Search prompts (powered by Gemini). AI tools were primarily used as a writing aid: to polish some passages, check the language, and check consistency throughout the paper. A secondary application was in certain numerical computations appearing in some proofs, where the human authors gave precise prompts to optimise certain functions and compute the resulting bounds, much as a numerical computing tool would be used in this context. All text and all technical claims, AI-assisted or not, were ultimately reviewed and validated by the authors, and they take full responsibility for the content of this paper.

References

AB12. Jean-Philippe Aumasson and Daniel J. Bernstein. Siphash: A fast short-input PRF. In Steven D. Galbraith and Mridul Nandi, editors, *Progress in Cryptology - INDOCRYPT 2012, 13th International Conference on Cryptology in India, Kolkata, India, December 9-12, 2012. Proceedings*, volume 7668 of *Lecture Notes in Computer Science*, pages 489–508. Springer, 2012. doi:10.1007/978-3-642-34931-7_28. *(CITED ON P. 1.)*

Bab02. László Babai. The fourier transform and equations over finite abelian groups. Online Lecture Notes, 2002. URL: https://people.cs.uchicago.edu/~laci/reu02/fourier.pdf. *(CITED ON P. 8.)*

BBC+98. Robert Beals, Harry Buhrman, Richard Cleve, Michele Mosca, and Ronald de Wolf. Quantum lower bounds by polynomials. In *39th Annual Symposium on Foundations of Computer Science, FOCS 1998, Palo Alto, California, USA, November 8-11, 1998*, pages 352–361. IEEE Computer Society, 1998. doi:10.1109/SFCS.1998.743485. *(CITED ON P. 3.)*

Ber97. Daniel J. Bernstein. SURF: simple unpredictable random function, 1997. URL: http://cr.yp.to/papers/surf.pdf. *(CITED ON P. 1.)*

BI99. M. Bellare and R. Impagliazzo. A tool for obtaining tighter security analyses of pseudorandom function based constructions, with applications to PRP to PRF conversion. Cryptology ePrint Archive, Paper 1999/024, 1999. URL: https://eprint.iacr.org/1999/024. *(CITED ON P. 2.)*

BIL+21. Subhadeep Banik, Takanori Isobe, Fukang Liu, Kazuhiko Minematsu, and Kosei Sakamoto. Orthros: A low-latency PRF. *IACR Trans. Symmetric Cryptol.*, 2021(1):37–77, 2021. doi:10.46586/TOSC.V2021.I1.37-77. *(CITED ON P. 1.)*

BKR94. Mihir Bellare, Joe Kilian, and Phillip Rogaway. The Security of Cipher Block Chaining. In *CRYPTO 1994*, volume 839 of *LNCS*, pages 341–358. Springer, 1994. doi:10.1007/3-540-48658-5_32. *(CITED ON P. 1.)*

BKR98. Mihir Bellare, Ted Krovetz, and Phillip Rogaway. Luby-rackoff backwards: Increasing security by making block ciphers non-invertible. In Kaisa Nyberg, editor, *Advances in Cryptology - EUROCRYPT '98, International Conference on the Theory and Application of Cryptographic Techniques, Espoo, Finland, May 31 - June 4, 1998, Proceeding*, volume 1403 of *Lecture Notes in Computer Science*, pages 266–280. Springer, 1998. doi:10.1007/BFB0054132. *(CITED ON PP. 1 AND 2.)*

BLNPS21. Xavier Bonnetain, Gaetan Leurent, Maria Naya-Plasencia, and Andre Schrottenloher. Quantum linearization attacks. In Mehdi Tibouchi and Huaxiong Wang, editors, *Advances in Cryptology – ASIACRYPT 2021, Part I*, volume 13090 of *Lecture Notes in Computer Science*, pages 422–452, Cham, 2021. Springer. doi:10.1007/978-3-030-92062-3_15. *(CITED ON PP. 13 AND 31.)*

BN18. Srimanta Bhattacharya and Mridul Nandi. Revisiting variable output length XOR pseudorandom function. *IACR Transactions on Symmetric Cryptology*, 2018(1):314–335, 2018. doi:10.13154/tosc.v2018.i1.314-335. *(CITED ON P. 2.)*

Che22. Lily Chen. Recommendation for key derivation using pseudorandom functions. Special Publication 800-108r1-upd1, NIST, U.S. Department of Commerce, 2022. doi:10.6028/NIST.SP.800-108r1-upd1. *(CITED ON P. 1.)*

CKLL22. Wonseok Choi, Hwigyeom Kim, Jooyoung Lee, and Yeongmin Lee. Multi-user security of the sum of truncated random permutations. In Shweta Agrawal and Dongdai Lin, editors, *Advances in Cryptology - ASIACRYPT 2022 - 28th International Conference on the Theory and Application of Cryptology and Information Security, Taipei, Taiwan, December 5-9, 2022, Proceedings, Part II*, volume 13792 of *Lecture Notes in Computer Science*, pages 682–710. Springer, 2022. doi:10.1007/978-3-031-22966-4_23. *(CITED ON PP. 2, 3, AND 16.)*

CLP15. Benoît Cogliati, Rodolphe Lampe, and Jacques Patarin. The indistinguishability of the XOR of k permutations. In Carlos Cid and Christian Rechberger, editors, *Fast Software Encryption – FSE 2014*, volume 8540 of *Lecture Notes in Computer Science*, pages 285–302, Berlin, Heidelberg, 2015. Springer. doi:10.1007/978-3-662-46706-0_15. *(CITED ON P. 2.)*

CN08. Donghoon Chang and Mridul Nandi. A Short Proof of the PRP/PRF Switching Lemma. Cryptology ePrint Archive, Report 2008/078, 2008. URL: http://eprint.iacr.org/2008/078. *(CITED ON P. 1.)*

CP20. Benoît Cogliati and Jacques Patarin. Mirror Theory: A simple proof of the $P_i \oplus P_j$ Theorem with $\xi_{\max} = 2$. Cryptology ePrint Archive, Paper 2020/734, 2020. URL: https://eprint.iacr.org/2020/734. *(CITED ON P. 2.)*

CS16. Benoît Cogliati and Yannick Seurin. EWCDM: an efficient, beyond-birthday secure, nonce-misuse resistant MAC. In Matthew Robshaw and Jonathan Katz, editors, *Advances in Cryptology - CRYPTO 2016 - 36th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2016, Proceedings, Part I*, volume 9814 of *Lecture Notes in Computer Science*, pages 121–149. Springer, 2016. doi:10.1007/978-3-662-53018-4_5. *(CITED ON P. 1.)*

DHT17. Wei Dai, Viet Tung Hoang, and Stefano Tessaro. Information-theoretic indistinguishability via the chi-squared method. In Jonathan Katz and Hovav Shacham, editors, *Advances in Cryptology - CRYPTO 2017 - 37th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 20-24, 2017, Proceedings, Part III*, volume 10403 of *Lecture Notes in Computer Science*, pages 497–523. Springer, 2017. doi:10.1007/978-3-319-63697-9_17. *(CITED ON P. 2.)*

Din24. Itai Dinur. Tight indistinguishability bounds for the XOR of independent random permutations by fourier analysis. In Marc Joye and Gregor Leander, editors, *Advances in Cryptology - EUROCRYPT 2024 - 43rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zurich, Switzerland, May 26-30, 2024, Proceedings, Part I*, volume 14651 of *Lecture Notes in Computer Science*, pages 33–62. Springer, 2024. doi:10.1007/978-3-031-58716-0_2. *(CITED ON PP. 2, 5, 10, AND 24.)*

Din25. Itai Dinur. Combining outputs of a random permutation: New constructions and tight security bounds by fourier analysis. In Serge Fehr and Pierre-Alain Fouque, editors, *Advances in Cryptology - EUROCRYPT 2025 - 44th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Madrid, Spain, May 4-8, 2025, Proceedings, Part I*, volume 15601 of *Lecture Notes in Computer Science*, pages 244–273. Springer, 2025. doi:10.1007/978-3-031-91107-1_9. *(CITED ON PP. 2, 5, 17, 31, 32, 33, AND 36.)*

DNS22. Avijit Dutta, Mridul Nandi, and Abishanka Saha. Proof of mirror theory for $\xi_{\max} = 2$. *IEEE Trans. Inf. Theory*, 68(9):6218–6232, 2022. doi:10.1109/TIT.2022.3171178. *(CITED ON P. 2.)*

Ebe17. Sean Eberhard. More on additive triples of bijections, 2017. URL: https://arxiv.org/abs/1704.02407, arXiv:1704.02407. *(CITED ON PP. 2, 4, 5, 10, 11, 23, AND 32.)*

FGL+24. Antonio Flórez-Gutiérrez, Lorenzo Grassi, Gregor Leander, Ferdinand Sibleyras, and Yosuke Todo. General practical cryptanalysis of the sum of round-reduced block ciphers and ZIP-AES. In Kai-Min Chung and Yu Sasaki, editors, *Advances in Cryptology - ASIACRYPT 2024 - 30th International Conference on the Theory and Application of Cryptology and Information Security, Kolkata, India, December 9-13, 2024, Proceedings, Part IX*, volume 15492 of *Lecture Notes in Computer Science*, pages 280–311. Springer, 2024. doi:10.1007/978-981-96-0947-5_10. *(CITED ON P. 1.)*

Fre77. David Freedman. A Remark on the Difference between Sampling with and without Replacement. *J. American Stat. Association*, 72(359):681–681, 1977. doi:10.1080/01621459.1977.10480637. *(CITED ON P. 1.)*

HWKS98. Chris Hall, David A. Wagner, John Kelsey, and Bruce Schneier. Building prfs from prps. In Hugo Krawczyk, editor, *Advances in Cryptology - CRYPTO '98, 18th Annual International Cryptology Conference,*

Santa Barbara, California, USA, August 23-27, 1998, Proceedings, volume 1462 of *Lecture Notes in Computer Science*, pages 370–389. Springer, 1998. doi:10.1007/BFB0055742. *(CITED ON P. 2.)*

IMV16. Tetsu Iwata, Bart Mennink, and Damian Vizár. CENC is optimally secure. Cryptology ePrint Archive, Paper 2016/1087, 2016. URL: https://eprint.iacr.org/2016/1087. *(CITED ON P. 2.)*

Iwa06. Tetsu Iwata. New blockcipher modes of operation with beyond the birthday bound security. In Matthew J. B. Robshaw, editor, *Fast Software Encryption, 13th International Workshop, FSE 2006, Graz, Austria, March 15-17, 2006, Revised Selected Papers*, volume 4047 of *Lecture Notes in Computer Science*, pages 310–327. Springer, 2006. doi:10.1007/11799313_20. *(CITED ON PP. 2 AND 17.)*

Luc00. Stefan Lucks. The sum of prps is a secure PRF. In Bart Preneel, editor, *Advances in Cryptology - EUROCRYPT 2000, International Conference on the Theory and Application of Cryptographic Techniques, Bruges, Belgium, May 14-18, 2000, Proceeding*, volume 1807 of *Lecture Notes in Computer Science*, pages 470–484. Springer, 2000. doi:10.1007/3-540-45539-6_34. *(CITED ON P. 2.)*

MN17a. Bart Mennink and Samuel Neves. Encrypted Davies–Meyer and its dual: Towards optimal security using mirror theory. In Jonathan Katz and Hovav Shacham, editors, *Advances in Cryptology – CRYPTO 2017, Part III*, volume 10403 of *Lecture Notes in Computer Science*, pages 556–583, Cham, 2017. Springer. doi:10.1007/978-3-319-63697-9_19. *(CITED ON P. 1.)*

MN17b. Bart Mennink and Samuel Neves. Optimal prfs from blockcipher designs. *IACR Trans. Symmetric Cryptol.*, 2017(3):228–252, 2017. doi:10.13154/TOSC.V2017.I3.228-252. *(CITED ON P. 1.)*

MP15. Bart Mennink and Bart Preneel. On the XOR of multiple random permutations. In Stanislaw Jarecki and Dieter Gollmann, editors, *Applied Cryptography and Network Security – ACNS 2015*, volume 9092 of *Lecture Notes in Computer Science*, pages 619–634, Cham, 2015. Springer. doi:10.1007/978-3-319-28166-7_30. *(CITED ON P. 2.)*

MS17. Bart Mennink and Alan Szepieniec. XOR of PRPs in a quantum world. In Tanja Lange and Tsuyoshi Takagi, editors, *Post-Quantum Cryptography – PQCrypto 2017*, volume 10346 of *Lecture Notes in Computer Science*, pages 367–383, Cham, 2017. Springer. doi:10.1007/978-3-319-59879-6_21. *(CITED ON P. 3.)*

Nan24. Mridul Nandi. Improving tightness gap of GGM construction and its applications. In Sourav Mukhopadhyay and Pantelimon Stanica, editors, *Progress in Cryptology - INDOCRYPT 2024 - 25th International Conference on Cryptology in India, Chennai, India, December 18-21, 2024, Proceedings, Part I*, volume 15495 of *Lecture Notes in Computer Science*, pages 28–50. Springer, 2024. doi:10.1007/978-3-031-80308-6_2. *(CITED ON P. 1.)*

O'D14. Ryan O'Donnell. *Analysis of Boolean Functions*. Cambridge University Press, Cambridge, 2014. doi:10.1017/CBO9781139814782. *(CITED ON P. 5.)*

Pat08. Jacques Patarin. A proof of security in o(2n) for the xor of two random permutations. In Reihaneh Safavi-Naini, editor, *Information Theoretic Security, Third International Conference, ICITS 2008, Calgary, Canada, August 10-13, 2008, Proceedings*, volume 5155 of *Lecture Notes in Computer Science*, pages 232–248. Springer, 2008. doi:10.1007/978-3-540-85093-9_22. *(CITED ON P. 2.)*

Pat10. Jacques Patarin. Introduction to mirror theory: Analysis of systems of linear equalities and linear non equalities for cryptography. Cryptology ePrint Archive, Paper 2010/287, 2010. URL: https://eprint.iacr.org/2010/287. *(CITED ON P. 2.)*

Pat13. Jacques Patarin. Generic attacks for the XOR of k random permutations. In Michael J. Jacobson Jr., Michael E. Locasto, Payman Mohassel, and Reihaneh Safavi-Naini, editors, *Applied Cryptography and Network Security – ACNS 2013*, volume 7954 of *Lecture Notes in Computer Science*, pages 154–169, Berlin, Heidelberg, 2013. Springer. doi:10.1007/978-3-642-38980-1_10. *(CITED ON P. 2.)*

Pat16. Jacques Patarin. Mirror theory and cryptography. Cryptology ePrint Archive, Paper 2016/702, 2016. URL: https://eprint.iacr.org/2016/702. *(CITED ON P. 2.)*

Rao66. J. N. K. Rao. On the comparison of sampling with and without replacement. *Revue de l'Institut International de Statistique / Review of the International Statistical Institute*, 34(2):125–138, 1966. doi:10.2307/1401762. *(CITED ON P. 1.)*

RS06. Phillip Rogaway and Thomas Shrimpton. A provable-security treatment of the key-wrap problem. In Serge Vaudenay, editor, *Advances in Cryptology - EUROCRYPT 2006, 25th Annual International Conference on the Theory and Applications of Cryptographic Techniques, St. Petersburg, Russia, May 28 - June 1, 2006, Proceedings*, volume 4004 of *Lecture Notes in Computer Science*, pages 373–390. Springer, 2006. doi:10.1007/11761679_23. *(CITED ON P. 1.)*

Ter99. Audrey Terras. *Fourier Analysis on Finite Groups and Applications*, volume 43 of *London Mathematical Society Student Texts*. Cambridge University Press, 1999. doi:10.1017/CBO9780511626265. *(CITED ON P. 8.)*

Zha12a. Mark Zhandry. How to construct quantum random functions. In *53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2012, New Brunswick, NJ, USA, October 20-23, 2012*, pages 679–687. IEEE Computer Society, 2012. doi:10.1109/FOCS.2012.37. *(CITED ON PP. 4 AND 28.)*

Zha12b. Mark Zhandry. Secure identity-based encryption in the quantum random oracle model. In Reihaneh Safavi-Naini and Ran Canetti, editors, *Advances in Cryptology - CRYPTO 2012 - 32nd Annual Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2012. Proceedings*, volume 7417 of *Lecture Notes in Computer Science*, pages 758–775. Springer, 2012. doi:10.1007/978-3-642-32009-5_44. *(CITED ON PP. 3 AND 14.)*

A Injection Coefficient Estimates

The following coefficient bound is due to [Ebe17, Lemma 4.1], written in our density normalisation.

Proposition 6. *Let G be a finite abelian group of order $N \geq 2$. Then, for all $1 \leq t \leq \lfloor N/2 \rfloor$,*

$$\Lambda_t \leq \binom{N}{t}^{-1/2}.$$

Proof. From [Ebe17, Lemma 4.1],

$$|\hat{1}_S(\chi)| \leq \binom{N}{t}^{-1/2} \frac{N!}{N^N} \qquad (|\chi| = t \leq N/2).$$

Multiplying by $N^N/N!$ gives the corresponding density coefficient, and (8) then gives the stated bound on Λ_t. □

A.1 Small Supports

At support size one, every nonprincipal character has zero coefficient, so $\Theta_1 = \Lambda_1 = 0$. For nonprincipal $\chi_1, \chi_2 \in \widehat{G}$,

$$N(N-1)\widehat{\mu_{I_2}}(\chi_1, \chi_2) = \sum_{y_1 \neq y_2} \overline{\chi_1(y_1)\chi_2(y_2)} = -\sum_{y \in G} \overline{(\chi_1\chi_2)(y)}.$$

The unrestricted sum over G^2 is zero, since each character is nonprincipal. Subtracting the equal-output assignments gives the last equality. By character orthogonality,

$$\widehat{\mu_{I_2}}(\chi_1, \chi_2) = \begin{cases} -1/(N-1), & \chi_1\chi_2 = \hat{\mathbf{1}}, \\ 0, & \text{otherwise.} \end{cases}$$

There are $N-1$ such ordered pairs, one for each choice of χ_1, since $\chi_2 = \chi_1^{-1}$ is then forced. Hence $\Lambda_2 = 1/(N-1)$ and $\Theta_2 = 1/(N-1)$.

For $N \geq 3$ and three nonprincipal characters, inclusion–exclusion removes the three possible pairwise equalities. The unrestricted sum and each single-pair sum vanish, since an unrestricted nonprincipal character sums to zero. The all-equal term has coefficient two, so

$$(N)_3\widehat{\mu_{I_3}}(\chi_1, \chi_2, \chi_3) = 2\sum_{y \in G} \overline{(\chi_1\chi_2\chi_3)(y)}.$$

Consequently,

$$\widehat{\mu_{I_3}}(\chi_1, \chi_2, \chi_3) = \begin{cases} 2/((N-1)(N-2)), & \chi_1\chi_2\chi_3 = \hat{\mathbf{1}}, \\ 0, & \text{otherwise.} \end{cases}$$

There are $(N-1)(N-2)$ such ordered triples: choose $\chi_1 \neq \hat{\mathbf{1}}$, then $\chi_2 \notin \{\hat{\mathbf{1}}, \chi_1^{-1}\}$, and finally $\chi_3 = (\chi_1\chi_2)^{-1}$. Thus

$$\Lambda_3 = \frac{2}{(N-1)(N-2)}, \qquad \Theta_3 = \frac{4}{(N-1)(N-2)}.$$

B Proof of Lemma 3

We retain the notation $\Theta_t = \Theta_{t,2}$ and Λ_t for the actual fixed-support moment and maximum. Throughout this appendix, $N \geq 1024$, and $3 \leq T \leq 8N/15$. All logarithms in this appendix are natural.

Exact second moments. Parseval gives

$$\sum_{\alpha\in\widehat{G}^j} |\widehat{\mu_{I_j}}(\alpha)|^2 = \|\mu_{I_j}\|_2^2 = \frac{N^j}{(N)_j}.$$

Separating these characters by their support, and using (8), yields

$$\sum_{t=0}^{j} \binom{j}{t} \Theta_t = \frac{N^j}{(N)_j},$$

where we take $\Theta_0 = 1$. Binomial inversion therefore gives the exact identity

$$\Theta_t = \sum_{j=0}^{t} (-1)^{t-j} \binom{t}{j} \frac{N^j}{(N)_j}. \tag{31}$$

This is the expression obtained in [Din24, Proposition 20] in the binary setting. The preceding calculation also makes clear that it depends only on the cardinality N.

Let V be an auxiliary real random variable with density, relative to Lebesgue measure,

$$p_V(v) = \frac{v^N e^{-v}}{N!} \qquad (v > 0), \qquad Z = \frac{N}{V} - 1.$$

For $0 \le j \le N$, direct integration gives

$$\mathbf{E}\left[\left(\frac{N}{V}\right)^j\right] = \frac{N^j (N-j)!}{N!} = \frac{N^j}{(N)_j}.$$

Consequently, the binomial-inversion identity (31) gives

$$\Theta_t = \mathbf{E}[Z^t] \qquad (0 \le t \le N). \tag{32}$$

This variable only represents the numerical moment sequence; it is not part of the oracle distribution. Its even moments allow us to interpolate between support sizes using Hölder's inequality.

We also use the generating function

$$F_N(z) = (1-z)^N \exp\left(\frac{Nz}{1-z}\right), \qquad [z^t]F_N(z) = \binom{N}{t}\Theta_t \quad (0 \le t \le N). \tag{33}$$

Indeed, substituting (31) and interchanging the finite sums yields

$$\sum_{t=0}^{N} \binom{N}{t} \Theta_t z^t = (1-z)^N \sum_{j=0}^{N} \frac{1}{j!} \left(\frac{Nz}{1-z}\right)^j.$$

The coefficients through degree N agree with those of F_N. Moreover,

$$F_N(z) = \exp\left(N \sum_{j=2}^{\infty} \frac{j-1}{j} z^j\right)$$

has non-negative coefficients. Hence, for every $0 < \rho < 1$, $[z^t]F_N(z)\,\rho^t \le F_N(\rho)$, and therefore $[z^t]F_N(z) \le \rho^{-t} F_N(\rho)$. Differentiating gives $(1-z)^2 F_N'(z) = NzF_N(z)$, and comparison of coefficients gives

$$\Theta_{t+1} = \frac{t}{N-t}(2\Theta_t + \Theta_{t-1}) \quad (1 \le t < N), \qquad \Theta_0 = 1, \quad \Theta_1 = 0.$$

In particular,

$$\begin{aligned}
\Theta_4 &= \frac{3(N+6)}{(N-1)(N-2)(N-3)}, \\
\Theta_5 &= \frac{8(5N+12)}{(N-1)(N-2)(N-3)(N-4)}, \\
\Theta_6 &= \frac{5(3N^2+86N+120)}{(N-1)(N-2)(N-3)(N-4)(N-5)}.
\end{aligned} \tag{34}$$

Extending the maximum beyond $N/2$. To control the remaining levels, we use Eberhard's coefficient-merging identity. For $2 \le t \le N$ and nonprincipal $\chi_1, \dots, \chi_t \in \widehat{G}$, write

$$\chi^{(i)} = (\chi_1, \dots, \chi_i\chi_t, \dots, \chi_{t-1}).$$

Then using (9)

$$\widehat{\mu_{I_t}}(\chi_1, \dots, \chi_t) = -\frac{1}{N-t+1} \sum_{i=1}^{t-1} \widehat{\mu_{I_{t-1}}}(\chi^{(i)}).$$

The merged coordinate $\chi_i\chi_t$ is either nonprincipal or principal, so the resulting character lies at level $t-1$ or $t-2$. Taking absolute values and using (8) gives

$$\Lambda_t \le \frac{t-1}{N-t+1} \max\{\Lambda_{t-1}, \Lambda_{t-2}\} \qquad (3 \le t \le N). \tag{35}$$

Put $m_0 = \lfloor N/2 \rfloor$ and $C = \binom{N}{m_0}$. We apply this recurrence for $t \ge m_0 + 1$, starting from Proposition 6, which gives $\Lambda_t \le \binom{N}{t}^{-1/2}$ for $1 \le t \le N/2$. We claim that

$$\Lambda_t \le \frac{\sqrt{C}}{\binom{N}{t}} \qquad (m_0 \le t \le N). \tag{36}$$

At $t = m_0$, this is exactly the sparse bound. At $t = m_0 + 1$, if $N = 2m_0$, the recurrence factor is one and

$$\binom{N}{m_0 - 1}^{-1/2} = \frac{1}{\sqrt{C}} \sqrt{\frac{m_0 + 1}{m_0}} \le \frac{\sqrt{C}}{\binom{N}{m_0+1}}.$$

If $N = 2m_0 + 1$, the factor is $m_0/(m_0 + 1)$ and

$$\frac{m_0}{m_0+1} \binom{N}{m_0 - 1}^{-1/2} = \frac{1}{\sqrt{C}} \frac{\sqrt{m_0(m_0+2)}}{m_0+1} \le \frac{1}{\sqrt{C}} = \frac{\sqrt{C}}{\binom{N}{m_0+1}}.$$

For $t \ge m_0 + 2$, the binomial coefficients are non-increasing, so the induction step follows from

$$\frac{t-1}{N-t+1} \frac{1}{\binom{N}{t-1}} = \frac{t-1}{t} \frac{1}{\binom{N}{t}} \le \frac{1}{\binom{N}{t}}.$$

This proves (36) and extends the maximum bound beyond $N/2$.

The levels up to $N/2$. Put $L = 2\lfloor N/4 \rfloor$, so L is even and is either m_0 or $m_0 - 1$. Since the largest binomial coefficient is at least $2^N/(N+1)$,

$$\binom{N}{L} \ge \frac{2^N}{2(N+1)}.$$

Using (33) at $z = 2/5$ therefore gives

$$\Theta_L \le 2(N+1)\beta^N, \qquad \beta = \frac{3}{10} e^{2/3} \sqrt{\frac{5}{2}} < e^{-1/13}.$$

Also, $\Theta_6 \ge 15/N^3$ by (34). For $N \ge 1024$,

$$\Theta_L \le 2(N+1)e^{-N/13} \le \frac{15}{N^3} e^{-(L-6)/10} \le \Theta_6 e^{-(L-6)/10}. \tag{37}$$

For the middle inequality, it suffices, using $L \le N/2$, to check

$$\frac{7N}{260} - \log\left(\frac{2(N+1)N^3}{15}\right) + \frac{3}{5} > 0.$$

The left-hand side is greater than 2.45727 at $N = 1024$, and its derivative is $7/260 - 1/(N+1) - 3/N > 0$ in the required range.

Both 6 and L are even. Hölder's inequality applied to (32), followed by (37), gives, for $6 \le t \le L$,

$$\Theta_t \le \mathbf{E}[|Z|^t] \le \Theta_6^{(L-t)/(L-6)} \Theta_L^{(t-6)/(L-6)} \le \Theta_6 e^{-(t-6)/10}.$$

Hence

$$\sum_{t=4}^{L} \Theta_t \le \Theta_4 + \Theta_5 + \frac{\Theta_6}{1 - e^{-1/10}} < \Theta_4 + \Theta_5 + 11\Theta_6. \tag{38}$$

For $4 \le t \le L$, the sparse maximum gives

$$\binom{N}{t} \Theta_{t,4} \le \binom{N}{t} \Lambda_t^2 \Theta_t \le \Theta_t.$$

The support-three contribution is given exactly by Appendix A:

$$\binom{N}{3} \Theta_{3,4} = \frac{8N}{3(N-1)^2(N-2)^2}.$$

The remaining levels. For $L < t \le T$, we have $t \ge m_0$. Using (36) and then (33) at $z = 2/5$,

$$\binom{N}{t} \Theta_{t,4} \le \frac{C}{\binom{N}{t}} \Theta_t \le \frac{2^N}{\binom{N}{t}^2} \left(\frac{3}{5}\right)^N e^{2N/3} \left(\frac{5}{2}\right)^t.$$

For $x = t/N$, put $h(x) = -x \log x - (1-x)\log(1-x)$. The binomial probability with parameters N, x is largest at t, and is at least $1/(N+1)$ there. Thus

$$\binom{N}{t} \ge \frac{e^{Nh(x)}}{N+1}.$$

It follows that

$$\binom{N}{t} \Theta_{t,4} \le (N+1)^2 e^{NJ(x)}, \qquad J(x) = \log \frac{6}{5} + \frac{2}{3} + x \log \frac{5}{2} - 2h(x).$$

Here $x \ge m_0/N > 0.49$, and $J'(x) = \log(5/2) - 2\log((1-x)/x) > 0$ throughout the required interval. A direct evaluation gives

$$J(8/15) < -0.04417000 < -\frac{1}{23}.$$

There are fewer than N remaining levels, so

$$\sum_{L<t\le T} \binom{N}{t} \Theta_{t,4} \le N(N+1)^2 e^{-N/23}. \tag{39}$$

This also covers the possible odd level just below $N/2$.

Combining (38), (39), and the support-three contribution gives

$$\sum_{t=3}^{T} \binom{N}{t} \Theta_{t,4} < \frac{8N}{3(N-1)^2(N-2)^2} + \Theta_4 + \Theta_5 + 11\Theta_6 + N(N+1)^2 e^{-N/23}. \tag{40}$$

After multiplication by N^2, each rational term decreases with N: write it as a polynomial with non-negative coefficients in $1/N$, divided by products of $1 - i/N$. The exponential term decreases as well, since

$$\frac{d}{dN} \log\left(N^3 (N+1)^2 e^{-N/23}\right) = \frac{3}{N} + \frac{2}{N+1} - \frac{1}{23} < 0.$$

At $N = 1024$, the normalised right-hand side of (40) is less than $3.245641 < 13/4$. This proves (12).

A common maximum for all these levels. At support three,

$$\Lambda_3^2\binom{N}{4} = \frac{N(N-3)}{6(N-1)(N-2)} < 1.$$

For $4 \le t \le m_0$, we have

$$\Lambda_t \le \binom{N}{t}^{-1/2} \le \binom{N}{4}^{-1/2}.$$

For $m_0 < t \le 8N/15$, the entropy estimate and (36) give

$$\Lambda_t \le (N+1)\exp\left(N\left(\frac{1}{2}\log 2 - h(t/N)\right)\right) \le (N+1)e^{-N/3} < N^{-2} < \binom{N}{4}^{-1/2}.$$

Indeed, $h(8/15) - \frac{1}{2}\log 2 > 1/3$, and $(N+1)e^{-N/3} < N^{-2}$ holds at $N = 1024$ and persists since $2/N + 1/(N+1) - 1/3 < 0$. This proves (13).

Now let m be any integer with $T \le m \le N$. For $3 \le t \le T \le m \le N$, the ratio $\binom{m}{t}/\binom{N}{t}$ is at most $\binom{m}{3}/\binom{N}{3}$. Also, $\Theta_{t,2k} \le \Lambda_t^{2k-4}\Theta_{t,4}$. Thus the two estimates just proved give

$$\sum_{t=3}^{T}\binom{m}{t}\Theta_{t,2k} \le \frac{\binom{m}{3}}{\binom{N}{3}}\binom{N}{4}^{-(k-2)}\sum_{t=3}^{T}\binom{N}{t}\Theta_{t,4} \le \frac{\binom{m}{3}}{\binom{N}{3}}\frac{13}{4N^2}\binom{N}{4}^{-(k-2)}.$$

This proves (14). □

C Query-dependent Quantum Bounds for Sums of Permutations

Let $F\colon \mathcal{X} \to H$ have density φ and uniform one-point marginals, and let $\mathsf{g}\colon \mathcal{X} \to H$ be uniform. For a q-query quantum distinguisher $\mathcal{A}$ with acceptance function a, Lemmas 1 and 2, with $r = 2$, give

$$\big\|\,|F\rangle - |\mathsf{g}\rangle\big\|_{\mathcal{A}} \le |\mathbf{E}[a\varphi_{=2}]| + \frac{1}{2}\|R_q\|_2, \qquad R_q = \sum_{t=3}^{2q}\varphi_{=t}, \tag{41}$$

for $q \ge 1$. We apply this bound to the unprojected and projected sums. Their level-two components follow from Lemma 4; the common injection estimate (14) controls the higher levels.

C.1 A Bound for the Centred Collision Count

We first bound the correlation of a quantum acceptance function with the centred collision count. We state it for an arbitrary finite input set and a finite abelian output group, so that it also applies to truncation.

Lemma 6. *Let $\mathcal{X}$ be a finite set, let H be a finite abelian group of order M, and define, for $f : \mathcal{X} \to H$,*

$$K_H(f) := \sum_{\{x,y\}\subseteq\mathcal{X}}\Big(M\mathbf{1}_{\{f(x)=f(y)\}} - 1\Big),$$

where the sum is over unordered pairs of distinct inputs, and the expectation below is with respect to $\mathsf{U}_{H^{\mathcal{X}}}$. If a is the acceptance probability of a q-query quantum algorithm, then

$$|\mathbf{E}_{\mathsf{U}}[aK_H]| \le \frac{\pi^2}{6}(2q-1)^3 = O(q^3) \qquad (q \ge 1).$$

For $q = 0$, the expectation is zero. The implicit constant is independent of $\mathcal{X}$ and H.

Notice that K_H has mean zero.

Proof. If $q = 0$, $M = 1$, or $|\mathcal{X}| < 2$, the claim follows from $\mathbf{E}[K_H] = 0$ or $K_H = 0$. Assume otherwise. For a positive integer r, sample independent uniform random functions $b\colon \mathcal{X} \to [r]$ and $c\colon [r] \to H$, and set $F_r = c \circ b$. For $z = 1/r$, let $p(z) = \mathbf{E}[a(F_r)]$. We first make the small-range bound concrete. By Lemma 2, the Fourier expansion of a uses supports of size at most $2q$. On a support of size $s \geq 1$, the labels assigned by b induce a partition. A fixed partition with $j \geq 1$ blocks has probability

$$\frac{(r)_j}{r^s} = z^{s-j} \prod_{i=0}^{j-1} (1 - iz),$$

which has degree at most $s - 1$, since the factor at $i = 0$ is one. The conditional character expectation depends on the partition but not on r. Thus p is a polynomial of degree at most $2q - 1$, with $p(0) = \mathbf{E}_{\mathsf{U}}[a]$ and $0 \leq p(1/r) \leq 1$ for every positive integer r. Its coefficients are real, since its values at these real points are real. Zhandry's polynomial bound [Zha12a, Theorem B.1, with $\Delta = 1$] yields

$$|\mathbf{E}[a(F_r)] - \mathbf{E}_{\mathsf{U}}[a]| \leq \frac{\pi^2 (2q-1)^3}{6r} \leq C_{\mathrm{SR}} \frac{q^3}{r}, \qquad C_{\mathrm{SR}} = \frac{4\pi^2}{3}. \tag{42}$$

Here the support cutoff applies directly to the addition oracle on H.

Let $D = |\mathcal{X}|$, and let ϱ_r be the density of F_r on $H^{\mathcal{X}}$. As $r \to \infty$, the probability that b is injective is

$$\frac{(r)_D}{r^D} = 1 - \frac{\binom{D}{2}}{r} + O_D(r^{-2}).$$

For each fixed pair $\{x, y\}$, the probability that this is the only collision of b is

$$\frac{(r)_{D-1}}{r^D} = \frac{1}{r} + O_D(r^{-2}).$$

Conditioned on the first event, the values of c at the distinct labels chosen by b are independent and uniform, so F_r is a uniform random function. Conditioned on the second event, the outputs at x and y are equal, and all other outputs remain independent and uniform. The density of this distribution is $M 1_{\{f(x)=f(y)\}}$. All other collision patterns have total probability $O_D(r^{-2})$. Consequently,

$$\varrho_r = 1 + \frac{K_H}{r} + O_{D,M}(r^{-2}),$$

where the error is uniform on the finite space $H^{\mathcal{X}}$, with $\mathcal{X}$ and H fixed. Thus,

$$\lim_{r \to \infty} r \left(\mathbf{E}[a(F_r)] - \mathbf{E}_{\mathsf{U}}[a] \right) = \mathbf{E}_{\mathsf{U}}[a K_H].$$

Multiplying (42) by r and taking the limit proves the claim. The limit is taken with $\mathcal{X}$, H and the algorithm fixed. Although the vanishing error can depend on D, M, the surviving constant C_{SR} does not. □

C.2 Proof of Theorem 3

Fix $k \geq 2$, and let $\varphi_k = \mu_S^{*k}$ be the density of the complete table of π_+^k. Its one-point marginals are uniform. Since the two-point marginals determine the level-two projection, Lemma 4 gives

$$h_{2,k} := (\varphi_k)_{=2} = \eta_k K_G, \qquad \eta_k = \left(-\frac{1}{N-1} \right)^k. \tag{43}$$

Thus Lemma 6 implies, for every q-query acceptance function a,

$$|\mathbf{E}[a h_{2,k}]| \leq \frac{\pi^2 (2q-1)^3}{6(N-1)^k}.$$

For $q = 1$, there is no remainder, proving the one-query bounds.

Suppose $N \geq 1024$ and $2 \leq q \leq \sqrt{N}$. By Parseval's identity, convolution-multiplication duality, and (14) with $m = N$ and $T = 2q$,

$$\|R_{q,k}\|_2^2 = \sum_{t=3}^{2q} \binom{N}{t} \Theta_{t,2k} \leq \frac{13}{4N^2} \binom{N}{4}^{-(k-2)}, \qquad R_{q,k} = \sum_{t=3}^{2q} (\varphi_k)_{=t}.$$

The remainder in (41) is therefore $O_k(N^{-2k+3})$. For $k = 2$, this gives $O(q^3/N^2 + 1/N)$; for $k \geq 3$, it gives $O_k(q^3/N^k)$. The finitely many orders $N < 1024$ are absorbed into the implicit constants.

Corollary 2 gives the uniform bound $O_k(N^{-(k-3/2)})$ throughout $2q < N$. For $q > \sqrt{N}$, this is no larger than q^3/N^k. Combining the local and uniform bounds proves the theorem. □

Keeping the constants in the preceding calculation gives the following bound for $q < 4N/15$.

Corollary 7. *Let G be a finite abelian group of order $N \geq 1024$, fix $k \geq 2$, and let $\mathsf{g}\colon G \to G$ be uniform. For $2 \leq q \leq 4N/15$,*

$$\left\| |\pi_+^k\rangle - |\mathsf{g}\rangle \right\|_{(q)} \leq \min \left\{ \sqrt{\frac{N}{8(N-1)^{2k-2}} + \frac{13}{16N^2}\binom{N}{4}^{-(k-2)}}, \right. \\ \left. \frac{\pi^2(2q-1)^3}{6(N-1)^k} + \frac{\sqrt{13}}{4N}\binom{N}{4}^{-(k-2)/2} \right\}. \tag{44}$$

For $q = 1$, the advantage is at most $\pi^2/(6(N-1)^k)$; for $q = 0$, it is zero.

Proof. The first bound is Corollary 1. For the second, the two-input identity (43), which also holds for $k = 2$, and Lemma 6 give

$$|\mathbf{E}_{\mathsf{U}}[a h_{2,k}]| \leq \frac{\pi^2(2q-1)^3}{6(N-1)^k}.$$

For the remaining levels, (14) gives

$$\|R_{q,k}\|_2^2 = \sum_{t=3}^{2q} \binom{N}{t} \Theta_{t,2k} \leq \frac{13}{4N^2}\binom{N}{4}^{-(k-2)}.$$

Apply (41), retaining its factor $1/2$ on the remainder. For $q = 1$, the remainder is empty, and for $q = 0$ the oracles are not queried. □

C.3 Proof of Theorem 6

The same decomposition gives a bound that retains the output-size dependence for the construction from Section 5.1. The additional observation is that the proof of Proposition 5 retains the rank of the span of the character coordinates. Keeping this rank information gives a finer estimate on the remainder than applying the proposition directly.

Throughout this subsection, take $G = \mathbb{F}_2^n$, $H = \mathbb{F}_2^m$, $N = 2^n$, and $M = 2^m$, with $0 \leq m \leq n$. Let $\tau\colon G \to H$ be surjective and linear.

Two permutations. For $M = 1$, both oracles are the unique function into the trivial group. Assume $M \geq 2$. As before, the finitely many orders $N < 1024$ can be absorbed into the asymptotic constant. Fix a q-query distinguisher $\mathcal{A}$ with acceptance probability a. Let ν be the density of $\tau\pi_+^2$ on H^G, and set $W = \tau^\top(H)$. As in the proof of Theorem 5,

$$\widehat{\nu}(\eta) = \widehat{\mu_S}(\tau^\top \eta)^2. \tag{45}$$

Here $\tau^\top$ acts coordinatewise and is injective, so it preserves support size. Write

$$\Pi_{\leq 2q}(\nu - 1) = h_2 + R_q, \qquad h_2 = \nu_{=2}, \qquad R_q = \sum_{t=3}^{2q} \nu_{=t}.$$

The pushforward identity (15) gives

$$h_2(f) = \frac{K_H(f)}{(N-1)^2}, \qquad |\mathbf{E}_{\mathsf{U}}[a h_2]| \leq O(q^3/N^2),$$

where the second inequality follows from Lemma 6.

For $q = 1$, $R_q = 0$. For $2 \leq q \leq \sqrt{N}$, Parseval and (45) give

$$\|R_q\|_2^2 = \sum_{\substack{\chi \in W^G \\ 3 \leq |\chi| \leq 2q}} |\widehat{\mu_S}(\chi)|^4. \tag{46}$$

The subspace-averaging identity (17), restricted to the support levels in (46), gives

$$\|R_q\|_2^2 = \sum_{3\le|\chi|\le 2q} p_{n,m}(r(\chi))|\widehat{\mu_{\mathcal{S}}}(\chi)|^4. \tag{47}$$

Here $p_{n,m}$ is the same probability function as in (19). The restriction to these levels is allowed because the diagonal action of $\mathrm{GL}(n,2)$ preserves support size. In particular,

$$p_{n,m}(1) = \frac{M-1}{N-1}, \qquad p_{n,m}(r) \le p_{n,m}(2) = \frac{(M-1)(M-2)}{(N-1)(N-2)} \quad (r \ge 2).$$

A rank-one tuple over $\mathbb{F}_2$ has the same nonzero character at every position in its support. Its permutation coefficient vanishes when the support size is odd, by translating all outputs by a vector on which this character takes value -1. There are $(N-1)\binom{N}{t}$ rank-one tuples of support size t. For $4 \le t \le 2q \le 2\sqrt{N} < N/2$, Proposition 6 gives $\Lambda_t \le \binom{N}{t}^{-1/2}$. Thus,

$$\begin{aligned}\sum_{\substack{3\le|\chi|\le 2q\\ r(\chi)=1}} |\widehat{\mu_{\mathcal{S}}}(\chi)|^4 &\le (N-1)\sum_{t=4}^{2q}\binom{N}{t}^{-1}\\ &= O(N^{-3}).\end{aligned} \tag{48}$$

For the last step, the successive ratio of $\binom{N}{t}^{-1}$ is $(t+1)/(N-t) = O(N^{-1/2})$ throughout $4 \le t < 2q \le 2\sqrt{N}$. Here $N \ge 1024$ suffices.

Applying (48) to the rank-one part of (47), and (14), with $k = 2$ and $m = N$, to the remaining part, gives

$$\begin{aligned}\|R_q\|_2^2 &\le p_{n,m}(1)O(N^{-3}) + p_{n,m}(2)O(N^{-2})\\ &\le O\left(\frac{M}{N^4} + \frac{M^2}{N^4}\right) = O(M^2/N^4).\end{aligned}$$

Hence $\|R_q\|_2 = O(M/N^2)$. Now, (41) yields

$$\left\||\tau\pi_+^2\rangle - |\mathsf{g}\rangle\right\|_{(q)} \le O\left(\frac{q^3+M}{N^2}\right) \qquad (q \le \sqrt{N}). \tag{49}$$

Again, $R_q = 0$ for $q = 1$.

Finally, Corollary 3 with $k = 2$ gives

$$\left\||\tau\pi_+^2\rangle - |\mathsf{g}\rangle\right\|_{(q)} \le O(\sqrt{M}/N) \qquad (2q < N).$$

Combining this with (49) proves (21). For $q > \sqrt{N}$, the minimum selects the uniform bound, since $q^3/N^2 > N^{-1/2} \ge \sqrt{M}/N$.

For $M^{1/3} \le q \le \sqrt{N}$, the local bound (49) is $O(q^3/N^2)$. In particular, for a fixed-size output group, this holds throughout $1 \le q \le \sqrt{N}$, with the implicit constant allowed to depend on that fixed size. The improvement from $1/N$ to M/N^2 in the remainder is a consequence of separating rank-one characters from the characters of rank at least two.

Three or more permutations. The case $M = 1$ is immediate, and the finitely many orders $N < 1024$ can be absorbed into the constant depending on k. Fix a q-query distinguisher $\mathcal{A}$ with acceptance probability a, and let ν_k be the density of $\tau\pi_+^k$ on H^G. For $q \le \sqrt{N}$, split its low-support part into $h_{2,k} + R_{q,k}$. By (15),

$$h_{2,k}(f) = \frac{(-1)^k}{(N-1)^k} K_H(f),$$

and hence Lemma 6 gives a contribution $O_k(q^3/N^k)$.

If $q = 1$, the remainder is zero. Otherwise, the pullback formula (45) and the convolution identity show that each coefficient of the k-sum is the corresponding coefficient of the two-sum multiplied by $\widehat{\mu_S}(\chi)^{k-2}$. As above, Lemma 3 gives

$$\max_{3\le|\chi|\le 2q} |\widehat{\mu_S}(\chi)| \le \binom{N}{4}^{-1/2} = O(N^{-2}) \qquad (q \le \sqrt{N}).$$

Together with the estimate $\|R_{q,2}\|_2 = O(M/N^2)$ proved above, this gives

$$\|R_{q,k}\|_2 \le O_k(N^{-2k+4})\frac{M}{N^2} = O_k\left(\frac{M}{N^{2k-2}}\right).$$

Since $M \le N$ and $k \ge 3$, this is at most $O_k(N^{-k})$, and therefore at most $O_k(q^3/N^k)$. This proves the local bound $O_k(q^3/N^k)$ for $q \le \sqrt{N}$. Corollary 3 gives the uniform bound $O_k(\sqrt{M}/N^{k-1})$. For $q > \sqrt{N}$, the uniform bound is no larger than $N^{-(k-3/2)} < q^3/N^k$, and the result follows. □

Remark 3. The query-dependent refinement uses the exact centred-collision form of the support-two component for sums of independent permutations. For LXoP this component has a different linear structure, so the same refinement does not follow from the preceding argument.

D A Deutsch-Type One-Query Algorithm

We give the standard two-point phase-query algorithm with an arbitrary nonzero output mask; see also Bonnetain et al. [BLNPS21]. Let $G = \mathbb{F}_2^n$ and let

$$U_f\,|u\rangle_X\,|v\rangle_Y = |u\rangle_X\,|v \oplus f(u)\rangle_Y$$

be the standard quantum oracle.

For distinct $x, x' \in G$ and nonzero $\lambda \in G$, there exists a one-query quantum algorithm that returns $b_\lambda(x,x') = \langle \lambda, f(x) \oplus f(x')\rangle$ with certainty.

The algorithm uses a one-qubit register B and two n-qubit registers X and Y.

Algorithm.
1. Prepare $|0\rangle_B\,|0^n\rangle_X\,|\lambda\rangle_Y$.
2. Apply H to B and $H^{\otimes n}$ to Y.
3. XOR x into X if $B = 0$, and XOR x' into X if $B = 1$.
4. Apply U_f once to X and Y.
5. Undo Step 3.
6. Apply H to B, measure it, and return the outcome.

Adding $z \in G$ to $H^{\otimes n}\,|\lambda\rangle$ contributes the phase $(-1)^{\langle\lambda,z\rangle}$. Hence, after the oracle query and Step 5, register B is in the state

$$\frac{(-1)^{\langle\lambda,f(x)\rangle}\,|0\rangle + (-1)^{\langle\lambda,f(x')\rangle}\,|1\rangle}{\sqrt{2}}$$

$$= (-1)^{\langle\lambda,f(x)\rangle}\frac{|0\rangle + (-1)^{b_\lambda(x,x')}\,|1\rangle}{\sqrt{2}} = (-1)^{\langle\lambda,f(x)\rangle}H\,|b_\lambda(x,x')\rangle\,.$$

Since $H^2 = I$, the final Hadamard transform leaves B in $|b_\lambda(x,x')\rangle$ up to a global phase. Measuring B returns the required value with certainty. Only Step 4 queries U_f. □

E The Recursive Moment Argument for LXoP

Throughout this appendix, $G = \mathbb{F}_2^n$ and $N = 2^n$. We prove Lemma 5. The argument follows Dinur's recursive moment bound [Din25], beginning with (9) and keeping track of a family of masks when bounding its second Fourier moment. We first prove the following more precise estimate for $t \ge 2$.

Proposition 7. *Fix $w \geq 1$, let σ be w-admissible, and let $t \geq 2$ satisfy $(w+1)t \leq N/8$. Put $d_t = \lceil t/2 \rceil$. Then*

$$L_t \leq 2^{d_t+(w+1)t} \left(\frac{(w+1)t}{N-(w+1)t} \right)^{t+d_t}. \tag{50}$$

We use Dinur's fixed-support second-moment estimate [Din25, Lemma 2],

$$\Theta_k \leq \left(\frac{k}{N-k} \right)^{k/2}, \qquad 1 \leq k \leq N/2. \tag{51}$$

E.1 From Coefficient Merging to the Moment Bound

Let S be a finite set and $T\colon S \to G^r$, with $1 \leq r \leq N$, be injective. Suppose the masks $T(\alpha)$ have the same support, of size k_0, and fix an integer $d \geq 0$ with $2d < k_0 \leq N/8$. At each nonempty internal node, the primary coordinate is chosen from the common support using only the initial data and the previous merge indices and zero/nonzero outcomes. We assume that, for every secondary choice and each nonempty outcome, the removed entry is uniquely determined by the child mask and these known indices and outcomes. We derive the bound under this assumption, and verify it for the LXoP family in Appendix E.2.

A current mask is a tuple $\beta \in G^r$. For distinct indices j, i in its support, define the merged mask by

$$\beta_\ell^{(j,i)} = \begin{cases} 0, & \ell = j, \\ \beta_i + \beta_j, & \ell = i, \\ \beta_\ell, & \ell \notin \{i, j\}. \end{cases}$$

We keep the zero coordinates rather than delete them, so the positions retain their original names. This does not affect the Fourier coefficient, by (8). If the support of β is J, with $|J| = k \geq 2$, Equation (9), applied to the nonzero entries and with the primary entry placed last, gives

$$\widehat{\mu_{I_r}}(\beta) = -\frac{1}{N-k+1} \sum_{i \in J \setminus \{j\}} \widehat{\mu_{I_r}}(\beta^{(j,i)}).$$

This is exactly the identity used in [Din25, Proposition 8], attributed there to [Ebe17, Section 4]. Taking absolute values gives the maximum recurrence used in Appendix B. For the present purpose, we instead square the identity and apply Cauchy–Schwarz:

$$|\widehat{\mu_{I_r}}(\beta)|^2 \leq \frac{k-1}{(N-k+1)^2} \sum_{i \in J \setminus \{j\}} |\widehat{\mu_{I_r}}(\beta^{(j,i)})|^2. \tag{52}$$

This is the one-step estimate of [Din25, Proposition 9]. The distinction is that we will sum it over the structured family, rather than replace each term by a maximum coefficient.

A recursion node. At a node v, let $S_v \subseteq S$ be the set of original masks still under consideration, and let T_v map them to their current masks. Write J_v for their common support and $k_v = |J_v|$. At the root, these are S, T, and the initial support of size k_0. Choose one primary index $j \in J_v$ for this entire family, using the initial data and the merge history as above. Before depth d, we have $k_v \geq k_0 - 2(d-1) > 2$, so a secondary index is available. For each secondary index $i \in J_v \setminus \{j\}$, put

$$T_{v,i}(\alpha) = T_v(\alpha)^{(j,i)},$$

and split S_v into

$$S_{v,i,0} = \{\alpha \in S_v : T_{v,i}(\alpha)_i = 0\},$$
$$S_{v,i,1} = \{\alpha \in S_v : T_{v,i}(\alpha)_i \neq 0\}.$$

The zero branch has common support $J_v \setminus \{j, i\}$, while the nonzero branch has common support $J_v \setminus \{j\}$. Thus the support remains fixed within each child, even though the two children lie at different Fourier levels. Summing (52) gives

$$\begin{aligned} &\sum_{\alpha \in S_v} |\widehat{\mu_{I_r}}(T_v(\alpha))|^2 \\ &\quad \leq \frac{k_v - 1}{(N-k_v+1)^2} \sum_{i \in J_v \setminus \{j\}} \sum_{b \in \{0,1\}} \sum_{\alpha \in S_{v,i,b}} |\widehat{\mu_{I_r}}(T_{v,i}(\alpha))|^2. \end{aligned} \tag{53}$$

For a fixed i, the two branches partition S_v. Different choices of i need not give disjoint images; each such choice is already a separate term in (53).

Why recovery is needed. By the recovery assumption, the removed primary entry β_j is determined by the child mask and the known branch. The secondary value is then recovered from

$$\beta_i = (\beta_i + \beta_j) + \beta_j,$$

and the other coordinates have not changed. Hence the full parent mask is determined. Starting from the injective map T, this proves inductively that the map at each node is injective. This is Dinur's applicability condition [Din25, Proposition 11]. It cannot be dropped, since initial injectivity alone does not imply injectivity after a merge. For example, all masks of the form (γ, γ) collapse to $(0, 0)$ when their two coordinates are merged.

At a leaf v, trim the common zero coordinates. The resulting masks are distinct elements of $(G \setminus \{0\})^{k_v}$, so

$$\sum_{\alpha \in S_v} |\widehat{\mu_{\mathcal{I}_r}}(T_v(\alpha))|^2 \le \Theta_{k_v}. \tag{54}$$

No factor $\binom{r}{k_v}$ is needed here, because the leaf has one fixed support. Without recovery, a trimmed mask could occur repeatedly on the left, and the comparison with Θ_{k_v} would not follow.

Summing the recursion. We stop at depth d and apply the fixed-support estimate (51) at each leaf. After d merges, every nonempty leaf has

$$k_0 - 2d \le k_v \le k_0 - d.$$

In particular, $k_v > 0$ by $2d < k_0$, and $k_v \le k_0 \le N/8$. Since $k_0/(N - k_0) < 1$, the leaf contribution is at most

$$\Theta_{k_v} \le \left(\frac{k_v}{N - k_v}\right)^{k_v/2} \le \left(\frac{k_0}{N - k_0}\right)^{k_0/2 - d}.$$

Each node has at most $2(k_v - 1) \le 2k_0$ children, and the factor preceding the child sums in (53) is at most $k_0/(N - k_0)^2$. There are therefore at most $(2k_0)^d$ nonempty leaves, with exactly d such factors along each root-to-leaf path. Iterating (53) and then applying (54) gives

$$\begin{aligned} \sum_{\alpha \in S} |\widehat{\mu_{\mathcal{I}_r}}(T(\alpha))|^2 &\le (2k_0)^d \left(\frac{k_0}{(N - k_0)^2}\right)^d \left(\frac{k_0}{N - k_0}\right)^{k_0/2 - d} \\ &= 2^d \left(\frac{k_0}{N - k_0}\right)^{k_0/2 + d}. \end{aligned} \tag{55}$$

For $d = 0$, the same calculation is just the leaf estimate. This is the weaker of the two bounds in [Din25, Lemma 3], which is sufficient for the general-w analysis. Relative to the direct bound $\Theta_{k_0} \le (k_0/(N - k_0))^{k_0/2}$, the right-hand side gains a factor $(2k_0/(N - k_0))^d$. The purpose of the block relations is to make this recursion possible for sufficiently many steps.

E.2 Recovering Labels from an Untouched Block

We retain $C_j = A^{w-j}$ from (27). At a fixed recursion node, choose a primary coordinate in an untouched block, meaning that none of its entries has been selected as primary or secondary in the previous merges. Suppressing the block index, write $\beta = (\beta_0, \dots, \beta_w)$ for its entries before the merge, and β'_ℓ for its entries after the merge. The indices in the following two calculations are local to this block. The relation

$$\sum_{\ell=0}^{w} C_\ell \beta_\ell = 0$$

is available because the block is untouched. We do not assume that every block still satisfies this relation after a merge. In fact, the merged masks need not remain in the image of $T_{\sigma,w}$.

The secondary coordinate is in another block. Only the primary coordinate j changes within the chosen block. Thus $\beta'_\ell = \beta_\ell$ for $\ell \neq j$, and

$$\beta_j = C_j^{-1} \sum_{\ell \neq j} C_\ell \beta'_\ell.$$

The secondary value in the other block is then recovered by adding β_j to its post-merge value. No relation in that other block is needed.

The secondary coordinate is in the same block. Let $i \neq j$ be its index. Since $\beta'_i = \beta_i + \beta_j$, substitution in the block relation gives

$$(C_i + C_j)\beta_j = C_i \beta'_i + \sum_{\ell \notin \{i,j\}} C_\ell \beta'_\ell.$$

Now

$$C_i + C_j = A^{w-\max\{i,j\}}(I + A^{|i-j|})$$

is invertible: $A = \sigma^\top$ is invertible, and $I + A^{|i-j|} = (I + \sigma^{|i-j|})^\top$ is invertible by w-admissibility. Consequently,

$$\beta_j = (C_i + C_j)^{-1} \left(C_i \beta'_i + \sum_{\ell \notin \{i,j\}} C_\ell \beta'_\ell \right), \qquad \beta_i = \beta'_i + \beta_j.$$

This also covers $\beta'_i = 0$. The zero value and its coordinate position are retained, and do not remove any information needed by the recovery formula. These formulas determine the parent mask within each branch, establishing the recovery assumption used to derive (55).

An example with two output blocks. For $w = 2$, one block satisfies

$$A^2\beta_0 + A\beta_1 + \beta_2 = 0.$$

If coordinate 0 is primary and coordinate 1 is secondary, the child block is $(0, \gamma, \delta)$, where $\gamma = \beta_0 + \beta_1$ and $\delta = \beta_2$. The parent is recovered by

$$\beta_0 = (A^2 + A)^{-1}(A\gamma + \delta), \qquad \beta_1 = \gamma + \beta_0, \qquad \beta_2 = \delta.$$

The formula applies equally when $\gamma = 0$. Here $A^2 + A = A(I + A)$ is invertible. This is precisely where the linear combining map is needed: it lets us reverse a merge that would otherwise lose one character label.

Proof of Proposition 7. Fix $t \geq 2$ with $(w+1)t \leq N/8$, and put $r = (w+1)t$ and $d_t = \lceil t/2 \rceil$. Let

$$S_t = \{\alpha \in (G^w)^t : \alpha^{(i)} \neq 0 \text{ for all } i \in [t]\}.$$

Partition S_t according to the exact support of $T_{\sigma,w}(\alpha)$. There are at most 2^r support patterns. Fix a nonempty part S with injection support $J_0 \subseteq [r]$, and let T be the restriction of $T_{\sigma,w}$ to S. The map T is injective, and all its masks have the same support, of size $k_0 = |J_0|$. Each nonzero construction block gives between 2 and $w+1$ nonzero injection entries, so $2t \leq k_0 \leq (w+1)t \leq N/8$. Also $2d_t \leq t + 1 < 2t \leq k_0$, since $t \geq 2$.

Each merge touches at most two original blocks. At depth $\ell < d_t$, at least $t - 2\ell > 0$ blocks have not been touched. Their original entries are still present, and each such block has a nonzero entry. Choose the first untouched block and its first nonzero coordinate as primary. This choice is the same for all masks at the node and is determined by the initial support and the previous merge indices and outcomes, not by the values of the entries. The preceding recovery formulas apply for every secondary choice, including when the merged entry is zero. Thus each child map remains injective, and the recursion can be continued to depth d_t. No primary selection is needed at a leaf.

Applying (55) to this fixed part gives

$$\begin{aligned} \sum_{\alpha \in S} \left| \widehat{\mu_{I_r}}(T_{\sigma,w}(\alpha)) \right|^2 &\leq 2^{d_t} \left(\frac{k_0}{N - k_0} \right)^{k_0/2 + d_t} \\ &\leq 2^{d_t} \left(\frac{(w+1)t}{N - (w+1)t} \right)^{t + d_t}. \end{aligned}$$

For the last inequality, the ratio is at most $((w+1)t)/(N-(w+1)t) < 1$, while $k_0/2 \geq t$. Summing over at most $2^{(w+1)t}$ support patterns and using the pullback equality in (26) proves (50). To obtain (29), note that $d_t \leq 3t/4$, $t + d_t \geq 3t/2$, and $N - (w+1)t \geq 7N/8$. Since $(w+1)t/(N-(w+1)t) < 1$, the extra factors are absorbed into a constant C_w raised to the power t. □

E.3 The One-Input Moment

We give the remaining calculation for (28). Put $a = w + 1$ and first assume $N \geq 8a$. For this single construction block, suppress the block index and partition the nonzero masks $\alpha \in G^w$ according to the support of their pullback $\beta = T_{\sigma,w}(\alpha)$. Each support has size between 2 and a. Translation invariance of the injection distribution, as recorded in the preliminaries, gives

$$\widehat{\mu_{I_a}}(\beta) = 0 \quad \text{unless} \quad \sum_{j=0}^{w} \beta_j = 0.$$

For a pullback of support size two, at positions i, j, this requires $\beta_i = \beta_j$. The block relation would then give

$$(C_i + C_j)\beta_i = 0,$$

contradicting w-admissibility and $\beta_i \neq 0$. Thus these coefficients vanish.

Consider next a fixed support of size three, at positions i, j, ℓ. A nonzero coefficient requires $\beta_\ell = \beta_i + \beta_j$. Substitution in (27) gives

$$(C_i + C_\ell)\beta_i + (C_j + C_\ell)\beta_j = 0.$$

Both matrices are invertible. After choosing the nonzero value β_i, the values β_j and β_ℓ are determined. Hence there are at most $N - 1$ contributing masks for this support. By Appendix A, their nonzero coefficient has magnitude $2/((N-1)(N-2))$. Their total second moment is therefore at most

$$(N-1)\left(\frac{2}{(N-1)(N-2)}\right)^2 = \frac{4}{(N-1)(N-2)^2} = O(N^{-3}).$$

The injectivity of $T_{\sigma,w}$ ensures that each such injection mask corresponds to at most one construction mask.

For a fixed pullback support of size $k_0 \geq 4$, apply the recursive estimate (55) with $r = a$ and $d = 1$. Indeed, $2 < k_0 \leq a \leq N/8$, and the sole original block is untouched before that merge, so the recovery argument of Appendix E.2 applies. The contribution is at most

$$2\left(\frac{k_0}{N-k_0}\right)^{k_0/2+1} = O_w(N^{-3}).$$

Since w is fixed, there are only $O_w(1)$ support patterns. Summing their contributions gives $L_1 = O_w(N^{-3})$. The finitely many $N < 8a$ are absorbed into the implicit constant. All these estimates are uniform over the w-admissible maps under consideration.

F Proofs of Concrete Bounds for Dinur's LXoP

We prove the two bounds in Corollary 6.

F.1 One-Block Output

The case $q = 0$ is immediate; assume $q \geq 1$. Let ξ denote the density of the complete $\pi_+^{\sigma,1}$ function on its actual domain, of size $D = N/2$. By input symmetry and the definition of $\xi_{n,t}^{(\sigma,1)}$,

$$\sum_{|\alpha|=t} |\widehat{\xi}(\alpha)|^2 = \binom{D}{t} \sum_{\alpha \in (G\setminus\{0\})^t} |\widehat{\xi_{n,t}^{(\sigma,1)}}(\alpha)|^2.$$

The $t = 1$ term is zero by the support-two calculation in Appendix E.3, as also shown in [Din25, Proposition 15]. For $t \geq 2$, [Din25, Proposition 16] gives

$$\sum_{\alpha \in (G\setminus\{0\})^t} |\widehat{\xi_{n,t}^{(\sigma,1)}}(\alpha)|^2 \leq \frac{2^{3t/2+3c_t} t^{t+2c_t} (t-2c_t)^{t/2-c_t}}{(N-2t)^{3t/2+c_t}},$$

where $c_t = 0$ for even t and $c_t = 1/2$ for odd t. Proposition 16 applies for $2 \leq t \leq N/16$, which contains every level used below. Suppose $2 \leq t \leq N/32$. Using $D = N/2$, $\binom{D}{t} \leq (eD/t)^t$, $N - 2t \geq 15N/16$, and $c_t \leq 1/2$, we obtain

$$\sum_{|\alpha|=t} |\widehat{\xi}(\alpha)|^2 \leq 2^{3/2} t \left(e\sqrt{2} \left(\frac{16}{15}\right)^{3/2} \sqrt{\frac{t}{N}} \right)^t.$$

Since $e^2(16/15)^3 < 9$ and $t/N \leq 1/32$,

$$\sum_{|\alpha|=t} |\widehat{\xi}(\alpha)|^2 \leq \frac{51}{N} t^2 \left(\frac{3}{4}\right)^{t-2}.$$

Therefore, for $q \leq N/64$,

$$\sum_{1 \leq |\alpha| \leq 2q} |\widehat{\xi}(\alpha)|^2 \leq \frac{51}{N} \sum_{t=2}^{\infty} t^2 \left(\frac{3}{4}\right)^{t-2} = \frac{7548}{N}.$$

Theorem 1 then yields the first bound in Corollary 6. □

F.2 Two-Block Output

The case $q = 0$ is immediate; assume $q \geq 1$. Let ξ denote the density of the complete $\pi_+^{\sigma,2}$ function, whose actual domain has size $D = N/4$. In the one-input calculation of Appendix E.3, exactly the $N - 1$ construction masks (γ, γ) with $\gamma \neq 0$ contribute. Each has coefficient $2/((N-1)(N-2))$. Thus, as in [Din25, Lemma 8],

$$D \sum_{\alpha \in G^2\setminus\{0\}} \left| \widehat{\xi_{n,1}^{(\sigma,2)}}(\alpha) \right|^2 = \frac{N}{(N-1)(N-2)^2}. \tag{56}$$

For $2 \leq t \leq N/32$, Proposition 19 of the same work, together with input symmetry, gives

$$\sum_{|\alpha|=t} |\widehat{\xi}(\alpha)|^2 \leq \binom{D}{t} 2^{7t/2+3c_t} \left(\frac{t}{N-2t}\right)^{3t/2+c_t} =: A_t,$$

where $c_t = 0$ for even t and $c_t = 1/2$ for odd t. Since $2q \leq N/192$, this estimate applies at every required level.

For two consecutive levels of the same parity with $4 \leq t$ and $t + 2 \leq 2q$, a direct calculation gives

$$\begin{aligned} \frac{A_{t+2}}{A_t} &= 2^7 \frac{(D-t)(D-t-1)}{(t+1)(t+2)} \left(\frac{t+2}{N-2t-4}\right)^3 \\ &\quad \cdot \left(\frac{(t+2)(N-2t)}{t(N-2t-4)}\right)^{3t/2+c_t} < 0.855. \end{aligned} \tag{57}$$

We give some details for the uniform numerical bound. For fixed t, the displayed expression decreases with N when $N \geq 192(t+2)$. Substituting $N = 192(t+2)$ and using $c_t \leq 1/2$ bounds it by

$$\frac{128}{190^3} \frac{(47t+96)(47t+95)}{(t+1)(t+2)} \left(1 + \frac{192}{95t}\right)^{3t/2+1/2} < \frac{128 \cdot 47^2}{190^3} e^{288/95} < 0.855.$$

The middle inequality holds for $t \geq 4$, using $\log(1+x) \leq x - x^2/(2(1+x))$. This proves the bound uniformly in N and t.

We next evaluate the first few levels. For $N \geq 2048$, their scaled upper bounds decrease with N, and their values at $N = 2048$ give

$$NA_2 < 32.126, \qquad NA_3 < 1.277, \qquad NA_4 < 5.397, \qquad NA_5 < 0.406.$$

The right-hand side of (56) is smaller than $0.001/N$. Using the ratio 0.855 in (57), separately for the even and odd tails, yields

$$\begin{aligned}\sum_{1\leq|\alpha|\leq 2q} |\widehat{\xi}(\alpha)|^2 &\leq \frac{N}{(N-1)(N-2)^2} + \sum_{t=2}^{2q} A_t \\ &< \frac{1}{N}\left(0.001 + 32.126 + 1.277 + \frac{5.397 + 0.406}{1 - 0.855}\right) < \frac{74}{N}.\end{aligned}$$

Missing low levels when q is small can be added as positive upper bounds. For $N = 1024$, we have $q \leq 2$, and direct evaluation of the levels $t \leq 4$ gives total second Fourier moment less than $46/N < 74/N$. Theorem 1 now proves the second bound in Corollary 6. □